

# Veritas Storage Foundation™ and High Availability Solutions HA and Disaster Recovery Solutions Guide for Enterprise Vault

Windows Server 2003  
Windows Server 2008

5.1 Service Pack 1



# Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Enterprise Vault

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

Product version: 5.1. Service Pack 1

Document version: 5.1.SP1.0

## Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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

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

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

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

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

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

# Technical Support

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

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

<http://www.symantec.com/business/support/index.jsp>

## Contacting Technical Support

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

[http://www.symantec.com/business/support/contact\\_techsupp\\_static.jsp](http://www.symantec.com/business/support/contact_techsupp_static.jsp)

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

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

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

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

## Licensing and registration

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

<http://customercare.symantec.com>

## Customer service

Customer service information is available at the following URL:

<http://customercare.symantec.com>

Customer Service is available to assist with the following types of issues:

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

## Maintenance agreement resources

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

|                                 |                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------|
| Asia-Pacific and Japan          | <a href="mailto:customercare_apac@symantec.com">customercare_apac@symantec.com</a> |
| Europe, Middle-East, and Africa | <a href="mailto:semea@symantec.com">semea@symantec.com</a>                         |
| North America and Latin America | <a href="mailto:supportsolutions@symantec.com">supportsolutions@symantec.com</a>   |

## Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

|                                  |                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Symantec Early Warning Solutions | These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.                                                                                                                                                                                                                                             |
| Managed Security Services        | These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.                                                                                                                                                                                                                                                            |
| Consulting Services              | Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources. |
| Educational Services             | Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.                                                                                                                                                                                                                                           |

To access more information about Enterprise services, please visit our Web site at the following URL:

[www.symantec.com](http://www.symantec.com)

Select your country or language from the site index.

# Contents

## Section 1 Introduction and Concepts

### Chapter 1 Introducing Veritas Storage Foundation and High Availability Solutions for Enterprise Vault

|                                                         |    |
|---------------------------------------------------------|----|
| About clustering solutions with SFW HA .....            | 14 |
| About high availability .....                           | 14 |
| How a high availability solution works .....            | 14 |
| About replication .....                                 | 15 |
| About disaster recovery .....                           | 15 |
| What you can do with a disaster recovery solution ..... | 16 |
| Typical disaster recovery configuration .....           | 16 |
| Where to get more information .....                     | 17 |

## Section 2 Configuration Workflows

### Chapter 2 Configuration workflows for Enterprise Vault

|                                                         |    |
|---------------------------------------------------------|----|
| About using the workflow tables .....                   | 23 |
| High availability (HA) configuration (New Server) ..... | 24 |
| Disaster recovery configuration .....                   | 27 |
| DR configuration tasks: Primary site .....              | 27 |
| DR configuration tasks: Secondary site .....            | 30 |

### Chapter 3 Using the Solutions Configuration Center

|                                                          |    |
|----------------------------------------------------------|----|
| About the Solutions Configuration Center .....           | 35 |
| Starting the Configuration Center .....                  | 36 |
| Available options from the Configuration Center .....    | 37 |
| About running the Configuration Center wizards .....     | 43 |
| Following the workflow in the Configuration Center ..... | 44 |
| Solutions wizard logs .....                              | 47 |

## Section 3 Requirements and Planning

### Chapter 4 Requirements and planning for your HA and DR configurations

|                                                                                             |    |
|---------------------------------------------------------------------------------------------|----|
| Reviewing the requirements .....                                                            | 52 |
| Disk space requirements .....                                                               | 52 |
| Requirements for Veritas Storage Foundation High Availability for<br>Windows (SFW HA) ..... | 52 |
| Supported Enterprise Vault versions .....                                                   | 53 |
| System requirements for SFW HA .....                                                        | 54 |
| Network requirements for SFW HA .....                                                       | 54 |
| Permission requirements for SFW HA .....                                                    | 55 |
| Additional requirements for SFW HA .....                                                    | 55 |
| Best practices for SFW HA .....                                                             | 56 |
| Reviewing the HA configuration .....                                                        | 57 |
| Following the HA workflow in the Solutions Configuration Center .....                       | 59 |
| Reviewing the disaster recovery configuration .....                                         | 59 |
| IP addresses for disaster recovery configuration .....                                      | 61 |
| Supported disaster recovery configurations for service group<br>dependencies .....          | 62 |

## Section 4 Deployment

### Chapter 5 Installing and configuring SFW HA

|                                                                        |    |
|------------------------------------------------------------------------|----|
| Configuring the storage hardware and network .....                     | 66 |
| Installing Veritas Storage Foundation HA for Windows .....             | 68 |
| Setting Windows driver signing options .....                           | 68 |
| Installing Storage Foundation HA for Windows .....                     | 70 |
| Resetting the driver signing options .....                             | 73 |
| Configuring cluster disk groups and volumes for Enterprise Vault ..... | 74 |
| About cluster disk groups and volumes .....                            | 74 |
| Prerequisites for configuring cluster disk groups and volumes .....    | 75 |
| Considerations for volumes for a VVR configuration .....               | 75 |
| Sample disk group and volume configuration .....                       | 76 |
| Viewing the available disk storage .....                               | 76 |
| Creating a cluster disk group .....                                    | 77 |
| Creating volumes .....                                                 | 78 |

|           |                                                                     |     |
|-----------|---------------------------------------------------------------------|-----|
|           | About managing disk groups and volumes .....                        | 81  |
|           | Importing a disk group and mounting a volume .....                  | 81  |
|           | Unmounting a volume and deporting a disk group .....                | 81  |
|           | Adding drive letters to mount the volumes .....                     | 82  |
|           | Deporting the cluster disk group .....                              | 83  |
|           | Configuring the cluster .....                                       | 84  |
|           | Configuring Web console .....                                       | 96  |
|           | Configuring notification .....                                      | 97  |
| Chapter 6 | Installing and configuring Enterprise Vault for failover            |     |
|           | Installing Enterprise Vault .....                                   | 102 |
|           | Configuring the Enterprise Vault service group .....                | 102 |
|           | Before you configure an EV service group .....                      | 102 |
|           | Creating an EV service group .....                                  | 103 |
|           | Configuring Enterprise Vault Server in a cluster environment .....  | 107 |
|           | Setting service group dependencies for high availability .....      | 108 |
|           | Verifying the Enterprise Vault cluster configuration .....          | 108 |
|           | Setting up Enterprise Vault .....                                   | 110 |
|           | Considerations when modifying an EV service group .....             | 110 |
| Chapter 7 | Configuring disaster recovery for Enterprise Vault                  |     |
|           | Tasks for configuring disaster recovery for Enterprise Vault .....  | 114 |
|           | Verifying your primary site configuration .....                     | 117 |
|           | Guidelines for installing SFW HA and configuring the cluster on the |     |
|           | secondary site .....                                                | 118 |
|           | Setting up security for VVR .....                                   | 119 |
|           | Assigning user privileges (secure clusters only) .....              | 122 |
|           | Configuring disaster recovery with the DR wizard .....              | 123 |
|           | Cloning the storage on the secondary site using the DR wizard ..... | 126 |
|           | Installing and configuring Enterprise Vault on the                  |     |
|           | secondary site .....                                                | 129 |
|           | Configuring VVR replication and global clustering .....             | 132 |
|           | Setting service group dependencies for disaster recovery .....      | 138 |
|           | Verifying the disaster recovery configuration .....                 | 139 |
|           | Establishing secure communication within the                        |     |
|           | global cluster (optional) .....                                     | 140 |
|           | Adding multiple DR sites (optional) .....                           | 142 |
|           | Recovery procedures for service group dependencies .....            | 143 |
| Index     |                                                                     | 147 |



## Introduction and Concepts

This section contains the following chapters:

- [Introducing Veritas Storage Foundation and High Availability Solutions for Enterprise Vault](#)



# Introducing Veritas Storage Foundation and High Availability Solutions for Enterprise Vault

This chapter contains the following topics:

- [About clustering solutions with SFW HA](#)
- [About high availability](#)
- [How a high availability solution works](#)
- [About replication](#)
- [About disaster recovery](#)
- [What you can do with a disaster recovery solution](#)
- [Typical disaster recovery configuration](#)
- [Where to get more information](#)

## About clustering solutions with SFW HA

Veritas Storage Foundation HA for Windows (SFW HA) provides 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 Veritas 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 installing and configuring clustered environments using Veritas Storage Foundation HA for Windows (SFW HA). SFW HA includes Veritas Storage Foundation for Windows and Veritas Cluster Server.

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

## How a high availability solution works

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

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

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

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

## About replication

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

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

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

For more information on VVR refer to the *Veritas 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. Veritas Storage Foundation HA for Windows (SFW HA) provides the capability for implementing disaster recovery.

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

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

## What you can do with a disaster recovery solution

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

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

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


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

## Typical disaster recovery configuration

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

[Figure 1-1](#) 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.

## Where to get more information

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

**Table 1-1** shows the available Veritas Storage Foundation and High Availability Solutions solutions guides for SQL Server 2008. Separate guides are available

for earlier versions of SQL, for Microsoft Exchange, and for additional application solutions.

**Table 1-1** SFW HA solutions guides for SQL Server 2008

| Title                                                                                                                             | Description                                                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2008</i> | Solutions for SQL Server 2008 and Veritas Cluster Server clustering with Veritas Storage Foundation HA for Windows <ul style="list-style-type: none"><li>■ High availability (HA)</li><li>■ Campus clusters</li><li>■ Replicated data clusters</li><li>■ Disaster recovery (DR) with Veritas Volume Replicator or hardware array replication</li></ul> |
| <i>Veritas Storage Foundation and High Availability Solutions Microsoft Clustering Solutions Guide for Microsoft SQL 2008</i>     | Solutions for SQL Server 2008 and Microsoft clustering with Veritas Storage Foundation for Windows: <ul style="list-style-type: none"><li>■ High availability (HA)</li><li>■ Campus clusters</li><li>■ Disaster recovery (DR) with Veritas Volume Replicator</li></ul>                                                                                 |
| <i>Veritas Storage Foundation and High Availability Solutions Quick Recovery Solutions Guide for Microsoft SQL 2008</i>           | Quick Recovery solutions for SQL Server 2008 using either Veritas Storage Foundation for Windows or Veritas Storage Foundation HA for Windows.                                                                                                                                                                                                         |

[Table 1-2](#) shows the available Veritas Storage Foundation and High Availability Solutions solutions guides for SQL Server 2000/2005.

Table 1-2 SFW HA solutions guides for SQL Server 2000/2005

| Title                                                                                                                                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2000/2005</i>      | <p>Solutions for SQL Server 2000/2005 and Veritas Cluster Server clustering with Veritas Storage Foundation HA for Windows</p> <ul style="list-style-type: none"><li>■ High availability (HA)</li><li>■ Campus clusters</li><li>■ Replicated data clusters</li><li>■ Disaster recovery (DR) with Veritas Volume Replicator or hardware array replication</li></ul>                                                                                |
| <i>Veritas Storage Foundation and High Availability Solutions Microsoft Clustering and Quick Recovery Solutions Guide for Microsoft SQL</i> | <p>Quick Recovery solutions for SQL Server 2000 and 2005 using either Veritas Storage Foundation for Windows or Veritas Storage Foundation HA for Windows.</p> <p>Solutions for SQL Server 2000/2005 and Microsoft clustering with Veritas Storage Foundation for Windows:</p> <ul style="list-style-type: none"><li>■ High availability (HA)</li><li>■ Campus clusters</li><li>■ Disaster recovery (DR) with Veritas Volume Replicator</li></ul> |



## Configuration Workflows

This section contains the following chapters:

- [Using the Solutions Configuration Center](#)
- [Configuration workflows for Enterprise Vault](#)



# Configuration workflows for Enterprise Vault

This chapter contains the following topics:

- [About using the workflow tables](#)
- [High availability \(HA\) configuration \(New Server\)](#)
- [Disaster recovery configuration](#)

## About using the workflow tables

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

Separate workflow tables are provided for HA and DR configurations. Use the appropriate workflow table as a guideline to perform the installation and configuration.

Symantec recommends using the Solutions Configuration Center as a guide for installing and configuring SFW HA for Enterprise Vault.

See [“About the Solutions Configuration Center”](#) on page 35.

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

For example, in using the Solutions Configuration Center to set up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration. Likewise, in this guide, you first refer to the High Availability

workflow to set up high availability. You then continue with the disaster recovery workflow.

## High availability (HA) configuration (New Server)

[Table 2-1](#) outlines the high-level objectives and the tasks to complete each objective for an Active-Passive configuration.

**Note:** Symantec 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. See [“Where to get more information”](#) on page 17.

**Table 2-1** Enterprise Vault Server: Active-Passive configuration tasks

| Action                                     | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Verify hardware and software requirements  | See <a href="#">“Reviewing the requirements”</a> on page 52.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Review the HA configuration                | <ul style="list-style-type: none"><li>■ Understand active-passive configuration</li></ul> See <a href="#">“Reviewing the HA configuration”</a> on page 57.                                                                                                                                                                                                                                                                                                                             |
| Configure the storage hardware and network | <ul style="list-style-type: none"><li>■ Set up the storage hardware for a cluster environment</li><li>■ Verify the DNS entries for the systems on which Enterprise Vault Server will be installed</li></ul> See <a href="#">“Configuring the storage hardware and network”</a> on page 66.                                                                                                                                                                                             |
| Install SFW HA                             | <ul style="list-style-type: none"><li>■ Verify the driver signing option for the system</li><li>■ Install Veritas Storage Foundation HA for Windows</li><li>■ Select the option to install the Enterprise Vault (EV) Cluster Setup Wizard</li><li>■ To configure SQL Server for high availability, select the option to install Veritas Cluster Server Agent for Microsoft SQL Server</li></ul> See <a href="#">“Installing Veritas Storage Foundation HA for Windows”</a> on page 68. |

**Table 2-1** Enterprise Vault Server: Active-Passive configuration tasks

| Action                                                        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure disk groups and volumes for Enterprise Vault Server | <ul style="list-style-type: none"> <li>■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA)</li> <li>■ Create dynamic volumes for the MSMQ data, registry replication data, and EV services data</li> </ul> <p>See “<a href="#">Configuring cluster disk groups and volumes for Enterprise Vault</a>” on page 74.</p>                                                                                                                |
| Configure VCS cluster                                         | <p>If the cluster has not already been configured for SQL Server:</p> <ul style="list-style-type: none"> <li>■ Verify static IP addresses and name resolution configured for each node</li> <li>■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster</li> </ul> <p>See “<a href="#">Configuring the cluster</a>” on page 84.</p>                                                           |
| Install Enterprise Vault on the cluster nodes                 | <ul style="list-style-type: none"> <li>■ 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.</li> </ul> <p>Refer to the Enterprise Vault documentation for installation instructions</p>                                                                                                                                                            |
| Create an Enterprise Vault service group                      | <ul style="list-style-type: none"> <li>■ Ensure that you have met the prerequisites</li> <li>■ Ensure that the disk group and volumes for the various Enterprise Vault components are mounted on the first node</li> <li>■ Create a EV service group using the Enterprise Vault Cluster Setup Wizard</li> <li>■ Bring the EV service group online on the first node</li> </ul> <p>See “<a href="#">Configuring the Enterprise Vault service group</a>” on page 102</p> |

**Table 2-1** Enterprise Vault Server: Active-Passive configuration tasks

| Action                                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure Enterprise Vault for the cluster environment on the first node       | <ul style="list-style-type: none"> <li>■ Launch the Enterprise Vault Configuration Wizard on the first node</li> <li>■ Choose the option to create a new Enterprise Vault server with cluster support</li> <li>■ Complete running the wizard on the first node</li> </ul> <p>See <a href="#">“Configuring Enterprise Vault Server in a cluster environment”</a> on page 107.</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"> <li>■ Bring the EV service group online on the first node</li> <li>■ Launch the Enterprise Vault Configuration Wizard on the second node</li> <li>■ Choose the option to add the node as a failover node for an existing clustered server</li> <li>■ Complete running the wizard on the second node</li> <li>■ Repeat these steps for any additional nodes in the EV cluster</li> </ul> <p>See <a href="#">“Configuring Enterprise Vault Server in a cluster environment”</a> on page 107.</p> <p>Refer to the Enterprise Vault documentation for more information.</p> |
| Perform additional configuration steps for Enterprise Vault                    | <p>See <a href="#">“Setting up Enterprise Vault”</a> on page 110.</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 <a href="#">“Verifying the Enterprise Vault cluster configuration”</a> on page 108.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Verify the HA configuration                                                    | <p>Test failover between nodes</p> <p>See <a href="#">“Verifying the Enterprise Vault cluster configuration”</a> on page 108</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

# Disaster recovery configuration

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

**Note:** Symantec 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. See [“Where to get more information”](#) on page 17.

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

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

See [“About the Solutions Configuration Center”](#) on page 35.

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

## DR configuration tasks: Primary site

[Table 2-2](#) outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the primary site.

**Table 2-2**            Configuring the primary site for disaster recovery

| Action                                     | Description                                                                                                                                                                                                                                            |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Verify hardware and software prerequisites | See <a href="#">“Reviewing the requirements”</a> on page 52.<br><br><b>Note:</b> If the DR site is on a different network segment, ensure that you allocate two IP addresses for the virtual server, one for the primary site and one for the DR site. |
| Understand the configuration               | Understand the DR configuration<br><br>See <a href="#">“Reviewing the disaster recovery configuration”</a> on page 59.                                                                                                                                 |

**Table 2-2** Configuring the primary site for disaster recovery (Continued)

| Action                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the storage hardware and network                     | <p>For all nodes in the cluster:</p> <ul style="list-style-type: none"> <li>■ Set up the storage hardware for a cluster environment</li> <li>■ Verify the DNS entries for the systems on which EV will be installed</li> </ul> <p>See <a href="#">“Configuring the storage hardware and network”</a> on page 66.</p>                                                                                                                                                                                                                                                                                                                                                                  |
| Install SFW HA                                                 | <ul style="list-style-type: none"> <li>■ Verify the driver signing option for the system</li> <li>■ Install Veritas Storage Foundation for Windows HA on all nodes that will become part of the cluster</li> <li>■ Select the option to install the Global Cluster Option (GCO).</li> <li>■ Select the option to install VVR replication</li> <li>■ Select the option to install the Enterprise Vault (EV) Cluster Setup Wizard</li> <li>■ To configure SQL Server for high availability, select the option to install Veritas Cluster Server Agent for Microsoft SQL Server</li> </ul> <p>See <a href="#">“Installing Veritas Storage Foundation HA for Windows”</a> on page 68.</p> |
| Configure the cluster                                          | <p>If the cluster has not already been configured for SQL Server:</p> <ul style="list-style-type: none"> <li>■ Verify static IP addresses and name resolution configured for each node</li> <li>■ Configure cluster components using the Veritas Cluster Server Configuration Wizard (VCW)</li> <li>■ Set up secure communication for the cluster</li> </ul> <p>See <a href="#">“Configuring the cluster”</a> on page 84.</p>                                                                                                                                                                                                                                                         |
| Configure cluster disk groups and volumes for Enterprise Vault | <ul style="list-style-type: none"> <li>■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA)</li> <li>■ Create dynamic volumes for the MSMQ data, registry replication data, and EV services data</li> </ul> <p>See <a href="#">“Configuring cluster disk groups and volumes for Enterprise Vault”</a> on page 74.</p>                                                                                                                                                                                                                                                                                                                               |

**Table 2-2** Configuring the primary site for disaster recovery (Continued)

| Action                                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Install Enterprise Vault on the cluster nodes                                  | <ul style="list-style-type: none"> <li>■ 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.</li> </ul> <p>Refer to the Enterprise Vault documentation for installation instructions</p>                                                                                                                                                                                                                                                                                  |
| Create an Enterprise Vault service group                                       | <ul style="list-style-type: none"> <li>■ Ensure that you have met the prerequisites</li> <li>■ Ensure that the disk group and volumes for the various Enterprise Vault components are mounted on the first node</li> <li>■ Create a EV service group using the Enterprise Vault Cluster Setup Wizard</li> <li>■ Bring the EV service group online on the first node</li> </ul> <p>See <a href="#">“Configuring the Enterprise Vault service group”</a> on page 102.</p>                                                                                                                                    |
| Configure Enterprise Vault for the cluster environment on the first node       | <ul style="list-style-type: none"> <li>■ Launch the Enterprise Vault Configuration Wizard on the first node</li> <li>■ Choose the option to create a new Enterprise Vault server with cluster support</li> <li>■ Complete running the wizard on the first node</li> </ul> <p>See <a href="#">“Configuring Enterprise Vault Server in a cluster environment”</a> on page 107.</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"> <li>■ Bring the EV service group online on the first node</li> <li>■ Launch the Enterprise Vault Configuration Wizard on the second node</li> <li>■ Choose the option to add the node as a failover node for an existing clustered server</li> <li>■ Complete running the wizard on the second node</li> <li>■ Repeat these steps for any additional nodes in the EV cluster</li> </ul> <p>See <a href="#">“Configuring Enterprise Vault Server in a cluster environment”</a> on page 107.</p> <p>Refer to the Enterprise Vault documentation for more information.</p> |

**Table 2-2**                    Configuring the primary site for disaster recovery (Continued)

| Action                                                          | Description                                                                                                                                     |
|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Perform additional configuration steps for Enterprise Vault     | See <a href="#">“Setting up Enterprise Vault”</a> on page 110.<br>Refer to the Enterprise Vault documentation for more information.             |
| (Optional) Configure the appropriate service group dependencies | Configure the appropriate service group dependencies<br>See <a href="#">“Verifying the Enterprise Vault cluster configuration”</a> on page 108. |
| Verify the primary site configuration                           | Test failover between nodes on the primary site.<br>See <a href="#">“Verifying the Enterprise Vault cluster configuration”</a> on page 108.     |

DR configuration tasks: Secondary site

[Table 2-3](#) outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the secondary site.

**Table 2-3**                    Configuring the secondary site for disaster recovery

| Action                                                                                     | Description                                                                                                                                                                                                                        |
|--------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Install SFW HA and configure the cluster on the secondary site                             | <b>Caution:</b> Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.<br><br>See <a href="#">“Verifying your primary site configuration”</a> on page 117. |
| 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<br><br>See <a href="#">“Verifying your primary site configuration”</a> on page 117.                |
| Set up security for VVR                                                                    | Ensure that you have completed setting up VVR security before running the DR wizard.<br><br>See <a href="#">“Setting up security for VVR”</a> on page 119.                                                                         |
| (Secure cluster only) Assign user privileges                                               | For a secure cluster only, assign user privileges<br><br>See <a href="#">“Assigning user privileges (secure clusters only)”</a> on page 122.                                                                                       |

**Table 2-3** Configuring the secondary site for disaster recovery (Continued)

| Action                                                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Start running the DR wizard                                              | <ul style="list-style-type: none"> <li>■ Review prerequisites for the DR wizard</li> <li>■ 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</li> </ul> <p>See <a href="#">“Configuring disaster recovery with the DR wizard”</a> on page 123.</p>                                                                                                              |
| Clone the storage configuration                                          | <p>Clone the storage configuration on the secondary site using the DR wizard</p> <p>See <a href="#">“Cloning the storage on the secondary site using the DR wizard”</a> on page 126.</p>                                                                                                                                                                                                                                                                                                            |
| Install Enterprise Vault on the cluster nodes                            | <ul style="list-style-type: none"> <li>■ 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.</li> </ul> <p>Refer to the Enterprise Vault documentation for installation instructions</p>                                                                                                                                                                           |
| Create an Enterprise Vault service group                                 | <ul style="list-style-type: none"> <li>■ Ensure that you have met the prerequisites</li> <li>■ Ensure that the disk group and volumes for the various Enterprise Vault components are mounted on the first node</li> <li>■ 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</li> </ul> <p>See <a href="#">“Installing and configuring Enterprise Vault on the secondary site”</a> on page 129.</p> |
| Configure Enterprise Vault for the cluster environment on the first node | <ul style="list-style-type: none"> <li>■ Launch the Enterprise Vault Configuration Wizard on the first node</li> <li>■ Choose the option to create a new Enterprise Vault server with cluster support</li> <li>■ Complete running the wizard on the first node</li> </ul> <p>See <a href="#">“Installing and configuring Enterprise Vault on the secondary site”</a> on page 129.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>                                      |

**Table 2-3** Configuring the secondary site for disaster recovery (Continued)

| Action                                                                         | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure Enterprise Vault for the cluster environment on any additional nodes | <ul style="list-style-type: none"> <li>■ Bring the EV service group online on the first node</li> <li>■ Launch the Enterprise Vault Configuration Wizard on the second node</li> <li>■ Choose the option to add the node as a failover node for an existing clustered server</li> <li>■ Complete running the wizard on the second node</li> <li>■ Repeat these steps for any additional nodes in the EV cluster</li> </ul> <p>See <a href="#">“Installing and configuring Enterprise Vault on the secondary site”</a> on page 129.</p> <p>Refer to the Enterprise Vault documentation for more information.</p> |
| Perform additional configuration steps for Enterprise Vault                    | <p>See <a href="#">“Setting up Enterprise Vault”</a> on page 110.</p> <p>Refer to the Enterprise Vault documentation for more information.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Configure replication and global clustering                                    | <ul style="list-style-type: none"> <li>■ Use the DR wizard to configure VVR replication and global clustering</li> </ul> <p>See <a href="#">“Configuring VVR replication and global clustering”</a> on page 132.</p>                                                                                                                                                                                                                                                                                                                                                                                            |
| (Optional) Configure the appropriate service group dependencies                | <p>Configure the appropriate service group dependencies</p> <p>See <a href="#">“Setting service group dependencies for disaster recovery”</a> on page 138.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Verify the disaster recover configuration                                      | <p>Verify that the secondary site has been fully configured for disaster recovery</p> <p>See <a href="#">“Verifying the disaster recovery configuration”</a> on page 139.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| (Optional) Add secure communication                                            | <p>Add secure communication between local clusters within the global cluster (optional task)</p> <p>See <a href="#">“Establishing secure communication within the global cluster (optional)”</a> on page 140.</p>                                                                                                                                                                                                                                                                                                                                                                                               |
| (Optional) Add additional DR sites                                             | <p>Optionally, add additional DR sites to a VVR environment</p> <p>See <a href="#">“Adding multiple DR sites (optional)”</a> on page 142.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

Table 2-3                      Configuring the secondary site for disaster recovery (Continued)

| Action                                             | Description                                                                                                                                                                                                                                                      |
|----------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Handling service group dependencies after failover | <p>If your environment includes dependent service groups, review the considerations for bringing the service groups online after failover to the secondary site</p> <p>See <a href="#">“Recovery procedures for service group dependencies”</a> on page 143.</p> |



# Using the Solutions Configuration Center

This chapter covers the following topics:

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

## About the Solutions Configuration Center

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

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

You can use the Configuration Center and its wizards to set up your environment for any combination of the following solutions:

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

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

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

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

## Starting the Configuration Center

You can start the Configuration Center in the following ways:

- Click **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.
- Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- Click **Start > Run** and type **scc**.

## Available options from the Configuration Center

On the Applications tab, the Solutions Configuration Center is context-sensitive to the application. For example, the Solution Guides listed in the right pane match the selected application.

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

[Figure 3-1](#) shows the choices available on a 32-bit system when you click Solutions for Microsoft Exchange.

**Figure 3-1** Solutions Configuration Center for Microsoft Exchange




Figure 3-2 shows the choices available when you click Solutions for Microsoft SQL Server.

Figure 3-2 Solutions Configuration Center for Microsoft SQL Server




Figure 3-3 shows the choices available when you click Solutions for Enterprise Vault Server.

**Figure 3-3** Solutions Configuration Center for Enterprise Vault Server





Figure 3-4 shows the choices available when you click Solutions for Additional Applications.

Figure 3-4 Solutions Configuration Center for additional applications



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

Figure 3-5 shows one of the steps for implementing high availability for Exchange.

Figure 3-5 Context-sensitive step for Exchange




Figure 3-6 shows one of the steps for implementing high availability for SQL Server.

**Figure 3-6** Context-sensitive step for SQL Server



Figure 3-7 shows one of the steps for implementing high availability for Enterprise Vault Server.

**Figure 3-7** Context-sensitive step for Enterprise Vault Server





Figure 3-8 shows one of the steps for implementing high availability for additional applications.

**Figure 3-8** Context-sensitive step for additional applications



## About running the Configuration Center wizards

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

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

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

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

|                               |                                                                                                                                                                             |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| New Dynamic Disk Group Wizard | Launched from the Veritas Enterprise Administrator console                                                                                                                  |
| New Volume Wizard             | Launched from the Veritas Enterprise Administrator console                                                                                                                  |
| Exchange Setup Wizard         | <p>Installs and configures Exchange for the high availability environment</p> <p>If Exchange is already installed, refer to the documentation for further instructions.</p> |
| Exchange Configuration Wizard | Configures the service group for Exchange high availability                                                                                                                 |

|                                       |                                                                                                                                                                                                 |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SQL Server Configuration Wizard       | Configures the service group for SQL Server 2000 or SQL Server 2005 high availability<br><br>You must first install SQL Server on each node according to the instructions in the documentation. |
| SQL Server 2008 Configuration Wizard  | Configures the service group for SQL Server 2008 high availability<br><br>You must first install SQL Server on each node according to the instructions in the documentation.                    |
| Enterprise Vault Cluster Setup Wizard | Configures the service group for Enterprise Vault Server high availability.                                                                                                                     |
| MSDTC Wizard                          | Configures an MSDTC Server service group for SQL Server 2000, 2005, or 2008 environments.                                                                                                       |
| MSMQ Configuration Wizard             | Configures a Microsoft Message Queuing (MSMQ) service group.                                                                                                                                    |

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

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

## Following the workflow in the Configuration Center

During the multi-step High Availability Configuration workflow, you may find it helpful to run an SFW HA client on another system and leave the Configuration Center open on that system. In this way, you can see what step comes next, drill

down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format. When setting up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration.

Figure 3-9 shows the high-level overview of the workflow steps for configuring high availability for Exchange from the Solutions Configuration Center.

**Figure 3-9** Workflow for configuring Exchange high availability




Figure 3-10 shows the high-level overview of the workflow steps for configuring high availability for SQL Server from the Solutions Configuration Center.

**Figure 3-10** Workflow for configuring SQL Server high availability




Figure 3-11 shows the high-level overview of the workflow steps for configuring high availability for Enterprise Vault Server from the Solutions Configuration Center.

**Figure 3-11** Workflow for configuring high availability for Enterprise Vault Server





Figure 3-12 shows the high-level overview of the workflow steps for configuring high availability for additional applications from the Solutions Configuration Center.

**Figure 3-12** Workflow for configuring high availability for additional applications



## Solutions wizard logs

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

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

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

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

For Windows Server 2003:

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

For Windows Server 2008:

```
C:\ProgramData\Veritas\winsolutions\log
```



## Requirements and Planning

This section contains the following chapter:

- [Requirements and planning for your HA and DR configurations](#)



# Requirements and planning for your HA and DR configurations

This chapter contains the following topics:

- [Reviewing the requirements](#)
- [Reviewing the HA configuration](#)
- [Following the HA workflow in the Solutions Configuration Center](#)
- [Reviewing the disaster recovery configuration](#)

# Reviewing the requirements

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

## Disk space requirements

For normal operation, all installations require an additional 50 MB of disk space. [Table 4-1](#) estimates disk space requirements for SFW HA.

**Table 4-1** Disk space requirements

| Installation options                        | Install directory/drive |
|---------------------------------------------|-------------------------|
| SFW HA + all options<br>+ client components | 1564 MB                 |
| SFW HA + all options                        | 1197 MB                 |
| Client components                           | 528 MB                  |

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

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

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW 5.1 Service Pack 1 Hardware Compatibility List to confirm supported hardware:  
<http://entsupport.symantec.com/docs/302144>
- Review the Enterprise Vault Server versions supported with Veritas Storage Foundation High Availability for Windows (SFW HA).
- When installing SFW HA for a Disaster Recovery configuration, ensure that you select the Global Clustering Option and depending on your replication solution select Veritas Volume Replicator or a hardware replication agent.
- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the option to install Veritas Volume Replicator.

## Supported Enterprise Vault versions

The following versions of Enterprise Vault are tested and supported with SFW HA 5.1 SP1:

Enterprise Vault 8.0 SP1 and SP2

Review the SFW HA 5.1 Service Pack 1 Software Compatibility List for earlier supported versions.

## System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 1 GB of RAM per server for SFW HA.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.  
If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA). See the *Veritas Storage Foundation Administrator's Guide* for more information.
- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Symantec recommends three NICs.  
See "[Best practices for SFW HA](#)" on page 56.
- NIC teaming is not supported for the private network.
- All servers must have the same system architecture, run the same operating system, and be at the same service pack (SP) level.

## Network requirements for SFW HA

SFW HA has the following network requirements:

- Install SFW HA on servers in a Windows Server 2003 or Windows Server 2008 domain.
- Do not install SFW HA on servers that are assigned the role of a Domain Controller. Configuring a cluster on a domain controller is not supported.
- Ensure that your firewall settings allow access to ports used by SFW HA wizards and services. For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Static IP addresses for the following purposes:
  - One static IP address available 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, the Cluster Management Console (web console), or the Global Cluster Option. The same IP address may be used for all options.
- For VVR replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For VVR replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.  
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS 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 *Veritas Cluster Server Bundled Agents Reference Guide*.

## Permission requirements for SFW HA

The following permissions are required:

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

## Additional requirements for SFW HA

Please review the following additional requirements:

- Installation media for all products and third-party applications.
- Licenses for all products and third-party applications.

- You must install the operating system in the same path on all systems. For example, if you install Windows Server 2003 on C:\WINDOWS of one node, installations on all other nodes must be on C:\WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- 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.

## Best practices for SFW HA


Symantec recommends that you perform the following tasks:

- Verify that you have three network adapters (two NICs exclusively for the private network and one for the public network).  
When using only two NICs, lower the priority of one NIC and use the low-priority NIC for public and private communication.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the private network.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.
- Although you can use a single node cluster as the primary and secondary zones, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the vxclus UseSystemBus ON command. This is applicable for a Replicated Data Cluster configuration.

## Reviewing the HA configuration

In a typical example of 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. [Figure 4-1](#) illustrates a typical Active-Passive configuration.

**Figure 4-1** 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.)

Symantec 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. See [“Where to get more information”](#) on page 17.

### Sample Active-Passive configuration

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

Table 4-5 describes the objects created and used during the installation and configuration using sample names.

Table 4-5            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, you should obtain all required IP addresses before beginning 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 Web Console or the notification service, you require a cluster IP address. The cluster IP address is also used by the Global Cluster Option for disaster recovery.

## Following the HA workflow in the Solutions Configuration Center

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

[Figure 4-2](#) shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

**Figure 4-2** Configuration steps in the Solutions Configuration Center




See [“Following the workflow in the Configuration Center”](#) on page 44.

## Reviewing the disaster recovery configuration

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

[Figure 4-3](#) illustrates a typical Active-Passive disaster recovery configuration.

**Figure 4-3** 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 VVR for setting up the Replicated Volume Group (RVG). The application data is stored on the volumes that are under the control of the RVG.

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.

Table 4-6 describes the objects created and used during the installation and configuration.

**Table 4-6** 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_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                                                                                               |
| 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 (VVR 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, you should obtain all required IP addresses before beginning 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 *Veritas Cluster Server Administrator's Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

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

The wizard clones dependent service groups as global groups.

## Deployment

This section contains the following chapters:

- [Installing and configuring SFW HA](#)
- [Installing and configuring Enterprise Vault for failover](#)
- [Configuring disaster recovery for Enterprise Vault](#)



# Installing and configuring SFW HA

This chapter contains the following topics:

- [Configuring the storage hardware and network](#)
- [Installing Veritas Storage Foundation HA for Windows](#)
- [Configuring cluster disk groups and volumes for Enterprise Vault](#)
- [About managing disk groups and volumes](#)
- [Configuring the cluster](#)

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

### To verify the DNS settings and binding order

- 1 From the Control Panel, access the Network Connections window.
- 2 Ensure the public network adapter is the first bound adapter as follows:
  - From the Advanced menu, click **Advanced Settings**.
  - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the Connections list. If necessary, use the arrow button to move the adapter to the top of the list.
- 3 Ensure that DNS name resolution is enabled. Make sure that you use the public network adapter, and not those configured for the private network. Do the following:
  - In the Network Connections window, double-click the adapter for the public network to access its properties.
  - In the Public Status dialog box, on the General tab, click **Properties**.

- In the Public Properties dialog box, on the General tab, select the **Internet Protocol (TCP/IP)** check box and click **Properties**.
- Select the **Use the following DNS server addresses** option and verify the correct value for the IP address of the DNS server.
- Click **Advanced**.
- In the DNS tab, make sure the **Register this connection's address in DNS** check box is selected. Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.

# Installing Veritas Storage Foundation HA for Windows

The product installer enables you to install the software for Veritas Storage Foundation HA for Windows. The installer automatically installs Veritas Storage Foundation for Windows and Veritas Cluster Server. You must select the option to install the Enterprise Vault (EV) Cluster Setup Wizard. For a disaster recovery configuration, select the option to install GCO and VVR.

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

## Setting Windows driver signing options

Some drivers provided by Symantec may not be signed by Microsoft. Depending upon your installation options, these unsigned drivers may stop your installation.

When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

The following table describes the product installer behavior on local and remote systems when installing options with unsigned drivers on Windows Server 2003.

**Table 5-1** Installation behavior with unsigned drivers

| Driver Signing Setting | Installation behavior on the local system  | Installation behavior on remote systems                                                                                             |
|------------------------|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------|
| Ignore                 | Always allowed                             | Always allowed                                                                                                                      |
| Warn                   | Warning message, user interaction required | Installation proceeds. The user must log on locally to the remote system to respond to the dialog box to complete the installation. |
| Block                  | Never allowed                              | Never allowed                                                                                                                       |

On local systems set the driver signing option to either Ignore or Warn. On remote systems set the option to Ignore in order to allow the installation to proceed without user interaction.

---

**Note:** Windows Server 2008 does not allow you to change the driver signing option. Symantec provides an installation option to install a Symantec Trusted certificate to allow installation of Symantec drivers that are not certified by Microsoft. For details, see “[Installing Symantec Trusted certificate for unsigned drivers](#)” on page 69.

---

#### To change the driver signing options on each system

- 1 Log on locally to the system.
- 2 Open the Control Panel and click **System**.
- 3 Click the **Hardware** tab and click **Driver Signing**.
- 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
- 5 Click **OK**.
- 6 Repeat for each computer.

If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

### Installing Symantec Trusted certificate for unsigned drivers

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

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

---

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

---

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

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

## Installing Storage Foundation HA for Windows

Install Veritas Storage Foundation HA for Windows.

### To install the product

- 1 Insert the DVD containing the installation software into your system's disk drive or download the installation software from the Symantec website.
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The Select Product screen appears.
- 3 Review the links on the Select Product screen.  
Links on this screen access Late Breaking News, the Configuration Checker, as well as begin the process to install Storage Foundation HA for Windows. Click on **Read Late Breaking News** for the latest information about updates, patches, and software issues regarding this release.
- 4 Click **Storage Foundation HA 5.1 SP1 for Windows**.
- 5 Do one of the following:
  - Click **Complete/Custom** to begin installation.
  - Click the **Administrative Console** link to install only the client components.
- 6 Review the Welcome message and the listed prerequisites. Ensure that any of the listed prerequisites applicable to your installation are met prior to proceeding.  
Click **Next**.
- 7 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click the radio button for **I AGREE TO the terms of the license agreement**, and then click **Next**.
- 8 Enter the product license key before adding license keys for features. Enter the license key in the top field and click **Add**.  
If you do not have a license key, click **Next** to use the default evaluation license key. This license key is valid for a limited evaluation period only.
- 9 Repeat for additional license keys. Click **Next**
  - To remove a license key, click the key to select it and click **Remove**.
  - To see the license key's details, click the key.
- 10 Select the appropriate SFW product options and click **Next**. Be sure to select the following as appropriate for your installation:

|                                            |                                                               |
|--------------------------------------------|---------------------------------------------------------------|
| Enterprise Vault (EV) Cluster Setup Wizard | Required to configure high availability for Enterprise Vault. |
|--------------------------------------------|---------------------------------------------------------------|

|                           |                                                                                                                                                                                                                                                                                       |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Client                    | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration. |
| Global Cluster Option     | Required for a disaster recovery configuration only.                                                                                                                                                                                                                                  |
| Veritas Volume Replicator | If you plan to use VVR for replication, select the option to install VVR.                                                                                                                                                                                                             |

11 Select the following for the installation and click **Next**.

|              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Domain       | <p>Select a domain from the list.</p> <p>Depending on domain and network size, speed, and activity, the domain and computer lists can take some time to populate.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Computer     | <p>To add a computer for installation, select it from the Computer list or type the computer's name in the Computer field. Then click <b>Add</b>.</p> <p>To remove a computer after adding it, click the name in the Selected computers for installation field and click <b>Remove</b>.</p> <p>Click a computer's name to see its description.</p> <p>When installing the software on multiple computers in a single installation using the product installer, all computers must have the same platform type (for example, x86 or x64). However, the computers can have different Windows operating systems. For example, you can install the software on multiple computers at once running Windows 2003 and Windows 2008.</p> |
| Install Path | <p>Optionally, change the installation path.</p> <ul style="list-style-type: none"> <li>■ To change the path, select a computer in the Selected computers for installation field, type the new path, and click <b>Change</b>.</li> <li>■ To restore the default path, select a computer and click <b>Default</b>.</li> </ul> <p>The default path is:<br/>C:\Program Files\Veritas</p> <p>For 64-bit installations, the default path is:<br/>C:\Program Files (x86)\Veritas</p>                                                                                                                                                                                                                                                   |

- 12 When the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the target computers. In this situation, after the installer displays an error message, enter the host names or the IP addresses of the missing computers manually.
- 13 The installer checks the prerequisites for the selected computers and displays the results. Review the information and click **Next**.  
If a computer fails validation, address the issue, and repeat the validation. Click the computer in the list to display information about the failure. Click **Validate Again** to begin the validation process again.
- 14 Depending upon your earlier product installer selections and operating system, you may receive one or more of the following messages.

|                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>If you are using multiple paths and selected a specific DSM, you receive the Veritas Dynamic Multi-pathing warning:</p> | <p>The time to install the Veritas Dynamic Multi-pathing MPIO feature depends on the number of physical paths connected during the installation. To reduce installation time for this feature, Symantec recommends only one physical path connection during installation. After the installation completes, reconnect additional physical paths before rebooting the system.</p> |
|----------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

If applicable to your installation, perform the above procedure.

|                                                                                                                                         |                                                                                                                                                                    |
|-----------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>If you are using multiple paths and selected a specific DSM on a Windows Server 2008 machine, you receive an additional message:</p> | <p>On Windows Server 2008, the Microsoft Multipath input/output (Microsoft MPIO) feature must be enabled before installing DMP Device Specific Modules (DSMs).</p> |
|-----------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|

If applicable to your installation, perform the above procedure.

|                                                                                                                                                     |                                                                                                                                                                                                                                                             |
|-----------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>When installing Veritas Storage Foundation for Windows (Server Components) with the MSCS option selected, you receive the following message:</p> | <p>When installing Veritas Storage Foundation 5.1 for Windows (Server Components) with the MSCS option, you may want MSCS Quorum Arbitration Time (Min. and Max) to be adjusted to ensure optimal functionality with Veritas dynamic volumes with MSCS.</p> |
|-----------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

For additional information, see the *Storage Foundation for Windows Administrator Guide* for details.

If applicable to your installation, perform the above procedure.

- 15 When finished reviewing the message or messages, click **OK**.

- 16 The Summary screen appears displaying an Install report.  
Review the information in the Install report. Click **Back** to make changes, if necessary. Click **Install** if information is validated.
- 17 The Installation Status screen displays status messages and the progress of the installation.  
If an installation fails, click **Next** to review the report and address the reason for failure. You may have to either repair the installation or uninstall and re-install.
- 18 When the installation completes, review the summary screen and click **Next**.
- 19 If you are installing on remote nodes, click **Reboot**. Note that you cannot reboot the local node now, and that failed nodes are unchecked by default. Click the check box next to the remote nodes that you want to reboot.
- 20 When the nodes have finished rebooting successfully, the Reboot Status shows Online and the **Next** button is available. Click **Next**.
- 21 Review the log files and click **Finish**.
- 22 Click **Yes** to reboot the local node.

## Resetting the driver signing options

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

### To reset the driver signing options

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

## Configuring cluster disk groups and volumes for Enterprise Vault

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

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

- [About cluster disk groups and volumes](#)
- [Prerequisites for configuring cluster disk groups and volumes](#)
- [Considerations for volumes for a VVR configuration](#)
- [Sample disk group and volume configuration](#)

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

- [Viewing the available disk storage](#)
- [Creating a cluster disk group](#)
- [Creating volumes](#)

### 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 *Veritas Storage Foundation Administrator's Guide* for more information.

---

## Prerequisites for configuring cluster disk groups and volumes

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

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

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 disaster recovery configuration using Veritas Volume Replicator, read the following topic:

- [Considerations for volumes for a VVR configuration](#)

## Considerations for volumes for a VVR configuration

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

- VVR does not support the following types of volumes:
  - SFW (software) RAID 5 volumes
  - Volumes with the Dirty Region Log (DRL)

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

## Sample disk group and volume configuration

For 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 Open the VEA console (launch from the Solutions Configuration Center, or by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**). 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.

### To create a dynamic (cluster) disk group

---

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

---

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.  
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:
  - In the **Group Name** field, enter the name of the disk group, for example, EVDG.
  - Click the checkbox for **Create cluster group**.
  - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.  
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

---

**Note:** For Windows Server 2003, Windows Disk Management Compatible Dynamic Disk Group creates a disk group that is compatible with the disk groups created with Windows Disk Management and with earlier versions of Volume Manager for Windows products.

---

- Click **Next**.

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

## Creating volumes

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

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

- [Considerations for volumes for a VVR configuration](#)

---

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


---

### To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.  
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.  
You can right-click the disk group you have just created.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.
  - Automatic disk selection is the default setting and is recommended for campus clusters. SFW automatically selects the disks based on the following criteria:
    - Their port assignment (disks with two different ports are selected). Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.
    - Amount of available space on the disks. SFW will pick two disks (one from each array) with the most space.

- To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list.
  - You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
  - Click **Next**.
- 7 Specify the volume attributes.
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
  - Provide a size for the volume. If you click the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
  - Select a layout type.  
For campus clusters, select either **Concatenated** or **Striped**.
  - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.  
For campus clusters, if you select **Striped**, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
  - To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.  
For campus clusters, you select the **Mirrored** checkbox for either layout type.
  - In the Mirror Info area, select the appropriate mirroring options.  
For campus clusters, in the **Mirror Info** area, after selecting the **Mirrored** checkbox, click **Mirror across** and select **Enclosures** from the drop-down list.
  - Verify that **Enable logging** is not selected.
  - Click **Next**.
- 8 Assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
  - To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.

- If creating a Replicator Log volume for Veritas Volume Replicator, select **Do not assign a drive letter**.



- 9 Click **Next**.
- 10 Create an NTFS file system.
  - Make sure the **Format this volume** checkbox is checked and click **NTFS**.
  - For a VVR configuration, for the Replicator Log volume only, clear the **Format this volume** check box.
  - Select an allocation size or accept the default.
  - The file system label is optional. SFW makes the volume name the file system label.
  - Select **Perform a quick format** if you want to save time.
  - Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
  - Click **Next**.
- 11 Click **Finish** to create the new volume.
- 12 Repeat these steps to create additional volumes.  
Create the cluster disk group and volumes on the first node of the cluster only.

## About managing disk groups and volumes

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

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

### Importing a disk group and mounting a volume

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

#### To import a disk group

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

#### To mount a volume

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

### Unmounting a volume and deporting a disk group

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

#### To unmount a volume and deport the dynamic disk group


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

## Adding drive letters to mount the volumes

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

#### To add a drive letter or path to a volume

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



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder.

- *To assign a drive letter*  
Select the **Assign a Drive Letter** option and select a drive letter from the drop-down list.
- *To mount the volume as a folder*  
Select the **Mount as an empty NTFS folder** option and click **Browse** to locate an empty folder on the shared disk.

---

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

---

- 5 Click **OK**.

## Deporting the cluster disk group

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

## Configuring the cluster

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

Complete the following tasks before creating a cluster:


- Verify that each node uses static IP addresses and that name resolution is configured for each node.
- Verify that you have the required privileges.  
See [“Reviewing the requirements”](#) on page 52.

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

### To configure a VCS cluster

- 1 Start the VCS Cluster Configuration Wizard.  
Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.
- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.

- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.




Do one of the following:

- To discover information about all systems and users in the domain:
  - Clear the **Specify systems and users manually** check box.
  - Click **Next**.Proceed to [step 8](#) on page 86.
- To specify systems and user names manually (recommended for large domains):
  - Check the **Specify systems and users manually** check box. Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
  - Click **Next**.  
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 86. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**.  
Do not specify systems that are part of another cluster.  
Proceed to [step 8](#) on page 86.

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



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

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

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

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

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

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


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

The screenshot shows the 'Veritas Cluster Server Configuration Wizard' window, specifically the 'Cluster Details' panel. The left sidebar contains a list of steps: Domain Selection, VCV Options, Cluster Details (highlighted), Cluster Selection, Validate Systems, Edit Options, NIC Selection, Service Account, Security, Summary, and Finish. The main area of the wizard is titled 'Cluster Details' and includes the instruction: 'Enter necessary details to create the new cluster'. Below this, there are three input fields: 'Cluster Name' with the text 'MYCLUSTER', 'Cluster ID' with the value '2', and 'Operating System' with the value 'Windows 2003 (x86)'. A section titled 'Select the systems to create the cluster.' contains a checkbox labeled 'Select all systems' which is checked. Below this is a list box titled 'Available Systems' containing two entries: 'SYSTEM1' and 'SYSTEM2', both of which are checked. At the bottom of the main area, it states 'Total number of systems selected to create the cluster : 2' and 'Click "Next" to continue.' The bottom of the window has three buttons: 'Back', 'Next', and 'Cancel'.

|                  |                                                                                                                                                                                                                                                                                                                                                                            |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cluster Name     | Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.                                                                                                                                                                                                                                                               |
| Cluster ID       | Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.<br><br><b>Caution:</b> If you chose to specify systems and users manually in <a href="#">step 4</a> or if you share a private network between more than one domain, make sure that the cluster ID is unique. |
| Operating System | From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.                                                  |

- Available Systems
- Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.
- The wizard discovers the network interface cards (NICs) on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

- 10
- The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.  
If a system is not validated, review the message associated with the failure and restart the wizard after rectifying the problem.  
If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 90.
- 11
- On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:
  - To configure the VCS private network over the ethernet, complete the following steps:




- Select the check boxes next to the two NICs to be assigned to the private network.

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

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

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:




- Check the **Configure LLT over UDP** check box.
- Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to

65535. The default ports numbers are 50000 and 50001 respectively.

- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.  
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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


- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.  
This account does not require Domain Administrator privileges.



The screenshot shows the 'Veritas Cluster Server Configuration Wizard' window, specifically the 'VCS Helper Service User Account' panel. The title bar reads 'Veritas Cluster Server Configuration Wizard'. The panel title is 'VCS Helper Service User Account' with the subtitle 'Specify a user account for the VCS Helper service'. On the left, a navigation pane lists steps: Domain Selection, VCS Options, Cluster Details, Cluster Selection, Validate Systems, Edit Options, NIC Selection, Service Account (selected), Security, Summary, and Finish. The main area contains instructions: 'Specify a user account in the domain \'vcscdm.com\' for the VCS Helper service. The service will run in the context of the specified user on all nodes in the cluster.' There are two radio buttons: 'Existing User' (selected) and 'New user'. Under 'Existing User', there is a 'Specify User:' label and a dropdown menu showing 'UserA'. Under 'New user', there is a 'Create New User:' label and an empty text field. At the bottom, there are 'Back', 'Next', and 'Cancel' buttons. A 'Click \'Next\' to continue.' instruction is also present.

Specify a domain user as follows:

- To specify an existing user, do one of the following:
    - Click **Existing user** and select a user name from the drop-down list
    - If you chose not to retrieve the list of users in [step 4](#) on page 85, type the user name in the **Specify User** field, and then click **Next**.
  - To specify a new user, click **New user** and type a valid user name in the Create New User field, and then click **Next**.  
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
  - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13 On the Configure Security Service Option panel, specify the security options for the cluster and then click **Next**.  
Do one of the following:
- To use the single sign-on feature, complete the following steps:



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

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

- If you know the name of the system that will serve as the root broker, click **Specify an existing root broker or a node in this cluster to configure as root broker**, type the system name, and then click **Next**.  
If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.  
If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard then configures all nodes in the cluster as authentication brokers.
- If you want to search the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.
- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.  
For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.  
To search for all Windows Server 2003 systems, select **Operating System** from the first drop-down list, **is (exactly)** from the second drop-down list, type **\*2003\*** in the adjacent field, click **Add** and then click **OK**.

Table 5-2 contains some more examples of search criteria.

**Table 5-2** Search criteria examples


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

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

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

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

- To use a VCS user privilege, complete the following steps:



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS Encrypt utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password. After the cluster is configured, you can use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.
- Click **Next**.

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

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

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

- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

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

---


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

---

You are not required to configure the Cluster Management Console (Single Cluster Mode) or Web Console, for this HA environment. Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on VCS Cluster Management Console (Single Cluster Mode), and the Notification resource.

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

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




- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console. See [“Configuring Web console”](#) on page 96.
- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients. See [“Configuring notification”](#) on page 97.

## Configuring Web console

This section describes steps to configure the VCS Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.

### To configure the Web console

- 1 On the Web Console Network Selection panel, specify the network information for the Web Console resources and click **Next**.




- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
  - If you choose to configure a new IP address, type the IP address and associated subnet mask.
  - Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.
  - 3 If you chose to configure a Notifier resource, proceed to: [“Configuring notification”](#) on page 97. Otherwise, click **Finish** to exit the wizard.

## Configuring notification

This section describes steps to configure notification.


To configure notification

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




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

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



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




The screenshot shows the 'Notifier SMTP Configuration' window of the Veritas Cluster Server Configuration Wizard. The window title is 'Veritas Cluster Server Configuration Wizard' and the subtitle is 'Notifier SMTP Configuration'. Below the subtitle is the instruction 'Specify information about SMTP recipients'. On the left, a navigation pane shows the following steps: 'Domain Selection' (checked), 'VCW Options' (checked), 'Create Cluster' (checked), 'Select Components' (checked), 'Configure' (selected), 'Summary', and 'Finish'. The main area contains a text box for 'SMTP Server Name / IP' with the value 'SMTPServer'. Below this is the instruction 'Enter SMTP recipients and select a severity level for each recipient.' followed by a table with two columns: 'Recipients' and 'Severity'. The first row of the table contains 'admin@example.com' and 'Information'. Below the table are two buttons: '+' and '-'. To the left of these buttons is the text 'Click '+' to add a recipient. Click '-' to remove a recipient.' Below the buttons is the text 'Click 'Next' to continue.' At the bottom right are three buttons: 'Back', 'Next', and 'Cancel'.

| Recipients        | Severity    |
|-------------------|-------------|
| admin@example.com | Information |
|                   |             |
|                   |             |
|                   |             |
|                   |             |

- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity to the recipient.
- Click + to add fields; click - to remove a field.

- 4 On the Notifier Network Card Selection panel, specify the network information and click **Next**.



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

# Installing and configuring Enterprise Vault for failover

This chapter contains 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.

### 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 SFW HA 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 Daemon (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 following to the Firewall Exceptions list:
  - Port 14150 or the VCS Command Server service,  
`%vcs_home%\bin\CmdServer.exe`.  
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
  - Port 14141  
For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- 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 81.

## 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 EV Cluster Setup Wizard.  
Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center** to start the Solutions Configuration Center (SCC). 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> |

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

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Virtual Server Name  | <p>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 16 characters.</p> <p>You will need to specify the same name when running the Enterprise Vault Configuration Wizard.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| Virtual IP Address   | Type a unique virtual IP address for the EV virtual server.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Subnet Mask          | Type the subnet to which the virtual server IP belongs.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Advanced Settings    | <p>Click <b>Advanced Settings</b> to specify additional details for the Lanman resource.</p> <p>On the Lanman Advanced Configuration dialog box, complete the following:</p> <div><div>1</div><div>Check <b>Active Directory Update required</b> check box to enable the Lanman resource to update the Active Directory with the virtual name.</div></div> <div><div>2</div><div><p>In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format <b>CN=containername,DC=domainname,DC=com</b>.</p><p>To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box.</p><p>By default, the Lanman resource adds the virtual server to the default container "Computers."</p></div></div> <div><div>3</div><div><p>Click <b>OK</b>.</p><p>The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.</p></div></div> |
| Adapter Display Name | <p>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.</p> <p>For each system in the cluster, select the public network adapter name. Verify that you select the adapters assigned to the public network, not the private.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

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



#### MSMQ Directory

Type a path or click ... (ellipsis button) to browse for the directory. All MSMQ data is stored at this location.

**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.

#### Replication Directory

Type a path or click ... (ellipsis button) to browse for the registry replication directory. This directory contains the list of registry keys to replicate.

Symantec recommends that you configure the MSMQ and registry replication directories on different volumes.

- 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 and 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.

|            |                                                                                                                                                                                                                                                                                                                                                                              |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 | <p>Displays the attributes and their configured values, for a resource selected in the Resources list.</p>                                                                                                                                                                                                                                                                   |

- 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**.

You can now proceed to configuring EV using the EV Configuration Wizard. See [“Configuring Enterprise Vault Server in a cluster environment”](#) on page 107.

# 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 Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. 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 EVSQL, enter virtualsvr\EVSQL.
  - 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 Veritas Cluster Server, Symantec 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 *Veritas Cluster Server Administrator's Guide*.

## Verifying the Enterprise Vault cluster configuration

Failover simulation is an important part of configuration testing.

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

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

### To switch service groups

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
  - Click **Switch To**, and click the appropriate node from the menu.
  - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.  
If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

### To shut down an active cluster node

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

# 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

Table 6-1 lists the items that you can modify in the EV service group. For more information on modifying service groups, see the *Veritas Cluster Server Administrator's Guide*.

**Table 6-1** 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. |
| 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 that 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.



# Configuring disaster recovery for Enterprise Vault

This chapter contains the following topics:

- [Tasks for configuring disaster recovery for Enterprise Vault](#)
- [Verifying your primary site configuration](#)
- [Guidelines for installing SFW HA and configuring the cluster on the secondary site](#)
- [Setting up security for VVR](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [Configuring disaster recovery with the DR wizard](#)
- [Cloning the storage on the secondary site using the DR wizard](#)
- [Installing and configuring Enterprise Vault on the secondary site](#)
- [Configuring VVR replication and global clustering](#)
- [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](#)

# Tasks for configuring disaster recovery for Enterprise Vault

After setting up an SFW HA high availability environment for Enterprise Vault on a primary site, you can create a secondary or “failover” site for disaster recovery.

For Enterprise Vault, the Disaster Recovery (DR) wizard helps you to clone the storage from the primary site to the secondary site. The DR wizard also helps you set up replication and the global cluster (GCO option).

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

See “[About the Solutions Configuration Center](#)” on page 35.

**Note:** Symantec 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. See “[Where to get more information](#)” on page 17.

**Table 7-1**                      Configuring the secondary site for disaster recovery

| Action                                                                                     | Description                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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.<br>See “ <a href="#">Verifying your primary site configuration</a> ” on page 117.                                                                                                 |
| Install SFW HA and configure the cluster on the secondary site                             | <b>Caution:</b> Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.<br><br>See “ <a href="#">Guidelines for installing SFW HA and configuring the cluster on the secondary site</a> ” on page 118. |
| Set up security for VVR                                                                    | Set up security for VVR.<br><br>See “ <a href="#">Setting up security for VVR</a> ” on page 119.                                                                                                                                                                              |
| (Secure cluster only) Assign user privileges                                               | For a secure cluster only, assign user privileges<br><br>See “ <a href="#">Assigning user privileges (secure clusters only)</a> ” on page 122.                                                                                                                                |

**Table 7-1** Configuring the secondary site for disaster recovery (Continued)

| Action                                                          | Description                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Start the DR wizard to begin disaster recovery configuration    | <ul style="list-style-type: none"> <li>■ Review prerequisites for the DR wizard</li> <li>■ Start the DR wizard and make the initial selections: selecting a primary site system, the service group, the secondary site system, and the replication method</li> </ul> <p>See <a href="#">“Configuring disaster recovery with the DR wizard”</a> on page 123.</p> |
| Clone the storage configuration                                 | <p>Clone the storage configuration on the secondary site using the DR wizard</p> <p>See <a href="#">“Cloning the storage on the secondary site using the DR wizard”</a> on page 126.</p>                                                                                                                                                                        |
| Install Enterprise Vault on the cluster nodes                   | See <a href="#">“Installing and configuring Enterprise Vault on the secondary site”</a> on page 129                                                                                                                                                                                                                                                             |
| Create the EV service group configuration on the secondary site | <p>The EV service group must be offline on the primary site before you configure the EV service group on the secondary site.</p> <p>You must use the same service group name and virtual server name as on the primary site.</p> <p>See <a href="#">“Installing and configuring Enterprise Vault on the secondary site”</a> on page 129.</p>                    |
| Configure the Enterprise Vault server in a cluster environment  | See <a href="#">“Installing and configuring Enterprise Vault on the secondary site”</a> on page 129.                                                                                                                                                                                                                                                            |
| Configure VVR replication and global clustering                 | <p>Use the DR wizard to configure replication and global clustering</p> <p>See <a href="#">“Configuring VVR replication and global clustering”</a> on page 132.</p>                                                                                                                                                                                             |
| Verify the disaster recovery configuration                      | <p>Verify that the secondary site has been fully configured for disaster recovery</p> <p>See <a href="#">“Verifying the disaster recovery configuration”</a> on page 139.</p>                                                                                                                                                                                   |
| (Optional) Add secure communication                             | <p>Add secure communication between local clusters within the global cluster (optional task)</p> <p>See <a href="#">“Establishing secure communication within the global cluster (optional)”</a> on page 140.</p>                                                                                                                                               |

**Table 7-1**                      Configuring the secondary site for disaster recovery (Continued)

| Action                                             | Description                                                                                                                                                                                                                                                 |
|----------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| (Optional) Add additional DR sites                 | Optionally, add additional DR sites to a VVR environment<br><br>See “ <a href="#">Adding multiple DR sites (optional)</a> ” on page 142.                                                                                                                    |
| 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<br><br>See “ <a href="#">Recovery procedures for service group dependencies</a> ” on page 143. |

## 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 SFW HA and configuring the cluster on the secondary site

Use the following guidelines for installing SFW HA and configuring the cluster on the secondary site.

**Note:** Symantec 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. See [“Configuring the storage hardware and network”](#) on page 66.
- Ensure that when installing SFW HA you install the appropriate disaster recovery options at both the primary and secondary sites. Be sure to select the following installation options as appropriate for your environment:

|                                       |                                                                                                                                                                                                                                                                                       |
|---------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enterprise Vault Cluster Setup Wizard | Required to configure the EV service group.                                                                                                                                                                                                                                           |
| Client                                | Required to install VCS Cluster Manager (Java console) and Veritas Enterprise Administrator console, which are used during configuring high availability. Also required to install the Solutions Configuration Center which provides information and wizards to assist configuration. |
| Global Cluster Option                 | Required for a disaster recovery configuration.                                                                                                                                                                                                                                       |
| Veritas Volume Replicator             | Required for VVR replication.                                                                                                                                                                                                                                                         |

- For more information see the *SFW HA Installation and Upgrade Guide*.
- Configure the cluster with the VCS Cluster Configuration Wizard (VCW). Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster. See [“Configuring the cluster”](#) on page 84.

**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 security for VVR

If you are using Veritas Volume Replicator (VVR) replication, you must configure the VxSAS service on all cluster nodes. For a disaster recovery environment, you configure the service on all nodes on both the primary and secondary sites.

---

**Note:** Symantec 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 procedure may already be complete.

---

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

The procedure has these prerequisites:

- You must be logged on with administrative privileges on the server for the wizard to be launched.
- The account you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- Make sure that the hosts on which you want to configure the VxSAS service are accessible from the local host.

---

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

---

### To configure the VxSAS service

- 1 To launch the wizard, select **Start > All Programs > Symantec > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or run `vxscfg.exe` from the command prompt of the required machine.  
Read the information provided on the Welcome page and click **Next**.
- 2 Complete the Account Information panel as follows:

|                                  |                                        |
|----------------------------------|----------------------------------------|
| Account name<br>(domain\account) | Enter the administrative account name. |
|----------------------------------|----------------------------------------|

Password

Specify a password.

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

- 3
- On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

Selecting domains


The Available domains pane lists all the domains that are present in the Windows network neighborhood.  
  
Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.

Adding a domain

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

Click **Next**.

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



#### Selecting hosts

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

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

#### Adding a host

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

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

- 5 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.  
Click **Back** to change any information you had provided earlier.
- 6 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 *Veritas Cluster Server Administrator's Guide*.

### To assign user privileges at the primary site

- 1 Set the configuration to read/write mode:  
`haconf -makerw`
- 2 Add the user. Specify the name in the format `username@domain`.  
`hauser -add user [-priv <Administrator|Operator>]`
- 3 Modify the attribute of the service group to add the user. Specify the 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 VVR 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:

- SFW HA is installed and a cluster is configured at the secondary site. Ensure that the name assigned to the secondary site cluster is different than the name assigned to the primary site cluster.
- 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.
- 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 incoming and outgoing TCP requests on port 7419.

**Note:** The DR wizard does not support VVR configurations that include a Bunker secondary site.

In addition, see the following replication prerequisite:

- [“Setting up security for VVR”](#) on page 119

**To start configuring disaster recovery with the DR wizard**

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. Expand the Solutions for 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.<br><br>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. |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Click **Next**.

- 4 In the Service Group Selection panel, select the Enterprise Vault service group for which you want to configure the storage and replication. The wizard can configure only one level of dependency. In a VVR environment, the wizard configures a dependency for the RVG service group, so no other dependency is supported. The panel lists only service groups that contain a MountV resource. Click **Next**.
- 5 In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.

Click **Next**.

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

Configure Veritas Volume Replicator (VVR) and the Global Cluster Option (GCO)

Select this option if you want to configure VVR replication.

Select this option even if you plan to configure VVR replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a VVR environment.

The wizard verifies each configuration task and recognizes if a task has been completed successfully.

You cannot mix replication methods. That is, if your primary site is using array-based replication, and you select the VVR option, the wizard will warn you that you cannot use VVR replication for the disaster recovery site.

Configure EMC SRDF and the Global Cluster Option (GCO)

Select this replication option if you want to configure the settings for the VCS EMC SRDF agent. All disks used for the service group on the primary site must belong to an EMC SRDF array.

Select this option even if you plan to configure EMC SRDF replication or the GCO option manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.

Configure Hitachi TrueCopy and the Global Cluster Option (GCO)

Select this replication option if you want to configure the settings for the VCS Hitachi TrueCopy agent. All disks used for the service group on the primary site must belong to a Hitachi TrueCopy array.

Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.

|                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the Global Cluster Option (GCO) only | <p>If you select this option, the DR wizard does not configure any replication settings. It configures the global cluster option.</p> <p>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.</p> <p>If you select the GCO only option, the DR wizard sets up the storage and service group configuration on the secondary site for an array-based hardware replication environment. Therefore, you cannot use this option to clone the storage and service group for a VVR replication environment.</p> |
|------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Click **Next**.

- 7 Continue with cloning the storage.  
See [“Cloning the storage on the secondary site using the DR wizard”](#) on page 126.

## Cloning the storage on the secondary site using the DR wizard

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at the secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, the wizard can correct this and then complete the configuration.

If you have not yet started the wizard, see the following topic for the wizard prerequisites before continuing with the storage cloning procedure:

- [“Configuring disaster recovery with the DR wizard”](#) on page 123.

### To clone the storage configuration from the primary site to the secondary site (VVR replication method)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system. In the Replication Options panel, select the VVR replication method and click **Next**.
- 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 | <p>Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary.</p> <ul style="list-style-type: none"> <li>■ If the volume does not exist, a new volume will be created.</li> <li>■ If the volume exists but is of a smaller size than that on the primary, the volume will be expanded to the required size.</li> <li>■ If the volume is of a greater size than that on the primary, the volume will be recreated using the appropriate size.</li> <li>■ If the volume is the same as that on the primary, the message indicates that the volumes are identical and no action is required.</li> </ul> |

The summary view shows the following:

|                                             |                                                                                                                                       |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------|
| Disk groups that do not exist               | Displays the names of any disk groups that exist on the primary but do not exist on the secondary.                                    |
| Existing disk groups that need modification | Displays the names of any disk groups on the secondary that need to be modified to match the primary.                                 |
| Free disks present on secondary             | Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information. |

If the panel displays a message indicating that the available disks are inadequate to clone the primary site configuration on the secondary, you can free some disks on the secondary or add more storage. Then click **Refresh/Validate** to have the wizard update its information about the secondary storage configuration.

You continue with the wizard to provide information for the recommended actions. Before proceeding to the service group configuration, the wizard ensures that the configuration of the disk groups and volumes for the service group is the same at the primary and secondary site.

Click **Next**.

- 3
- In the Disk Selection for Storage Cloning panel, for each of the disk groups that does not exist or is not same as the corresponding disk group at the primary site, select disks that the wizard can use to create the respective disk groups at the secondary site.

Selecting Disks

For each of the disk groups that needs to be created, select the required disks from the Available Disks pane. Either double-click on the host name or the >> option to move the hosts into the Selected disks pane.

Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures.

Click **Next**.

- 4
- In the Volume Layout for Secondary Site Storage panel, complete the requested information:

Disk Group

Displays the disk group name to which the volume belongs.

Volume (Volume Size)

Displays the name and the size of the volume, corresponding to that on the primary, that needs to be created on the secondary.

Available Disks

Select the disks on which you want the wizard to create the volumes. From the Available Disks pane, either double-click on the disk name or the >> option to move the disks into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group.

Select disks for each unavailable volume that you want to clone on to the secondary.

Layout

By default, the same layout as the one specified for the primary volume is selected. Click **Edit** to change the layout to suit your specific requirements.

Selected Disks

Displays the list of disks that have been moved in from the Available Disks pane.

View Primary Layout Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 5 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.
- 6 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 7 In the Storage Cloning Configuration Result screen, view the results and click **Next**.
- 8 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

See “[Installing and configuring Enterprise Vault on the secondary site](#)” on page 129.

## 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 81.
- 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.  
 Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. 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 102.
- 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).  
 To launch the wizard, click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**. 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 *EVSQ*L, 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 VVR replication and global clustering

You use the DR wizard to configure VVR replication and global clustering.

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

- Ensure that Veritas Volume Replicator is installed at the primary and secondary site.
- Ensure that Global Cluster Option (GCO) is installed at the primary and secondary site. One static IP address must be available per site for configuring GCO.
- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that you configure a VCS user with the same name and privileges in each cluster.
- Ensure that VVR Security Service (VxSAS) is configured at the primary and secondary site.  
See [“Setting up security for VVR”](#) on page 119.
- Ensure that the storage has been cloned on the secondary site.  
See [“Cloning the storage on the secondary site using the DR wizard”](#) on page 126.
- Ensure that Enterprise Vault is installed and configured on the secondary site.  
See [“Installing and configuring Enterprise Vault on the secondary site”](#) on page 129.
- Bring the EV service group offline at the secondary site and online at the primary site.

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

### To configure VVR replication and GCO

- 1 Verify that the application server service group is online at the primary site and the appropriate disk groups are imported at the secondary site.
- 2 Launch the wizard and proceed to the Replication Setup panel as follows:
  - Start the DR Configuration Wizard from the Solutions Configuration Center. Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.  
Expand the Solutions for Enterprise Vault tab and click **Disaster Recovery Configuration > Configure replication and the global cluster option (GCO) > Disaster Recovery Configuration Wizard**.

- On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
  - On the Replication Methods panel, click **Configure VVR and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.

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 Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**. If not, click **Cancel** and restart the wizard after meeting the requirements.
- 4 In the Replication Settings for Replicated Volume Group panel, specify the requested information. If you are adding a DR site to an existing DR configuration, fields that must match the existing settings, such as the RVG or RDS name, are dimmed so that you cannot change them.

|                      |                                                                                                                                                                                                                                   |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disk Group           | The left column lists the disk groups. By design, an RVG is created for each disk group.                                                                                                                                          |
| RVG Name             | Displays the default RVG name. If required, change this to a name of your choice.                                                                                                                                                 |
| RDS Name             | Displays the default Replicated Data Set (RDS) name. If required, change this to a name of your choice.                                                                                                                           |
| Available Volumes    | Displays the list of available volumes that have not been selected to be a part of the RVG. <p>Either double-click on the volume name or use the &gt; option to move the volumes into the Selected RVG Volumes pane.</p>          |
| Selected RVG Volumes | Displays the list of volumes that have been selected to be a part of the RVG. <p>To remove a selected volume, either double-click the volume name or use the &lt; option to move the volumes into the Available Volumes pane.</p> |

- Primary SRL

If you did not create a Replicator Log volume on the primary site, click **Create New** on the drop-down menu. On the New Volume dialog box, specify the name, size, and disk.

Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.
- Secondary SRL

If you did not create a Replicator Log volume on the primary site, click **Create New** on the drop-down menu. On the New Volume dialog box, specify the same name and size as you specified for the primary SRL.

Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.
- Start Replication after the wizard completes

Select this check box to start replication automatically after the wizard completes the necessary configurations.

Once replication is configured and running, deselecting the checkbox does not stop replication.

- Click **Advanced Settings** to specify some additional replication properties. The options on the dialog box are described column-wise, from left to right; refer to the *Veritas Volume Replicator Administrator's Guide* for additional information on VVR replication options.

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             |
|                       | Mbps                 | Initial Synchronization: | Auto Synchronous |

OKCancel

- Replication Mode
- Select the required mode of replication; **Synchronous**, **Asynchronous**, or **Synchronous Override**. The default is synchronous override.

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Log Protection       | <p>Select the appropriate log protection from the list.</p> <p>The <b>AutoDCM</b> 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.</p> <p>The <b>Off</b> option disables Replicator Log Overflow protection.</p> <p>The <b>Override</b> 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.</p> <p>If the Secondary becomes inactive due to disconnection or administrative action then Replicator log protection is disabled, and the Replicator Log overflows.</p> <p>The <b>Fail</b> 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.</p> |
| 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, VVR replication uses the maximum available bandwidth. You can select <b>Specify</b> to specify a bandwidth limit.</p> <p>The default unit is Mega bits 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 <b>Off</b>.</p> <p>When this option is selected the <b>High Mark Value</b> and the <b>Low Mark Value</b> are disabled. Select the <b>Fail</b> or <b>Override</b> option to enable Latency protection.</p> <p>This <b>Override</b> option behaves like the <b>Off</b> option when the Secondary is disconnected and behaves like the <b>Fail</b> option when the Secondary is connected.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| High Mark Value         | <p>This option is enabled only when Latency Protection is set to <b>Override</b> or <b>Fail</b>. 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 <b>Override</b> or <b>Fail</b>. When the updates in the Replicator log reach the <b>High Mark Value</b>, 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 <b>Low Mark Value</b>. The default is 9950.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Initial Synchronization | <p>If you are doing an initial setup, then use the <b>Auto Synchronous</b> option to synchronize the secondary site and start replication. This is the default.</p> <p>When this option is selected, VVR by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.</p> <p>If you want to use the <b>Synchronize from Checkpoint</b> method then you must first create a checkpoint.</p> <p>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 <b>Synchronize from Checkpoint</b> option to start replication from the checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.</p> |

To apply changes to advanced settings, click **OK**. On the Replication Settings for Replicated Volume Group panel click **Next**.

- 5 In the Replication Attribute Settings panel, specify required replication attribute information for the cluster at the primary and secondary site. Click the arrow icon to expand an RVG row and display the replication attribute fields. If you are configuring an additional secondary site (multiple DR sites), some fields are disabled.

|             |                                                                                                                                                             |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disk Group  | Displays the list of disk groups that have been configured.                                                                                                 |
| RVG Name    | Displays the Replicated Volume Groups corresponding to the disk groups.                                                                                     |
| IP Address  | Enter replication IPs that will be used for replication, one for the primary site and another for the secondary site.                                       |
| Subnet Mask | Enter the subnet mask for the system at the primary site and the secondary site.                                                                            |
| Public NIC  | Select the public NIC from the drop-down list for the system at the primary and secondary site.                                                             |
| Copy        | Enables you to copy the above network settings to any additional RVGs that are listed on this screen. If there is only one RVG, this option does not apply. |

After specifying the replication attributes for each of the RVGs, click **Next**.

- 6 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

|                       |                                                                                                                                                                                           |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Use existing settings | Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists. |
| Resource Name         | Select the existing WAC resource name from the resource name list box.                                                                                                                    |
| Create new settings   | Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.                                                                                       |
| IP Address            | Enter a virtual IP for the WAC resource.                                                                                                                                                  |

|                               |                                                                                                                                                                                                                                                                                                                                                           |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Subnet Mask                   | Enter the subnet mask for the system at the primary site and the secondary site.                                                                                                                                                                                                                                                                          |
| Public NIC                    | Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.                                                                                                                                                                                                                                           |
| Start GCO after configuration | <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. If you have a printer installed, you can click the printer icon at the bottom of the scrollable list to print the settings.  
Click **Next** to implement the settings.
- 8
- In the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for global clustering. If a task could not be completed successfully, it is marked with an (x) symbol. For any critical errors, the wizard displays an error message. For less critical errors, the Information column displays a brief description about the task failure and the next screen displays additional information on what action you can take to remedy it. Click **Next**.
- 9
- In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

# Setting service group dependencies for disaster recovery

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

In the disaster recovery environment with VVR 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 more information on setting service group dependencies, see the *Veritas 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 VVR replication configuration. This includes ensuring that the RVGs have been created at primary and secondary with the correct volume inclusion, replication mode, Replicator Log configuration, and any specified advanced options.
- Confirm that the replication state matches what was specified during configuration. If specified to start immediately, ensure that it is started. If specified to start later, ensure that it is stopped.
- Ensure that the VVR RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.  
See also "[Setting service group dependencies for disaster recovery](#)" on page 138.
- 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 configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you chose to start replication manually in the DR wizard, to avoid replicating large amounts of data over the network the first time, then you will need to start the process necessary to synchronize from checkpoint. This typically consists of
  - starting a VVR replication checkpoint
  - performing a block level backup
  - ending the VVR replication checkpoint
  - restoring the block level backup at the DR site
  - starting replication from the VVR replication checkpointTo learn more about the process of starting replication from a checkpoint, refer to the *Veritas Volume Replicator Administrator's Guide*.
- Do not attempt a wide area failover until data has been replicated and the state is consistent and up to date.

## Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.
- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

#### To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat step 1 to step 3 for the additional clusters in the global cluster.

#### To add the -secure option to the StartProgram resource

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View**, and then **Properties** view.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add -secure switch to the path of the executable Scalar Value.  
 For example:  
`"C:\Program Files\Veritas\Cluster Server\bin\wac.exe"  
 -secure`
- 5 Repeat step 4 for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the **Save and Close Configuration** icon in the tool bar.
- 8 Repeat step 1 to step 7 for each cluster in the global cluster.

#### To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.  
 Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster. The complete syntax of the command is:

```
vssat setuptrust --broker <host:port> --securitylevel  
<low|medium|high> [--hashfile <filename> | --hash <root  
hash in hex>]
```

For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2:

from RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low
```

from RB2, type:

```
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the ClusterService-Proc (wac) resource online on all clusters

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the Process agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat step 1 and step 2 for the additional clusters in the global cluster.

## Adding multiple DR sites (optional)

In a Veritas Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. Veritas Cluster Server supports up to four DR sites. In other replication environments, additional DR sites require manual configuration.

Run the DR wizard and on the Secondary System selection panel, select the new site.

Before you start the wizard on the task of configuring replication and global clustering, ensure that the cluster service group is online at the existing primary and secondary sites. This enables the wizard to configure GCO not only between the selected primary site and the new secondary site but also between the new site and the earlier configured secondary site. Otherwise, the wizard displays a warning message after the global clustering task.

When configuring the VVR replication settings with the wizard for the additional site, fields that must match existing settings are dimmed so that you cannot change them. For example, you cannot change the RVG name or RVG layout on the Replication Settings panel. Similarly, on the Global Cluster Settings panel, GCO has already been configured at the primary site, so the primary site fields are dimmed.

# Recovery procedures for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

See “[Supported disaster recovery configurations for service group dependencies](#)” on page 62.

The procedure and requirements for bringing service group dependencies online at the secondary site depends on their configuration: soft, firm, or hard.

In general, if a child or parent remains online at the primary site, you take it offline before you bring the child and parent service groups online in the correct order on the secondary site.

An exception is the RVG service group, used for VVR replication, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

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

Table 7-2      Online, local, soft dependency link

| Failure condition             | Results                                                                                | Action required                                                                                                         |
|-------------------------------|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------|
| The child service group fails | ■ The parent remains online on the primary site.                                       | 1 Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online.         |
|                               | ■ An alert notification at the secondary site occurs for the child service group only. | 2 Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent). |
|                               | ■ The RVG group remains online.                                                        |                                                                                                                         |

Table 7-2                      Online, local, soft dependency link

| Failure condition              | Results                                                                   | Action required                                                                                                |
|--------------------------------|---------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------|
| The parent service group fails | ■ The child remains online on the primary site.                           | 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. |
|                                | ■ An alert notification at the secondary site occurs for the parent only. | 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).         |
|                                | ■ The RVG group remains online.                                           |                                                                                                                |

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 7-3                      Online, local, firm dependency link

| Failure condition              | Results                                                                                | Action required                                                                                                                                             |
|--------------------------------|----------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| The child service group fails  | ■ The parent goes offline on the primary site.                                         | Secondary site: Bring the service groups online in the appropriate order (child first, then parent).<br><br>Leave the RVG group online at the primary site. |
|                                | ■ An alert notification at the secondary site occurs for the child service group only. |                                                                                                                                                             |
|                                | ■ The RVG group remains online.                                                        |                                                                                                                                                             |
| The parent service group fails | ■ The child remains online on the primary site.                                        | 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online.                                              |
|                                | ■ An alert notification at the secondary site occurs for the parent only.              | 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).                                                      |
|                                | ■ The RVG group remains online.                                                        |                                                                                                                                                             |

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

Table 7-4                      Online, local, hard dependency link

| Failure condition              | Results                                                                                | Action required                                                                                                |
|--------------------------------|----------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------|
| The child service group fails  | ■ The parent goes offline on the primary site.                                         | Secondary site: Bring the service groups online in the appropriate order (child first, then parent).           |
|                                | ■ An alert notification at the secondary site occurs for the child service group only. | Do not take the RVG group offline at the primary site.                                                         |
|                                | ■ The RVG group remains online.                                                        |                                                                                                                |
| The parent service group fails | ■ The child remains online on the primary site.                                        | 1 Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. |
|                                | ■ An alert notification at the secondary site occurs for the parent only.              | 2 Secondary site: Bring the service groups online in the appropriate order (child first, then parent).         |
|                                | ■ The RVG group remains online.                                                        |                                                                                                                |



# Index

## A

- active/passive configuration  
    illustration 57

## C

- cloning for DR
  - secondary storage (VVR replication) 126
- cluster
  - configure LLT over ethernet 88
  - configure LLT over UDP 89
  - configuring network and storage 66
- clusters
  - assigning user privileges 122
  - configuring the cluster 84
  - configuring the hardware and network 66
  - verifying the HA failover configuration 108
  - verifying the primary site configuration for DR 117
- configuration overview
  - disaster recovery 59
  - high availability 57
- configure
  - LLT over ethernet 88
  - LLT over UDP using VCW 89
- configure cluster
  - ethernet 88
  - UDP 89

## D

- disaster recovery (DR)
  - cloning secondary storage (VVR replication) 126
  - deploying for Enterprise Vault 113
  - DR wizard overview 123
  - DR wizard requirements 123
  - illustrated 16
  - IP addresses 61
  - multiple sites 142
  - typical configuration 16
  - verifying HA configuration at primary site 117

- disk groups
  - cloning for secondary site (VVR replication) 126
  - creating 77
  - deporting 83
  - overview 74
  - preconditions 75
  - sample configuration 76
- DR wizard
  - cloning secondary storage VVR replication 126
  - overview 123
  - requirements 123
- drive letters
  - adding to mount volumes 82
- driver signing options
  - resetting 73

## E

- Enterprise Vault
  - DR configuration overview 59
  - HA configuration overview 57
  - HA sample configuration 58
  - installing on secondary (DR) site 129
  - supported software 53
- Enterprise Vault Cluster Setup Wizard 102
- Enterprise Vault Configuration Wizard 107
- EV service group
  - creating 103
  - modifying 110
  - prerequisites 102

## G

- Global Cluster Option (GCO)
  - prerequisites 138
  - secure configuration 140

**H**

- hardware configuration for a cluster 66
- high availability (HA)
  - defined 14
  - verifying the failover 108

**I**

- installing Enterprise Vault Server
  - secondary site 129
- installing SFW HA
  - HA 68
- IP addresses
  - active/passive configuration 58
  - disaster recovery configuration 61

**L**

- LLT over ethernet
  - configuring using VCW 88
- LLT over UDP
  - configuring using VCW 89

**M**

- multiple DR sites 142

**N**

- network configuration for the cluster 66

**P**

- permissions requirements 55
- prerequisites
  - SFW HA 52
- primary host, defined 15
- primary site
  - verifying the cluster configuration 117

**R**

- replication
  - configuring for VVR with DR wizard 132
  - defined 15
- requirements
  - permissions 55
- requirements, additional for SFW HA 55
- requirements, network 54
- requirements, system 54

- resetting
  - driver signing options 73

**S**

- sample configurations
  - Enterprise Vault
    - HA active/passive 58
  - secondary host, defined 15
- secure clusters
  - assigning user privileges 122
- secure GCO, establishing 140
- Security Services
  - configuring 91
- service group
  - configuring for Enterprise Vault 102
- service group dependencies for DR 138
- service group dependencies for HA 108
- service groups
  - dependencies 62, 143
- setting bandwidth
  - using RDS wizard 135
- SFW HA
  - additional requirements 55
  - best practices 56
  - network requirements 54
  - system requirements 54
- SFW HA installation 68
- software requirements
  - Enterprise Vault 53
- Solutions Configuration Center
  - context sensitivity 37
  - overview 35
  - running wizards remotely 43
  - starting 36
  - wizard descriptions 43
  - workflow for active/active configuration 59
- storage cloning with the DR wizard
  - for VVR replication 126
- storage hardware configuration 66

**U**

- user privilege assignment 122

**V**

- VCS
  - configuring the cluster 84
- VCS Configuration Wizard 84

## volumes

- adding drive letters 82
- considerations for VVR and disaster recovery 75
- creating on a cluster disk group 78
- preconditions on a cluster disk group 75
- sample configuration 76

## VVR

- configuration diagram 60
- configuring replication with DR wizard 132
- VxSAS 119

## VxSAS 119

