

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

Windows Server 2008 (x64),
Windows Server 2008 R2 (x64)

6.0.1

December 2012



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

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

Product version: 6.0.1

Document version: 6.0.1 Rev 0

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate 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. See the Third-party Legal Notices document for this product, which is available online or included in the base release media.

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

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

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

by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation

350 Ellis Street

Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

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

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan	customercare_apac@symantec.com
Europe, Middle-East, and Africa	semea@symantec.com
North America and Latin America	supportsolutions@symantec.com

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Section 1 Introduction and Concepts

Chapter 1 Introducing Veritas Storage Foundation and High Availability Solutions for SharePoint Server 2010

About clustering solutions with SFW HA	14
About high availability	14
How a high availability solution works	15
About replication	15
About disaster recovery	16
What you can do with a disaster recovery solution	16
Typical disaster recovery configuration	16
About high availability support for SharePoint Server 2010	18
About disaster recovery support for SharePoint Server 2010	18
About quick recovery support for SharePoint Server 2010	19
Where to get more information	20

Chapter 2 Introducing the VCS agent for SharePoint Server 2010

About the VCS agent for Microsoft SharePoint Server 2010	24
SharePoint Server 2010 agent functions	24
SharePoint Server 2010 agent state definitions	25
SharePoint Server 2010 agent resource type definition	25
SharePoint Server 2010 agent attribute definitions	26

Section 2 Configuration Workflows

Chapter 3 Configuration workflows for SharePoint Server 2010

About using the workflow tables	33
High availability (HA) configuration	34
Disaster recovery configuration	36

Chapter 4 Using the Solutions Configuration Center

About the Solutions Configuration Center	40
Starting the Solutions Configuration Center	40

Available options from the Configuration Center	40
How to launch the Configuration Center wizards	42
Using the Configuration Center from remote systems	42
Solutions wizard logs	43
Following the workflow in the Configuration Center	43

Section 3 Requirements and Planning

Chapter 5 Requirements and planning for your HA and DR configurations

Reviewing the requirements	48
Disk space requirements	48
Requirements for Veritas Storage Foundation High Availability for Windows (SFW HA)	48
Supported SharePoint 2010 versions	49
System requirements for SFW HA	50
Network requirements for SFW HA	50
IPv6 support	51
Permission requirements for SFW HA	52
Additional requirements for SFW HA	52
Best practices for SFW HA	53
Reviewing the HA configuration	54
Sample SharePoint Server 2010 HA configuration	56
Following the HA workflow in the Solutions Configuration Center	57
Reviewing the disaster recovery configuration	58

Section 4 Deployment

Chapter 6 Installing and configuring SFW HA

Configuring the storage hardware and network	64
Installing Veritas Storage Foundation HA for Windows	66
Installing SFW HA server components using the product installer	67
Applying the selected installation and product options to multiple systems	73
Installing the client components using the product installer	73
Configuring the cluster	76
Configuring notification	85
Adding nodes to an existing cluster	87

Chapter 7 Installing and configuring SharePoint Server 2010 for high availability

Installing and configuring SharePoint Server	95
Configuring 64-bit Perl for SharePoint 2010	96
Configuring SharePoint Server service groups	96
Before you configure a SharePoint service group	97
Creating a SharePoint service group	98
About service groups for SharePoint Search	100
Verifying the SharePoint cluster configuration	101
Considerations when modifying a SharePoint service group	102
Chapter 8	Configuring disaster recovery for SharePoint Server 2010
Tasks for configuring disaster recovery for SharePoint Server 2010	105
Configuring the SQL Server service group for DR	
in the SharePoint environment	107
Updating the SQL Server IP address	108
Updating the IP address for web requests	109
Requirements	109
Customizing the DNS update settings for the web servers	110
Configuring a resource for the web servers	112
Configuring the secondary site for SharePoint disaster recovery	115
Installing SFW HA and configuring the cluster on the	
secondary site	115
Installing the SharePoint servers on the secondary site	115
Configuring the SharePoint service groups on the secondary site	116
Verifying the service group configuration	116
Index	117

Introduction and Concepts

This section contains the following chapters:

- [Introducing Veritas Storage Foundation and High Availability Solutions for SharePoint Server 2010](#)
- [Introducing the VCS agent for SharePoint Server 2010](#)

Introducing Veritas Storage Foundation and High Availability Solutions for SharePoint Server 2010

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](#)
- [About high availability support for SharePoint Server 2010](#)
- [Where to get more information](#)

About clustering solutions with SFW HA

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

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

This guide describes the high availability and disaster recovery solutions for SharePoint Server 2010.

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.

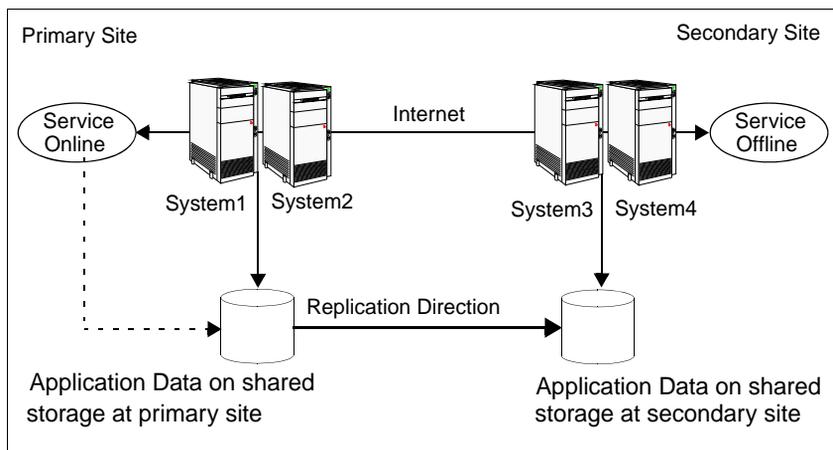
Strategically planning a DR solution provides businesses with affordable ways to meet their service level agreements, comply with government regulations, and minimize their business risks.

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.

About high availability support for SharePoint Server 2010

The high availability (HA) solution for SharePoint Server 2010 is a combination of monitoring and recovery support for SharePoint 2010 applications and high availability support for SQL Server databases used by SharePoint Server 2010.

The SharePoint 2010 high availability configuration components are as follows:

- VCS provides a new agent for SharePoint 2010 that performs the task of managing the SharePoint 2010 Web Applications, Service Applications, and services configured in the server farm. Depending on the configuration, the agent monitors, starts, and stops the SharePoint components in the cluster.
- SharePoint 2010 Web Applications are configured in a VCS parallel service group. A parallel service group runs simultaneously on multiple nodes in a cluster. The parallel service group manages the Web Applications configured in the farm. The state of the parallel service group represents the state of the Web Applications configured in the farm. If a Web Application becomes unavailable, the agent attempts to restart the application in the farm.
- SharePoint 2010 Service Applications and services are configured in a separate service group that is created locally on each cluster node. The service group manages the components configured on the local node only. If any of the components become unavailable, the agent attempts to restart the component on the local node.
- The VCS SQL Server database agents are used to configure high availability for the SharePoint databases. The agents monitor the health of the SharePoint databases as well as underlying resources and hardware. If a failure occurs, predefined actions bring up SQL Server on another node in the cluster.

About disaster recovery support for SharePoint Server 2010

Disaster recovery (DR) support for SharePoint Server 2010 involves configuring service groups for the SharePoint Web and Application servers at the primary and secondary sites and configuring DR for the SharePoint databases using the VCS DR solution for SQL Server.

After you have configured a primary site for high availability, you can set up a secondary site to create a wide area disaster recovery environment. Wide area disaster recovery uses a global cluster to enable SQL Server to failover between clusters at geographically-dispersed sites.

You can configure SharePoint Web and Application servers on the secondary site to allow for running applications and services on the secondary site if the

primary site fails. After completing the configuration, you will be able to efficiently bring your application and web services and data online at an alternate site in the event of a catastrophic failure at your primary production site.

Note: For configuring DR for SharePoint 2010, the SharePoint servers at the primary site and the secondary site must belong to the same SharePoint farm.

About quick recovery support for SharePoint Server 2010

Quick recovery (QR) solution for SharePoint Server 2010 involves scheduling and creating snapshot copies of production volumes of the SQL database. Configuring QR requires using the SFW FlashSnap technology along with Microsoft Volume Shadow Copy Services (VSS) framework to quiesce the database and ensure a persistent snapshot of the production data.

Use the FlashSnap solution to take snapshots of the SharePoint 2010 Web Applications data, Service Applications data, and the farm configuration data. You create a VSS snapshot from the SQL cluster node that hosts the SharePoint Server components data. You use the VSS snapshot wizard to take snapshots of the volumes associated with the SQL databases.

Refer to the *SFW Administrator's Guide* for more details.

For more information on quick recovery for SQL Server, refer to the Quick Recovery Solutions Guides.

Where to get more information

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

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

Table 1-1 SFW HA solutions guides for SQL Server

Title	Description
<i>Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2005</i>	Solutions for SQL Server 2005 and Veritas Cluster Server clustering with Veritas Storage Foundation HA for Windows <ul style="list-style-type: none"> ■ High availability (HA) ■ Campus clusters ■ Replicated data clusters ■ Disaster recovery (DR) with Veritas Volume Replicator or hardware array replication
<i>Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2008 and 2008 R2</i>	Solutions for SQL Server 2008, SQL Server 2008 R2, and Veritas Cluster Server clustering with Veritas Storage Foundation HA for Windows <ul style="list-style-type: none"> ■ High availability (HA) ■ Campus clusters ■ Replicated data clusters ■ Disaster recovery (DR) with Veritas Volume Replicator or hardware array replication
<i>Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2012</i>	Solutions for SQL Server 2012 and Veritas Cluster Server clustering with Veritas Storage Foundation HA for Windows <ul style="list-style-type: none"> ■ High availability (HA) ■ Campus clusters ■ Replicated data clusters ■ Disaster recovery (DR) with Veritas Volume Replicator or hardware array replication
<i>Veritas Storage Foundation and High Availability Solutions Microsoft Clustering Solutions Guide for Microsoft SQL 2005, 2008, 2008 R2, and 2012</i>	Solutions for SQL Server and Microsoft clustering with Veritas Storage Foundation for Windows: <ul style="list-style-type: none"> ■ High availability (HA) ■ Campus clusters ■ Disaster recovery (DR) with Veritas Volume Replicator

Table 1-1 SFW HA solutions guides for SQL Server

Title	Description
<i>Veritas Storage Foundation and High Availability Solutions Quick Recovery Solutions Guide for Microsoft SQL 2005, 2008, 2008 R2, and 2012</i>	Quick Recovery solutions for SQL Server 2005, 2008, and 2008 R2 using either Veritas Storage Foundation for Windows or Veritas Storage Foundation HA for Windows.

Introducing the VCS agent for SharePoint Server 2010

This chapter contains the following topics:

- [About the VCS agent for Microsoft SharePoint Server 2010](#)
- [SharePoint Server 2010 agent functions](#)
- [SharePoint Server 2010 agent state definitions](#)
- [SharePoint Server 2010 agent resource type definition](#)
- [SharePoint Server 2010 agent attribute definitions](#)

About the VCS agent for Microsoft SharePoint Server 2010

The VCS application agent for Microsoft SharePoint Server manages SharePoint Server 2010 Service Applications, Web Applications, and services in a VCS cluster. The agent provides monitoring support in making a SharePoint Server 2010 applications highly available in a VCS environment.

Depending on the configuration, the agent performs the following operations:

- monitors and starts the configured SharePoint services
- monitors the configured Web Applications, brings them online, and takes them offline
- monitors the configured Service Applications, brings them online, and takes them offline

If any of the configured SharePoint component fails or is unavailable, the agent attempts to start the component on the local node. If the components fails to start, the agent declares the resource as faulted.

SharePoint Server 2010 agent functions

Agent functions include the following:

Online	Starts the configured Web Applications, Service Applications, or services.
Offline	Stops the configured Web Applications and Service Applications. The agent also stops monitoring the configured services on the node.
Monitor	Verifies the status of the configured Web Application, Service Application or service. If the components are running, the agent reports the resource as ONLINE. If any of the components are not running, the agent reports the resource as FAULTED.
Clean	Forcibly stops the configured Web Applications and Service Applications. The agent also stops monitoring the configured services on the node.

SharePoint Server 2010 agent state definitions

Agent state definitions are as follows:

ONLINE	Indicates that the configured Web Applications, Service Applications, or services are running on the cluster node.
OFFLINE	Indicates that the configured Web Applications and Service Applications are stopped on the cluster node. It also indicates that the monitoring for the services is also stopped.
FAULTED	Indicates that the agent is unable to start the configured Web Applications, Service Applications, or services on the cluster node.
UNKNOWN	Indicates that the agent is unable to determine the status of the configured SharePoint components on the cluster node.

SharePoint Server 2010 agent resource type definition

The resource type represents the VCS configuration definition of the agent and specifies how the agent is defined in the cluster configuration file, main.cf.

The SharePoint Server 2010 agent is represented by the SharePointServer resource type.

```

type SharePointServer (
  static i18nstr ArgList[] = { AppType,
  AppName, AppPoolMon, FarmAdminAccount, FarmAdminPassword,
  ServiceIDList, StopSPSService }
  str AppType
  i18nstr AppName
  str AppPoolMon = NONE
  i18nstr FarmAdminAccount
  str FarmAdminPassword
  i18nstr ServiceIDList[]
  boolean StopSPSService = 0
)

```

SharePoint Server 2010 agent attribute definitions

Review the tables of required and optional attributes to familiarize yourself with the agent attributes for a SharePointServer resource type. This information will assist you during the agent configuration.

Table 2-1 SharePoint Server 2010 agent required attributes

Required Attributes	Type and Dimension	Definition
AppType	string-scalar	<p>Defines whether the agent is configured to monitor a SharePoint Web Application, Service Application, or service.</p> <p>This attribute can take one of the following values:</p> <ul style="list-style-type: none">■ WebApp■ ServiceApp■ SPSService <p>The default value is WebApp.</p> <p>If this attribute value is set to WebApp or ServiceApp, then you must specify a value for the AppName attribute.</p> <p>If this attribute value is set to SPSService, the AppName attribute value is ignored.</p>

Table 2-1 SharePoint Server 2010 agent required attributes (Continued)

Required Attributes	Type and Dimension	Definition
AppPoolMon	string-scalar	<p>Defines the monitoring modes for the application pool associated with the Web site being monitored.</p> <p>Configure this attribute only if AppType attribute value is set to WebApp and IIS is configured to run in the Worker Process Isolation mode.</p> <p>The attribute can take one of the following values:</p> <ul style="list-style-type: none"> ■ NONE: Indicates that the agent does not monitor the application pool associated with the Web site. ■ DEFAULT: Indicates that the agent monitors the root application pool associated with the Web site. If this attribute is set, the agent starts, stops, and monitors the root application pool associated with the Web site. If the root application pool is stopped externally, the service group faults; the agent then attempts to start the root application pool. ■ ALL: Indicates that the agent starts all the application pools associated with the Web site, but monitors and stops the root application pool only. If any application pool is stopped externally, the service group faults; the agent then attempts to start the application pool. <p>The default value is NONE.</p>

Table 2-1 SharePoint Server 2010 agent required attributes (Continued)

Required Attributes	Type and Dimension	Definition
ServiceIDList	string-vector	<p>Defines the service IDs of the SharePoint services that are managed by the agent. This attribute is always local, that is, it is different for each cluster node.</p> <p>This attribute can take the following values:</p> <ul style="list-style-type: none"> ■ If AppType attribute value is set to WebApp, specify the service ID of the Microsoft SharePoint Foundation Web Application service. ■ If AppType attribute value is set to ServiceApp, specify the service ID of the service on which the Service Application depends. ■ If AppType attribute value is set to SPSService, specify the service IDs of the SharePoint services. <p>Note: If you are configuring this attribute manually, use the VCS hadiscover command or the SharePoint server cmdlets to retrieve the service IDs.</p>

Table 2-2 SharePoint Server 2010 agent optional attributes

Optional Attribute	Type and Dimension	Definition
AppName	string-scalar	<p>The name of the SharePoint Web Application or Service Application that is managed by the agent. The value of this attribute depends on the value of the AppType attribute.</p> <p>This attribute can take the following values:</p> <ul style="list-style-type: none"> ■ If AppType attribute value is set as WebApp, specify the Web Application name. ■ If AppType attribute value is set as ServiceApp, specify the application pool ID for the SharePoint Service Application. <p>Note: This attribute is ignored if AppType attribute value is set as SPSService.</p>

Table 2-2 SharePoint Server 2010 agent optional attributes (Continued)

Optional Attribute	Type and Dimension	Definition
FarmAdminAccount	string-scalar	<p>A user account that has the SharePoint Server Farm Admin privileges.</p> <p>User name can be of the form <i>username@domain.com</i>, <i>domain\username</i>, or <i>domain.com\username</i>.</p> <p>The agent uses the Farm Admin user account context to manage the services specified in the ServiceIDList attribute value.</p>
FarmAdminPassword	string-scalar	<p>The password of the user specified in the FarmAdminAccount attribute value.</p> <p>The password is stored in the VCS configuration in an encrypted form.</p>
StopSPSService	boolean	<p>When a resource in the VCS cluster is taken offline:</p> <ul style="list-style-type: none"> ■ If the value of this attribute is set to true, the agent stops all the SharePoint services in its ServiceIDList. ■ If the value of this attribute is set to false, the agent does not change the state of the SharePoint services in its ServiceIDList, but it stops monitoring the services. <p>The default value of this attribute is false.</p>

Configuration Workflows

This section contains the following chapters:

- [Configuration workflows for SharePoint Server 2010](#)
- [Using the Solutions Configuration Center](#)

Configuration workflows for SharePoint Server 2010

This chapter contains the following topics:

- [About using the workflow tables](#)
- [High availability \(HA\) configuration](#)
- [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 SharePoint Server.

See [“About the Solutions Configuration Center”](#) on page 40.

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

Table 3-1 outlines the high-level objectives and the tasks to complete each objective for a high availability configuration.

Note: Symantec recommends as a best practice to configure SQL Server for high availability before configuring SharePoint Server 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 20.

Table 3-1 SharePoint Server: High availability configuration tasks

Action	Description
Verify hardware and software requirements	See “Reviewing the requirements” on page 48.
Review the HA configuration	<ul style="list-style-type: none"> ■ Understand active-passive configuration See “Reviewing the HA configuration” on page 54.
Configure the storage hardware and network	<ul style="list-style-type: none"> ■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which SharePoint Server will be installed See “Configuring the storage hardware and network” on page 64.
Install SFW HA	<ul style="list-style-type: none"> ■ Install Veritas Storage Foundation HA for Windows See “Installing Veritas Storage Foundation HA for Windows” on page 66.

Table 3-1 SharePoint Server: High availability configuration tasks (Continued)

Action	Description
Configure VCS cluster	<p>You can include both SQL Server and SharePoint Server systems in the same cluster if they use the same operating system platform.</p> <p>If you are configuring SharePoint Server in a separate cluster, perform the following tasks:</p> <ul style="list-style-type: none"> ■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster <p>See “Configuring the cluster” on page 76.</p> <p>If you are adding the SharePoint systems to the existing SQL Server cluster, perform the following task:</p> <ul style="list-style-type: none"> ■ Run the VCS Cluster Configuration Wizard (VCW) to add the nodes <p>See “Adding nodes to an existing cluster” on page 87.</p>
Install SharePoint Server on the cluster nodes	<ul style="list-style-type: none"> ■ Install and configure Microsoft SharePoint Server on each cluster node and configure the farm. While installing, select the Complete installation mode. The Stand-alone install mode is not supported. <p>Refer to the SharePoint Server documentation for installation instructions</p>
Create SharePoint Server service groups	<ul style="list-style-type: none"> ■ Launch the VCS SharePoint Server Configuration Wizard on a node on which SharePoint is installed and configured to create SharePoint service groups <p>See “Configuring SharePoint Server service groups” on page 96.</p>
Verify the HA configuration	<p>Test failover between nodes</p> <p>See “Verifying the SharePoint cluster configuration” on page 101.</p>

Disaster recovery configuration

For configuring disaster recovery, you first begin by configuring the primary site for high availability.

See [“High availability \(HA\) configuration”](#) on page 34.

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

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

Table 3-2 Configuring the secondary site for disaster recovery

Action	Description
Configure SQL Server for disaster recovery at the secondary site	For the steps for configuring SQL Server for high availability and disaster recovery, refer to the HA and DR solutions guide for the desired SQL Server version.
Modify the SQL Server service group on the primary and secondary site	Edit the SQL Server service group on both the primary and secondary site to allow updating the NLB details if a disaster recovery failover occurs. See “Configuring the SQL Server service group for DR in the SharePoint environment” on page 107.
Verify that SharePoint has been configured for high availability at the primary site	Verify that SharePoint has been configured for high availability at the primary site. See “Verifying the SharePoint cluster configuration” on page 101.
Install SFW HA and configure the cluster on the secondary site	Install SFW HA on the SharePoint server systems on the secondary site. You can optionally use the same SFW HA cluster for both SQL Server and SharePoint Server if all systems use the same operating system platform. Otherwise, create a separate cluster for SharePoint. See “Configuring the secondary site for SharePoint disaster recovery” on page 115.

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

Action	Description
Install SharePoint on the cluster nodes on the secondary site	<p>Install Microsoft SharePoint Server on the SharePoint servers on the secondary site. Run the Microsoft SharePoint Products Configuration wizard to add the servers to the existing primary site farm. Choose the option to connect to an existing server farm.</p> <p>Note: You do not need to configure the same number of web servers or service applications on the secondary site as on the primary site. However, you should provide for all required services.</p>
Create the SharePoint service groups on the secondary site	<p>Configure the SharePoint Server service groups for the secondary site</p> <p>The VCS SharePoint Server Configuration Wizard helps you create SharePoint Server service groups.</p> <p>See “Configuring SharePoint Server service groups” on page 96.</p>
Verify the disaster recovery configuration	<p>In the Veritas Cluster Server Java console, ensure that you can bring the SharePoint service groups online and offline.</p>

Using the Solutions Configuration Center

This chapter covers the following topics:

- [About the Solutions Configuration Center](#)
- [Starting the Solutions Configuration Center](#)
- [Available options from the Configuration Center](#)
- [How to launch the Configuration Center wizards](#)
- [Using the Configuration Center from remote systems](#)
- [Solutions wizard logs](#)
- [Following the workflow in the Configuration Center](#)

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 Veritas Storage Foundation High Availability for Windows (SFW HA) environment. The Configuration Center provides solutions for the following applications:

- Microsoft SQL Server 2008, 2008 R2, and 2012
- Additional applications

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

- High availability at a single site for a new installation
- High availability at a single site for an existing server
- Campus cluster disaster recovery, including the following:
 - Campus cluster using SFW HA
 - Campus cluster using Microsoft clustering
- Wide area disaster recovery involving multiple sites
- Quick Recovery for on-host recovery from logical errors in application data
- Fire drill to test the fault readiness of a disaster recovery environment

Starting the Solutions Configuration Center

You can start the Solutions Configuration Center in the following ways:

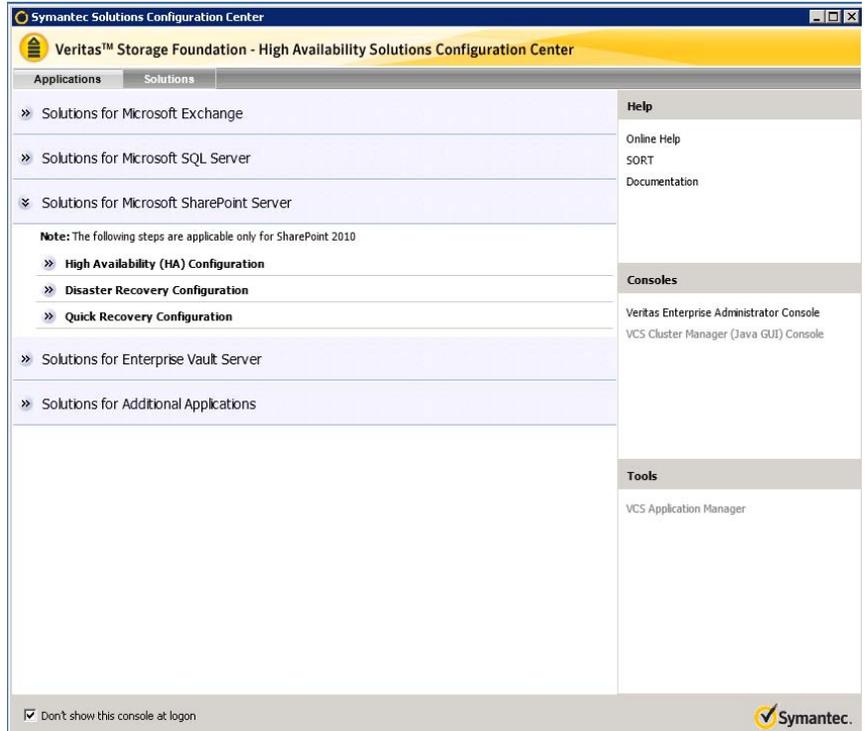
- Click **Start > All Programs > Symantec > Veritas Storage Foundation > Solutions Configuration Center**.
- Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center**.
- Click **Start > Run** and type **scc**.

Available options from the Configuration Center

On the Applications tab, the Solutions Configuration Center is context-sensitive to the application. For example, the solutions displayed when you click the application name are those available for that application. The steps that are shown when you click on a solution are customized for that application.

Figure 4-1 shows the solutions available when you click Solutions for Microsoft SharePoint Server.

Figure 4-1 Solutions Configuration Center for SharePoint Server 2010



How to launch the Configuration Center wizards

The Solutions Configuration Center provides two ways to access wizards:

Applications tab	Lists solutions by application. Provides step-by-step configuration instructions that include buttons to launch the appropriate wizard for each step.
Solutions tab	(For advanced users) Lists wizards by solution, without additional instructions, as follows: <ul style="list-style-type: none">■ High Availability Configuration Wizards■ Disaster Recovery Configuration Wizards■ Quick Recovery Configuration Wizards■ Fire Drill Configuration Wizards

You can go directly to a particular wizard.

Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows 2012 Server Core systems.

Using the Configuration Center from remote systems

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

Disaster Recovery Configuration Wizard	Configures wide area disaster recovery, including cloning storage, cloning service groups, and configuring the global cluster. Also can configure Veritas Volume Replicator (VVR) replication or configure the VCS resource for EMC SRDF and Hitachi TrueCopy array-based hardware replication. Requires first configuring high availability on the primary site. To configure IPv6 settings, the wizard must be launched from a system on which the IPv6 stack is installed.
Fire Drill Wizard	Sets up a fire drill to test disaster recovery. Requires configuring disaster recovery first. To configure IPv6 settings, the wizard must be launched from a system on which the IPv6 stack is installed.
Quick Recovery Configuration Wizard	Schedules preparation of snapshot mirrors and schedules the Quick Recovery snapshots.

VCS Configuration Wizard Sets up the VCS cluster.

VVR Security Service Configuration Wizard Configures the VVR security service.

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

New Dynamic Disk Group Wizard Launched from the Veritas Enterprise Administrator console.

New Volume Wizard Launched from the Veritas Enterprise Administrator console.

MSMQ Configuration Wizard Configures a Microsoft Message Queuing (MSMQ) service group.

SharePoint 2010 Configuration Wizard Configures SharePoint Server 2010 service groups
You can run the wizard from any SFW HA cluster node where SharePoint Server is installed and configured.

SFW Configuration Utility for Hyper-V Live Migration Support Configure SFW for Microsoft Hyper-V Live Migration support on the selected systems.

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 path:

C:\ProgramData\Veritas\winsolutions\log

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 Solutions

Configuration Center open on that system. In this way, you can see what step comes next, drill down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format.

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

Figure 4-2 shows the high-level overview of the workflow steps for configuring high availability for SharePoint Server 2010 from the Solutions Configuration Center.

Figure 4-2 Workflow for configuring high availability for SharePoint Server 2010



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](#)
- [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

The following table estimates disk space requirements for SFW HA.

Table 5-1 Disk space requirements

Installation options	Required disk space
SFW HA + all options	832 MB
Client components	916 MB

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

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

- Review the general installation requirements for SFW HA in the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- Review the SFW HA 6.0.1 Hardware Compatibility List to confirm supported hardware:
<http://www.symantec.com/docs/TECH152806>
- Review the SFW HA 6.0.1 Software Compatibility List to confirm supported software:
<http://www.symantec.com/docs/TECH153742>

Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows 2012 Server Core systems.

- Review the SharePoint 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 VCS Global Clustering Option and if required for your replication solution, select the SFW Veritas Volume Replicator option.

- When installing SFW HA for a Replicated Data Cluster configuration, ensure that you select the VCS option to install Veritas Volume Replicator.

Supported SharePoint 2010 versions

This release of SFW HA provides support for SharePoint Server 2010 RTM and SP1 on the following operating systems:

- Windows Server 2008 with Service Pack 2 or later, x64
- Windows Server 2008 R2

See the Microsoft documentation for details on required SQL Server database versions supported with SharePoint Server 2010.

System requirements for SFW HA

Systems must meet the following requirements for SFW HA:

- Memory must be a minimum 4 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 53.
- NIC teaming is not supported for the VCS private network.

Network requirements for SFW HA

SFW HA has the following network requirements:

- 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 are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated. Static IP addresses are required for the following purposes:
 - A minimum of one static IP address for each physical node in the cluster.
 - One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.

- For 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.
- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.
- If Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).

IPv6 support

For IPv6 networks, the following is supported:

Types of addresses	<p>The following types of IPv6 addresses are supported:</p> <ul style="list-style-type: none"> ■ Unicast addresses Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	<p>LLT over UDP is supported on both IPv4 and IPv6.</p> <p>You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.</p>

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

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

Note: Support is limited to mixed mode (IPv4 and IPv6) network configurations only; a pure IPv6 environment is currently not supported.

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.
- If User Access Control (UAC) is enabled on Windows Server systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, `C:\Program Files\Veritas`). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.

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

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 VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- This is applicable for a Replicated Data Cluster configuration.
You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the vxclus UseSystemBus ON command.

Reviewing the HA configuration

Symantec recommends as a best practice to configure SQL Server for high availability before configuring SharePoint Server.

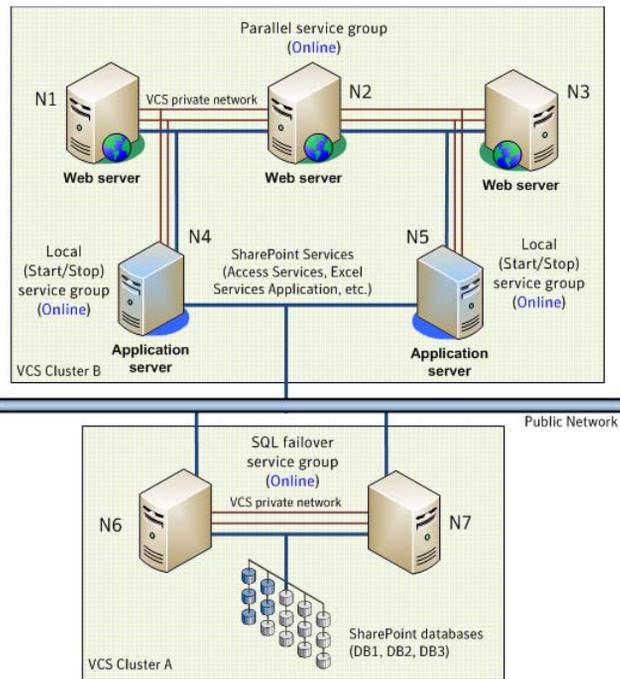
Configuring SQL Server for high availability is covered in the SQL Server solutions guides. See [“Where to get more information”](#) on page 20.

In a typical example of a SharePoint Server 2010 high availability environment, SharePoint Web Applications and Service Applications are configured on a separate set of cluster nodes. A VCS parallel service group manages the Web Applications residing on the Web servers and local service groups manage the application servers. The SharePoint databases are made highly available using the VCS SQL Server service group. The databases reside on shared storage that is accessible from all the SharePoint server nodes in the cluster.

[Figure 5-1](#) illustrates a typical SharePoint Server 2010 configuration. The SharePoint farm layout is as follows:

- Nodes N1, N2, and N3 are the Web front end servers
- Nodes N4 and N5 are the application servers
- Nodes N6 and N7 host the SharePoint SQL databases

Figure 5-1 SharePoint Server 2010 high availability configuration



The graphic displays SQL and SharePoint Servers in different clusters. However, if the SharePoint Servers and SQL Servers are using the same operating system and platform, you can configure both SQL and SharePoint nodes in the same cluster.

The SharePoint Web Applications are configured in a parallel service group that is online on Nodes N1, N2, and N3. The application servers host SharePoint services such as Access Services and Excel Services that are used by the Web servers. These application services are configured in local service groups on nodes N4 and N5 separately. If any of the configured Web or Service applications become unavailable, the SharePoint agent attempts to restart those components in the cluster. If the component fails to come online, the agent declares the resource as faulted.

The databases are made highly available by the SQL service group that is configured on nodes N6 and N7. The databases are configured on the shared storage. The SQL virtual server is online on node N6. All client requests are handled by node N6. N7 waits in a warm standby state as a backup node, prepared to begin handling client requests if N6 becomes unavailable. If N6 fails, N7 becomes the active node and the SQL virtual server comes online on N7.

From the user’s perspective there will be a small delay as the backup node comes online, but the interruption in effective service is minimized.

Sample SharePoint Server 2010 HA configuration

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

[Table 5-2](#) shows a sample SharePoint configuration. If you plan to take snapshots of SharePoint components using FlashSnap, you must ensure that the SharePoint database components are configured on volumes on shared storage.

Table 5-2 Sample SharePoint Server 2010 HA configuration objects

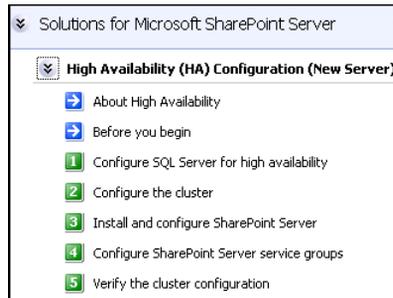
Name	Object
N1, N2, N3, N4, N5	SharePoint Server nodes
N6, N7	SQL Server nodes
SharePoint_Config-WebApplication1	Name of the parallel service group configured for the SharePoint Web Applications.
SharePoint_Config-N4-ServiceApp1	Names of the local service groups configured for the SharePoint Service Applications or services.
SharePoint_Config-N5-ServiceApp2	
INST1	SQL Server instance name
INST1_DG	cluster disk group
INST1-VS	SQL Server virtual server name
INST1_SG	SQL Server service group name
INST1_DB1_VOL	Volume for SQL Server database
INST1_DB1_LOG	Volume for SQL Server database logs

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 SharePoint Server.

[Figure 5-2](#) shows the workflow under the High Availability (HA) Configuration in the Solutions Configuration Center.

Figure 5-2 Configuration steps in the Solutions Configuration Center



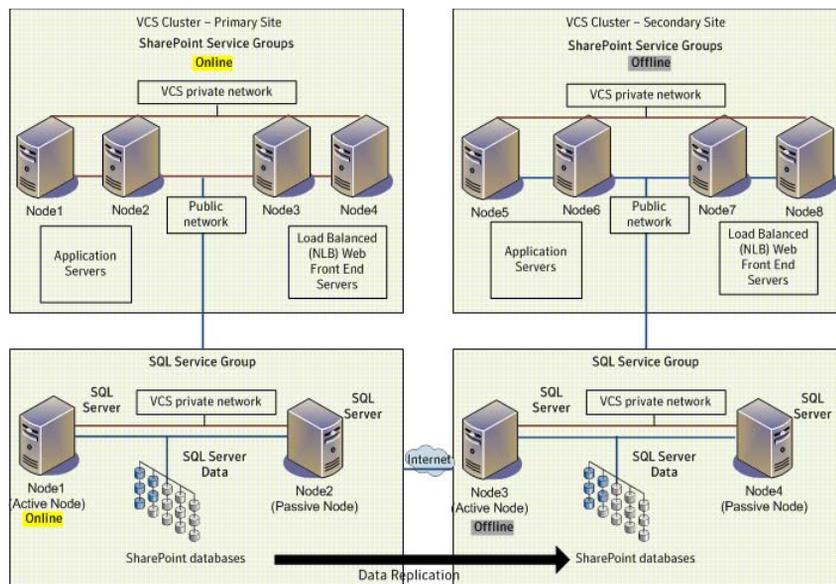
Reviewing the disaster recovery configuration

You configure SQL Server for disaster recovery before configuring SharePoint Server.

Configuring SQL Server for disaster recovery is covered in the SQL Server solutions guides. See [“Where to get more information”](#) on page 20.

Figure 5-3 shows an example SharePoint Server disaster recovery configuration.

Figure 5-3 Example SharePoint Server disaster recovery configuration



The example configuration for SharePoint disaster recovery shows SharePoint configured in a separate cluster from SQL Server. However, you can optionally configure SharePoint Server in the same cluster as SQL Server if all systems use the same operating system.

In the example setup, there are eight SharePoint servers, four for the primary site and four for the secondary site. This is an example only; any supported farm configuration can be used. The SharePoint nodes will form two separate clusters, one at the primary site and one at the secondary site.

Note: You do not need to configure the same number of SharePoint web servers or application servers on the secondary site as on the primary site. However, you should provide for all required services to be available on the secondary site.

The sample setup for SQL Server 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. Disaster recovery configuration for SQL Server configures a global cluster with replication of the databases from the primary to the secondary site.

If the SQL Server primary site fails, the replicated SQL Server databases on the secondary site come online, along with SQL Server. In addition, the SharePoint Servers on the secondary site will automatically start responding to clients.

If the SharePoint Servers fail on the primary site, but SQL Server remains online on the primary site, you would need to manually switch the SQL Server service group to the secondary site. This would be necessary for the secondary site SharePoint servers to respond to clients.

Deployment

This section contains the following chapters:

- [Installing and configuring SFW HA](#)
- [Installing and configuring SharePoint Server 2010 for high availability](#)
- [Configuring disaster recovery for SharePoint Server 2010](#)

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 the cluster](#)
- [Adding nodes to an existing 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

You can install SFW HA using either the product installer wizard or the command line interface (CLI). Before you begin to install the product, ensure that you have reviewed and performed the required preinstallation and planning tasks.

Note: The Windows MPIO feature must be enabled for Windows Server 2012 before you proceed with SFW HA 6.0.1 installation.

During the installation you can choose to separately install the server components or the client components. If you choose to install the server components, the following components are installed by default.

Table 6-1 List of options installed by default

Option	Description
Client components	Installs the VCS Java Console on the same node where the server components are installed.
High Availability Hardware Replication Agents	<ul style="list-style-type: none"> <li data-bbox="625 907 1129 1020">■ Veritas Cluster Server Hardware Replication Agent for EMC MirrorView Enables VCS to manage MirrorView replicated devices. <li data-bbox="625 1029 1129 1112">■ Veritas Cluster Server Hardware Replication Agent for EMC SRDF Enables VCS to manage SRDF replicated devices. <li data-bbox="625 1121 1129 1229">■ Veritas Cluster Server Hardware Replication Agent for EMC SRDFSTAR Enables VCS to manage SRDFSTAR replicated devices. <li data-bbox="625 1237 1129 1350">■ Veritas Cluster Server Hardware Replication Agent for Hitachi TrueCopy Enables VCS to manage TrueCopy replicated devices. <li data-bbox="625 1359 1129 1472">■ Veritas Cluster Server Hardware Replication Agent for MetroMirror Enables VCS to manage MetroMirror replicated devices.
High Availability Database Agents	<ul style="list-style-type: none"> <li data-bbox="625 1498 1129 1576">■ Veritas Cluster Server Database Agent for SQL Server—installs the VCS agent for SQL Server 2008, 2008 R2, and 2012

Note: The high availability agents that get installed with the product software are also available in the form of an agent pack. The agent pack is released on a quarterly basis. The agent pack includes support for new applications as well as fixes and enhancements to existing agents. You can install the agent pack on an existing SFW HA installation. Refer to the Symantec Operations Readiness Tools (SORT) website for information on the latest agent pack availability. <https://sort.symantec.com>

Refer to the agent-specific configuration guide for more details about the application agents.

To install the server components,

See “[Installing SFW HA server components using the product installer](#)” on page 67.

To install the client components,

See “[Installing the client components using the product installer](#)” on page 73.

Installing SFW HA server components using the product installer

Use the following procedure to install SFW HA server components using the product installer.

Ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress.

During installation, some product features are installed by default and others must be selected, as follows:

- For a disaster recovery configuration, select the VCS global cluster (GCO) option.
- If you plan to use Veritas Volume Replicator (VVR) for replication, select the SFW VVR option.
- If you plan to use the Fast Failover feature, select the VCS Fast Failover option.

To install SFW HA server components using the product installer

- 1 Insert the disc containing the installation software into your system’s disk drive or download the installation software from the Symantec website. <https://fileconnect.symantec.com>
- 2 Allow the autorun feature to start the installation or double-click **Setup.exe**. The CD browser appears.

Note: If you are installing the software using the product software disc, the CD browser displays the installation options for all the products specified earlier. However, if you are downloading the installation package from the Symantec website, the CD browser displays the installation options only for the product to be installed.

3 Click to download the required contents.

Note: The client components are installed by default along with the server components. However, on a server core machine, the client components will not be installed.

Storage Foundation HA 6.0.1 for Windows	Click to install the server and client components for SFW HA.
---	---

Late Breaking News	Click to access the latest information about updates, patches, and software issues regarding this release.
--------------------	--

Windows Data Collector	Click to verify that your configuration meets all pertinent software and hardware requirements.
------------------------	---

SORT	Click to access the Symantec Operations Readiness Tools site. In addition to the product download you can also download the custom reports about your computer and Symantec enterprise products, a checklist providing configuration recommendations, and system and patch requirements to install or upgrade your software.
------	---

Browse Content	Click to view the software disc contents.
----------------	---

Technical Support	Click to contact Symantec Technical Support.
-------------------	--

- 4 On the Welcome panel, review the list of prerequisites and click **Next**.
- 5 On the License panel, read the license terms, select **I accept the terms of License Agreement**, and then click **Next**.
- 6 On the System Selection panel, select the systems and the desired Installation and Product options.
You can select the systems in one of the following ways:
- In the System Name or IP text box, manually type the system name or its IP address and click **Add**.

Note: The wizard does not support the internet protocol version 6. To add the systems having internet protocol version 6, you must type the system name.

The local host is populated by default.

- Alternatively, browse to select the systems.

The systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**.

Once you add or select a system, the wizard performs certain validation checks and notes the details in the Verification Details box. To review the details, select the desired system.

To select the installation and product options, perform the following tasks on each of the selected systems.

Note: To apply the selection to multiple systems, select the system for which you have selected the installation and product options and then click **Apply to multiple systems**.

See [“Applying the selected installation and product options to multiple systems”](#) on page 73.

- By default the wizard uses %ProgramFiles%\Veritas as the installation directory. To customize the installation directory, click **Browse** and select the desired location. Click **OK**.
- Select the required license type from the **License key** drop-down list.

Note: The default license type is “Keyless”.

If you select the "Keyless" license type, all the available product options are displayed and are selected by default.

If you select "User entered license key" as your license type, the License Details panel appears by default. On the License Details panel, enter the license key and then click **Add**. You can add multiple licenses for the various product options you want to use.

The wizard validates the entered license keys and displays the relevant error if the validation fails. After the validation is complete, click **OK**.

- From the list of product options, select the appropriate options to be installed. The options differ depending on your product and environment.

Storage
Foundation
Options

- Veritas Volume Replicator (VVR): Replicates data across multiple sites for disaster recovery.
- FlashSnap: Allows you to create and maintain split-mirror, persistent snapshots of volumes and application components. FlashSnap supports VSS based snapshots to provide application data in a consistent state after the application is restored.
- Replace Disk Management Snap-in with SFW VEA GUI: Replaces the Disk Management Snap-in in the Windows Computer Management console and the Server Manager console with the Veritas Enterprise Administrator GUI for Windows Server 2012.

- DMP Device Specific Modules
- 3PARDATA (V3PARAA)
 - Compellent array (VCOMPLNT)
 - Dell EqualLogic array (VEQLOGIC)
 - EMC Clarion (VEMCLLAR)
 - EMC Symmetrix/DMX (VEMCSYMM)
 - EMC VPLEX array (VEMCVPLX)
 - FUJITSU ETERNUS 2000 array (VFUJITSUAA)
 - Hitachi 95xx-AMS-WM (VHDSAP)
 - Hitachi TagmaStore/HP XP (VHDSAA)
 - HP 2000 array (VHPMSA2)
 - HP EVA-MSA (VHPEVA)
 - HUAWEI S5300/S2300 array (VHUAWEIAP)
 - IBM DS AP (VIBMAPDS)
 - IBM DS4000/SUN 6000 (VENGAP)
 - IBM DS6000 (VIBMAP)
 - IBM DS8000/ESS (VIBMAADS)
 - IBM XiV Storage System (VXIV)
 - NETAPP (VNETAPP)
 - NEXSAN SATA/SAS Beast, E60/E18 array (VNEXSAN)
 - PILLAR (VPILLAR)
 - Sun array (VSUN)
 - XioTech array (VXIOTECH)

Symantec maintains a Hardware Compatibility List (HCL). The HCL provides information on HBAs and firmware that have been tested with each supported array. Check the HCL for details about your hardware before installing or using DMP DSMs.

The HCL is located at:

<http://www.symantec.com/docs/TECH152806>

Note: Do not use a DMP DSM together with a third-party DSM for the same array. Only one DSM at a time can claim the LUNs in an array. According to Microsoft Multipath I/O (MPIO) documentation, if multiple DSMs are installed, the Microsoft MPIO framework contacts each DSM to determine which is appropriate to handle a device. There is no particular order in which the MPIO framework contacts the DSMs. The first DSM to claim ownership of the device is associated with that device. Other DSMs cannot claim an already claimed device. Therefore, to ensure that the DMP DSM claims the LUNs of an array, no other DSM should be installed for that same array.

Veritas Cluster
Server Options

- Global Cluster Option (GCO): Enables you to link the clusters located in different geographies. This provides wide-area failover and disaster recovery.
- Fast Failover: Improves the failover time taken by storage resources during the service group failovers, in a clustered environment. Fast failover is particularly noticeable in clusters having multiple storage stacks configured, typically over 20 disk groups and over 150 volumes.

7 On the System Selection panel, click **Next**.

Note that the wizard fails to proceed with the installation unless all the selected systems have passed the validation checks and are ready for installation. If the validation checks have failed on any of the systems, review the details and rectify the issue. Before you choose to proceed with the installation, select the system and click **Re-verify** to re-initiate the validation checks for this system.

8 On the Pre-install Summary panel, review the summary and click **Next**.

Note that the **Automatically reboot systems after installer completes operation** check box is selected by default. This will reboot all the selected remote systems immediately after the installation is complete on the respective system. If you do not want the wizard to initiate this auto reboot, clear the selection of **Automatically reboot systems after installer completes operation** check box.

9 On the Installation panel, review the progress of installation and click **Next** after the installation is complete.

If an installation is not successful on any of the systems, the status screen shows a failed installation.

10 On the Post-install Summary panel, review the installation result and click **Next**.

If the installation has failed on any of the systems, refer to the log file for details. You may have to re-install the software.

11 On the Finish panel, click **Finish**.

If you chose to initiate the auto reboot, a confirmation message to reboot the local system appears. Click **Yes** to reboot immediately or **No** to reboot later.

If you did not choose to initiate the auto reboot, ensure that you manually reboot these systems.

This completes the product installation.

Applying the selected installation and product options to multiple systems

The following procedure gives details on applying options to multiple systems.

To apply the selected installation and product options to multiple systems

- 1 Click on any one of the selected systems and select the desired installation and product options.
- 2 Click **Apply to multiple systems**.
- 3 On the Apply Installation Options panel, select the installation options to be applied and then select the desired systems. Click **OK**.

Installing the client components using the product installer

Use the following procedure to install SFW HA client components using the product installer.

Note: Client components cannot be installed on server core systems.

Before you begin the installation, ensure that there are no parallel installations, live updates, or Microsoft Windows updates in progress on the system where you want to install the client components.

To install SFW HA client components using the product installer

- 1 Insert the software disk containing the installation package into your system's disc drive or download the installation package from the following Symantec Web site.
<https://fileconnect.symantec.com>
- 2 Allow the autorun feature to start the installation or double-click Setup.exe. The CD browser appears.
- 3 Click to download the required contents.

Veritas Storage Foundation HA 6.0.1 for Windows	Click to install the server or client components for SFW HA.
Late Breaking News	Click to access the latest information about updates, patches, and software issues regarding this release.
Windows Data Collector	Click to verify that your configuration meets all pertinent software and hardware requirements.

SORT	Click to access the Symantec Operations Readiness Tools site. In addition to the product download you can also download the custom reports about your computer and Symantec enterprise products, a checklist providing configuration recommendations, and system and patch requirements to install or upgrade your software.
Browse Content	Click to view the software disc contents.
Technical Support	Click to contact Symantec Technical Support.

- 4 On the Welcome panel, review the list of prerequisites and click **Next**.
- 5 On the License Agreement panel, read the license terms, select **I accept the terms of License Agreement**, and then click **Next**.
- 6 On the System Selection panel, select the systems and the installation directory.

You can select the systems in one of the following ways:

- In the System Name or IP text box, manually type the system name or its IP address and click **Add**.

Note: The wizard does not support the internet protocol version 6. To add the systems having IPv6, you must type the system name.

Local host is populated by default.

- Alternatively, browse to select the systems.
The systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**.

Once you add or select a system, the wizard performs certain validation checks and notes the details in the Verification Details box. To review the details, select the desired system.

By default the wizard uses %ProgramFiles%\Veritas as the installation directory. To customize the installation directory, click **Browse** and select the desired location. Click **OK**.

To apply the customized directory to multiple systems, click **Apply to multiple systems**. On the Apply Installation Options panel, select the systems to apply the customized directory. Click **OK**.

- 7 On the System Selection panel, click **Next**.

Note that the wizard fails to proceed with the installation unless all the selected systems have passed the validation checks and are ready for installation. If the validation checks have failed on any of the system, review the details and rectify the issue. Before you choose to proceed with the installation select the system and click **Re-verify** to re-initiate the validation checks for this system.

- 8 On the Pre-install Summary panel, review the summary and click **Next**.
- 9 On the Installation panel, review the progress of installation and click **Next** after the installation is complete.
If an installation is not successful on any of the systems, the status screen shows a failed installation.
- 10 On the Post-install Summary panel, review the installation result and click **Next**.
If the installation has failed on any of the system, refer to the log file for details. You may have to re-install the software.
- 11 On the Finish panel, click **Finish**.
This completes the installation of the client components.

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:

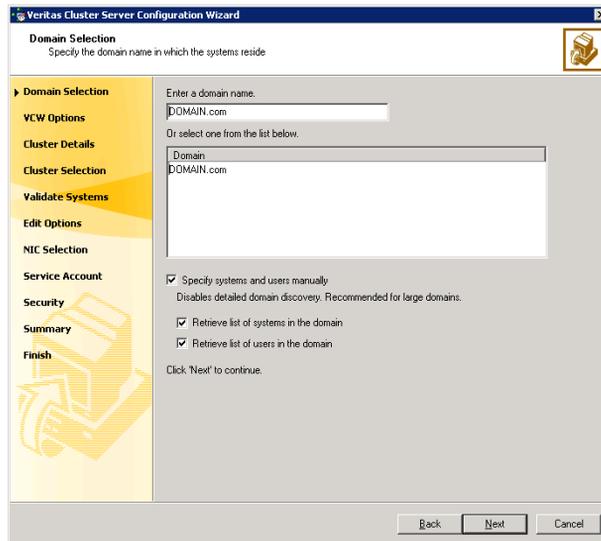
- Verify that each node uses static IP addresses and that name resolution is configured for each node.
- Verify that you have the required privileges.
See [“Reviewing the requirements”](#) on page 48.

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**.

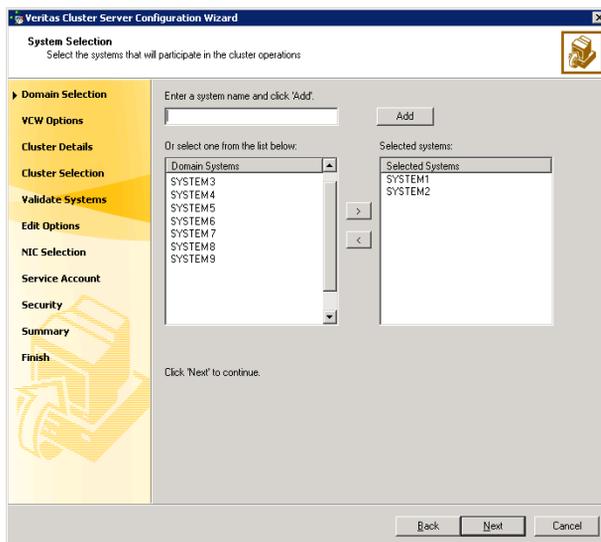
- 4 On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.



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 78.
 - 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 78. 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 78.

- 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**.

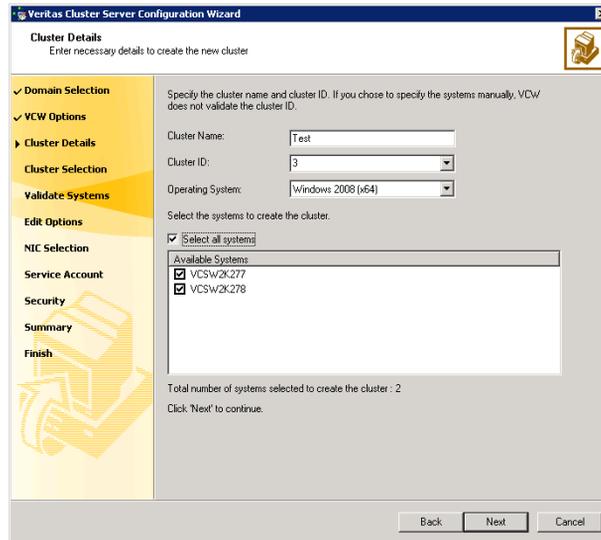
Click on a system name to see the validation details. If you wish to include a rejected system, rectify the error based on the reason for rejection and then run the wizard again.

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.

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

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



- | | |
|-------------------|---|
| Cluster Name | Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 65535.

Caution: If you chose to specify systems and users manually in step 4 or if you share a private network between more than one domain, make sure that the cluster ID is unique. |
| Operating System | From the drop-down list select the operating system. All the systems running Windows Server 2012 are displayed. |
| 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.

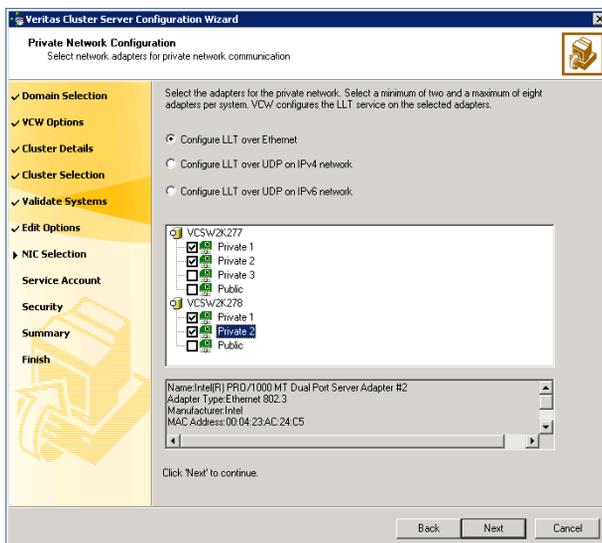
If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 82.

- 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 ethernet or over the User Datagram Protocol (UDP) layer using IPv4 or IPv6 network.

Do one of the following:

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



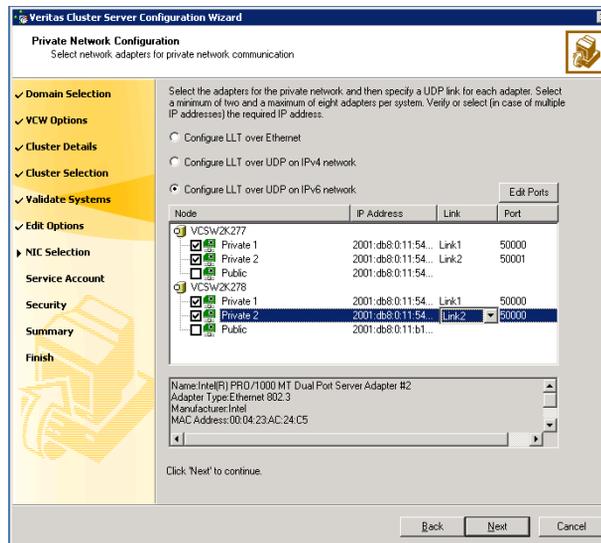
- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one of the NICs and use the low-priority NIC for both public and private communication.
- If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.

To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.

- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.

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

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



- Select **Configure LLT over UDP on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use.

The IPv6 option is disabled if the network does not support IPv6.

- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links.

Symantec recommends reserving at least two adapters exclusively for the VCS private network.

- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and

choose the desired IP address from the drop-down list. In case of IPv4, each IP address can be in a different subnet.

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

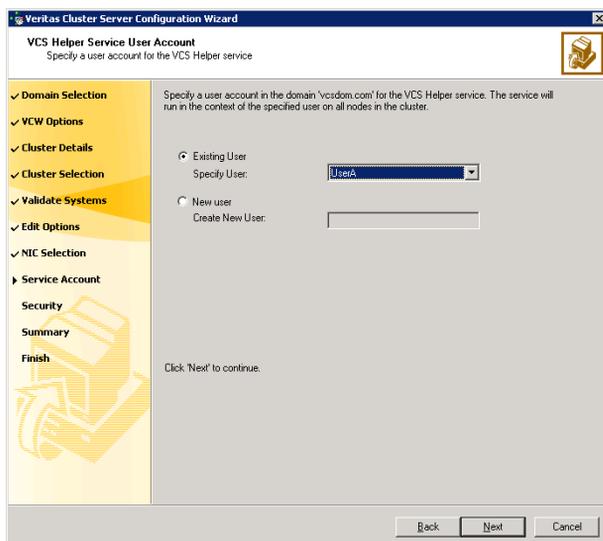
- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.

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

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

- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network.

This account does not require Domain Administrator privileges.



Specify a domain user as follows:

- To specify an existing user, do one of the following:
 - Click **Existing user** and select a user name from the drop-down list

- If you chose not to retrieve the list of users in [step 4](#) on page 77, 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 communications and then click **Next**.
Do one of the following:
- To use VCS cluster user privileges, click **Use VCS User Privileges** and then type a user name and password.
The wizard configures this user as a VCS Cluster Administrator. In this mode, communication between cluster nodes and clients, including the Cluster Manager (Java Console), occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password.
The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.
Symantec recommends that you specify a new user name and password.
 - To configure a secure cluster using the single sign-on feature, click **Use Single Sign-on**.
In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode.
The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.
- 14 Review the summary information on the Summary panel, and click **Configure**.

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

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

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

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

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

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.

- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** check box to configure notification of important events to designated recipients.
See “[Configuring notification](#)” on page 85.
 - Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.
Refer to the *Veritas Cluster Server Administrator’s Guide* for details on configuring GCO using the cluster configuration wizard.

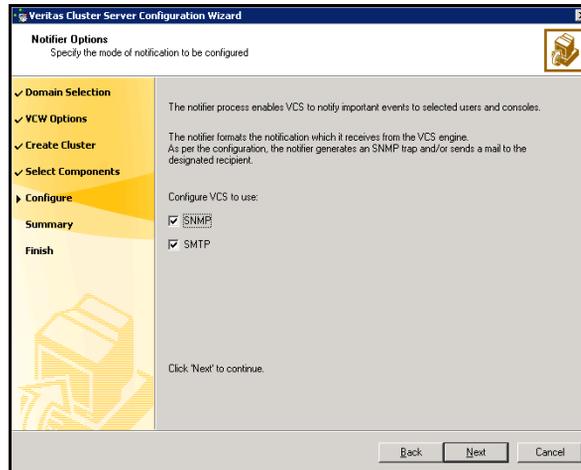
You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters discuss how to use the Disaster Recovery wizard to configure the GCO option.

Configuring notification

This section describes steps to configure notification.

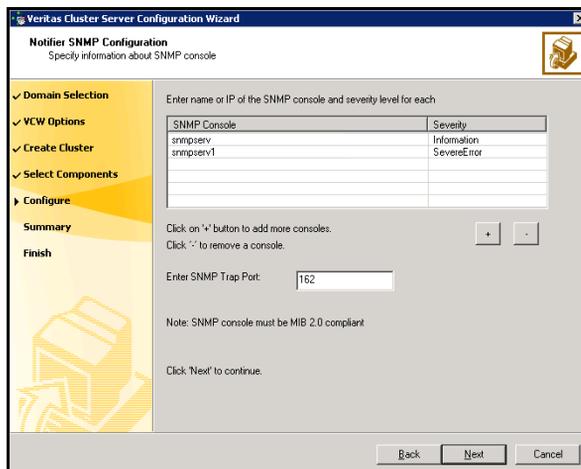
To configure notification

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

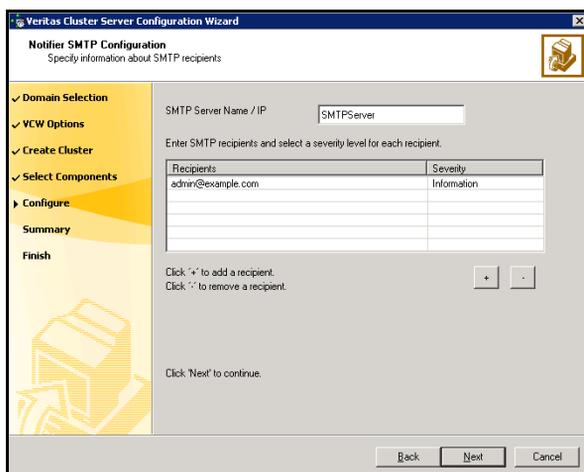


You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

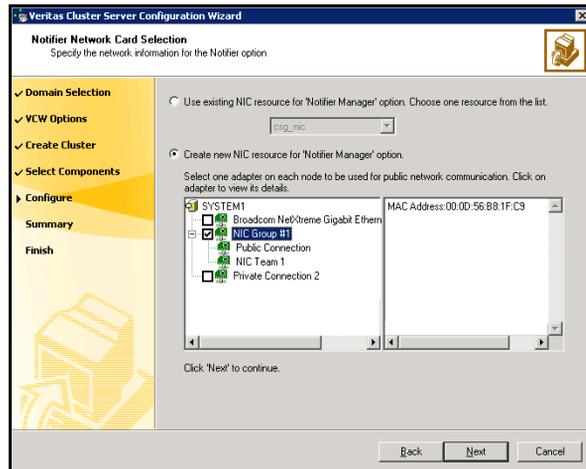


- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
 - Click the corresponding field in the Severity column and select a severity level for the console.
 - Click '+' to add a field; click '-' to remove a field.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



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

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



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

Adding nodes to an existing cluster

Optionally, you can configure both the SQL Server and SharePoint Server systems in the same SFW HA cluster if all systems use the same operating system and platform. If you have an existing SQL Server cluster and want to add the SharePoint systems to it, you can use this procedure.

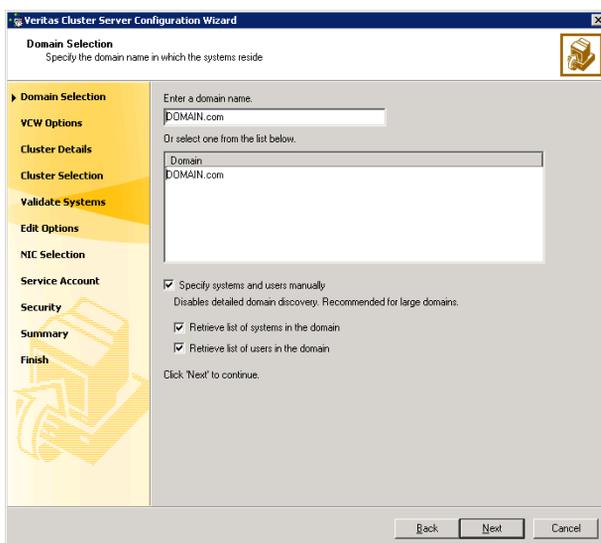
To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.

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

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

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



Do one of the following:

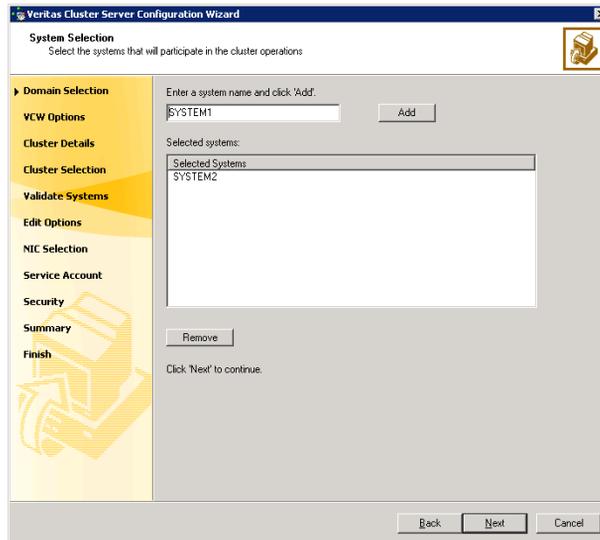
- To discover information about all the systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.

Proceed to [step 8](#) on page 91.

- 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**.

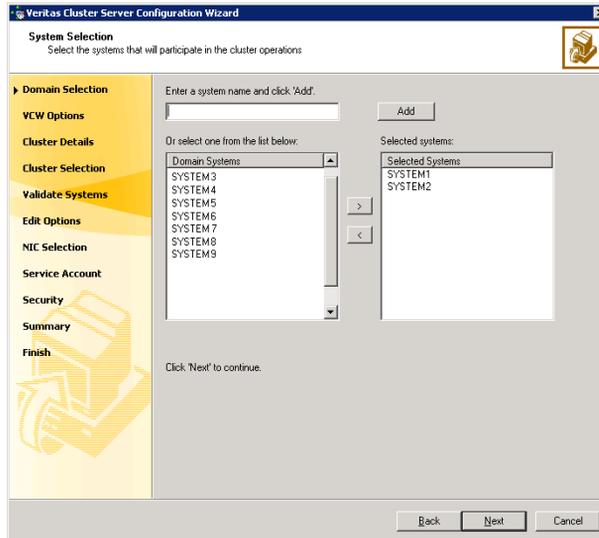
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 90. Otherwise proceed to the next step.

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



- Type the name of a node in the cluster and click **Add**.
 - Type the name of the system to be added to the cluster and click **Add**.
- If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.
- Proceed to [step 8](#) on page 91.

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



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

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

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

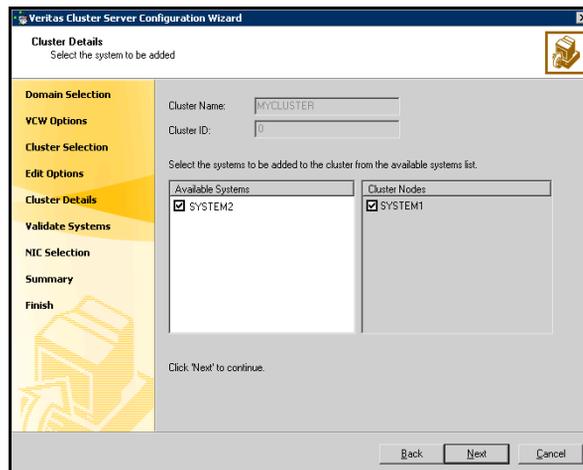
A system can be rejected for any of the following reasons:

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

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

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

The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges, that is when the cluster configuration does not use the Symantec Product Authentication Service for secure cluster communication.
- 11 On the Cluster Details panel, check the check boxes next to the systems to be added to the cluster and click **Next**.



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

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

If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.
- 13 On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**. How you configure the VCS private network communication depends on

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

Do one of the following:

- To configure the VCS private network over ethernet, complete the following steps:
 - Select the check boxes next to the two NICs to be assigned to the private network.
Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.
The wizard configures the LLT service (over ethernet) on the selected network adapters.
- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Symantec recommends reserving at least two NICs exclusively for the VCS private network.
 - Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. In case of IPv4, each IP address can be in a different subnet.

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

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

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

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

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

- 15 Specify the password for the user in whose context the VCS Helper service runs.
- 16 Review the summary information and click **Add**.
The wizard starts running commands to configure the VCS components on the node. In case of a secure cluster, the wizard also configures the VCS Authentication Service on the new node that is being added.
- 17 After all commands have been successfully run, click **Finish**.

Installing and configuring SharePoint Server 2010 for high availability

This chapter contains the following topics:

- [Installing and configuring SharePoint Server](#)
- [Configuring 64-bit Perl for SharePoint 2010](#)
- [Configuring SharePoint Server service groups](#)
- [Verifying the SharePoint cluster configuration](#)
- [Considerations when modifying a SharePoint service group](#)

Installing and configuring SharePoint Server

Install and configure Microsoft SharePoint Server on all the nodes that will be part of the SharePoint Server service group and configure the farm.

Note the following before you proceed:

- Symantec recommends that you first configure SQL Server for high availability before configuring SharePoint Server 2010.
- While installing SharePoint Server, ensure that you select **Server Farm** installation and then select **Complete** Server Type installation (Microsoft SharePoint Server 2010 installer > Server Type tab).

Note: The **Stand-alone** Server Type installation is not supported.

- VCS does not require you to install the SharePoint Server 2010 components on shared storage. You can install SharePoint to the local system disks.

- During configuration, for the database server name for the farm configuration database, specify the SQL Server that you configured for high availability earlier.

For installation and configuration instructions, see the Microsoft SharePoint Server documentation.

Note: For Perl scripts related to the SharePoint 2010 solution to work properly, use 64-bit Perl instead of the default version that is provided with the VCS installation. See [“Configuring 64-bit Perl for SharePoint 2010”](#) on page 96.

Configuring 64-bit Perl for SharePoint 2010

Perl scripts are used to update DNS entries pertaining to the Network Load Balancer (NLB) name for the SharePoint 2010 DR solution. The scripts fail to execute when VERITAS Perl is used, because it is a 32-bit version of Perl.

By default, when a 32-bit process tries to access a 64-bit component in the `C:\Windows\System32` directory, Windows File System Redirector redirects it to the `C:\Windows\SysWOW64` directory. When `C:\Windows\System32\dnscommand.exe` is used, Windows is unable to locate the 64-bit file, because `dnscommand.exe` is only installed in the `System32` directory.

To configure 64-bit Perl

- 1 Install a 64-bit version of Perl.
- 2 Copy the `ag_i18n_inc.pm` file from `%vcs_home%\VRTSPerl\lib` to `64-bit_Perl_Install\lib`.
- 3 Make sure that you add the file path of 64-bit Perl to the process resource attribute in the SQL Server service group for configuring the web servers. For further information about process resources, see [“Configuring the SQL Server service group for DR in the SharePoint environment”](#) on page 107.

Configuring SharePoint Server service groups

Configuring the SharePoint Server service group involves the following tasks:

- creating a parallel service group for the SharePoint Web Applications running on the front-end Web servers
- creating service groups for SharePoint Service Applications or services locally on the application servers

Use the VCS SharePoint Server Configuration Wizard to create the required service groups and its resources and define the attribute values for the configured resources.

Note the following before you proceed:

- The wizard discovers the Web Applications, Service Applications, and services in the farm where the local node resides and then configures them in the service groups.
- The wizard automatically configures all the discovered SharePoint applications and services configured in the local cluster farm. You cannot choose applications or services for the service group configuration. If you do not want to configure an application or a service, host it on a server outside the local cluster.
- The wizard has a single workflow that performs service group creation as well as modification tasks. If you wish to add or remove a SharePoint component from the configuration, you must run the wizard again. If you run the wizard after configuring the SharePoint service groups, the wizard modifies the existing service group configuration. The wizard rediscovers the SharePoint configuration in the farm and then adds or removes resources depending on the changes made. For example, if you add a node to the server farm, the wizard adds the required resources and service groups to the configuration. If an application is removed from the server farm, the wizard removes the corresponding resources from the service group and also updates the VCS configuration.
- If you have configured the Web Applications and Service Applications in different clusters, then you must run the configuration wizard once from a node in each cluster.
- After configuring the SharePoint service groups, you can add custom resources such as IP or NIC to monitor the network availability of the cluster nodes in the configuration. You can add these resources manually from the Cluster Manager (Java Console). If you run the wizard again, these custom resources are ignored.

Before you configure a SharePoint service group

Before you configure a SharePoint service group, do the following:

- Verify that you have configured a cluster using the VCS Cluster Configuration Wizard (VCW).
- Verify that you have installed and configured SharePoint Server 2010 on all the nodes that will be part of the SharePoint service groups.

- Ensure that the SharePoint Server 2010 Timer service is running on all the nodes that will be part of the SharePoint service groups.
- Ensure that the Veritas Command Server service is running on all the nodes that will be part of the SharePoint service groups.
- Verify that the Veritas High Availability Daemon (HAD) is running on the system from where you run the VCS SharePoint Server Configuration Wizard.
- Ensure that you have VCS Cluster Administrator privileges. This privilege is required to configure service groups.
- Ensure that the logged-on user has SharePoint Server 2010 Farm Administrator privileges on the SharePoint Server.
- Ensure that you run the wizard from a node where SharePoint Server 2010 is installed and configured.
- 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*.

Creating a SharePoint service group

Complete the following steps to create a service group for SharePoint Server.

To create the SharePoint Server service group

- 1 Launch the VCS SharePoint Server Configuration Wizard.
Click **Start > All Programs > Symantec > Veritas Cluster Server > Solutions Configuration Center** to start the Solutions Configuration Center (SCC).
Expand the Solutions for SharePoint Server tab and click **High Availability (HA) Configuration > Configure SharePoint Server Service Groups > SharePoint 2010 Configuration Wizard**.
- 2 Review the information in the Welcome panel and click **Next**.

- 3 On the Farm Admin User Details panel, specify the SharePoint Farm Admin user credentials and then click **Next**.

Farm Name	Displays the name of the farm configuration database where the node resides.
Farm Admin User Name	<p>Specify a user account that has Farm Admin privileges in the SharePoint farm where the current node resides.</p> <p>Click the ellipsis button to launch the Windows Select User dialog box and then specify the appropriate user account.</p> <p>The Farm Admin user account is used to manage the SharePoint applications and services configured in the SharePoint service groups in the cluster.</p>
Password	<p>Type the password of the user account specified in the Farm Admin User Name field.</p> <p>The wizard stores the user password in the VCS configuration in an encrypted form.</p>

- 4 On the Web Applications Details panel, review the list of Web Applications that the wizard discovers in the farm and then click **Next**.
The wizard configures these Web Applications in a parallel service group. The wizard configures only those components that are part of the local cluster.
- 5 On the Service Applications Details panel, review the list of Service Applications and services that the wizard discovers in the farm and then click **Next**.
The wizard configures these Service Applications and services in a local service group on each node. The wizard configures only those components that are part of the local cluster.

- 6 On the Service Groups Summary panel, review the service group configuration, edit the service group and resource names if required, and then click **Next**.

Resources Displays a list of configured service groups and its resources. The wizard assigns unique names to service group and resources.

- For parallel service groups, the wizard uses the following naming convention:
FarmConfigurationDatabaseName-WebApplications
- For local service groups, the wizard uses the following naming convention:
FarmConfigurationDatabasename-NodeName-ServiceApps

You can edit resource names only in the create mode. You cannot modify names of service groups and resources that already exist in the configuration.

To edit a 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.

Attributes Displays the attributes and their configured values, for a resource selected in the Resources list.

- 7 Click **Yes** on the message that informs that the wizard will run commands to modify the service group configuration. The wizard starts running commands to create the service groups. Various messages indicate the status of these commands.
- 8 On the completion panel, check **Bring the service group online** check box to bring the SharePoint service groups online in the cluster, and then click **Finish**.
This completes the SharePoint service group configuration.

About service groups for SharePoint Search

The VCS SharePoint Server Configuration Wizard no longer discovers services and service applications related to SharePoint Search. A new solution to monitor such services and service applications is provided in Agent Pack Q4 2012 for 6.0.1.

The procedure to configure a SharePoint Search service application for disaster recovery has also changed.

For information about the new solution, see the *Veritas Cluster Server Agent for Microsoft Sharepoint 2010 Search Service Application Configuration Guide*.

For information about the agent pack, see the *Veritas Cluster Server 6.0.1 Agent Pack Readme*.

If you have any SharePoint service groups configured in your setup, you need to run the SharePoint 2010 Configuration Wizard again. This removes any existing resources that might be configured for Search-related services and applications in the existing service groups.

Note: Make sure that you do this after you apply Agent Pack Q4 2012 for 6.0.1 and before you configure service groups for Search-related services or applications.

Verifying the SharePoint cluster configuration

Failover simulation is an important part of configuration testing. To verify the configuration in the cluster, you can take the service groups offline, or manually stop the configured applications on the active cluster node.

You can also simulate a local cluster failover for the SQL databases configured in the VCS SQL Server service group. Refer to the VCS SQL documentation for instructions.

Use Veritas Cluster Manager (Java Console) to perform all the service groups operations.

To take the service groups offline and bring them online

- 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 **Offline** and then choose the local system.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the node.

If there is more than one service group, you must repeat this step until all the service groups are offline.
- 2 Verify that the applications and services configured in the service groups are in the stopped state.
- 3 To start all the stopped services, bring all the services groups online on the node.

To manually stop the configured applications and services

- 1 To verify that the SharePoint applications and services are properly configured with VCS, manually stop these components either from the SharePoint Central Administration console or from the IIS Manager.
- 2 From the IIS Manager, in the Connections pane on the left, select a configured Web site and then in the Actions pane on the right, click **Stop**. The status of the Web Site will show as stopped.
- 3 In the Cluster Manager (Java Console) the corresponding service group resource state may temporarily show as faulted as the SharePoint agent attempts to start the stopped application.
- 4 When the resource comes online, refresh the IIS Manager view to verify that the IIS site is in the started state.

Considerations when modifying a SharePoint service group

Note the following while modifying SharePoint service groups:

- The wizard has a single workflow that performs service group creation as well as modification tasks. If you wish to add or remove a SharePoint component from the configuration, you must run the wizard again. If you run the wizard after configuring the SharePoint service groups, the wizard modifies the existing service group configuration. The wizard rediscovers the SharePoint configuration in the farm and then adds or removes resources depending on the changes made. For example, if you add a node to the server farm, the wizard adds the required resources and service groups to the configuration. If an application is removed from the server farm, the wizard removes the corresponding resources from the service group and also updates the VCS configuration.
- 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.
- The wizard automatically configures all the discovered SharePoint applications and services configured in the local cluster farm. You cannot choose applications or services for the service group configuration. If you do not want an application or a service to be part of the configuration, host it on a server outside the local cluster.

- When you run the wizard after configuring the SharePoint service groups, the wizard ignores any custom resources that you may have added to the service groups. If you wish to add, remove, or modify those custom resources, you must do so manually. The wizard does not provide any options to modify custom resources.
- 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 VCS SharePoint Server Configuration Wizard on this node to configure the SharePoint services for it.

Configuring disaster recovery for SharePoint Server 2010

This chapter contains the following topics:

- [Tasks for configuring disaster recovery for SharePoint Server 2010](#)
- [Configuring the SQL Server service group for DR in the SharePoint environment](#)
- [Configuring the secondary site for SharePoint disaster recovery](#)

Tasks for configuring disaster recovery for SharePoint Server 2010

After setting up an SFW HA high availability environment for a SharePoint Server 2010 farm on a primary site, you can create a secondary or “failover” site for disaster recovery.

In addition to configuring DR for the SQL Server components of the SharePoint farm, you can configure DR for SharePoint applications and services.

The following table lists the main tasks and sequence for configuring SharePoint applications and services for DR on the secondary site.

Table 8-1 Configuring the secondary site for disaster recovery

Action	Description
Configure SQL Server for disaster recovery at the secondary site	For the steps for configuring SQL Server for high availability and disaster recovery, refer to the HA and DR solutions guide for the desired SQL Server version.
Modify the SQL Server service group on the primary and secondary site	<p>Edit the SQL Server service group on both the primary and secondary site to allow updating the NLB details if a disaster recovery failover occurs.</p> <p>See “Configuring the SQL Server service group for DR in the SharePoint environment” on page 107.</p>
Verify that SharePoint has been configured for high availability at the primary site	<p>Verify that SharePoint has been configured for high availability at the primary site.</p> <p>See “Verifying the SharePoint cluster configuration” on page 101.</p>
Install SFW HA and configure the cluster on the secondary site	<p>Install SFW HA on the SharePoint server systems on the secondary site. Ensure that you select the option to install the Veritas Cluster Server Application Agent for SharePoint Server 2010.</p> <p>You can optionally use the same SFW HA cluster for both SQL Server and SharePoint Server if all systems use the same operating system platform. Otherwise, create a separate cluster for SharePoint.</p> <p>See “Configuring the secondary site for SharePoint disaster recovery” on page 115.</p>
Install SharePoint on the cluster nodes on the secondary site	<p>Install Microsoft SharePoint Server on the SharePoint servers on the secondary site. Run the Microsoft SharePoint Products Configuration wizard to add the servers to the existing primary site farm. Choose the option to connect to an existing server farm.</p> <p>Note: You do not need to configure the same number of web servers or service applications on the secondary site as on the primary site. However, you should provide for all required services.</p>

Table 8-1 Configuring the secondary site for disaster recovery (Continued)

Action	Description
Create the SharePoint service groups on the secondary site	Configure the SharePoint Server service groups for the secondary site The VCS SharePoint Server Configuration Wizard helps you create SharePoint Server service groups. See “Configuring SharePoint Server service groups” on page 96.
Verify the disaster recovery configuration	In the Veritas Cluster Server Java console, ensure that you can bring the SharePoint service groups online and offline.

Configuring the SQL Server service group for DR in the SharePoint environment

To create the VCS SQL Server service group on the primary site, follow the instructions in the SQL Server solutions guide, as follows:

- *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2005*
- *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2008 and 2008 R2*
- *Veritas Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL 2012*

After creating the SQL Server service group, you edit the default configuration of the service group to automate updating the Network Load Balancing (NLB) details when you switch between sites.

The following provide additional details:

- Edit the service group to change the attribute settings of the VCS Lanman agent resource.
See [“Updating the SQL Server IP address”](#) on page 108.
- Optionally, depending on your environment, edit the service group to add a process resource that implements a VCS script to update the NLB details of the SharePoint farm. You must customize the script configuration settings file separately for each site.
See [“Updating the IP address for web requests”](#) on page 109.

Updating the SQL Server IP address

You configure the VCS Lanman agent to update the DNS server with the virtual IP address for the SQL Server instance that is being brought online. The Lanman agent resource is created automatically as part of the SQL Server service group. However, you need to edit the default Lanman settings.

You must specify the following attribute settings for the Lanman agent, at a minimum:

DNSUpdate	True	This setting causes the update of the SQL Server IP address on the DNS server.
DNSCriticalForOnline	True	The server will not be able to come online if the DNS update is not successful.
DNSOptions	PurgeDuplicate	Removes duplicate DNS entries from the DNS servers.

More information on Lanman agent settings is provided in the agent documentation.

See *Cluster Server Bundled Agents Reference Guide*.

The procedure shows how to edit the Lanman resource of an existing SQL Server service group from the VCS Cluster Manager Java Console. You do this after you create the service group on the primary site and again on the secondary site after creating the service group there.

To configure the Lanman agent resource to update the SQL Server IP address

- 1 Start the Cluster Manager Java Console, log on to the cluster, and open the Cluster Explorer window (click anywhere in the active Cluster Monitor panel).
- 2 In the Cluster Explorer configuration tree, expand the SQL Server service group and expand **Lanman**.
- 3 Under Lanman, right-click the resource icon (labeled with the service group name and the "-Lanman" suffix) and click **View > Properties View**.
- 4 Expand the Properties View window as necessary to see all attributes under Type Specific Attributes.
- 5 Edit the following attribute settings by locating the row containing the setting, clicking the Edit icon in that row, and editing the setting as follows in the Edit Attribute dialog box. Leave Global (the default) enabled to apply

the attribute to all nodes in the cluster. If initially prompted to switch to read/write mode, click **Yes**.

DNSUpdateRequired	Check DNSUpdateRequired and click OK .
DNSCriticalForOnline	Check DNSCriticalForOnline and click OK .
DNSOptions	Under Vector Values, click the plus icon to display the list, select PurgeDuplicate and click OK .

- 6 If your site uses additional DNS servers, edit the setting for AdditionalDNSServers to specify the IP addresses.
- 7 In the Cluster Explorer window, click **File > Save Configuration**, and then click **File > Close Configuration**.
- 8 If you are configuring a resource for the web servers, continue with that procedure; otherwise, log off the cluster and exit the Cluster Manager. See [“Configuring a resource for the web servers”](#) on page 112.

Updating the IP address for web requests

You can configure VCS to update the DNS server with a site-specific IP address for the SharePoint NLB. This update occurs as part of the process of bringing the SQL Server service group online.

To automate this, you configure a VCS process resource as part of the SQL Server service group. You configure the resource after you create the service group on the primary site and you repeat the procedure on the service group that you create on the secondary site.

See [“Configuring a resource for the web servers”](#) on page 112.

The process resource uses Perl scripts. The scripts read information from a configuration settings file that you must customize separately for each site. See [“Customizing the DNS update settings for the web servers”](#) on page 110.

Requirements

The DNS update script files are available in the following directory:

```
%VCS_HOME%\bin\SQLServer2008
```

The files consist of the following:

- dnsupdate-online.pl
- dnsupdate-offline.pl
- dnsupdate-monitor.pl

■ `dnsupdate-settings.txt`

You customize the settings file for your environment. You need two copies of the settings file, one with settings for the primary site and one with settings for the secondary site.

See [“Customizing the DNS update settings for the web servers”](#) on page 110.

After customizing the settings file for each site, place the script files and the appropriate settings file for the site in a location where they are available from the cluster nodes. Since you specify the file names and locations as part of the service group process resource, you can choose the file names and locations. To avoid editing the service group again on the secondary site, you must use the same names and locations on both sites.

Warning: Do not place the settings file on a replicated volume. Otherwise, the active site’s settings file would overwrite the passive site’s settings file during replication.

In addition, the scripts require the `Dnscmd.exe` command line tool. `Dnscmd.exe` is installed as part of the Windows Server 2008 DNS Server Tools feature.

The scripts log to the engine log. The name of the log is `engine_A.txt`.

Customizing the DNS update settings for the web servers

You customize the settings file `dnsupdate-settings.txt` with the values required by the script used to update the DNS server. For each keyword (in brackets) you enter a value.

[Table 8-2](#) describes the contents of the settings file.

Table 8-2 DNS update settings file

Keyword	Value	Notes
[web alias]	The web server (or NLB) name	Same in both setting files

Table 8-2 DNS update settings file

Keyword	Value	Notes
[local ip]	Comma delimited pair of IP addresses: IP address for the web server or NLB on this site, IP address for the DNS server to be updated Example: 192.168.1.2, 192.168.10.10	When editing the primary site settings file, the local IP is that of the primary site web server or NLB. For the secondary site file, the local IP is that of the secondary site web server or NLB. If you have additional IP addresses for additional web servers or DNS servers, enter them as a comma delimited pair on separate lines.
[remote ip]	Comma delimited pair of IP addresses: IP address for the web server or NLB on the remote site, IP address of the DNS server to be updated Example: 192.168.1.1, 192.168.10.10	When editing the primary site settings file, the remote IP is that of the secondary site web server or NLB. For the secondary site file, the remote IP is that of the primary site web server or NLB. The DNS server to be updated is the one that manages the IP address for the web server or NLB. If you have additional IP addresses for additional web servers or DNS servers, enter them as a comma delimited pair on on separate lines.
[dns command]	Path to the location of DNScmd.exe Example: \\Windows\System32\dnscommand.exe	By default, on Windows Server 2008, the script will look for DNScmd.exe in \\Windows\System32\dnscommand.exe on the drive where SFW HA is installed, unless you specify another value.

Table 8-2 DNS update settings file

Keyword	Value	Notes
[domain name]	Fully qualified domain of the web server Example: symantecdomain.com	Same in both settings files
[nslookup command]	Full path for nslookup.exe Example: \\Windows\System32\nslookup.exe	By default, the script will look for nslookup.exe on the drive where SFW HA is installed in the default directory shown, unless you specify another value.

Configuring a resource for the web servers

You can add a process resource to the SQL Server service group to enable switching to the web servers at the site where the SQL Server service group is brought online. The process resource executes a Perl script to update the DNS server IP address for the web servers.

You add the process resource after you create the service group on the primary site. After you create the service group on the secondary site, you add the process resource to that service group as well.

The procedure shows how to add a resource using the Java Console. You can also use other methods, as described in the VCS documentation.

See *Veritas Cluster Server Administrator's Guide*.

Verify that the Perl executable, the scripts, and the customized settings file is available from the systems on which the service group is configured.

In addition, ensure that DNScmd.exe is installed to the same drive as the SFW HA application.

To configure a resource for the web servers

- 1 Start the Cluster Manager Java Console, log on to the cluster, and open the Cluster Explorer window (click anywhere in the active Cluster Monitor panel).
- 2 In the Cluster Explorer configuration tree, right-click the name of the SQL service group and click **Add Resource**. If prompted to switch to read-write mode, click **Yes**.
- 3 In the Add Resource dialog box, specify a name for the resource and in the Resource Type list, click **Process**.

4 Edit the following process resource attributes:

StartProgram	<p>The full path names of the following, in the order shown, separated by spaces:</p> <ul style="list-style-type: none"> ■ The Perl script executable ■ The dnsupdate-online script ■ The script settings file <p>Example:</p> <pre>c:\Program Files\Veritas\VRTSPerl\bin\perl.exe c:\bin\dnsupdate-online.pl c:\bin\dnsupdate-settings.txt</pre>
StopProgram	<p>The full path names of the following, in the order shown, separated by spaces:</p> <ul style="list-style-type: none"> ■ The Perl script executable ■ The dnsupdate-offline script ■ The script settings file <p>Example:</p> <pre>c:\Program Files\Veritas\VRTSPerl\bin\perl.exe c:\bin\dnsupdate-offline.pl c:\bin\dnsupdate-settings.txt</pre>
MonitorProgram	<p>The full path names of the following, in the order shown, separated by spaces:</p> <ul style="list-style-type: none"> ■ The Perl script executable ■ The dnsupdate-monitor script ■ The script settings file <p>Example:</p> <pre>c:\Program Files\Veritas\VRTSPerl\bin\perl.exe c:\bin\dnsupdate-monitor.pl c:\bin\dnsupdate-settings.txt</pre>
UserName	The name of the user account to run the script. The account must have access and change rights to the DNS server.
Password	The password for the user account.
Domain	The domain name for that user account.

- 5** In the Add Resource dialog box, check **Enabled** and click **OK**.
- 6** In the Resource view, right-click the process resource you just created and click **Link**.
- 7** On the Link Resources dialog box, in the list of resources, select the name of the SQL Server resource and click **OK**.
- 8** In the Cluster Explorer window, click **File > Save Configuration**, and then click **File > Close Configuration**.

Configuring the secondary site for SharePoint disaster recovery

See the following topics:

- [“Installing SFW HA and configuring the cluster on the secondary site”](#) on page 115
- [“Installing the SharePoint servers on the secondary site”](#) on page 115
- [“Configuring the SharePoint service groups on the secondary site”](#) on page 116
- [“Verifying the service group configuration”](#) on page 116

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.

- Ensure that you have configured the SharePoint Server systems for the SFW HA cluster.
See [“Configuring the storage hardware and network”](#) on page 64.
- If you have not yet done so, install SFW HA on the SharePoint Server systems. Ensure that when installing SFW HA on the SharePoint systems, you select the option to install the Veritas Cluster Server Application Agent for SharePoint Server 2010.
- If both SQL Server and SharePoint Server systems use the same operating system platform, you can optionally use the same SFW HA cluster for both. In such a case, you can add the SharePoint Server systems to the existing SQL Server cluster on the secondary site. Otherwise, create a separate cluster for the SharePoint systems on the secondary site.
 - See the following:
 - [“Configuring the cluster”](#) on page 76
 - [“Adding nodes to an existing cluster”](#) on page 87

Installing the SharePoint servers on the secondary site

When you install the SharePoint servers on the secondary site, ensure that you select the installation option that allows you to add the servers to the existing primary site farm. During configuration with the Microsoft SharePoint Products Configuration Wizard, on the Connect to a server farm panel, select the option to connect to an existing server farm.

You do not need to configure the same number of SharePoint web servers or application servers on the secondary site as on the primary site. However, you should provide for all required services to be available on the secondary site.

Configuring the SharePoint service groups on the secondary site

Run the VCS SharePoint Server Configuration Wizard from a SharePoint server system on the secondary site. Configure the SharePoint Server service groups for the secondary site using the same process as on the primary site. The SharePoint Server service groups can be online on both the primary and secondary site.

See [“Configuring SharePoint Server service groups”](#) on page 96.

Verifying the service group configuration

In the Veritas Cluster Server Java console, ensure that you can bring the SharePoint service groups online and offline.

For information on bringing service groups online and offline, see the *Veritas Cluster Server Administrator’s Guide*.

Index

A

- adding nodes to an existing cluster
 - HA 87
- agent functions
 - SharePoint Server 2010 agent 24
- agent state definition
 - SharePoint Server 2010 agent 25
- AppName attribute
 - SharePoint Server 2010 agent 28
- AppPoolMon attribute
 - SharePoint Server 2010 agent 27
- AppType attribute
 - SharePoint Server 2010 agent 26
- attributes
 - for SharePoint Server 2010 agent 26

C

- cluster
 - configure LLT over ethernet 80
 - configure LLT over UDP 81
 - configuring network and storage 64
- clusters
 - configuring the cluster 76
 - configuring the hardware and network 64
 - verifying the HA failover configuration 101
- configuration overview
 - disaster recovery 58
 - high availability 54
- configure
 - LLT over ethernet 80
 - LLT over UDP using VCW 81
- configure cluster
 - ethernet 80
 - UDP 81
- customizing settings file 110

D

- disaster recovery
 - configuring SQL Server 107
 - configuring VCS Lanman agent 108

- deployment process 105
 - SharePoint installation on secondary site 115
- disaster recovery (DR)
 - deploying for SharePoint 105
 - illustrated 16
 - typical configuration 16
- DNS configuration for DR 108
- DNS update script files 109
- DNS update settings for web servers 110
- dnsupdate-settings.txt 110

F

- FarmAdminAccount attribute
 - SharePoint Server 2010 agent 29
- FarmAdminPassword attribute
 - SharePoint Server 2010 agent 29
- functions
 - SharePoint Server 2010 agent 24

G

- GUI installation 67

H

- hardware configuration for a cluster 64
- high availability (HA)
 - defined 14
 - verifying the failover 101

I

- installing SFW HA
 - HA 66
- IP address update 108
- IPv6 support 51

L

- Lanman agent configuration for DR 108
- LLT over ethernet
 - configuring using VCW 80

LLT over UDP
 configuring using VCW 81

N

network configuration for the cluster 64
 NLB
 configuring VCS process resource for disaster recovery 109

P

permissions requirements 52
 prerequisites
 SFW HA 48
 primary host, defined 16

R

replication
 defined 15
 requirements
 permissions 52
 requirements, additional for SFW HA 52
 requirements, network 50
 requirements, system 50
 resource for updating DNS server IP address for web servers 112
 resource type
 SharePoint Server 2010 agent 25

S

script files for DNS update 109
 secondary host, defined 16
 secondary site
 SharePoint installation 115
 secure cluster 83
 Security Services
 configuring 83
 ServiceIDList attribute
 SharePoint Server 2010 agent 28
 SFW HA
 additional requirements 52
 best practices 53
 network requirements 50
 system requirements 50
 SFW HA installation 66
 SharePoint
 DR configuration overview 58
 HA configuration overview 54

installation and configuration 95
 installation on secondary site 115
 SharePoint Server 2010 agent
 attributes 26
 functions 24
 state definition 25
 type definition 25
 SharePoint Server 2010 agent attributes
 AppName 28
 AppPoolMon 27
 AppType 26
 FarmAdminAccount 29
 FarmAdminPassword 29
 ServiceIDList 28
 StopSPSService 29
 SharePoint Server Configuration Wizard 98
 SharePoint service group
 creating 98
 modifying 102
 prerequisites 97
 SharePoint web servers
 configuring IP address for disaster recovery 109
 Solutions Configuration Center
 context sensitivity 40
 overview 40
 running wizards remotely 42
 starting 40
 wizard descriptions 42
 workflow for active/active configuration 57
 SQL Server IP address update 108
 SQL Server service group
 resource for updating DNS IP address 112
 state definition
 SharePoint Server 2010 agent 25
 StopSPSService attribute
 SharePoint Server 2010 agent 29
 storage hardware configuration 64

T

type definition
 SharePoint Server 2010 agent 25

V

VCS
 configuring the cluster 76
 VCS Configuration Wizard 76
 VCS process resource for DR 109

W

web servers

 configuring process resource for DR 109

