

Storage Foundation and High Availability Solutions

7.3.1 HA and DR Solutions

Guide for Microsoft SharePoint 2010 - Windows

Last updated: 2017-11-05

Legal Notice

Copyright © 2017 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

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

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing Storage Foundation and High Availability Solutions for SharePoint 2010	8
	About clustering solutions with SFW HA	8
	About high availability	9
	How a high availability solution works	9
	About replication	10
	About disaster recovery	10
	What you can do with a disaster recovery solution	11
	Typical disaster recovery configuration	11
	About high availability support for SharePoint Server	12
	About disaster recovery support for SharePoint Server	13
	Before you begin configuring quick recovery for SharePoint Server 2010	13
	About quick recovery support for SharePoint Server	14
	About the SharePoint Search service application	14
Chapter 2	Introducing the VCS agent for SharePoint Server 2010	16
	About the VCS agent for Microsoft SharePoint Server 2010	16
	SharePoint Server agent functions	17
	SharePoint Server agent state definitions	17
	SharePoint Server agent resource type definition	18
	SharePoint Server agent attribute definitions	18
Chapter 3	Configuration workflows for SharePoint Server 2010	22
	Reviewing the HA configuration	22
	Sample SharePoint Server HA configuration	24
	Following the HA workflow in the Solutions Configuration Center	25
	Reviewing the disaster recovery configuration	25
	High availability (HA) configuration	27
	Disaster recovery configuration	29

	Notes and recommendations for cluster and application configuration	31
	IPv6 support	33
	Configuring the storage hardware and network	34
	Configuring the cluster using the Cluster Configuration Wizard	35
	Configuring notification	44
	Adding nodes to a cluster	47
Chapter 4	Using the Solutions Configuration Center	52
	About the Solutions Configuration Center	52
	Starting the Solutions Configuration Center	53
	Options in the Solutions Configuration Center	53
	About launching wizards from the Solutions Configuration Center	54
	Remote and local access to Solutions wizards	55
	Solutions wizards and logs	56
	Workflows in the Solutions Configuration Center	57
Chapter 5	Installing and configuring SharePoint Server 2010 for high availability	58
	About installing and configuring SharePoint	58
	Configuring 64-bit Perl for SharePoint	59
	About configuring SharePoint service groups	59
	Before you configure a SharePoint service group	60
	Creating a SharePoint service group	61
	About service groups for SharePoint Search	63
	Verifying the SharePoint cluster configuration	64
	Considerations when modifying a SharePoint service group	65
Chapter 6	Configuring disaster recovery for SharePoint Server 2010	67
	Tasks for configuring disaster recovery for SharePoint Server	67
	Configuring the SQL Server service group for DR in the SharePoint environment	69
	Updating the SQL Server IP address	69
	Updating the IP address for web requests	71
	Configuring the secondary site for SharePoint disaster recovery	75
	Installing InfoScale Enterprise and configuring the cluster on the secondary site	76
	About installing the SharePoint servers on the secondary site	76

	About configuring the SharePoint service groups on the secondary site	77
	Verifying the service group configuration	77
Chapter 7	Introducing the VCS agent for SharePoint Search Service Application	78
	About the VCS agent for SharePoint Search service application	78
	How the VCS agent makes SharePoint Search service application highly available	79
	VCS agent for SharePoint Search service application - functions	79
	VCS agent for SharePoint Search service application - state definitions	79
	Resource type definition	80
	Attribute definitions	80
	Sample configuration file	81
	Configuring the SharePoint Search Service Application service group	86
	Prerequisites for configuring a service group for a SharePoint Search service application	86
	Installing and configuring SharePoint Server 2010	86
	Changing the index location of the Crawl and Query components	87
	Configuring a service group for a SharePoint Search service application manually	90
	Configuring the service group for a Search service application using the wizard	98
	Verifying the application service group	100
	Configuring a Search service application for disaster recovery	102
	Administering the SharePoint Search Service Application service group	103
	About administering the application service group	103
	Modifying the application service group	103
	Deleting the application service group	104
Chapter 8	Troubleshooting	105
	About troubleshooting VCS agents	105
	Troubleshooting issues with SharePoint Search service application components	105
	Restoring the Crawl or Query component registry keys	106
	VCS logging	108

VCS Cluster Configuration Wizard (VCW) logs	109
Agent error messages and descriptions	109
VCS agent for SharePoint Search service application	110

Appendix A Using Veritas AppProtect for vSphere 111

About Just In Time Availability	112
Prerequisites	116
Setting up a plan	118
Deleting a plan	120
Managing a plan	120
Viewing the history tab	122
Limitations of Just In Time Availability	122
Getting started with Just In Time Availability	123
Supported operating systems and configurations	125
Viewing the properties	126
Log files	126
Plan states	127
Troubleshooting Just In Time Availability	129

Introducing Storage Foundation and High Availability Solutions for SharePoint 2010

This chapter includes 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](#)
- [About the SharePoint Search service application](#)

About clustering solutions with SFW HA

Storage Foundation and High Availability Solutions (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 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 Storage Foundation HA for Windows (SFW HA). SFW HA includes Storage Foundation for Windows and Cluster Server.

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

How a high availability solution works

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

Using 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 secondary site, and the application can be restarted at the secondary site.

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

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

About disaster recovery

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

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

The SharePoint high availability configuration components are as follows:

- VCS provides an agent for SharePoint that performs the task of managing the SharePoint 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 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 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 agents for 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

Disaster recovery (DR) support for SharePoint Server 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, the SharePoint servers at the primary site and the secondary site must belong to the same SharePoint farm.

Before you begin configuring quick recovery for SharePoint Server 2010

Review the following prerequisites:

- Ensure that your system hardware and software meets the requirements.

- Set up SQL Server for use with InfoScale Storage.
- Ensure that the SQL Server VSS Writer service is started.
- Ensure that you have disks with enough space to store the snapshot volumes. Each snapshot set requires the same amount of space as the original volumes.
- The vxsnap commands must be invoked on a local system. On Windows Server systems, all CLI commands must run in the command window in the **Run as Administrator** mode.

Refer to the SFW quick recovery solutions guides and the *SFW Administrator's Guide* for more details.

About quick recovery support for SharePoint Server

Quick recovery (QR) solution for SharePoint Server 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 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.

About the SharePoint Search service application

A SharePoint Search service application is used to create indexes and to service search requests. The application crawls the contents of websites and creates an index, which is used to serve the client requests for search.

A Search service application contains the following components:

- Admin
This component stores the configuration data and security descriptors. Each Search service application instance can only have one Admin component.
- Crawl (Crawl Server or Indexer)

This component crawls (accesses and catalogs) the associated website contents and creates the index. It propagates the index files to Query Server. After propagating all the index files, they are removed the Crawler.

- **Query (Query Server)**
This component is are responsible for serving search queries. A Query Server is a server that runs one or more Query components. Query Servers store the full or partial search index.

Introducing the VCS agent for SharePoint Server 2010

This chapter includes the following topics:

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

About the VCS agent for Microsoft SharePoint Server 2010

The VCS application agent for Microsoft SharePoint Server manages SharePoint Server Service Applications, Web Applications, and services in a VCS cluster. The agent provides monitoring support in making a SharePoint Server 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 restart the component on the local node. If the component fails to start, the agent declares the resource as faulted.

SharePoint Server 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 <code>FAULTED</code> .
Clean	Forcibly stops the configured Web Applications and Service Applications. The agent also stops monitoring the configured services on the node.

SharePoint Server 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 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 agent is represented by the SharePointServer resource type.

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

SharePoint Server 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 agent required attributes

Required Attributes	Definition
AppType	<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> <p>Type and Dimension: string-scalar</p>
AppPoolMon	<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 restart 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 restart the application pool. <p>The default value is NONE.</p> <p>Type and Dimension: string-scalar</p>

Table 2-1 SharePoint Server agent required attributes *(continued)*

Required Attributes	Definition
ServiceIDList	<p>Defines the service IDs of the SharePoint services that are managed by the agent. This attribute is always local.</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> <p>Type and Dimension: string-vector</p>

Table 2-2 SharePoint Server agent optional attributes

Optional Attribute	Definition
AppName	<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> <p>Type and Dimension: string-scalar</p>

Table 2-2 SharePoint Server agent optional attributes (*continued*)

Optional Attribute	Definition
FarmAdminAccount	<p>A user account that has the SharePoint Server Farm Admin privileges.</p> <p>User name can be of the form username@domain.com, domain\username, or domain.com\username.</p> <p>The agent uses the Farm Admin user account context to manage the services specified in the ServiceIDList attribute value.</p> <p>Type and Dimension: string-scalar</p>
FarmAdminPassword	<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> <p>Type and Dimension: string-scalar</p>
StopSPSService	<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> <p>Type and Dimension: boolean</p>

Configuration workflows for SharePoint Server 2010

This chapter includes the following topics:

- [Reviewing the HA configuration](#)
- [Reviewing the disaster recovery configuration](#)
- [High availability \(HA\) configuration](#)
- [Disaster recovery configuration](#)
- [Notes and recommendations for cluster and application configuration](#)
- [Configuring the storage hardware and network](#)
- [Configuring the cluster using the Cluster Configuration Wizard](#)

Reviewing the HA configuration

Veritas 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.

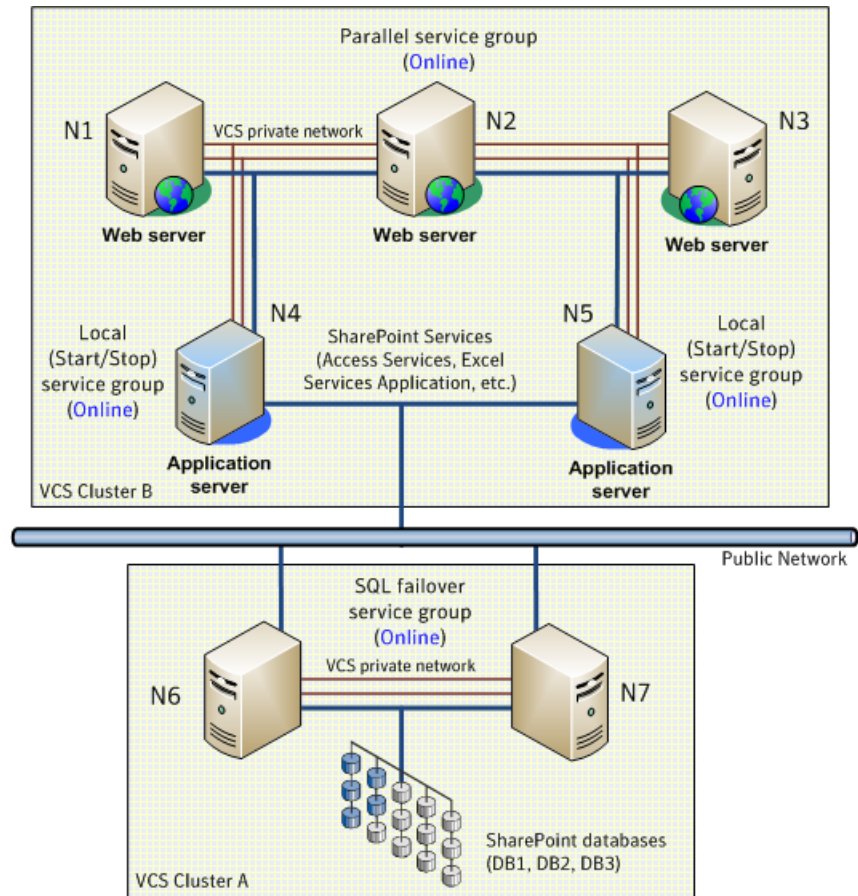
In a typical example of a SharePoint Server high availability environment, SharePoint Web Applications and Service Applications are configured on a separate set of cluster nodes. A VCS parallel service group manages the Web Applications residing on the Web servers and local service groups manage the application servers. The

SharePoint databases are made highly available using the SQL Server service group. The databases reside on shared storage that is accessible from all the SharePoint server nodes in the cluster.

The following figure illustrates a typical SharePoint Server 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 3-1 Typical SharePoint Server 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 HA configuration

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

The following table 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 3-1 Sample SharePoint Server 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 SharePoint_Config-N5-ServiceApp2	Names of the local service groups configured for the SharePoint Service Applications or services.
INST1	SQL Server instance name

Table 3-1 Sample SharePoint Server HA configuration objects (*continued*)


Name	Object
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 configuring a new Storage Foundation HA environment for SharePoint Server.

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

Figure 3-2 Configuration steps in the Solution 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.

The following figure shows an example SharePoint Server disaster recovery configuration.

Figure 3-3 Example SharePoint Server disaster recovery configuration

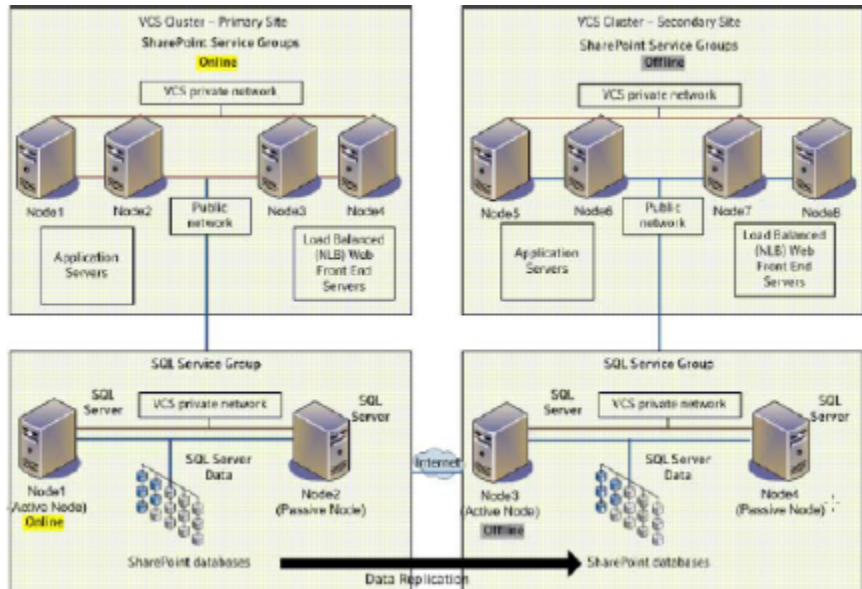


Table 3-2 Sample Disaster Recovery configuration objects

Object Name	Description
Primary site	
SYSTEM1 & SYSTEM2	first and second nodes of the primary site
CLUS1	separate SharePoint cluster, if not using the SQL Server cluster
SP_SG	SharePoint service group
Secondary site	
SYSTEM3 & SYSTEM4	First and second nodes of the secondary site
CLUS1	separate SharePoint cluster, if not using the SQL Server cluster
SP_SG	SharePoint service group

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.

High availability (HA) configuration

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

Note: Veritas 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.

Table 3-3 SharePoint Server: High availability configuration tasks

Action	Description
Verify hardware and software requirements	
Review the HA configuration	<ul style="list-style-type: none"> ■ Understand active-passive configuration. ■ Review the sample configuration. <p>See “Reviewing the HA configuration” on page 22.</p>

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

Action	Description
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. <p>See “Configuring the storage hardware and network” on page 34.</p>
Install the product	<p>Install InfoScale Enterprise.</p> <p>See <i>Veritas InfoScale Installation and Upgrade Guide</i></p>
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 using the Cluster Configuration Wizard” on page 35.</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.
Install SharePoint Server on the cluster nodes	<p>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> <p>Refer to the SharePoint Server documentation for installation instructions</p>

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

Action	Description
Create SharePoint Server service groups	Launch the VCS SharePoint Server Configuration Wizard on a node on which SharePoint is installed and configured to create SharePoint service groups. See “About configuring SharePoint service groups” on page 59.
Verify the HA configuration	Test failover between nodes. See “Verifying the SharePoint cluster configuration” on page 64.

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 27.

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.

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

Table 3-4 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 69.

Table 3-4 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
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 64.</p>
Install InfoScale Enterprise and configure the cluster on the secondary site	<p>Install InfoScale Enterprise on the SharePoint server systems on the secondary site. Ensure that you select the option to install the Cluster Server Application Agent for SharePoint Server 2010.</p> <p>See “Configuring the secondary site for SharePoint disaster recovery” on page 75.</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>
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 “About configuring SharePoint service groups” on page 59.</p>
Verify the disaster recovery configuration	<p>In the Cluster Server Java console, ensure that you can bring the SharePoint service groups online and offline.</p>

Notes and recommendations for cluster and application configuration

- Review the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:
<https://sort.veritas.com/documents>

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

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

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

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

- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.

If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA).

See the *Storage Foundation Administrator's Guide* for more information.

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Veritas recommends three NICs.
- NIC teaming is not supported for the VCS private network.
- Static IP addresses are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated.

Static IP addresses are required for the following purposes:

- A minimum of one static IP address for each physical node in the cluster.

- One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.
- For Volume Replicator replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For Volume Replicator replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in SFW HA because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
See the *Cluster Server Bundled Agents Reference Guide*.
- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.
- If Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.
- If User Access Control (UAC) is enabled on Windows systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, C:\Program Files\Veritas). As a workaround, an OS administrator user can

set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.

- For a Replicated Data Cluster, install only in a single domain.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- This is applicable for a Replicated Data Cluster configuration.
 This is applicable for a Replicated Data Cluster configuration. You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command.
- To configure a RDC cluster, you need to create virtual IP addresses for the following:
 - Application virtual server; this IP address should be the same on all nodes at the primary and secondary zones
 - Replication IP address for the primary zone
 - Replication IP address for the secondary zone
 Before you start deploying your environment, you should have these IP addresses available.

IPv6 support

For IPv6 networks, the following is supported:

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

VCS agents, wizards, and other components

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

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

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

Configuring the storage hardware and network

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

To configure the hardware

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

To verify the DNS settings and binding order

- 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.

Configuring the cluster using the Cluster Configuration Wizard

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

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

Note the following prerequisites before you proceed:

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

To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Veritas recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Veritas recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.

- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

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

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list. When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.
- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.
- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.
 Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.
- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

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

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

- Clear **Specify systems and users manually**.
- Click **Next**.
 Proceed to step [8](#).

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

- Select **Specify systems and users manually**.
 Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
 If you chose to retrieve the list of systems, proceed to step [6](#). Otherwise, proceed to the next step.

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

Do not specify systems that are part of another cluster.

Proceed to step [8](#).

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

Do not select systems that are part of another cluster.

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

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

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

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

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

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

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

Specify the cluster details as follows:

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

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

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

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

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


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

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

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

Do one of the following:

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



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

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

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

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

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

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

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

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

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

Specify the domain user details as follows:

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

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

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

Do one of the following:

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

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

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

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

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

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

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

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

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
See [“Configuring notification”](#) on page 44.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.
You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.

Configuring notification

This section describes steps to configure notification.

To configure notification

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

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

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

Veritas Cluster Server Configuration Wizard

Notifier SNMP Configuration
Specify information about SNMP console.

Domain Selection

Create Cluster

Select Components

Configure

Summary

Finish

Enter the name or the IP address of the SNMP console and then select the desired severity level.

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

Click on '+' button to add more consoles. + -

Click '-' to remove a console.

SNMP Trap Port:

Note: SNMP console must be MIB 2.0 compliant.

Click 'Next' to continue.

Back Next Cancel

Do the following:

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

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

Veritas Cluster Server Configuration Wizard

Notifier SMTP Configuration
Specify information about SMTP recipients.

Domain Selection

Create Cluster

Select Components

Configure

Summary

Finish

SMTP Server Name / IP

Enter SMTP recipients and select a severity level for each recipient.

Recipients	Severity
Click here to change the text..	Information

Click '+' to add a recipient.
Click '-' to remove a recipient.

Click 'Next' to continue.

VERITAS

Back Next Cancel

Do the following:

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

Do the following:

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

Adding nodes to a cluster

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.

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

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

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

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

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

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

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

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

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

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

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

Proceed to step 8.

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

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

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

- 7 The System Report panel displays the validation status, whether Accepted or Rejected, of all the systems you specified earlier.

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

- The system does not respond to a ping request.
- WMI access is disabled on the system.
- The wizard is unable to retrieve information about the system's architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.

- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**.

If you chose to specify the systems manually in step 4, only the clusters configured with the specified systems are displayed.

- 10 On the Edit Cluster Options panel, click **Add Nodes** and click **Next**.

In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges (a cluster that is not a secure cluster).

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

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

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

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

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

Do one of the following:

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

Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.

To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs

together. All NICs in a team have an identical MAC address. Veritas recommends that you do not select teamed NICs for the private network. The wizard configures the LLT service (over Ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, do the following:
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Veritas recommends reserving at least two NICs exclusively for the VCS private network. You could lower the priority of one NIC and use the low-priority NIC for both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. In case of IPv4, each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
 - For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

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

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

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

Using the Solutions Configuration Center

This chapter includes the following topics:

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

About the Solutions Configuration Center

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

The Solutions Configuration Center provides solutions for the following applications:

- Microsoft Exchange Server 2010
- Microsoft SQL Server 2012, 2014, and 2016
- Enterprise Vault Server (high availability and disaster recovery solutions)
- Microsoft SharePoint Server 2010 and 2013 (high availability, disaster recovery, and Quick Recovery solutions)
- Additional applications

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

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

Starting the Solutions Configuration Center

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

To start the Solutions Configuration Center

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

or

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

or

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

or

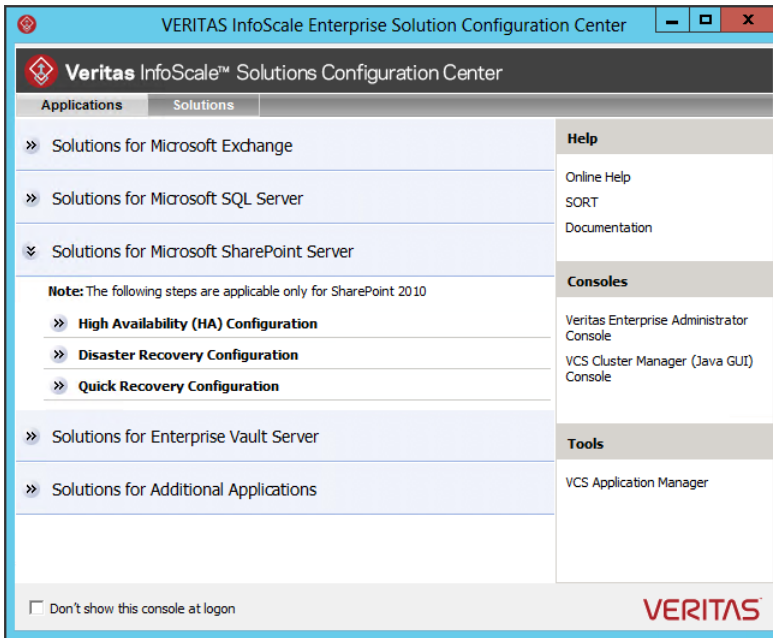
Navigate to the Apps menu and then click **SCC**.

Options in the Solutions Configuration Center

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

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

Figure 4-1 Solutions Configuration Center for SharePoint Server 2010



About launching wizards from the Solutions Configuration Center

The Solutions Configuration Center provides two ways to access wizards:

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

Solutions	<p>(For advanced users) Lists wizards by solution, without additional instructions, as follows:</p> <ul style="list-style-type: none"> ■ High Availability Configuration Wizards ■ Disaster Recovery Configuration Wizards ■ Quick Recovery Configuration Wizards ■ Fire Drill Configuration Wizards <p>You can go directly to a particular wizard.</p> <p>Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.</p> <p>The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems.</p> <p>Refer to the following Microsoft knowledge database article for more details:</p> <p>http://technet.microsoft.com/en-us/library/dd184075.aspx</p>
-----------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Remote and local access to Solutions wizards

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

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

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

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

New Dynamic Disk Group Wizard	Launched from the Veritas Enterprise Administrator console
New Volume Wizard	Launched from the Veritas Enterprise Administrator console
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.
MSMQ Configuration Wizard	Configures a Microsoft Message Queuing (MSMQ) service group
SFW Configuration Utility for Hyper-V Live Migration Support	Configures SFW for Microsoft Hyper-V Live Migration support on the selected systems

Solutions wizards and logs

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

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

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

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

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

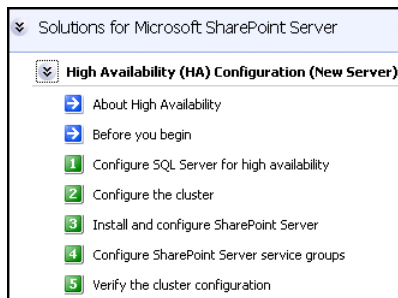
Workflows in the Solutions Configuration Center

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

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

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

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



Installing and configuring SharePoint Server 2010 for high availability

This chapter includes the following topics:

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

About installing and configuring SharePoint

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

Note the following before you proceed:

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

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

- VCS does not require you to install the SharePoint components on shared storage. You can install SharePoint on 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 documentation.

Note: For Perl scripts related to the SharePoint solution to work properly, use 64-bit Perl instead of the default version that is provided with the product installation.

See [“Configuring 64-bit Perl for SharePoint ”](#) on page 59.

Configuring 64-bit Perl for SharePoint

Perl scripts are used to update DNS entries pertaining to the Network Load Balancer (NLB) name for the SharePoint 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_dir\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 69.

About configuring SharePoint service groups

Configuring the SharePoint 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 SharePoint 2010 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 on all the nodes that will be part of the SharePoint service groups.

- Ensure that the SharePoint Server 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 Engine (HAD) is running on the system from where you run the VCS SharePoint Server 2010 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 Farm Administrator privileges on the SharePoint Server.
- Ensure that you run the wizard from a node where SharePoint Server is installed and configured.
- If you have configured a firewall, add the required ports and services to the Firewall Exception list.

For a detailed list of services and ports used by the product, refer to the *Veritas InfoScale 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 2010 Configuration Wizard.
 Launch Solutions Configuration Center from **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
 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 nodes reside.
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	<p>Displays a list of configured service groups and its resources. The wizard assigns unique names to service group and resources.</p> <ul style="list-style-type: none"> ■ 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 <p>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.</p> <p>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.</p>
Attributes	<p>Displays the attributes and their configured values, for a resource selected in the Resources list.</p>

- 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 Cluster Server Agent for Microsoft Sharepoint 2010 Search Service Application Configuration Guide.

For information about the agent pack, see the 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 Server databases configured in the SQL Server service group. Refer to the application-specific 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 includes the following topics:

- [Tasks for configuring disaster recovery for SharePoint Server](#)
- [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

After setting up an SFW HA high availability environment for a SharePoint Server 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 6-1 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 69.</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 64.</p>
Install the product and configure the cluster on the secondary site	<p>Install InfoScale Enterprise on the SharePoint server systems on the secondary site. Ensure that you select the option to install the 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 75.</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 6-1 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
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 “About configuring SharePoint service groups” on page 59.</p>
Verify the disaster recovery configuration	In the 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 SQL Server service group on the primary site, follow the instructions in the *Storage Foundation and High Availability Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL Server*.

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 69.
- 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 71.

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 73.

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 73.

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 72.

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 72.

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 6-2 DNS update settings file

Keyword	Value	Notes
[web alias]	The web server (or NLB) name	Same in both setting files
[local ip]	<p>Comma delimited pair of IP addresses:</p> <p>IP address for the web server or NLB on this site, IP address for the DNS server to be updated</p> <p>Example: 192.168.1.2, 192.168.10.10</p>	<p>When editing the primary site settings file, the local IP is that of the primary site web server or NLB.</p> <p>For the secondary site file, the local IP is that of the secondary site web server or NLB.</p> <p>If you have additional IP addresses for additional web servers or DNS servers, enter them as a comma delimited pair on separate lines.</p>

Table 6-2 DNS update settings file (*continued*)

Keyword	Value	Notes
[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\\dnscmd.exe	By default, on Windows Server 2008, the script will look for DNScmd.exe in \\Windows\\System32\\dnscmd.exe on the drive where the product is installed, unless you specify another value.
[domain name]	Fully qualified domain of the web server Example: veritasdomain.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 the product 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 *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

The full path names of the following, in the order shown, separated by spaces:

- The Perl script executable
- The dnsupdate-online script
- The script settings file

Example:

c:\Program Files\Veritas\VRTSPerl\bin\perl.exe

c:\bin\dnsupdate-online.pl c:\bin\dnsupdate-settings.txt

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> <p>c:\Program Files\Veritas\VRTSPer\bin\perl.exe c:\bin\dnsupdate-offline.pl c:\bin\dnsupdate-settings.txt</p>
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> <p>c:\Program Files\Veritas\VRTSPer\bin\perl.exe c:\bin\dnsupdate-monitor.pl c:\bin\dnsupdate-settings.txt</p>
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:

- See [“Installing InfoScale Enterprise and configuring the cluster on the secondary site”](#) on page 76.
- See [“About installing the SharePoint servers on the secondary site”](#) on page 76.
- See [“About configuring the SharePoint service groups on the secondary site”](#) on page 77.
- See [“Verifying the service group configuration”](#) on page 77.

Installing InfoScale Enterprise and configuring the cluster on the secondary site

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

- Ensure that you have configured the SharePoint Server systems for the SFW HA cluster.
 See [“Configuring the storage hardware and network”](#) on page 34.
- If you have not yet done so, install InfoScale Enterprise on the SharePoint Server systems. Ensure that when installing InfoScale Enterprise on the SharePoint systems, you select the option to install the 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:
 See [“Configuring the cluster using the Cluster Configuration Wizard”](#) on page 35.

About 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.

About configuring the SharePoint service groups on the secondary site

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

See [“About configuring SharePoint service groups”](#) on page 59.

Verifying the service group configuration

In the 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 Cluster Server Administrator's Guide.

Introducing the VCS agent for SharePoint Search Service Application

This chapter includes the following topics:

- [About the VCS agent for SharePoint Search service application](#)
- [Configuring the SharePoint Search Service Application service group](#)
- [Administering the SharePoint Search Service Application service group](#)

About the VCS agent for SharePoint Search service application

VCS application agents monitor specific resources within an enterprise application. The VCS agent for Microsoft SharePoint 2010 Search service application provides high availability for the Search service application in a VCS cluster.

The VCS agent monitors the 'SharePoint Server Search 14' service and the following application components:

- Admin
- Crawl
- Query

How the VCS agent makes SharePoint Search service application highly available

When the agent detects a system failure, the application service group switches to the next available system listed in the service group's SystemList attribute. VCS starts the configured application on the new system and brings the components online, thus ensuring high availability. For example, if the node on which the Search service application Admin component is running faults, the agent brings the Admin component online on the failover node.

If a configured component fails, the agent reports its status as UNKNOWN. An administrator can then intervene and troubleshoot the problem with the component.

Note: A Search service application with FAST Search is not supported.

VCS agent for SharePoint Search service application - functions

Agent functions include the following:

Online	Starts the configured application and its components.
Monitor	Verifies the status of the configured application components, and accordingly reports the status of the resources as follows: <ul style="list-style-type: none">■ If a component is Ready, the agent reports ONLINE.■ If a component is Disabled, the agent reports UNKNOWN.
Offline	Stops monitoring the Search service application and its components, but does not take the application offline.

VCS agent for SharePoint Search service application - state definitions

Agent state definitions are as follows:

ONLINE	Indicates that the configured application components are running on the cluster node.
OFFLINE	Indicates that the agent is not currently monitoring the application components.
FAULTED	Indicates that the agent is unable to start the configured application components on the cluster node.

UNKNOWN	Indicates that the agent is unable to determine the status of the configured application components on the cluster node. This state is also reported when the component's status is anything other than Enabled or Ready.
---------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Resource type definition

The resource definition is as follows:

```
type SharePointSearch (
    static i18nstr ArgList[] = { SPSFarmAdmin, Domain, Password,
                                SearchCompType, SearchAppName, ComponentID }
    i18nstr SPSFarmAdmin
    i18nstr Domain
    str Password
    str SearchCompType
    i18nstr SearchAppName
    str ComponentID
)
```

Attribute definitions

The following table describes the required attributes of the VCS agent for SharePoint Search service application.

Table 7-1 VCS agent for SharePoint Search service application - required attributes

Required attribute	Definition
SPSFarmAdmin	<p>A user account that has the SharePoint Server Farm Admin privileges.</p> <p>The user name can take one of the following forms:</p> <ul style="list-style-type: none"> username@domain.com domain\username domain.com\username <p>The agent uses the Farm Admin user account context to monitor the Search service application components.</p> <p>Type and dimension: string-scalar</p>
Domain	<p>The name of the domain to which the user specified in the SPSFarmAdmin attribute value belongs.</p> <p>Type and dimension: string-scalar</p>

Table 7-1 VCS agent for SharePoint Search service application - required attributes (*continued*)

Required attribute	Definition
Password	<p>The password of the user specified in the SPSFarmAdmin attribute value.</p> <p>The password is stored in the VCS configuration in an encrypted form.</p> <p>Type and dimension: string-scalar</p>
SearchCompType	<p>Defines whether the agent is configured to monitor the SharePoint Server Search 14 service or the Search service application components.</p> <p>This attribute can take one of the following values:</p> <ul style="list-style-type: none"> ■ SearchService: If you specify this value, you do not need to provide values for any other type-specific attribute. ■ Admin: If you specify this value, you need to provide values for all the other type-specific attributes, except ComponentID. ■ Crawl: If you specify this value, you need to provide values for all the other type-specific attributes. ■ Query: If you specify this value, you need to provide values for all the other type-specific attributes. <p>Type and dimension: string-scalar</p>
SearchAppName	<p>The name of the SharePoint Search service application that is managed by the agent.</p> <p>This attribute is ignored if the SearchCompType attribute value is set to SearchService.</p> <p>Type and dimension: string-scalar</p>
ComponentID	<p>An identifier that uniquely identifies the application component to be monitored.</p> <p>Type and dimension: string-scalar</p>

Sample configuration file

```
include "types.cf"

cluster SPSPri (
    SecureClus = 1
)
```

```

system SPS_R2_N1 (
)

system SPS_R2_N2 (
)

system SPS_R2_N3 (
)

group SPSPRI (
    SystemList = { SPS_R2_N1 = 0, SPS_R2_N2 = 1, SPS_R2_N3 = 2 }
)

FileShare SharePointSearch_VM_Query_Share (
    PathName = "\\QueryIndexComp"
    ShareName = cce65f70-1747-42b7-b877-bb82db4b0a68-query-1
    MountResName = SharePointSearch_VM_Query_MountV
    UserPermissions = { WSS_WPG = FULL_CONTROL }
    ShareComment = "Used by Microsoft Search Server 2010
to copy index files between servers"
)

MountV SharePointSearch_VM_Crawl_MountV (
    MountPath = "S:"
    VolumeName = CrawlVol
    VMDGResName = SharePointSearch_VM_Crawl_VMDg
    ForceUnmount = ALL
)

MountV SharePointSearch_VM_Query_MountV (
    MountPath = "Q:"
    VolumeName = IndexQueryVol
    VMDGResName = SharePointSearch_VM_Query_VMDg
    ForceUnmount = ALL
)

NIC SharePointSearch_VM_NIC (
    MACAddress @SPS_R2_N1 = 02-BF-0A-D9-3D-0B
    MACAddress @SPS_R2_N2 = 02-BF-0A-D9-3D-0B
    MACAddress @SPS_R2_N3 = 00-0C-29-64-B6-3F
)

RegRep SharePointSearch_VM_Crawl_RegRep (

```

```

MountResName = SharePointSearch_VM_Crawl_MountV
Keys = {
    "HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Office Server\\
    14.0\\Search\\Applications\\
    cce65f70-1747-42b7-b877-bb82db4b0a68-crawl-1"
    = "cce65f70-1747-42b7-b877-bb82db4b0a68-crawl-1_App.reg",
    "HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Office Server\\
    14.0\\Search\\Components\\
    cce65f70-1747-42b7-b877-bb82db4b0a68-crawl-1"
    = "cce65f70-1747-42b7-b877-bb82db4b0a68-crawl-1_Comp.reg"
}
)

RegRep SharePointSearch_VM_Query_RegRep (
    MountResName = SharePointSearch_VM_Query_MountV
    Keys = {
        "HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Office Server\\
        14.0\\Search\\Applications\\
        cce65f70-1747-42b7-b877-bb82db4b0a68-query-1"
        = "cce65f70-1747-42b7-b877-bb82db4b0a68-query-1_App.reg",
        "HKEY_LOCAL_MACHINE\\SOFTWARE\\Microsoft\\Office Server\\
        14.0\\Search\\Components\\
        cce65f70-1747-42b7-b877-bb82db4b0a68-query-1"
        = "cce65f70-1747-42b7-b877-bb82db4b0a68-query-1_Comp.reg"
    }
)

SharePointSearch SharePointSearch_VM_Admin (
    SPSFarmAdmin = Administrator
    Domain = IPV6
    Password = IWOUlWlQIoJOkOL
    SearchCompType = Admin
    SearchAppName = NewSearch
)

SharePointSearch SharePointSearch_VM_Crawl (
    SPSFarmAdmin = Administrator
    Domain = IPV6
    Password = aogMdoDiaGbgCgd
    SearchCompType = Crawl
    SearchAppName = NewSearch
    ComponentID = 422a3a73-9d3d-4dd1-b411-ec187975af7c
)

```

```
SharePointSearch SharePointSearch_VM_Query (
    SPSFarmAdmin = Administrator
    Domain = IPV6
    Password = ftlRitInflglHli
    SearchCompType = Query
    SearchAppName = NewSearch
    ComponentID = dcb7d24d-cadb-44ee-a84a-504f7bf88cce
)
```

```
SharePointSearch SharePointSearch_VM_SearchService (
    SearchCompType = SearchService
)
```

```
VMdG SharePointSearch_VM_Crawl_VMDg (
    DiskGroupName = SPSCrawlDG
    DGGuid = d4df4e41-2f73-43c7-b4bd-dac0c05d52aa
)
```

```
VMdG SharePointSearch_VM_Query_VMDg (
    DiskGroupName = SPSIndex
    DGGuid = 0447418c-bcda-4d74-b286-dd15ef958239
)
```

```
SharePointSearch_VM_Query_Share requires SharePointSearch_VM_Query_MountV
SharePointSearch_VM_Crawl_MountV requires SharePointSearch_VM_Crawl_VMDg
SharePointSearch_VM_Query_MountV requires SharePointSearch_VM_Query_VMDg
SharePointSearch_VM_Crawl_RegRep requires SharePointSearch_VM_Crawl
SharePointSearch_VM_Query_RegRep requires SharePointSearch_VM_Query
SharePointSearch_VM_Admin requires SharePointSearch_VM_NIC
SharePointSearch_VM_Crawl requires SharePointSearch_VM_Admin
SharePointSearch_VM_Crawl requires SharePointSearch_VM_Crawl_MountV
SharePointSearch_VM_Query requires SharePointSearch_VM_Admin
SharePointSearch_VM_Query requires SharePointSearch_VM_Query_Share
SharePointSearch_VM_SearchService requires SharePointSearch_VM_Crawl_RegRep
SharePointSearch_VM_SearchService requires SharePointSearch_VM_Query_RegRep
```

```
// resource dependency tree
//
// group SPSPRI
// {
//   SharePointSearch SharePointSearch_VM_SearchService
//   {
```

```
//      RegRep SharePointSearch_VM_Crawl_RegRep
//      {
//          SharePointSearch SharePointSearch_VM_Crawl
//          {
//              SharePointSearch SharePointSearch_VM_Admin
//              {
//                  NIC SharePointSearch_VM_NIC
//              }
//              MountV SharePointSearch_VM_Crawl_MountV
//              {
//                  VMDg SharePointSearch_VM_Crawl_VMDg
//              }
//          }
//      }
//      RegRep SharePointSearch_VM_Query_RegRep
//      {
//          SharePointSearch SharePointSearch_VM_Query
//          {
//              SharePointSearch SharePointSearch_VM_Admin
//              {
//                  NIC SharePointSearch_VM_NIC
//              }
//              FileShare SharePointSearch_VM_Query_Share
//              {
//                  MountV SharePointSearch_VM_Query_MountV
//                  {
//                      VMDg SharePointSearch_VM_Query_VMDg
//                  }
//              }
//          }
//      }
//      }
```

Configuring the SharePoint Search Service Application service group

Prerequisites for configuring a service group for a SharePoint Search service application

Before you configure the service group for a SharePoint Search service application, make sure that you have completed the following activities:

- Configure the cluster.
- Ensure that you have installed Agent Pack Q4 2012.
- Install Microsoft SharePoint 2010 and configure the shared storage to be used for high availability.
See [“Installing and configuring SharePoint Server 2010”](#) on page 86.
- Create the SharePoint Search service application (or edit an existing application) that you want to configure for high availability.
- Change the index location of the Crawl and Query components from the default location to that on the shared storage.
See [“Changing the index location of the Crawl and Query components”](#) on page 87.
- Ensure that the Search Query and Site Settings service is running on all the cluster nodes that will be part of the service group.
The Search Query and Site Settings service is an Internet Information Services (IIS) service. By default, this service runs on each server that includes a search query component. The service manages the query processing tasks, which include sending queries to one or more of the appropriate query components and building the results set. At least one instance of the service must be running to serve queries.
- Ensure that the SharePoint Server Search 14 Windows service is configured to run under the SharePoint FarmAdmin credentials.

Installing and configuring SharePoint Server 2010

Install Microsoft SharePoint Server 2010 and configure the farm.

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

Before you proceed, note the following:

- Veritas 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** installation (Microsoft SharePoint Server 2010 installer > **Server Type** tab).

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

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

Changing the index location of the Crawl and Query components

By default, the indexes of the SharePoint Search service application components are stored on the local disk. Before you create a corresponding application service group, you must change the index location of each Crawl and Query component to the shared storage.

To change the index location

- 1 From the Start menu, launch the SharePoint 2010 Central Administration web page.
- 2 In the Application Management section, click **Manage service applications**.
- 3 Select the required Search service application, and click the **Manage** button.

- 4 On the Search Administration page, click the **Modify** button for Search Application Topology.

Search Application Topology

Current 

Modify

Category	Server Name	Status
Admin		
Administration Component	SPS_R2_N2	Online
Crawl - VCSW2K284\NEWSQL\NewSearch_CrawlStoreDB_4127603b599f4696a8bb8b1d30daee83		
Crawl Component 0	SPS_R2_N2	Online
Databases		
Administration Database : NewSearch_DB_f94c45e9ed9d46c68af68b62da7ff0cb	VCSW2K284 \\NEWSQL	
Crawl Database : NewSearch_CrawlStoreDB_4127603b599f4696a8bb8b1d30daee83	VCSW2K284 \\NEWSQL	
Property Database : NewSearch_PropertyStoreDB_eebf2176ed3546d58b2797a431d6dd7c	VCSW2K284 \\NEWSQL	
Index Partition - 0 - VCSW2K284\NEWSQL\NewSearch_PropertyStoreDB_eebf2176ed3546d58b2797a431d6dd7c		
Query Component 0	SPS_R2_N2	Online



- 5 For each Crawl or Query component that you want to configure under VCS, select the component and click the **Edit Properties** menu to change its Index location.

Central Administration ► Manage Search Topology

Use this page to add, remove, or modify components of the topology of this Search Service Application.

Topology for Search Service Application: NewSearch

[Learn more about search topology](#)

New ▾

View ▾

Category	Server Name	Pending Changes
Admin		
Administration Component	SPS_R2_N2	
Crawl - VCSW2K284\NEWSQL\NewSearch_CrawlStoreDB_4127603b599f4696a8bb8b1d30daee83		
Crawl Component 0	SPS_R2_N2	
Databases		
Administration Database : NewSearch_DB_f94c45e9ed9d46c68af68b62da7ff0cb	VCSW2K284\NEWSQL	
Crawl Database : NewSearch_CrawlStoreDB_4127603b599f4696a8bb8b1d30daee83	VCSW2K284\NEWSQL	
Property Database : NewSearch_PropertyStoreDB_eebf2176ed3546d58b2797a431d6dd7c	VCSW2K284\NEWSQL	
Index Partition - 0 - VCSW2K284\NEWSQL\NewSearch_PropertyStoreDB_eebf2176ed3546d58b2797a431d6dd7c		
<div>Query Component 0 ▾</div> <div> <div>Edit Properties</div> <div>Add Mirror</div> <div>Delete</div> </div>	SPS_R2_N2	<div>Apply Topology Changes</div> <div>Cancel</div>

- 6 On the Edit Query Component dialog box, change the default value in the Location of Index field to the shared storage path.

The screenshot shows the 'Edit Query Component' dialog box. The 'Location of Index' text box is highlighted with a green rectangle, indicating the field to be modified. The text box contains the path 'Q:\QueryIndexComp'.

- 7 Click the **OK** button to close the dialog box.
- 8 Click the **Apply Topology Changes** button to save the changes.

Configuring a service group for a SharePoint Search service application manually

Use Veritas Cluster Manager (Java Console) to add a new service group for a SharePoint Search service application.

For information about using the Veritas Cluster Manager, see the *Cluster Server Administrator's Guide*.

To configure the application service group manually

- 1 Create a service group by providing the following values:
 - Service group name
 - Systems that will be part of the service group for the Search service application
 - Service group type
The service group should be of the type **Failover**.

- Service group template

In the SFW HA environment, select the **SharepointSearch-VMGroup** template.

In the SFW environment, select **SharepointSearch-NetAppGroup** template.

- 2 Launch the SharePoint 2010 Management Shell, and execute the following PowerShell script:

```
C:\Program Files\Veritas\Cluster  
Server\bin\SharePointSearch\SearchServiceAppDetails.ps1
```

Provide the application name as input. The script retrieves the details of the Search service application components and displays them on the screen. These details include the property values of the Admin, Crawl, and Query components. Some of these values are to be used as attribute values for the agent resources.

The following graphic depicts a sample output of the script.

```

Administrator: SharePoint 2010 Management Shell
PS C:\Program Files\Veritas\cluster server\bin\SharePointSearch> .\SearchService
AppDetails.ps1
Please enter the Search Application Name:: NewSearch

Getting Details about Search Application.....

Details of SharePoint Search application NewSearch
-----

NewSearch is Online
Admin Component is on server UCSW2K287
Location for Temp Index Storage for Admin Component is on server C:\Program Fil
es\Microsoft Office Servers\14.0\Data\Office Server\Applications
Admin DatabaseName : NewSearch_DB_f94c45e9ed9d46c68af68b62da7ff0cb

Crawl Components Details
-----

Number of Crawl Components : 1
Name of Crawl Components :
CrawlComponent ID : 422a3a73-9d3d-4dd1-b411-ec187975af7c
CrawlComponent Name : cce65f70-1747-42b7-b877-bb82db4b0a68-crawl-1
CrawlComponent Server is SPS_R2_N2
CrawlComponent IndexLocation is S:\CrawlStore
CrawlComponent State is Ready
Registry Entries to be set in RegRep Resource :
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office Server\14.0\Search\Applications\cce
65f70-1747-42b7-b877-bb82db4b0a68-crawl-1
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office Server\14.0\Search\Components\cce65
f70-1747-42b7-b877-bb82db4b0a68-crawl-1

Query Components Details
-----

Number of Query Components : 1
QueryComponent Name is cce65f70-1747-42b7-b877-bb82db4b0a68-query-1
QueryComponent ID is dcb7d24d-cadb-44ee-a84a-504f7bf88cce
QueryComponent Server is SPS_R2_N2
QueryComponent IndexLocation is Q:\QueryIndexComp
QueryComponent State is Ready
QueryComponent PropagationStatus is Idle
Registry Entries to be set in RegRep Resource :
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office Server\14.0\Search\Applications\cce
65f70-1747-42b7-b877-bb82db4b0a68-query-1
HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Office Server\14.0\Search\Components\cce65
f70-1747-42b7-b877-bb82db4b0a68-query-1
PS C:\Program Files\Veritas\cluster server\bin\SharePointSearch> _
    
```

- 3 In the Cluster Explorer, update the attributes of the storage resources (for example, VMDg and MountV in the SFW HA environment) and network resources with appropriate values.
- 4 Update the Query Share resource (for example, SharePointSearch_VM_Query_Share) attributes by copying over the values of the corresponding Query component properties as follows:
 - PathName = IndexLocation

Note: Do not include the drive letter when copying over this value.

- ShareName = Name
 - LanmanResName: Leave this attribute value empty.
- 5** Update the following attributes of the Admin, Crawl, and Query resources with the appropriate values:
- SPSFarmAdmin
 - Domain
 - Password
 - SearchAppName
 - ComponentID
- The component ID is not required for the Admin resource.
 For the Crawl or Query resource, copy the value of the corresponding Crawl or Query ComponentID property.
- 6** On the systems that are hosting the Crawl and Query components, back up the registry keys for those components. These registry keys are displayed in the output of the PowerShell script that you ran previously.

Note: It is important to back up these registry keys so that you can use them to restore the Search service application if the need arises.

- 7** In the Cluster Explorer, update the Keys attribute of all the RegRep resources by copying over the registry key names of the corresponding Crawl or Query component.


A sample key-value pair for the Query component is as follows:

HKLM\SOFTWARE\Microsoft\ Office Server\14.0\Search\Applications\ 3c8f51e3-ce82-45c9-bc99-81567c24beae-query-1	3c8f51e3-ce82-45c9-bc99-81567c24beae-query-1_App.reg
HKLM\SOFTWARE\Microsoft\ Office Server\14.0\Search\Components\ 3c8f51e3-ce82-45c9-bc99-81567c24beae-query-1	3c8f51e3-ce82-45c9-bc99-81567c24beae-query-1_Comp.reg

Note: Make sure that the values are unique by appending a string to indicate the key type.

- 8** Do not provide any attribute values for the SearchService resource.
- 9** Enable each resource in the service group.
 - When enabling the SearchService resource, a warning about its empty attributes is displayed. Ignore the warning and proceed.
 - When enabling the Query Share resource, a warning about the LanmanResName attribute is displayed. Ignore the warning and proceed.
- 10** Bring the service group online.

The following graphic depicts a sample service group configuration.




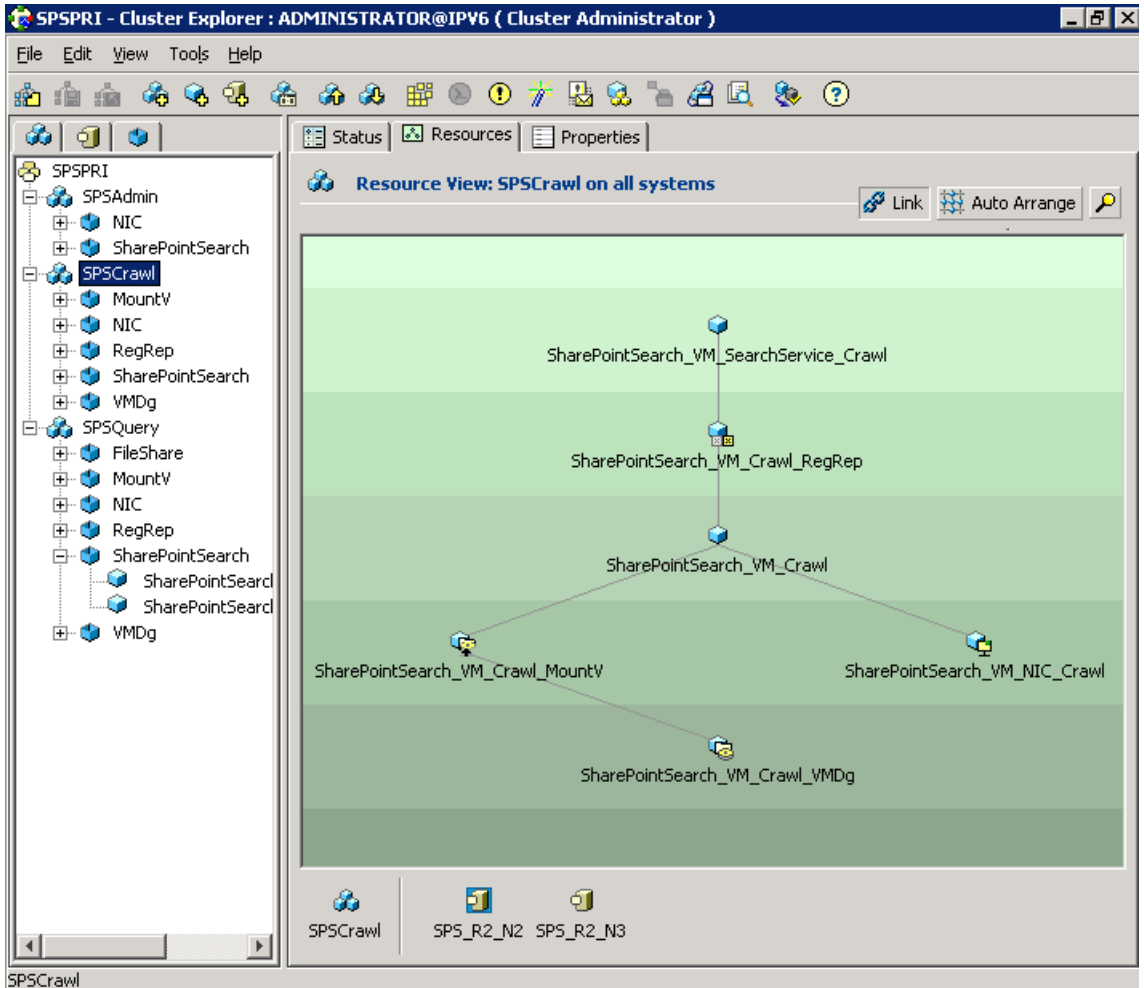
- 11** Optionally, you can split the service group depending on the Search service application components that you want to monitor. To do so, retain the resources required for that agent and delete the others.


Take the following requirements into consideration:

- For each service group pertaining to Crawl or Query component, make sure that the agent resource is a child of the RegRep resource, and the RegRep resource is a child of the SearchService resource.
- For the service group pertaining to the Admin component, make sure that the agent resource is a child of the SearchService resource.

- The following graphics depict a sample service group configuration for each Search service application component.







Configuring the service group for a Search service application using the wizard

The SharePoint Search Configuration Wizard allows you to create or delete a service group. However, you cannot modify an existing service group using the wizard. To modify a SharePoint Search service application, you must first delete any associated service groups. After modifying the application, you need to create the service groups again.

To configure the application service group using the wizard

- 1 Launch the wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SharePoint Search Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2 Review the prerequisites and other instructions on the Welcome panel, and click **Next**.
- 3 On the Group Selection panel, select **Create service group** and click **Next**.
- 4 On the System Selection panel, provide the following input:
 - Provide a name for the service group.
 - Select systems in the Available Cluster Systems list and click the left-arrow button.
Select one or more systems, according to the number of nodes on which to configure the service group.
To remove a system from the list, select it again and click the right-arrow button.
 - Specify the priority order of the systems by using the up or down arrow buttons in the Systems in Priority Order list.
 - To enable the service group to automatically come online on one of the systems, select the **Include selected systems in the service group's AutoStartList attribute** checkbox.

Click **Next**.

- 5 On the Farm Admin User panel, specify the password for the Farm Admin user.
The Farm Name and Farm Admin User Name are populated automatically.
- 6 On the Component Selection panel, specify the application and its components to be monitored as follows:
 - The **Search Application** drop-down list is populated with the Search service applications configured in the SharePoint farm. Select the application for which you want to configure a service group.
 - All the components of the selected application are displayed; select the checkboxes against the components that you wish to monitor using this service group.
Optionally, use the **Select all components** checkbox to indicate that all the components of the selected application should be monitored.

Note: Some component names might appear in bold with a different icon. The corresponding check boxes might appear selected, which indicates that those components are already being monitored as part of another service group.

You cannot deselect such components, but you can select the other components of the application for which resources have not been configured. No resources are created for components that are already part of another service group.

Click **Next**.

- 7 On the Network Configuration panel, review the selected systems and their network adapters, and click **Next**.

To change the network adapter for a system, use the drop-down box in the Adapter Display Name column.

- 8 On the Summary panel, review the service groups, resources, and attributes.

You can edit a service group or resource name by clicking the name and pressing **F2**.

Click **Next**.

- 9 When prompted to confirm the configuration changes that you are about to make, click **Yes**.

Along with the Search service application failover service group, a parallel service group named SharePointSearch-QueryProcessor is automatically created to monitor the Search Query and Site Settings service. If this parallel service group already exists in the cluster, any new nodes that are configured for the components that you selected in step 6 are added to its system list.

- 10 On the Finish panel, click **Finish** to exit the wizard.

To bring the newly created service group online immediately, select **Bring the service group online** before clicking **Finish**.

Verifying the application service group

This section provides steps to verify a service group configuration by bringing the service group online, taking it offline, and switching the service group to another cluster node.

You should perform a Site Search to verify whether SharePoint Search service application configured under VCS is working.

Bringing the service group online

Perform the following steps to bring the service group online from the VCS Java Console.

To bring a service group online

- 1 On the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

or

Select the cluster in the Cluster Explorer configuration tree, select the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Choose **Online**, and choose the appropriate system from the pop-up menu.
(Right-click > Online > *system_name*)

Taking the service group offline

Perform the following steps to take the service group offline from the VCS Java Console.

To take a service group offline

- 1 On the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

or

Select the cluster in the Cluster Explorer configuration tree, select the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Choose **Offline**, and choose the appropriate system from the pop-up menu.
(Right-click > Offline > *system_name*)

Switching the service group

The process of switching a service group involves taking it offline on its current system and bringing it online on another system.

To switch a service group

- 1 On the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

or

Select the cluster in the Cluster Explorer configuration tree, select the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Choose **Switch To**, and choose the appropriate system from the pop-up menu.
(Right-click > Switch To > *system_name*)

Disabling the service group

To disable the agent, you must change the service group for a Search service application to the OFFLINE state. You can switch the agent to another system.

To disable a service group

- 1 On the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

or

Select the cluster in the Cluster Explorer configuration tree, select the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Choose **Disable**, and choose the appropriate system from the pop-up menu. (Right-click > Disable > *system_name*)

Configuring a Search service application for disaster recovery

To configure a Search service application for disaster recovery, the cluster nodes on the primary and those on the secondary site must be part of the same SharePoint farm.

Use the Disaster Recovery Configuration Wizard (DR wizard) for additional applications. For more information, see the *Storage Foundation and High Availability Solutions Solutions Guide*.

If the service group for a Search service application does not contain a MountV resource, you cannot use the DR wizard to configure disaster recovery. For example, a service group that is only configured to monitor the Admin component of a Search service application does not contain a MountV resource. You need to manually configure disaster recovery for such a service group.

To configure disaster recovery for a service group that only monitors the Admin component

- 1 Make sure that the global clustering option (GCO) has been set up for the cluster on the primary as well as the secondary site.
- 2 Launch the Java GUI, and connect to the DR site.
- 3 Add a new failover service group with the same name as the one on the primary site.
- 4 Copy all the resources of the service group from the primary site and paste them into the service group at the secondary site.
- 5 At the secondary site, perform the following tasks sequentially:
 - Create dependency links between the resources similar to the service group at the primary site.

- Verify the attributes for all the resources.
- Edit the NIC resource attribute to provide the MAC addresses for each system.
- Enable all resources.
Ignore the warning that appears for the Search Service resource.

Note: Do not bring the service group online.

- 6 At the primary site, make the service group global.

On the Global Group Configuration Wizard, select the cluster on secondary site, and provide the remote cluster information when prompted.

Administering the SharePoint Search Service Application service group

About administering the application service group

This chapter describes the administrative tasks that you can perform on application service groups such as modifying the service group configuration and deleting a service group.

Modifying the application service group

To modify a SharePoint Search service application, you must first delete any associated service groups. After modifying the application, you need to create the service groups again.

To modify the service group

- 1 Delete the service group for the Search service application using the SharePoint Search Configuration wizard.
- 2 Remount the volumes that were taken offline by the wizard.
- 3 Make the desired changes to the Search service application.
- 4 Create the service group using one of the following methods:
 - Using the SharePoint Search Configuration wizard
See [“Configuring the service group for a Search service application using the wizard”](#) on page 98.
 - Using the service group templates

See [“Configuring a service group for a SharePoint Search service application manually”](#) on page 90.

Deleting the application service group

Deleting a service group for the Search service application is required in the following scenarios:

- When you want to modify the application
 After deleting a service group, you need to remount the volumes that were taken offline by the wizard.
- When you decide to stop monitoring the application
 When you delete a service group using the wizard, all the resources are taken offline. Any new search queries will fail, because the index locations are unavailable. To continue processing queries further, you need to remount the volumes that were taken offline by the wizard.

To delete a service group using the wizard

- 1 Launch the wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SharePoint Search Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2 On the Welcome panel, click **Next**.
- 3 On the Group Selection panel, select **Delete service group**.
 All the service groups that are configured in the cluster are listed.
 Select the service group that you want to delete, and click **Next**.
- 4 On the Summary panel, review the service groups, resources, and attributes that will be deleted, and click **Next**.
- 5 When prompted to confirm the configuration changes that you are about to make, click **Yes**.
 When the last service group for a Search service application is deleted from a cluster node, the node is removed from the system list of the SharePointSearch-QueryProcessor service group.
 After all the service groups for a Search service application in the cluster are deleted, the SharePointSearch-QueryProcessor service group is also deleted.
- 6 On the Finish panel, click **Finish** to exit the wizard.

Troubleshooting

This chapter includes the following topics:

- [About troubleshooting VCS agents](#)
- [Troubleshooting issues with SharePoint Search service application components](#)
- [VCS logging](#)
- [Agent error messages and descriptions](#)

About troubleshooting VCS agents

This chapter lists issues that you might encounter with the SharePoint Search service application components and the possible solutions. It also lists the error messages associated with the VCS agent for the application. Each message includes a description and a recommended solution, if applicable.

Troubleshooting issues with SharePoint Search service application components

You might encounter problems when using the VCS agent for a SharePoint Search service application. For example, the Crawl or Query component might be in the Not Responding state.

To troubleshoot when the Crawl or Query component is not responding

- 1 Launch the SharePoint 2010 Central Administration web page.
- 2 On the Search Administration page, view the details of the relevant Search service application.

- 3 Check the Status of the Crawl or Query component.
If the Status is "Not Responding", perform the following steps.
- 4 Launch the SharePoint 2010 Management Shell, and execute the following PowerShell script:

```
C:\Program Files\Veritas\Cluster Server\bin\SharePointSearch\SearchServiceAppDetails.ps1
```


Provide the application name as input. The script displays the details of the Search service application components.
- 5 Make sure that the storage for the index location on the Crawl Server or Query Server is accessible.
- 6 Restart the SharePoint Server Search 14 service.
- 7 Wait for a reasonable amount of time, and then verify that the component status is Online.
- 8 Ensure the following for the Query component only:
 - The ShareName attribute of the FileShare resource (for example, SharePointSearch_VM_Query_Share) has the correct value. Compare it with the Name property of the Query component in the output of the PowerShell script.
 - The FileShare resource is online.
- 9 If the component status is still not Online, then on the CrawlComponent or QueryComponent Server, perform one of the following tasks:
 - If you created the service group manually, restore the registry keys that you backed up before creating the service group or after modifying the Search service application.
See ["Configuring a service group for a SharePoint Search service application manually"](#) on page 90.
 - If you created the service group using the SharePoint Search Configuration Wizard, restore the registry keys that the wizard backed up.
See ["Restoring the Crawl or Query component registry keys"](#) on page 106.
- 10 Restart the SharePoint Server Search 14 service, and then check the component status again.

Restoring the Crawl or Query component registry keys

You might need to restore the registry keys for the Crawl or Query components if they stop responding.

To restore the registry keys backed up by the wizard

- 1 Open Windows Explorer on the node where you ran the SharePoint Search Configuration Wizard to create the service group.

- 2 Navigate to the
`%vcs_home%\bin\SharePointSearch\RegistryBackupForComponents` folder.

`%vcs_home%` is the folder where VCS is installed, for example:

```
C:\Program Files\Veritas\cluster server
```

- 3 Two files exist for each component in this folder, and they are named as follows:

```
applicationName_applicationGUID-componentName_timeStamp_keyType.reg
```

For example:

```
SearchApp_0c301859-27f9-4013-a992-822bd8be56e3-query-2  
_2012-12-07_14_33_41_Application.reg
```

and

```
SearchApp_0c301859-27f9-4013-a992-822bd8be56e3-query-2  
_2012-12-07_14_33_44_Component.reg
```

Identify the backup files that contain the appropriate application and component registry keys.

- 4 Copy these files to the node that currently owns the Crawl or Query component.

- 5 Run the following command:

```
SharePointSearchWizard.exe "registryFileNameIncludingAbsolutePath"
```

For example:

```
C:\>SharePointSearchWizard.exe "C:\Program Files\Veritas\  
cluster server\bin\SharePointSearch\RegistryBackupForComponents\  
SearchApp_0c301859-27f9-4013-a992-822bd8be56e3-query-2  
_2012-12-07_14_33_41_Application.reg"
```

- 6 A message box appears, informing you that the registry keys were restored successfully. Click **OK**.

VCS logging

VCS generates two error message logs: the engine logs and the agent logs. Log file names are appended by letters. Letter A indicates the first log file, B the second, C the third, and so on.

The agent log is located at %VCS_HOME%\log\agent_A.txt. The format of agent log messages is:

Timestamp (Year/MM/DD) | Mnemonic | Severity | UMI | Agent Type | Resource Name | Entry Point | Message Text

The following table describes the agent log message components and their descriptions.

Table 8-1 Log message components and their description

Log message component	Description
Timestamp	Denotes the date and time when the message was logged.
Mnemonic	Denotes which Veritas product logs the message. For Cluster Server, the mnemonic is 'VCS'.
Severity	<p>Denotes the severity of the message. Severity is classified into the following types:</p> <ul style="list-style-type: none"> ■ CRITICAL indicates a critical error within a VCS process. Contact Technical Support immediately. ■ ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action. ■ WARNING indicates a warning or error, but not an actual fault. ■ NOTE informs the user that VCS has initiated an action. ■ INFO informs the user of various state messages or comments. <p>Among these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.</p>
Unique Message ID (UMI)	<p>UMI is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the SharePoint agent would resemble: V-16-20083-107.</p> <p>Originator ID for all VCS products is 'V-16.'</p> <p>Category ID for SharePoint agent is 20083.</p> <p>Message ID is a unique number assigned to the message text.</p>

Table 8-1 Log message components and their description (*continued*)

Log message component	Description
Message Text	Denotes the actual message string.

You can view these message logs using Notepad or any text editor. All messages are logged to the engine and the agent logs. Messages of type CRITICAL and ERROR are written to the Windows event log.

A typical agent log resembles:

```
2012/09/20 07:53:51 VCS ERROR V-16-20083-107
SharePointSearch:SharePointSearch_VM_Admin:monitor:
Failed to open connection with the helper process. Error: 2.
```

VCS Cluster Configuration Wizard (VCW) logs

The VCS Cluster Configuration Wizard (VCW) log is located at %allUsersProfile%\Veritas\Cluster Server\vcw.log.

Here, %allusersprofile% is the file system directory containing application data for all users. A typical path is C:\ProgramData\.

The wizard log text is of the format *threadID | message Text*.

ThreadID is the ID of the thread initiated by the wizard and Message Text is the actual message generated by the wizard.

A typical wizard log resembles the following:

```
00000576-00000264: ExecMethod return 00000000.
00000576-00000110: CRegistry::Query for VCS License failed.
Error=0x00000000
00000576-00000264: ExecMethod return 00000000.
00000576-00000264: ExecMethod return 00000001.
00000576-00000127: QueryDWORDValue returned 0x00000001
00000576-00000132: CRegistry::Query for VxSS Root information failed.
Error=0x00000001
```

Agent error messages and descriptions

Each ERROR or WARNING message has a description and a recommended solution.

VCS agent for SharePoint Search service application

The following table lists the description of error messages and recommended solutions for the application agent.

Table 8-2 VCS agent for SharePoint Search service application - error messages

Message	Description
Failed to launch the process 'SPSearchHelperProcess'. Error = Error code.	Make sure that the correct values are provided for the following resource attributes: <ul style="list-style-type: none"> ■ SPSFarmAdmin ■ Domain ■ Password
Invalid Component Type	A resource goes into the UNKNOWN state. Make sure that you have selected one of the available values for the SearchCompType.
Failed to get status of Admin Component. Please ensure that SearchAppName attribute is specified correctly.	The resource corresponding to the Search Admin component goes into the UNKNOWN state. Ensure that the correct Search service application name is specified as the SearchAppName attribute value.
Crawl Component Not Found	Ensure that the ComponentID attribute value contains the correct ID of the Crawl component.
Query Component Not Found	Ensure that the ComponentID attribute value contains the correct ID of the Query component.
Search Service Application Not Found	The resource corresponding to the Search service application component goes into the UNKNOWN state. Ensure that the correct Search service application name is specified as the SearchAppName attribute value.
Attribute ComponentID is not specified for Crawl Component	Ensure that you have specified a value for the ComponentID attribute of the Crawl component.
Attribute ComponentID is not specified for Query Component	Ensure that you have specified a value for the ComponentID attribute of the Query component.

Using Veritas AppProtect for vSphere

This appendix includes the following topics:

- [About Just In Time Availability](#)
- [Prerequisites](#)
- [Setting up a plan](#)
- [Deleting a plan](#)
- [Managing a plan](#)
- [Viewing the history tab](#)
- [Limitations of Just In Time Availability](#)
- [Getting started with Just In Time Availability](#)
- [Supported operating systems and configurations](#)
- [Viewing the properties](#)
- [Log files](#)
- [Plan states](#)
- [Troubleshooting Just In Time Availability](#)

About Just In Time Availability

The Just In Time Availability solution provides increased availability to the applications on a single node InfoScale Availability cluster in VMware virtual environments.

Using the Just In Time Availability solution, you can create plans for:

- Planned Maintenance
- Unplanned Recovery

Planned Maintenance

In the event of planned maintenance, the Just In Time Availability solution enables you to clone a virtual machine, bring it online, and fail over the applications running on that virtual machine to the clone on the same ESX host. After the maintenance procedure is complete, you can fail back the applications to the original virtual machine. Besides failover and failback operations, you can delete a virtual machine clone, view the properties of the virtual machine and its clone, and so on.

Unplanned Recovery

When an application encounters an unexpected or unplanned failure on the original or primary virtual machine on the primary ESX host, the Just In Time Availability solution enables you to recover the application and bring it online using the unplanned recovery feature.

With **Unplanned Recovery Policies**, the Just In Time Availability solution enables you to set up recovery policies to mitigate unplanned failures that are encountered by an application. Just In Time Availability solution provides the following recovery policies; you may select one or all the recovery policies as per your need:

Unplanned Recovery Policies	Description
-----------------------------	-------------

Restart Application	
---------------------	--

	Just In Time Availability (JIT) solution attempts to restart the service group (SG), and bring the application online on the original virtual machine on primary ESX.
--	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------

	Maximum three retry attempts are permitted under this policy.
--	---------------------------------------------------------------

	Note: If all the three attempts fail, application continues to remain in faulted state or continues with the next policy as selected while creating a plan.
--	--------------------------------------------------------------------------------------------------------------------------------------------------------------------

Unplanned Recovery Policies	Description
Restart virtual machine (VM)	<p>Just In Time Availability (JIT) solution performs the following subsequent tasks:</p> <ul style="list-style-type: none">■ take the service group offline■ shut down the virtual machine■ power on the virtual machine■ bring the service group online on the original virtual machine on primary ESX <p>You are provided with Last attempt will be VM reset option to reset the virtual machine.</p> <p>By default, this checkbox is selected and the default retry attempt value is one. If you retain the default settings, then VM reset operation is performed on the virtual machine at the first attempt itself.</p> <p>Maximum three retry attempts are permitted for this operation.</p> <p>If you deselect the checkbox, then the virtual machine reset (VM Reset) operation is not performed.</p>
Restart VM on target ESX	<p>Using this policy, you can recover the faulted application on the virtual machine.</p> <p>In this policy, the original virtual machine is unregistered from the primary ESX; registered on the target ESX; and the faulted application is brought online on the target ESX.</p>
Restore VM on target ESX	<p>Using this policy, you can recover the faulted application on the virtual machine using a boot disk backup copy of the original virtual machine.</p> <p>In this policy, the original virtual machine is unregistered from the ESX and the boot disk backup copy of the original virtual machine is registered on target ESX. The faulted application is then brought online on the virtual machine.</p>

Unplanned Recovery Policies

Unplanned Failback

The **Unplanned Failback** operation lets you fail back the application from the boot disk backup copy of virtual machine on the target ESX to the original virtual machine on primary ESX.

If you have selected either **Restart VM on target ESX** or **Restore VM on target ESX** or both the recovery policies, you can perform the **Unplanned Failback** operation.

On the **Plans** tab, in the plans table list, right-click the virtual machine and click **Unplanned Failback**.

Note: **Unplanned Failback operation** operation is disabled and not available for the plans and the virtual machines which have **Restart Application** and **Restart VM** policies as the only selected options.

Based on the selected recovery policy for a plan, Just In Time Availability (JIT) solution performs the necessary operations in the sequential order.

For example, if you have selected **Restart Application** and **Restart VM** as the recovery policy, then in the event of unplanned application failure, first it performs tasks for **Restart Application** policy and if that fails, it moves to the next policy.

You may select one or all the recovery policies based on your requirement.

[Table A-1](#) lists the sequence of tasks that are performed for each Unplanned Recovery policy.

Table A-1 Tasks performed for each Unplanned Recovery policy

Unplanned Recovery Policy	Tasks Performed
Restart Application	◆ Make an attempt to restart the application.
Restart virtual machine (VM)	1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Power on the virtual machine 4 Brings the service group(s) online

Table A-1 Tasks performed for each Unplanned Recovery policy (*continued*)

Unplanned Recovery Policy	Tasks Performed
Restart VM on target ESX	<ol style="list-style-type: none"> 1 Takes the service group(s) offline 2 Shuts down the original virtual machine 3 Detaches the data disks from the original virtual machine 4 Unregisters the virtual machine from the primary ESX 5 Registers the original virtual machine on target ESX 6 Attaches the data disks back to the virtual machine 7 Power on the virtual machine 8 Brings the service group(s) online
Restore VM on target ESX	<ol style="list-style-type: none"> 1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Detaches the data disks from the virtual machine 4 Unregisters the original virtual machine from the target ESX 5 Registers the boot disk backup copy of the original virtual machine to the target ESX 6 Attaches the data disks back to the virtual machine 7 Power on the virtual machine 8 Brings the service group(s) online
Unplanned Failback	<ol style="list-style-type: none"> 1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Detaches the data disks from the virtual machine 4 Unregisters the virtual machine from the target ESX 5 Registers the virtual machine using the original boot disk backup copy to the primary ESX 6 Attaches the data disks to the virtual machine 7 Power on the virtual machine on primary ESX 8 Brings the service group(s) online on the virtual machine

Scheduler Settings

While creating a plan for unplanned recovery, with **Scheduler Settings**, you can set up a schedule for taking a back up of boot disk of all the virtual machines that are a part of the plan.

To use the Just In Time Availability solution, go to **vSphere Web Client > Home view > Veritas AppProtect**.

See [“Setting up a plan”](#) on page 118.

Prerequisites

Before getting started with Just In Time Availability, ensure that the following prerequisites are met:

- The Just In Time (JIT) solution feature cannot co-exist with VMware HA, VMware FT, and VMware DRS. This pre-requisite is applicable for **Unplanned Recovery** only.
- VIOM 7.2 version must be installed and configured using fully qualified domain name (FQDN) or IP.
- Make sure that you have the admin privileges for vCenter.
- VMware Tools must be installed and running on the guest virtual machine.
- VIOM Control Host add-on must be installed on VIOM server or machine.
- The virtual machines must be added in VIOM. The virtual machines, vSphere ESX servers, and VIOM must have the same Network Time Protocol (NTP) server configured.
- Make sure to specify VIOM Central Server FQDN or IP in the SNMP Settings of the vCenter Server.
- vCenter Server and VIOM must be configured using the same FQDN or IP address. Make sure that if FQDN is used to configure vCenter in VIOM Server that is used during the configuration.
- If raw disk mapping (RDM) disks are added to the virtual machine, then make sure that the virtual machine is in the physical compatibility mode. Veritas AppProtect does not support the virtual compatibility mode for RDM disks.
- For Microsoft Windows operating system, make sure that you have the Microsoft Windows product license key. The key is required to run the Sysprep utility, which enables customization of the Windows operating system for a clone operation.

- For RHEL7 and SUSE12 operating system, install the deployPkg plug-in file on the virtual machine.
For more information on installing the plug-in, see <https://kb.vmware.com/kb/2075048>
- Make sure that the InfoScale Availability service group is configured with one of the storage agents such as Mount, DiskGroup, LVMVolumeGroup, VMNSDg (for Windows), or DiskRes (for Windows), for the data disks. This configuration enables Veritas AppProtect to discover data disks for the applications. Also, ensure that the service group is online to determine data disk mapping.
- Virtual machines which have snapshots associated with them are not supported.
- Virtual machines with SCSI Bus Sharing are not supported.
- Make sure that the SNMP Traps are configured for the following from vCenter server to VIOM:
 - Registered virtual machine
 - Reconfigured virtual machine
 - Virtual machine which is getting cloned
- Make sure that the boot disk of VM's (vmdk) does not have spaces.
- For HA console add on upgrade from VIOM 7.1 to VIOM 7.2, refer *Veritas InfoScale Operations Manager 7.2 Add-ons User's Guide* for more details.
- Make sure to set the vSphere DRS Automation Level to manual, if you want to configure **Restart VM on target ESX** or **Restore VM on target ESX** policies for your plan.
- Ensure to update or edit the plan, when a virtual machine is migrated or if there are any modifications made to the settings of the virtual machines which are configured for that plan.
- Ensure to increase the tolerance limit of DiskRes resource to two, if you want to create a plan for unplanned recovery with **Restore VM on target ESX** as the unplanned recovery policy.

Note: This prerequisite is applicable for Windows operating system.

Setting up a plan

Plan is a template which involves a logical grouping of virtual machines so as to increase the availability of the application in the event of a planned failover and recovery of the application in the event of an unexpected application failure.

To set up a plan

- 1 Launch Veritas AppProtect from the **VMware vSphere Web Client > Home view > Veritas AppProtect** icon.

- 2 Click **Configure Plan**.

The **Plan Configuration** wizard appears.

- 3 Specify a unique **Plan Name** and **Description**, and then click **Next**.

The wizard validates the system details to ensure that all prerequisites are met.

- 4 Select the virtual machines that you want to include in the plan, review the host and operating system details, and then click **Next**.

The **Unplanned Recovery Settings** page appears.

- 5 On the **Unplanned Recovery Settings** page, you can configure the selected virtual machines for **Unplanned Recovery** as well.

Deselect the **Configure selected VMs for Unplanned Recovery as well** check box, if you do not want to include the selected virtual machines for unplanned recovery.

If you have selected the virtual machines for unplanned recovery, then set up the unplanned recovery policies as appropriate from the available options. You can set up policies to restart applications, restart virtual machines, restart virtual machine on target ESX, and restore a virtual machine on target ESX.

If you have selected **Restore VM on target ESX** as the unplanned recovery policy, then you can set up a schedule to create a boot disk back up copy of the virtual machine within the configured plan. You can set the frequency as daily, weekly, monthly, or manual as per your requirement.

After you have finished making necessary settings for Unplanned Recovery, Click **Next**.

- 6 The wizard validates the prerequisite attributes of the virtual machine and the ESX host, and adds the qualified virtual machines to the plan.

Click **Next** after the validation process completes.

- 7 In the **Disks** tab, you can view the selected application data disks. Just In Time Availability solution uses the selected data disks to perform detach-attach operation during a planned failover and unplanned recovery.

Note: If the disks are not auto-marked as selected to perform detach-attach operation, then first refresh the VIOM server and then the VCenter server in VIOM and then create a plan.

- 8 In the **Network Configuration** tab, specify the network interface configuration details for the cloned virtual machine. Make sure to specify at least one public interface and valid IP details.
- 9 In the **Unplanned Recovery Target** tab, specify the target ESX server to restore the virtual machine, and the target ESX port details.

Note: The **Unplanned Recovery Target** tab is visible only when **Restart VM on target ESX** or **Restore VM on target ESX** is selected.

- 10 In the **Windows Settings** tab, specify the domain name, Microsoft Windows product license key, domain user name, domain password, admin password, and time zone index.

Note: The **Windows Settings** tab is visible only when a Windows virtual machine is selected in the plan.

- 11 Click **Next**. The **Summary** wizard appears.
- 12 In the **Summary** wizard, review the plan details such as the plan name, unplanned recovery policies, schedule, and so on.

Deselect the **Start backup process on finish** checkbox if you do not want to initiate a backup process when the plan creation procedure is finished. This checkbox is selected by default.

Click **Create**. The plan is created and saved.

- 13 Click **Finish** to return to the plans tab and view the created plans.

See [“Managing a plan”](#) on page 120.

See [“Deleting a plan”](#) on page 120.

Deleting a plan

After you have finished performing failback operations from the clone to the primary virtual machine in case of planned maintenance and recovery operations in case of unplanned recovery, you may want to delete the plan.

To delete a plan

- 1 Launch **Veritas AppProtect** from the VMware vSphere Web Client Home view.
- 2 In the **Plans** tab, select the plan that you want to delete.
- 3 Click **Delete Plan**.

Note: The **Delete plan** icon is enabled only when the selected plan is in **Ready For Failover**, **Failed to Revert**, and **Failed to Failback** state.

Managing a plan

Planned Maintenance

After the maintenance plan is created, you can fail over the applications to the clone virtual machine and fail back the applications from the clone to the virtual machine. When the scheduled maintenance is complete, you can delete the cloned virtual machine or retain it for future use.

To perform failover, failback, revert, or delete clone operations, go to **Plans**, and select a plan. Based on the enabled operation, perform the following tasks:

To fail over the applications to the cloned virtual machine

- ◆ Click the **Failover** icon.

Just In Time Availability (JIT) performs the sequence of failover tasks, which includes taking the application offline, detaching the disks, cloning the virtual machine, attaching the disks, and so on.

To fail back the applications from the clone to the primary virtual machine

- ◆ Click the **Failback** icon.

Just In Time Availability (JIT) performs the sequence of failback tasks, which includes taking the application offline, detaching the disks, attaching the disks, and so on.

To revert a failover or a fallback operation

- ◆ Click the **Revert** icon.

If the failover or a fallback operation fails, the revert operation restores the applications on the virtual machine, and deletes the clone if created.

To delete a clone

- ◆ Click the **Delete Clone** icon.

After the fallback operation is complete, you can delete the clone. By default, the revert operation deletes the clone.

Note: Alternatively, right-click **Plan** in the **Plans** table on the **Plans** wizard to perform failover, fallback, revert, delete plan, and delete clone operations.

Unplanned Recovery

Once you have set up a plan for unplanned recovery during **Configure Plan** operation, based on the recovery policies selected for the plan, the application is recovered accordingly.

You can manage unplanned recovery policies settings by performing the following operations on the plan and its associated virtual machines.

Managing unplanned recovery settings

On the **Plans** tab, in the plans table which lists all the existing plans, navigate to the required plan and use the right-click option on the selected plan.

- **Edit:** Use this option to modify the configured plans settings such as adding or removing a virtual machine from the plan, and so on.
The same **Configuration Plan** wizard using which you had set up or configured a plan is displayed with pre-populated details.
See [“Setting up a plan”](#) on page 118.
- **Disable Unplanned Recovery:** Use this option to disable the Unplanned Recovery settings.
- **Enable Unplanned Recovery:** Use this option to enable the Unplanned Recovery settings.
- **Disable Scheduler:** Use this option to disable the scheduler settings.
- **Enable Scheduler:** Use this option to enable the scheduler settings.
- **Delete Plan:** Use this option to delete the created plan.
- **Properties:** Use this option to view the properties for unplanned recovery. It displays details such as the selected unplanned recovery policies and the

associated operations for the selected policies. It also provides information about the selected scheduler mode for performing boot disk back up operation for the selected virtual machines.

Managing virtual machines settings

On the **Plans** tab, in the plans table which lists all the existing plans and its associated virtual machines, navigate to the required virtual machine. Select the required virtual machine and use the right-click option on the selected virtual machine.

- **Remove VM From Plan:** Use this option to delete the virtual machine from the selected plan.
- **Create Clone Backup:** Use this option to create a boot disk backup copy of the virtual machine.
- **Unplanned Failback:** Use this option to fail back the application from the boot disk backup copy of the virtual machine on target ESX to the original virtual machine on primary ESX.

Note: This option is available only if you have set unplanned recovery policies as **Restart VM on target ESX** or **Restore VM on target ESX**.

- **Properties:** Use this option to view properties such as the last run time for backup operation, last successful backup attempt time and the target ESX details.

See [“Plan states”](#) on page 127.

Viewing the history tab

On the **History** tab, you can view the detailed summary of the operations that are performed on the virtual machine. The details include the plan name, virtual machine name, operation, the status of the operation, the start and the end time of the operation, and the description of the operation status.

To view the summary

- 1 Launch **Veritas AppProtect** from the VMware vSphere Web Client Home view.
- 2 Click the **History** tab.

Limitations of Just In Time Availability

The following limitations are applicable to Just In Time Availability:

- On a single ESX host only ten concurrent failover operations are supported. Across ESX hosts, twenty concurrent failover operations are supported.
- Linked mode vCenter is not supported.
- Only three backup operations per data store are active, the rest will be queued. Only five backup operations per ESX host are active, the rest will be queued.

See [“Supported operating systems and configurations”](#) on page 125.

Getting started with Just In Time Availability

You can access the Just In Time Availability solution from the **vSphere Web Client > Veritas AppProtect** interface.

The **Veritas AppProtect** is registered with Veritas InfoScale Operations Manager (VIOM), and is accessed from the **vSphere Web Client > Home** view.

[Figure A-1](#) describes the Veritas AppProtect interface in detail.

Figure A-1 Elements of the Veritas AppProtect interface

The screenshot displays the Veritas AppProtect interface within the vSphere Web Client. The interface is divided into several sections:

- Plans Tab:** Contains a table with columns: Plan Name, Virtual Machine, Status, Update Time, Clone IP, Last Backup Status, Validation Status, Unplanned Recovery, Scheduler, and Description. A row is selected with Plan Name 'a' and Virtual Machine 'iar730-07vm14', showing a 'Failed Over' status.
- Action Buttons:** Located below the table, including icons for failover, failback, revert, delete clone, and unplanned recovery.
- Unplanned Recovery Summary:** A detailed section showing the steps and status of the recovery process.

Step	Status	Started At	Completed At	Description
1 SG Restart	✗	Sep 8, 2016 1:12:27 PM	Sep 8, 2016 1:15:07 PM	Failure
VM Restart				
1 Offline SG on VM	✓	Sep 8, 2016 1:15:07 PM	Sep 8, 2016 1:15:14 PM	Success
2 Reset VM	✓	Sep 8, 2016 1:15:14 PM	Sep 8, 2016 1:16:16 PM	Success
3 Online SG on VM	✗	Sep 8, 2016 1:16:16 PM	Sep 8, 2016 1:18:31 PM	Failure
Restart VM on target ESX				
VM Restore				
Unplanned Failback				
- Diagnostic Information:** A link at the bottom right corner.

Table A-2 Elements of the Veritas AppProtect interface and the description

Label	Element	Description
1	Plans tab	<p>Enables setting up a plan for a planned failover and unplanned recovery.</p> <p>Displays the plan attributes, and the virtual machines that are added to the plan.</p> <p>Displays the status of virtual machines for unplanned recovery and schedule for virtual machine back up operation based on the criteria set while configuring or editing the plan.</p> <p>Shows the enabled or disabled failover, failback, delete clone, revert, delete plan, and properties operations icons based on the state of the selected plan for planned failover.</p>
2	History tab	Displays the status and the start and the end time of the specific operation performed on the created plans.
3	Configure Plan link	Opens the Plan Configuration wizard.
4	Plans table	Displays the attributes of the plan.
5	Failover icon	Fails over the applications from the original virtual machine to the clone.
6	Failback icon	Fails back the applications from the clone to the original virtual machine.
7	Delete Clone icon	Deletes the cloned virtual machine.
8	Revert State icon	Reverts the failed operation, restores the applications to the original virtual machines, and delete the clone virtual machines.
9	Delete Plan icon	Deletes the plan.
10	Properties icon	Displays the attributes of each virtual machine and the clone.

Table A-2 Elements of the Veritas AppProtect interface and the description
(continued)

Label	Element	Description
11	Operation-specific tabs	<p>Displays the sequence of the tasks that are performed for the selected operation.</p> <p>Based on the operation that is executed, the associate tab opens.</p> <p>For Planned Maintenance</p> <ul style="list-style-type: none"> 1 Failover 2 Failback 3 Revert 4 Delete Clone <p>For Unplanned Recovery</p> <ul style="list-style-type: none"> ◆ Unplanned Recovery Summary
12	Diagnostic information	Displays the logs that are reported for the Veritas AppProtect interface.

See “[Plan states](#)” on page 127.

Supported operating systems and configurations

Just In Time Availability supports the following operating systems:

- On Windows: Windows 2012, and Windows 2012 R2.
- On Linux: RHEL5.5, RHEL6, RHEL7, SUSE11, SUSE12.

Just In Time Availability supports the following configurations:

- Veritas Cluster Server (VCS) 6.0 or later, or InfoScale Availability 7.1 and later.
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) 7.1 and 7.2 version on the virtual machines.
For more information about VRTSsfmh, see the *Veritas InfoScale Operations Manager 7.2 User Guide*.
- Veritas InfoScale Operations Manager (VIOM) 7.2 as a central or managed server.
- VMware vSphere 5.5 Update 2, Update 3, or 6.0 and 6.0 Update 1 version.

Viewing the properties

Virtual Machine Properties

The **Virtual Machine Properties** window displays information about the virtual machine and its clone such as name, operating system, cluster name, service groups, DNS server, domain, IP addresses, and data disks.

To view the properties

- 1 On the **Plans** tab, select the virtual machine.
- 2 Click the **Properties** icon or right-click the virtual machine.

The **Virtual Machine Properties** window opens and displays the attributes of the virtual machine and its clone.

Plan Properties

The **Plan Properties** window displays information about the unplanned recovery policies selected; scheduler mode set; and the time when the last backup operation was run and was successful for a virtual machine.

To view properties for the plan

- 1 In the Plan Name table, select the plan.
- 2 Right-click the selected plan. A window with a list of options is displayed.
- 3 Click **Properties**

The **Plan Properties** window opens and displays the unplanned recovery policies selected and the schedule mode for virtual machine backup operation.

Log files

The following log files are helpful for resolving the issues that you may encounter while using Veritas AppProtect:

- Console related logs:

```
/var/opt/VRTSsfmcs/logs/*
```

These log files show console messages and are useful for debugging console issues.

- Operations logs:

```
/var/opt/VRTSsfmh/logs/vm_operations.log
```

This log file shows the messages pertinent to the Veritas AppProtect interface.

- VMware vSphere 6.0 logs:

C:\ProgramData\VMware\vCenterServer\logs\vsphere-client\logs*

These log files show the messages that are reported for the VMware vSphere Web Client version 6.0.

- VMware vSphere 5.5 U2 and U3 logs:

C:\ProgramData\VMware\vsphere Web Client\serviceability\logs*

These log files show the messages that are reported for the VMware vSphere Web Client version 5.5 U2 and U3.

- Veritas AppProtect interface logs:

The log file shows the logs that are reported for the Veritas AppProtect interface. To view the log files, on the **Planned Maintenance** tab or the **History** tab > **Diagnostic Information**.

Plan states

Based on the state of the plan, the operation icons are enabled and disabled on the **Plans** tab.

Table A-3 List of plan and operation states

Plan state	Failover	Failback	Revert	Delete clone	Delete Plan	Unplanned Failback	Create Clone backup	Properties
Ready For Failover	✓	–	–	✓ Note: Enabled when the selected maintenance plan has an associate clone.	✓ Note: Enabled when the selected maintenance plan does not have an associate clone.	–	✓	✓
Failed Over	–	✓	–	–	–	–	–	✓
Failed To Failover	–	–	✓	–	–	–	–	✓

Table A-3 List of plan and operation states (*continued*)

Plan state	Failover	Failback	Revert	Delete clone	Delete Plan	Unplanned Failback	Create Clone backup	Properties
Failed To Failback	–	–	✓	–	–	–	–	✓
Failed To Revert	–	–	✓	–	✓	–	–	✓
Unknown	–	–	✓	–	–	✓	–	✓
Failed To Delete Clone	–	–	–	✓	–	–	–	✓
Failover In Progress	–	–	–	–	–	–	–	✓
Failback In Progress	–	–	–	–	–	–	–	✓
Revert In Progress	–	–	–	–	–	–	–	✓
Delete Clone In Progress	–	–	–	–	–	–	–	✓
Application Faulted	–	–	–	–	–	–	–	✓
Failed To Restart VM	–	–	–	–	–	–	–	✓
Failed To Move VM	–	–	–	–	–	✓	–	✓
Failed To Restore VM	–	–	–	–	–	✓	–	✓
Unplanned	–	–	–	–	–	✓	✓	–
Unplanned Restored VM	–	–	–	–	–	✓	–	✓
Unplanned Failed to Failback	–	–	–	–	✓	–	–	–

Troubleshooting Just In Time Availability

Table A-4 lists the issues and the recommended solutions.

Table A-4 Issues and the corresponding resolutions

Issue	Recommended Solution
When setting up a maintenance plan, the registered virtual machine is not listed on the wizard.	To troubleshoot the issue, make sure the following: <ul style="list-style-type: none">■ ESX host on which the virtual machine resides, is connected to the vCenter.■ The virtual machine is added as a managed host to Management Server.■ On the virtual machine, at least one application is configured for monitoring, along with VCS.■ The virtual machine is registered in VIOM.■ VCS is configured on the virtual machine.■ The virtual machine does not contain RHEL7 and SUSE 12, which are not supported. Note: Windows 2012 R2 and 2008 R2 are supported. <ul style="list-style-type: none">■ VCS is configured with the service groups.
When setting up a maintenance plan, the listed virtual machine is not available for selection.	To troubleshoot the issue, make sure the following: <ul style="list-style-type: none">■ The virtual machine is not configured for Global Cluster option (GCO).■ Agents that support SAN are configured.
When Veritas AppProtect executes an operation, the timeout message is reported.	To troubleshoot the issue, perform the following: <ul style="list-style-type: none">■ If the failover or the failback operation fails, then click Planned Maintenance > Revert icon. Retry the operation.■ If the delete plan or the delete clone operation fails, then retry the operation.
The revert operation failed.	Manually revert the virtual machine to its original state.