Veritas™ Cluster Server Application Note: Disaster Recovery for Microsoft SharePoint® Portal Server 2003

Windows 2000, Windows Server 2003

5.0

11456159

symantec.

Veritas Cluster Server Application Note: Disaster Recovery for Microsoft SharePoint Portal Server 2003

Veritas Cluster Server 5.0

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Windows and Microsoft SharePoint are registered trademarks of Microsoft Corporation.

Licensing and registration

Veritas Cluster Server is a licensed product.

Technical support

For technical assistance, visit http://entsupport.symantec.com and select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Chapter 1

Planning for SharePoint Portal Server with VCS

This chapter covers the following topics:

What this application note covers

In a disaster recovery configuration, you set up a secondary site to provide data and services in the event of a disaster at the primary site.

This Application Note covers guidelines and requirements for implementing SharePoint Portal Server 2003 with Veritas Cluster Server (VCS) for disaster recovery.

VCS is a component of Veritas Storage Foundation HA for Windows (SFW HA).

Instructions on deploying SFW HA for disaster recovery with Microsoft SQL Server are covered in *Veritas Storage Foundation and HA Solutions HA and*

Disaster Recovery Solutions Guide for Microsoft SQL. You should use this Application Note as a supplement to that guide.

Software requirements

The following software is required for this solution:

Veritas Storage Foundation HA 5.0 for Windows, including: ■ Veritas Cluster Server Database Agent for Microsoft SQL ■ Veritas Volume Replicator ■ Global Clustering Option	For information on system and software requirements for SFW HA, see *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*.
One of the following: ■ Microsoft SQL Server 2000 and its supported operating systems ■ Microsoft SQL Server 2005 and its supported operating systems	For information on software versions supported with SFW HA, see *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*.
Microsoft Office SharePoint Portal Server 2003 SP2	For information on requirements, see the Microsoft SharePoint Portal Server 2003 documentation. If installing SharePoint in a SQL Server 2005 environment, see the Microsoft Knowledge Base article “How to install SharePoint Portal Server 2003 Service Pack 2 and Windows SharePoint Services Service Pack 2 in a deployment that uses SQL Server 2005 as the database back end”: http://support.microsoft.com/kb/917446

Implementation process overview

You can implement SharePoint Portal Server 2003 with VCS for disaster recovery in the following ways:

- Setting up a new installation of SFW HA and SQL Server with SharePoint Portal Server
- Converting an existing standalone SQL Server to an SFW HA environment with Sharepoint Portal Server

Deploying new SQL Server installation

Table 1-1 shows the process for a new installation of SFW HA, SQL Server, and SharePoint Portal Server.

Table 1-1 New installation of SQL Server and SharePoint Portal Server

Task	For more information
On the primary site, install and set up SFW HA and SQL Server for high availability.	See the *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL.*
On the primary site, install and configure the SharePoint servers.	See the Microsoft documentation for SharePoint Portal Server 2003. In addition, see "Configuring SharePoint Portal Server 2003" on page 30 of this Application Note.
On the primary site, edit the SQL Server service group for disaster recovery as covered in this Application Note.	See "Configuring the VCS SQL Server service group for the SharePoint environment" on page 30 of this Application Note.
On the secondary site, create a parallel SFW HA and SQL Server high availability environment.	See *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*
On the primary and secondary sites, configure SFW HA for disaster recovery.	See *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL.*
On the secondary site, install and configure the SharePoint servers.	See the Microsoft documentation for SharePoint Portal Server 2003. In addition, see "Configuring SharePoint Portal Server 2003" on page 30 of this Application Note.

Converting existing SQL Server installation

Table 1-2 shows the process for converting an existing installation of SQL Server 2000 and SharePoint Portal Server 2003.

Table 1-2 Converting existing stand-alone SQL 2000 server

Task	For more information
On the primary site, convert the stand-alone SQL Server to a clustered server in a Storage Foundation HA environment.	See *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*
On the primary site, edit the SQL Server service group for disaster recovery as covered in this Application Note.	See "Configuring the VCS SQL Server service group for the SharePoint environment" on page 30 of this Application Note.
On the secondary site, create a parallel SFW HA and SQL Server high availability environment.	See *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*
On the primary and secondary sites, configure SFW HA for disaster recovery.	See *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL.*
On the secondary site, install and configure the SharePoint servers.	See the Microsoft documentation for SharePoint Portal Server 2003. In addition, see "Configuring SharePoint Portal Server 2003" on page 30 of this Application Note.

Supported SharePoint Portal Server configurations

Table 1-3 shows the SharePoint Portal Server configurations supported for setting up SharePoint Portal Server with VCS for disaster recovery.

Note: VCS disaster recovery does not support a single server configuration in which both SQL Server and SharePoint are running on the same computer.

Table 1-3 SharePoint Portal Server configurations supported with VCS

Configuration	Description
Large Server Farm	■ One or more computers running SQL Server ■ Two or more front-end Web servers ■ Two or more search engines ■ One or more index management servers, one of which is the job server
Medium Server Farm	■ One or more computers running SQL Server ■ One or two front-end web servers with the search component enabled ■ One index management and job server
Small Server Farm	■ One or more computers running SQL Server ■ One computer running as the job server and running all of the following: the Web server, index component, and search component.

A shared services deployment is also supported for disaster recovery. In this configuration, index and search services provided by one server farm are used by a second server farm, in a parent and child relationship.

Primary and secondary site configuration

Table 1-4 shows how the secondary site configuration compares to the primary site configuration.

Table 1-4 Primary and secondary site configuration

Product	Configuration
SQL Server	Set up the SQL Server configuration on the secondary site the same as on the primary site. In other words, if the primary site has two SQL Server systems, one active, one passive, the secondary site has the same. See *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*.
SharePoint Portal Server	The number of computers and which components are installed on each does not need to match the primary site. All the SharePoint Portal Server components could be on multiple computers in the primary site, as on a large server farm configuration, and share the same computer on the secondary site, as in a small server farm configuration. You typically set up the SharePoint components on both primary and secondary sites as part of the same topology (server farm).

Network configuration requirements

You should be familiar with requirements for setting up the network configuration to support high availability and disaster recovery.

See *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*.

To summarize the network configuration requirements for the SQL Server and SharePoint Portal Server components:

- One domain that includes both primary and secondary sites
- One SQL Server virtual server name
  You configure a virtual server name for SQL Server as part of setting up high availability. You also assign an instance name. The SQL Server instance in both the primary and secondary sites is assigned the same virtual server name. The SharePoint Portal Server web server connection to the SQL database is configured by virtual server name and instance name. Since only one instance of the SQL application is running at one time, there is no host name collision.

- Two SQL Server IP addresses, one for the primary site and one for the secondary site
  The active and passive SQL servers on the same site share the same static IP address.
- Network Load Balancer (NLB) IP address or web server address
  Multiple SharePoint Portal Server web servers are typically set up on a network load balancer (NLB) cluster. You can use a hardware or software solution to implement the NLB cluster. You can implement a global NLB that includes both sites or separate local NLB clusters for each site.
  If using a local NLB, you can configure the NLB on the primary site with a different IP address than the NLB on the secondary site. This enables you to switch user requests to the IP address of the secondary site web servers if the primary site goes down.

Example configuration

Figure 1-1 shows an example disaster recovery configuration.

Figure 1-1 Example configuration

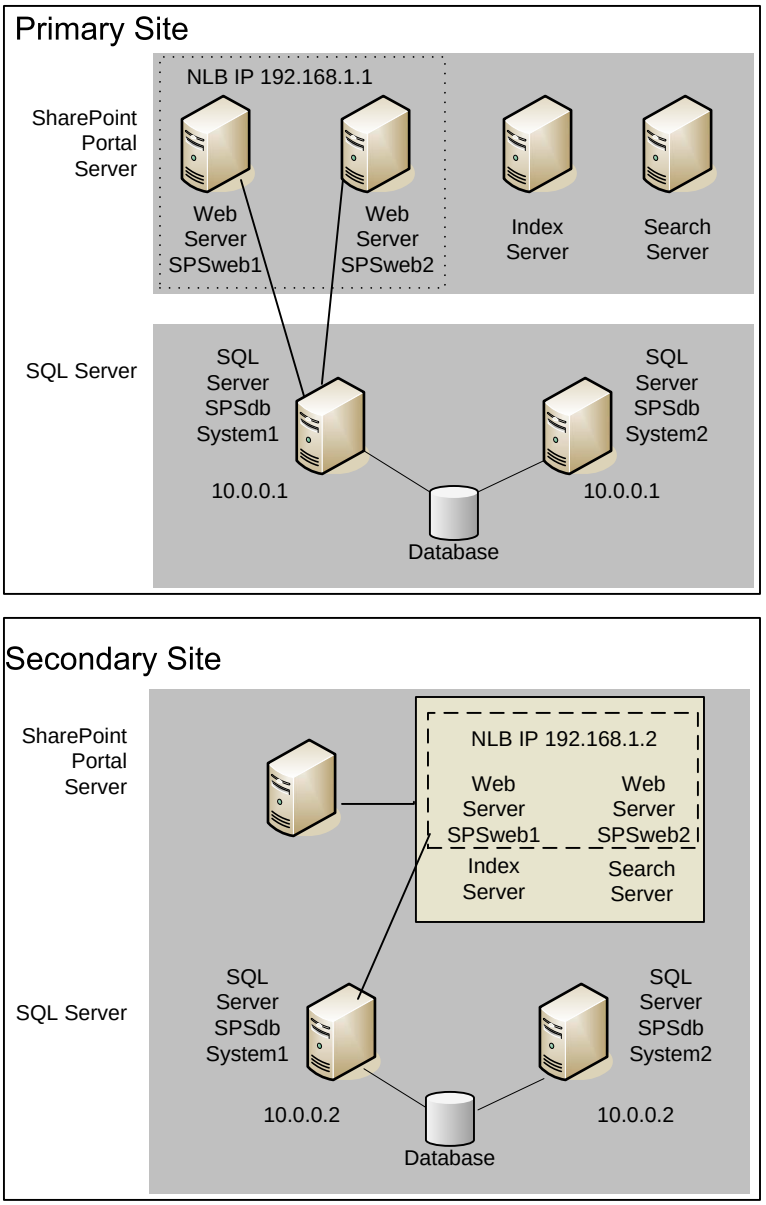

The illustrated example configuration is set up as follows:

Example primary site	An active/passive configuration of two machines running SQL Server (SYSTEM1 active and SYSTEM2 passive), with a virtual IP address of 10.0.0.1, both connected to the SQL database shared storage Four servers running SharePoint Portal Server 2003 components: ■ Two web servers using Microsoft Network Load Balancing (NLB) cluster service with an IP address of 192.168.1.1 The web servers connect to the active SQL Server instance. ■ An index management server ■ A search server
Example secondary site	An active/passive configuration of two machines running SQL Server (SYSTEM1 active and SYSTEM2 passive), with a virtual IP address of 10.0.0.2, both connected to the SQL database shared storage One server running the SharePoint Portal Server 2003 components: ■ Two web servers using Microsoft Network Load Balancing (NLB) cluster service with an IP address of 192.168.1.2 The web servers connect to the active SQL Server instance. ■ An index management server and a search server on the same machine

Replication considerations

VCS enables clustering and data replication for the SQL Server only. Any data stored locally on the SharePoint Portal Server web servers, index servers, and search servers rather than in the SQL database is not replicated. This non-replicated data includes the SharePoint index.

One way to provide search capabilities in the event of a disaster recovery scenario is to set up a scheduled propagation from the index server at the primary site to the search server at the secondary site. The search server at the secondary site can then handle the search requests if the primary site goes down.

For more information on index propagation, see the Microsoft documentation for SharePoint Portal Server 2003.

SharePoint server failover considerations

In a disaster recovery scenario, VCS brings the SQL Server service group online in the secondary site. VCS does not bring the SharePoint servers online. Therefore, you may prefer to maintain the secondary site SharePoint servers online but not in use until needed for disaster recovery.

You can configure VCS to perform a DNS update to switch user requests to the secondary site web servers if the primary site goes down.

See “DNS update considerations”.

DNS update considerations

When planning for disaster recovery in the SharePoint environment, you need to plan for updating IP addresses on the DNS server. You configure VCS to update the SQL Server virtual IP address on the DNS server when the remote site comes online. Optionally, if needed for your environment, you can configure VCS to update the web server/NLB virtual IP address on the DNS server.

About updating the SQL Server IP address

As part of the disaster recovery process of switching from the primary to secondary site, the DNS server must be updated with the site-specific virtual IP address for the SQL Server. Likewise, when switching back to the primary site, the DNS server must be updated again.

As an example, let’s assume that the SQL virtual server name is SPSdb. When the primary site is online, SPSdb is associated with the primary site virtual IP address, for example, 10.0.0.1. When the secondary site comes online, the DNS

server address list is updated so that SPSdb is associated with the IP address 10.0.0.2 (the secondary site virtual IP address).

You can configure VCS so that the update occurs automatically as part of the process of the SQL Server instance coming online.

See "Updating the SQL Server IP address" on page 31.

About updating the web server/NLB IP address

Multiple SharePoint web servers are typically set up on a network load balancer (NLB) cluster. You can use a hardware or software solution to implement the NLB cluster. You can implement a global NLB that includes both sites or a local NLB for each site.

If a local NLB is used, each NLB has a separate static IP address. When a site goes down in a disaster recovery scenario, user requests must be switched to the NLB at the secondary site. Therefore, in the DNS server, the virtual IP address associated with the NLB on the primary site must be updated with the virtual IP address for the NLB on the secondary site.

You can configure VCS so that the update occurs automatically as part of the process of the SQL Server instance coming online.

See "Updating the IP address for web requests" on page 32.

The web servers may take a few seconds to flush the cached IP address for the SQL Server and replace it with the new one. During that time a "cannot find content" message may be displayed in response to user requests.

Chapter 2

Configuring VCS with SharePoint Portal Server 2003

This chapter covers the following topics:

Tasks for a new installation of SQL Server

Setting up a new installation of SQL Server with SFW HA for high availability and disaster recovery is covered in the *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL*.

Some additional tasks are required to configure SFW HA for SharePoint, depending on your SharePoint environment.

The following table lists the full set of tasks so that you can use them as a checklist. The table also shows which tasks are specific to the SharePoint environment and where to find more information on the SharePoint related tasks and requirements.

Note: SFW HA provides a wizard that automates some disaster recovery configuration tasks on the secondary site. The task references in this table are based on configuring the secondary site manually rather than with the wizard.

Table 2-1 Tasks for a new installation of SQL Server

Objective	Tasks	Additional information for SharePoint
Reviewing the prerequisites	✔ Verifying hardware and software prerequisites	"Software requirements" on page 8 "Network configuration requirements" on page 12
Reviewing the SFW HA configuration	✔ Understanding active/passive configuration ✔ Reviewing the sample configuration	
Reviewing the SharePoint Portal Server configuration	✔ Reviewing SharePoint Portal Server configurations supported by SFW HA	"Supported SharePoint Portal Server configurations" on page 11
Configuring the network and storage on the primary site	✔ Setting up the storage hardware for a cluster environment ✔ Verifying the DNS entries for the systems on which SQL will be installed	

Table 2-1 Tasks for a new installation of SQL Server

Objective	Tasks	Additional information for SharePoint
Installing and configuring SFW HA on the primary site	✔ Verifying the driver signing option for the system ✔ Installing Veritas Storage Foundation HA for Windows (automatic installation) ✔ Selecting the option to install Veritas Cluster Server Enterprise Agent for Microsoft SQL Server and to install the GCO option	
Configuring cluster disk groups and volumes on the primary site	✔ Creating a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ✔ Creating dynamic volumes for the SQL system database, user databases and transaction logs using the Veritas Enterprise Administrator	
Configuring the cluster on the primary site	✔ Verifying static IP addresses and name resolution configured for each node ✔ Configuring cluster components using the Veritas Cluster Server Configuration Wizard ✔ Setting up secure communication for the cluster	
Installing and configuring SQL Server on the first node of the primary site	✔ Installing and configuring SQL Server ✔ Setting SQL Server services to manual start ✔ Configuring SQL services	

Table 2-1 Tasks for a new installation of SQL Server

Objective	Tasks	Additional information for SharePoint
Preparing to install SQL Server on the second or additional nodes of the primary site	✔ Stopping the SQL service ✔ Deporting the cluster disk group from the first node ✔ Importing the cluster disk group on an additional node ✔ Mounting the volumes (adding drive letters) ✔ Removing shared SQL files	
Installing and configuring SQL Server on the second or additional nodes of the primary site	✔ Installing SQL Server on the second or additional nodes	
Setting the internal name of the clustered instance on the primary site	✔ Setting the internal name of the clustered instance	
Configuring the VCS SQL Server service group on the primary site	✔ Creating a SQL Server service group using the VCS SQL Configuration wizard	
Creating a SQL Server user defined database on the primary site	✔ Creating volumes for a user-defined database and transaction log ✔ Creating a new user-defined database in SQL Server ✔ Adding resources for a user-defined database in VCS	
Verifying the SQL Server high availability configuration on the primary site	✔ Simulating failover ✔ Switching online nodes	

Table 2-1 Tasks for a new installation of SQL Server

Objective	Tasks	Additional information for SharePoint
Setting up the SharePoint servers on the primary site	✔ Installing and configuring SharePoint Portal Server 2003 on the secondary site	"Configuring SharePoint Portal Server 2003" on page 30
Modifying the SQL Server service group Lanman agent settings on the primary site	✔ Editing the attribute settings of the VCS Lanman agent resource to update the DNS server in a disaster recovery scenario	"Updating the SQL Server IP address" on page 31
Optionally, configuring the SQL Server service group to update the SharePoint web server/NLB IP address	✔ Customizing a VCS script configuration file for the primary site ✔ Editing the SQL service group to add a process resource for the script	"Updating the IP address for web requests" on page 32
Creating a parallel SFW HA environment on the secondary site	✔ Reviewing the prerequisites ✔ Reviewing the configuration ✔ Configuring the network and storage ✔ Installing SFW HA ✔ Configuring the cluster using the Veritas Cluster Server Configuration Wizard ✔ Configuring disk groups and volumes for SQL	
Installing and configuring SQL Server on the first node of the secondary site	✔ Installing and configuring SQL Server on the first node of the secondary site	

Table 2-1 Tasks for a new installation of SQL Server

Objective	Tasks	Additional information for SharePoint
Preparing to install SQL Server on additional nodes of the secondary site	✔ Stopping the SQL Service ✔ Deporting the cluster disk group from the first node ✔ Importing the cluster disk group on an additional node ✔ Adding drive letters ✔ Removing shared SQL files from the cluster disk group	
Installing SQL Server on additional nodes of the secondary site	✔ Installing SQL Server on additional nodes of the secondary site	
Setting the internal name of the clustered instance	✔ Setting the internal name of the clustered instance	
Creating the SQL service group configuration on the secondary site	✔ Creating the SQL service group configuration on the secondary site	
Configuring VVR and global clustering	✔ Configuring VVR components and global clustering	
Setting up the SharePoint servers on the secondary site	✔ Installing and configuring SharePoint Portal Server 2003 on the secondary site	"Configuring SharePoint Portal Server 2003" on page 30
Modifying the SQL Server service group Lanman agent settings on the secondary site	✔ Editing the attribute settings of the VCS Lanman agent resource to update the DNS server in a disaster recovery scenario	"Updating the SQL Server IP address" on page 31

Table 2-1 Tasks for a new installation of SQL Server

Objective	Tasks	Additional information for SharePoint
If using the scripts for updating web server IP address, installing and configuring the scripts for the secondary site	✔ Installing the script files in a location on the secondary site that matches the location on the primary site ✔ Customizing the script configuration file settings for the secondary site	"Updating the IP address for web requests" on page 32

Tasks for an existing installation of SQL Server

You can convert an existing standalone SQL Server 2000 site into an SFW HA high availability sites. Setting up high availability for an existing standalone SQL Server 2000 environment is covered in the *Veritas Storage Foundation and HA Solutions, HA and Disaster Recovery Solutions Guide for Microsoft SQL.*

A few additional tasks may be required to configure SFW HA for SharePoint, depending on your SharePoint environment.

Table 2-2 lists the full set of tasks so that you can use them as a checklist. The table also shows which tasks are specific to the SharePoint environment and where to find more information on SharePoint related tasks and requirements.

Note: SFW HA provides a wizard that automates some disaster recovery configuration tasks on the secondary site. The task references in this table are based on configuring the secondary site manually rather than with the wizard.

Table 2-2 Tasks for converting an existing SQL Server 2000 installation for SFW HA

Objective	Tasks	Additional information for SharePoint
Reviewing the prerequisites	✔ Verifying hardware and software prerequisites	"Software requirements" on page 8 "Network configuration requirements" on page 12

Table 2-2 Tasks for converting an existing SQL Server 2000 installation for SFW HA

Objective	Tasks	Additional information for SharePoint
Reviewing the SFW HA configuration	✔ Understanding active/passive configuration ✔ Reviewing the sample configuration	
Reviewing the SharePoint Portal Server configuration	✔ Reviewing SharePoint Portal Server configurations supported by SFW HA	"Supported SharePoint Portal Server configurations" on page 11
Configuring the network and storage	✔ Setting up the storage hardware for a cluster environmen ✔ Verifying the DNS entries for the systems on which SQL will be installed	
Preparing the standalone SQL Server	✔ Backing up existing data ✔ Setting SQL Server services to manual start	
Installing and configuring SFW HA	✔ Installing Veritas Storage Foundation HA for Windows (automatic installation) ✔ Configuring cluster disk groups and volumes ✔ Configuring the cluster using the Veritas Cluster Server Configuration Wizard	
Installing and configuring SQL Server on additional nodes	✔ Preparing for installation on additional nodes ✔ Installing SQL Server on additional nodes	

Table 2-2 Tasks for converting an existing SQL Server 2000 installation for SFW HA

Objective	Tasks	Additional information for SharePoint
Moving the existing SQL Server data files and user databases	✔ Stopping SQL Server service ✔ Modifying data file and user database locations	
Configuring the VCS SQL Server service group	✔ Creating a SQL Server service group using the VCS SQL Configuration wizard	
Configuring the SharePoint server connection to the SQL virtual server	✔ Configuring existing SQL Server clients to connect to the SQL virtual server name/instance name	"Configuring SharePoint Portal Server 2003" on page 30
Verifying the SQL Server high availability configuration on the primary site	✔ Simulating failover ✔ Switching online nodes	
Modifying the SQL Server service group Lanman agent settings on the primary site	✔ Editing the attribute settings of the VCS Lanman agent resource to update the DNS server in a disaster recovery scenario	"Updating the SQL Server IP address" on page 31
Optionally, configuring the SQL Server service group to update the SharePoint web server/NLB IP address	✔ Customizing a VCS script configuration file for the primary site ✔ Editing the SQL service group to add a process resource for the script	"Updating the IP address for web requests" on page 32

Table 2-2 Tasks for converting an existing SQL Server 2000 installation for SFW HA

Objective	Tasks	Additional information for SharePoint
Creating a parallel SFW HA environment on the secondary sitee	✔ Reviewing the prerequisites ✔ Reviewing the configuration ✔ Configuring the network and storage ✔ Installing SFW HA ✔ Configuring the cluster ✔ Configuring disk groups and volumes for SQ	
Installing and configuring SQL Server on the first node of the secondary site	✔ Installing and configuring SQL Server on the first node of the secondary site	
Preparing to install SQL Server on the second node of the secondary site	✔ Stopping the SQL Service ✔ Deporting the cluster disk group from the first node ✔ Importing the cluster disk group on an additional node ✔ Adding drive letters ✔ Removing shared SQL files from the cluster disk group	
Installing SQL Server on the second node of the secondary site	✔ Installing SQL Server on the second node of the secondary site	
Setting the internal name of the clustered instance	✔ Setting the internal name of the clustered instance	

Table 2-2 Tasks for converting an existing SQL Server 2000 installation for SFW HA

Objective	Tasks		Additional information for SharePoint
Creating the service group configuration on the secondary site	✔	Creating the SQL service group configuration on the secondary site using the VCS SQL Configuration wizard	
Configuring VVR and global clustering	✔	Configuring VVR components and global clustering	
Setting up the SharePoint servers on the secondary site	✔	Installing and configuring SharePoint Portal Server 2003 on the secondary site	"Configuring SharePoint Portal Server 2003" on page 30
Modifying the SQL Server service group Lanman agent settings on the secondary site	✔	Editing the attribute settings of the VCS Lanman agent resource to update the DNS server in a disaster recovery scenario	"Updating the SQL Server IP address" on page 31
	✔		
If using the scripts for updating web server IP address, installing and configuring the scripts for the secondary site	✔	Installing the script files in a location on the secondary site that matches the location on the primary site	"Updating the IP address for web requests" on page 32
	✔	Customizing the script configuration file settings for the secondary site	

Configuring SharePoint Portal Server 2003

For full information on installing SharePoint Portal Server 2003, see the Microsoft documentation for the product.

In addition, follow these guidelines when configuring SharePoint Portal Server 2003 on the primary and secondary sites:

- Installing the software
  If installing SharePoint in a SQL Server 2005 environment, see the Microsoft Knowledge Base article "How to install SharePoint Portal Server 2003 Service Pack 2 and Windows SharePoint Services Service Pack 2 in a deployment that uses SQL Server 2005 as the database back end": http://support.microsoft.com/kb/917446
- Setting up the SharePoint topology
  You configure all the SharePoint Portal Server components on the secondary site as part of the same server farm (topology) as the SharePoint components on the primary site.
- Configuring the SharePoint server connections to the database
  You configure the SharePoint servers to connect to the SQL database using the SQL virtual server name/instance.
  When you set up the SharePoint components at the secondary site, you configure the SharePoint servers at both sites to connect to the primary (active) site database.

Configuring the VCS SQL Server service group for the SharePoint environment

To create the VCS SQL Server service group on the primary site, follow the instructions in *Veritas Storage Foundation and HA Solutions HA and Disaster Recovery Solutions Guide for Microsoft SQL.*

See "Tasks for a new installation of SQL Server" on page 19 and "Tasks for an existing installation of SQL Server" on page 25.

After creating the service group, you edit the default configuration of the VCS SQL Server service group to automate updating IP addresses when you switch between sites. After creating the service group on the secondary site, you repeat the procedure for the secondary site.

See the following topics for details:

- Updating the SQL Server IP address
  You edit the service group to change the attribute settings of the VCS Lanman agent resource.

- Updating the IP address for web requests
  Optionally, depending on your environment, you edit the service group to add a process resource that implements a VCS script. You must customize the script configuration settings file separately for each site.

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 35.

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 Portal Server web servers or 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 35.

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 34.

See also "Example VCS configuration file entries (main.cf)" on page 37.

Requirements

To use the DNS update script files, download the required script and settings files from the Veritas Support Site:

http://entsupport.symantec.com/docs/283930

Download the following files:

- `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 34.

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 DNScmd.exe, which can be installed from the operating system Support Tools.

The scripts log to the engine log, `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.

See Table 2-3 for a description of the settings file contents.

Table 2-3 DNS update settings file

Keyword	Value	Notes
[web alias]	The web server (or NLB) name	Same in both setting files
[local ip]	Comma delimited pair of IP addresses: IP address for the web server or NLB on this site, IP address for the DNS server to be updated Example: 192.168.1.2, 192.168.10.10	When editing the primary site settings file, the local IP is that of the primary site web server or NLB. For the secondary site file, the local IP is that of the secondary site web server or NLB. If you have additional IP addresses for additional web servers or DNS servers, enter them as a comma delimited pair on separate lines.
[remote ip]	Comma delimited pair of IP addresses: IP address for the web server or NLB on the remote site, IP address of the DNS server to be updated Example: 192.168.1.1, 192.168.10.10	When editing the primary site settings file, the remote IP is that of the secondary site web server or NLB. For the secondary site file, the remote IP is that of the primary site web server or NLB. The DNS server to be updated is the one that manages the IP address for the web server or NLB. If you have additional IP addresses for additional web servers or DNS servers, enter them as a comma delimited pair on on separate lines.
[dns command]	Path to the location of DNScmd.exe Example: \Program Files\Support Tools	By default, the script will look for DNScmd.exe in \Program Files\Support Tools on the drive where SFW HA is installed, unless you specify another value.
[domain name]	Fully qualified domain of the web server Example: symantecdomain.com	Same in both settings files

Table 2-3 DNS update settings file

Keyword	Value	Notes
[nslookup command]	Full path for nslookup.exe Example: \Windows\System32\nslookup.exe	By default, the script will look for nslookup.exe on the drive where SFW HA is installed in the default directory shown, unless you specify another value.

Configuring a resource for the web servers

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

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

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

See *Veritas Cluster Server Administrator's Guide*.

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

In addition, ensure that DNScmd.exe is installed from the operating system Support Tools 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	The full path names of the following, in the order shown, separated by spaces: ■ The Perl script executable ■ The dnsupdate-offline script ■ The script settings file Example: c:\Program Files\Veritas\VRTSPerl\bin\perl.exe c:\bin\dnsupdate-offline.pl c:\bin\dnsupdate-settings.txt
MonitorProgram	The full path names of the following, in the order shown, separated by spaces: ■ The Perl script executable ■ The dnsupdate-monitor script ■ The script settings file Example: c:\Program Files\Veritas\VRTSPerl\bin\perl.exe c:\bin\dnsupdate-monitor.pl c:\bin\dnsupdate-settings.txt
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**.

Example VCS configuration file entries (main.cf)

The following is an example of the configuration entries created for SQL Server in the VCS `main.cf` file as a result of adding the process resource for the web server update script.

```
group SPS-SQL_Grp (
    SystemList = { Primary-Sys1 = 0, Primary-Sys2 = 1 }
    )

    IP SPS-SQL_Grp-IP (
        Address = "192.168.0.1"
        SubNetMask = "255.255.255.0"
        MACAddress @Primary-Sys1 = "00-11-33-55-77-99"
        MACAddress @Primary-Sys2 = "00-22-44-66-77-00"
        )

    Lanman SPS-SQL_Grp-Lanman (
        VirtualName = SPSDB
        IPResName = SPS-SQL_Grp-IP
        ADUpdateRequired = 1
        DNSCriticalForOnline = 1
        DNSOptions = { UpdateAll, PurgeDuplicate }
        )

    MSSearch SPS-SQL_Grp-MSSearch (
        AppName = "SQLServer$SPSDB"
        )

    MountV SPS-SQL_Grp-MountV (
        MountPath = "L:"
        VolumeName = DG1_Vol1
        VMDGResName = SPS-SQL_Grp-VMDg
        )

    NIC SPS-SQL_Grp-NIC (
        MACAddress @Primary-Sys1 = "00-11-33-55-77-99"
        MACAddress @Primary-Sys2 = "00-22-44-66-77-00"
        )

    RegRep SPS-SQL_Grp-RegRep-MSSQL (
        MountResName = SPS-SQL_Grp-MountV
        ReplicationDirectory = "\\RegRep\\SPS-SQL_Grp-RegRep-MSSQL"
        Keys = {
             "HKLM\\SOFTWARE\\Microsoft\\MSSQLServer\\Client" = "",
             "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\SPSDB" = "" }
        ExcludeKeys = {
```

```
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\SPSDB\\Setup",
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\SPSDB\\SQLServerAgent\\Subsystems",
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\SPSDB\\Tracking" }
        )

    RegRep SPS-SQL_Grp-RegRep-MSSearch (
        MountResName = SPS-SQL_Grp-MountV
        ReplicationDirectory =
"\\RegRep\\SPS-SQL_Grp-RegRep-MSSearch"
        Keys = {

"HKLM\\Software\\Microsoft\\Search\\1.0\\Applications\\SQLServer$SP
SDB" = "",

"HKLM\\Software\\Microsoft\\Search\\1.0\\CatalogNames\\SQLServer$SP
SDB" = "",

"HKLM\\Software\\Microsoft\\Search\\1.0\\Databases\\SQLServer$SPSDB
" = "",

"HKLM\\Software\\Microsoft\\Search\\1.0\\Gather\\SQLServer$SPSDB" =
"",
            "HKLM\\Software\\Microsoft\\Search\\1.0\\Gathering
Manager\\Applications\\SQLServer$SPSDB" = "",

"HKLM\\Software\\Microsoft\\Search\\1.0\\Indexer\\SQLServer$SPSDB"
= "" }
        )

    SQLServer2000 SPS-SQL_Grp-SQLServer2000 (
        Instance = SPSDB
        LanmanResName = SPS-SQL_Grp-Lanman
        MountResName = SPS-SQL_Grp-MountV
        DetailMonitor = 1
        Username = myuser
        Domain = mydomain.com
        Password = my_encrypted_password
        SQLFile = "c:\\Program Files\\Veritas\\cluster
server\\bin\\SQLServer2000\\sample_script.sql"
        )

    RVGPrimary SPS-SQL_Grp-RVGPrimary (
        RvgResourceName = SQL-RVG
        )

    Process SPS-SQL_Grp-Process (
        StartProgram =
"C:\\Progra~1\\Veritas\\VRTSPerl\\bin\\perl.exe
C:\\Bin\\prod_single_dns.pl C:\\Bin\\settings.txt"
```

```
        StopProgram =
"C:\\Progra~1\\Veritas\\VRTSPerl\\bin\\perl.exe
C:\\Bin\\dnsupdate-offline.pl C:\\Bin\\settings.txt"
        MonitorProgram =
"C:\\Progra~1\\Veritas\\VRTSPerl\\bin\\perl.exe
C:\\Bin\\dnsupdate-monitor.pl C:\\Bin\\settings.txt"
        UserName = myuser
        Password = my_encrypted_password
        Domain = "mydomain.com"
        )

    SPS-SQL_Grp-IP requires SPS-SQL_Grp-NIC
    SPS-SQL_Grp-Lanman requires SPS-SQL_Grp-IP
    SPS-SQL_Grp-MSSearch requires SPS-SQL_Grp-RegRep-MSSearch
    SPS-SQL_Grp-MSSearch requires SPS-SQL_Grp-SQLServer2000
    SPS-SQL_Grp-MountV requires SPS-SQL_Grp-RVGPrimary
    SPS-SQL_Grp-RegRep-MSSQL requires SPS-SQL_Grp-MountV
    SPS-SQL_Grp-RegRep-MSSearch requires SPS-SQL_Grp-MountV
    SPS-SQL_Grp-SQLServer2000 requires SPS-SQL_Grp-MountV
    SPS-SQL_Grp-SQLServer2000 requires SPS-SQL_Grp-RegRep-MSSQL
    SPS-SQL_Grp-SQLServer2000 requires SPS-SQL_Grp-Lanman
    SPS-SQL_Grp-Process requires SPS-SQL_Grp-SQLServer2000


    // resource dependency tree
    //
    //  group SPS-SQL_Grp
    //  {
    //  SPS-SQL_Grp-Process
    //      {
    //      MSSearch SPS-SQL_Grp-MSSearch
    //          {
    //          RegRep SPS-SQL_Grp-RegRep-MSSearch
    //              {
    //              MountV SPS-SQL_Grp-MountV
    //                  {
    //                  VMDg SPS-SQL_Grp-VMDg
    //                  }
    //              }
    //          SQLServer2000 SPS-SQL_Grp-SQLServer2000
    //              {
    //              MountV SPS-SQL_Grp-MountV
    //                  {
    //                  VMDg SPS-SQL_Grp-VMDg
    //                  }
    //              RegRep SPS-SQL_Grp-RegRep-MSSQL
    //                  {
    //                  MountV SPS-SQL_Grp-MountV
    //                      {
    //                      RVGPrimary SPS-SQL_Grp-RVGPrimary
    //                      }
    //                  }
```

```
//               Lanman SPS-SQL_Grp-Lanman
//                    {
//                    IP SPS-SQL_Grp-IP
//                        {
//                        NIC SPS-SQL_Grp-NIC
//                        }
//                    }
//               }
//          }
//     }
```