

Veritas™ Cluster Server for NetApp SnapMirror Installation and Configuration Guide for Microsoft SQL

Windows 2000, Windows Server 2003

5.0

Veritas Cluster Server for NetApp SnapMirror Installation and Configuration Guide for Microsoft SQL

Copyright © 2007 Symantec Corporation. All rights reserved.

VCS for NetApp SnapMirror 5.0

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

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

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

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

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 is a registered trademark of Microsoft Corporation.

Licensing and registration

VCS for NetApp SnapMirror is a licensed product. See the Licensing section in this guide for license installation instructions.

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	Introduction	
	About VCS enterprise agent for Network Appliance	13
	NetApp Filer agent	13
	NetApp SnapDrive agent	14
	NetApp SnapMirror agent	15
	About the VCS database agent for Microsoft SQL Server	16
	About the SQL Server 2000 agents	17
	Agent for SQL Server 2000	17
	Agent for MSSearch service	17
	About the SQL Server 2005 agents	18
	Agent for SQL Server 2005	18
	Agent for SQL Server 2005 agent service	18
	Agent for SQL Server 2005 analysis service	19
	About the agent for MSDTC service	19
	About SQL Server services	20
	MSSearch service	20
	Microsoft Data Transaction Coordinator (MSDTC) service	20
	Monitoring options	21
	How the agents make SQL Server highly available	21
	Local cluster configuration	21
	Disaster recovery configuration	22
	Running SQL Server in an active/active clustered environment	23
	Typical SQL Server 2000 configuration in a VCS cluster	24
	Typical SQL Server 2005 configuration in a VCS cluster	26
	Typical disaster recovery configuration	28
Chapter 2	Installing VCS for Network Appliance SnapMirror	
	Licensing	29
	Prerequisites	31
	Supported software	31
	Software requirements	32
	Configuration requirements	32
	System requirements	33
	Network requirements	33
	Permission requirements	34

Additional requirements	34
Best practices	34
Setting Windows driver signing options	35
Installing the software using the product installer	36
Installing the software from the command line	39
Example for setup.exe: client installation	41
Example for setup.exe: server installation	41
Configuring the cluster	42
Configuring Web Console	53
Configuring Notification	54
Configuring Wide-Area Connector process for global clusters	57

Chapter 3 Installing and configuring SQL Server

Configuring Microsoft iSCSI Initiator	59
Managing storage using Network Appliance Filer	61
Connecting virtual disks to the cluster node	62
Disconnecting virtual disks from the cluster nodes	62
Installing and configuring SQL Server 2000 on the first node	63
Prerequisites	63
Privileges requirements	63
Installing Microsoft SQL Server 2000 on the first node	64
Preparing to install SQL Server 2000 on the second node	67
Stopping the SQL Server 2000 Service	68
Disconnecting virtual disks from the first node	68
Connecting virtual disks to the second node	68
Renaming shared SQL Server 2000 files	68
Installing SQL Server 2000 on the second node	69
Installing SQL Server	69
Setting the internal name of the clustered instance	70
Configure VCS SQL service group.	72
Installing and configuring SQL Server 2005 on the first node	73
Installing Microsoft SQL Server 2005 on the first node	73
Preparing to install SQL Server 2005 on the second node	77
Stopping the SQL Server 2005 service	77
Disconnecting virtual disks from the first node	77
Connecting virtual disks to the second node	77
Renaming Shared SQL Server 2005 files	78
Installing SQL Server 2005 on the second node	79
Removing shared SQL Server 2005 files	79
Setting the internal name of the clustered instance	79
Configure VCS SQL service group.	82

Chapter 4	Configuring the VCS SQL Server service group	
	Prerequisites	84
	Configuring a SQL Server service group	85
	Running SnapManager for SQL	89
	Creating a SQL Server user-defined database	90
	Creating new volumes	90
	Creating a new SQL Server database	90
	Adding NetAppSnapDrive resource	91
	Verifying the service group configuration	93
	Bringing the service group online	93
	Taking the service group offline	93
	Switching the service group	94
	Administering a SQL Server service group	95
	Modifying a SQL service group configuration	95
	Deleting a SQL service group	96
Chapter 5	Configuring an MSDTC service group for High Availability	
	Prerequisites	97
	Reviewing the configuration	98
	Creating an MSDTC service group	101
	Creating an MSDTC client	103
	Verifying the installation	104
Chapter 6	Making a standalone SQL Server highly available	
	Prerequisites	106
	Reviewing the configuration	106
	Sample configuration	107
	Configuring the storage hardware and network	108
	Preparing the standalone SQL Server	110
	Backing up existing SQL data	110
	Setting SQL Server services to manual start	110
	Installing and configuring VCS for NetApp	111
	Installing and configuring SQL Server on additional nodes	111
	Moving the existing SQL Server data files and user databases	112
	Configuring the VCS SQL Server service group	113
	Creating a SQL Server user-defined database	113
	Verifying the installation	113

Chapter 7 Active/active configuration

Prerequisites	115
Reviewing the configuration	115
Sample configuration	116
Configuring the storage hardware and network	118
Installing VCS for NetApp SnapMirror and configuring the cluster	120
Configuring virtual disks and volumes for SQL Server	120
Installing and configuring the first instance of SQL Server	121
Configuring the VCS SQL Server service group for the first instance	122
Creating a SQL Server user-defined database	122
Repeating SQL Server installation for additional instances	123
Verifying the installation	124

Chapter 8 Disaster recovery configuration

What is a disaster recovery solution?	125
Why implement a disaster recovery solution?	125
Understanding replication	126
What needs to be protected in a SQL Server environment?	126
Typical disaster recovery configuration	127
Disaster recovery: New SQL Server 2000 installation	129
Reviewing the prerequisites	129
Reviewing the configuration	129
Installing VCS for NetApp	130
Configuring the cluster	130
Configuring storage on the NetApp Filer	130
Installing and configuring SQL Server 2000 at the primary site	130
Configuring the VCS SQL service group	131
Creating a parallel environment on the secondary site	132
Configuring DR components on the primary and secondary sites	133
Disaster recovery: New SQL Server 2005 installation	134
Reviewing the prerequisites	134
Reviewing the configuration	134
Installing VCS for NetApp	135
Configuring the cluster	135
Configuring storage on the NetApp Filer	135
Installing and configuring SQL Server 2005 at the primary site	135
Configuring the VCS SQL service group	136
Creating a parallel environment on the secondary site	137
Installing DR components on the primary and secondary sites	137
Configuring DR components	138
Configuring replication using Network Appliance SnapMirror	139


	Configuring SnapMirror resources at the primary site	140
	Configuring the Global Cluster Option for wide-area failover	142
	Prerequisites	142
	Linking clusters: Adding a remote cluster to a local cluster	143
	Converting a local service group to a global service group	144
	Bringing a global service group online	146
	Administering global service groups	147
	Deleting a remote cluster	148
Chapter 9	Removing VCS for NetApp SnapMirror	
	Prerequisites	153
	Removing the VCS configuration	153
	Uninstalling the agents	154
	Prerequisites	154
	Instructions	154
	License management	155
Chapter 10	Troubleshooting the agents	
	VCS logging	157
	VCW logs	158
	VCWsilent logs	159
	Network Appliance agents error messages	160
	Error messages and descriptions	162
	Agent for SQL Server 2000	162
	Agent for MSSearch	164
	Agent for MSDTC	165
	Agent for SQL Server 2005	166
	Agent for SQL Server 2005 Agent service	169
	Agent for SQL Server 2005 Analysis service	171
Appendix A	Resource type definitions	
	VCS enterprise agent for Network Appliance	174
	NetApp Filer agent	174
	NetApp SnapDrive agent	175
	NetApp SnapMirror agent	176
	SQL Server 2000 agents	178
	Agent for SQL Server 2000	178
	Agent for MSSearch service	182
	SQL Server 2005 agents	184
	Agent for SQL Server 2005	184
	Agent for SQL Server 2005 agent service	188
	Agent for SQL Server 2005 analysis service	190

	MSDTC agent	192
	Resource type definition	192
	Attribute definitions	192
	Sample dependency graph	193
Appendix B	Sample configurations	
	Sample configuration for a SQL Server 2000 cluster	195
	Sample configuration for a SQL Server 2005 cluster	200
Appendix C	Configuring the Symantec License Inventory Agent	
	About the Symantec License Inventory Manager	205
	When the Symantec License Inventory Agent is installed	206
	When the server and access points are installed	207
	What you can do with the agent	207
	How to remove the agent	207
Index		209

Introduction

The VCS database agent for Microsoft SQL Server provides high availability to SQL Server. The VCS enterprise agent for Network Appliance SnapMirror enables configuring Network Appliance filers in a VCS environment. Both the agents work together to provide high availability and disaster recovery to SQL Server in environments that use Network Appliance filers for shared storage. The agents also support disaster recovery configurations, set up using the VCS Global Cluster Option and Network Appliance SnapMirror for data replication. In a typical configuration, the agents are installed on each node in the cluster. The nodes are connected to the NetApp filers through a dedicated (private) storage network. VCS nodes are physically attached to the Network Appliance filer via an ethernet cable supporting iSCSI as the transport protocol.

Figure 1-1 Typical VCS configuration in a NetApp storage environment



This chapter provides an overview of the agents. For more information about the agents, including their VCS resource type definitions and attribute definitions, see “[Resource type definitions](#)” on page 173.

About VCS enterprise agent for Network Appliance

The VCS enterprise agent for Network Appliance provides failover support and recovery in environments employing Network Appliance filers for storage and SnapMirror for replication. The agent monitors and manages the state of replicated filer devices and ensures that at a time only one system has safe and exclusive access to the configured devices.

The agent can be used in local clusters, single VCS replicated data clusters, and multi-cluster environments set up using the VCS Global Cluster Option. The package contains three agents; each agent is described in subsequent sections:

- NetAppFiler agent—Monitors the state of the filer.
- NetAppSnapDrive agent—Connects and disconnects virtual disks (LUNs) using the iSCSI protocol.
- NetAppSnapMirror agent—Determines the role of the filer volumes with respect to replication and promotes a read-only snapmirrored volume to a read-write source volume during a wide-area failover.

NetApp Filer agent

The NetApp Filer agent monitors the state of the filer device. The agent is represented by the NetAppFiler resource type in VCS. NetAppFiler resources are persistent, meaning that they are not brought online or taken offline.

Agent operation

- Monitor—Verifies the state of the filer attached to the host by sending an ICMP ping command to the filer. If the filer does not respond, the agent reports the state of the filer as faulted.

NetApp SnapDrive agent

The NetApp SnapDrive agent monitors, connects, and disconnects filer volumes. The agent uses the iSCSI protocol.

Agent operations

- **Online**—Connects a virtual disk (LUN) using an iSCSI initiator. The agent presents the LUN as a locally-attached drive to the host. The agent also removes LUN-host mappings made before the online operation.
- **Offline**—Disconnects the virtual disk (LUN) from the host.
- **Monitor**—Verifies that the specified virtual disk (LUN) is connected to the host.
- **Open**—Verifies the iSCSI service is running on the filer. If the service is not running, the entry point starts the service.
- **Clean**—Attempts to forcibly disconnect a virtual disk (LUN).

NetApp SnapMirror agent

The NetApp SnapMirror agent monitors the replication state of filer devices. When a failover occurs, the agent reverses the direction of replication.

Agent operations

- **Online**—If the state of the local filer device is `SOURCE`, the agent creates a lock file to indicate that the resource can come online. This effectively makes the devices writable for the application.
If the state of the local filer is `TARGET`, the agent attempts to reverse the direction of replication by changing the state of the local filer to `SOURCE` and that of the original source to `SNAPMIRRORED`.
If the original source filer is down, the agent performs a mirror breakoff to enable local write access, if the filer is not already broken off.
If the original source returns to life, you must re-synchronize the data manually. The Online entry point touches a lock file if Read Write access is enabled successfully.
- **Offline**—Removes the lock file. The agent does not perform any filer operations because an offline entry point does not necessarily indicate an intention to give up the devices.
- **Monitor**—Verifies the lock file exists. If the lock file exists, the monitor entry point reports the status of the resource as online. If the lock file does not exist, the monitor entry point reports the status of the resource as offline.
- **Open**—Removes the lock file, thereby preventing potential concurrency violation if the group fails over to another node.

Note: The agent does not remove the lock file if the agent was started after an `hastop -force` command.

- **Clean**—Removes the lock file. No filer operations are performed since offlining this resource is no indication of a pending role swap.

About the VCS database agent for Microsoft SQL Server

The VCS database agent for Microsoft SQL Server provides high availability for Microsoft SQL Server 2000 and Microsoft SQL Server 2005 in a VCS cluster. The agent monitors Microsoft SQL Server RDBMS and its services on a VCS cluster to ensure high availability. The database agent also provides “active/active” support for SQL Server. VCS provides separate agents for SQL Server 2000 and SQL Server 2005.

Agents for SQL Server 2000

- Agent for SQL Server 2000 service. The agent provides high availability for SQL Server 2000 service.
- Agent for MSSearch service. The agent provides high availability for full-text search indices with a clustered SQL instance.

Agents for SQL Server 2005

- Agent for SQL Server 2005 service. The agent provides high availability for SQL Server 2005 service.
- Agent for SQL Server 2005 Agent service. The agent provides high availability for SQL Server 2005 agent service.
- Agent for SQL Server 2005 Analysis service. The agent provides high availability for SQL Server 2005 Analysis service.

The VCS database agent for Microsoft SQL Server also includes an MSDTC agent, which provides high availability for MSDTC service used in distributed transactions. The agent provides high availability for MSDTC service for SQL Server 2000 and SQL Server 2005.

About the SQL Server 2000 agents

The agents for SQL Server 2000 monitor specific resources within an enterprise application, determine the status of these resources, brings them online, and takes them offline.

Agent for SQL Server 2000

The agent brings the SQL Server 2000 service online, monitors the status, and takes it offline. The agent provides high availability for SQL Server 2000 in a clustered environment. The SQL Server 2000 agent monitors the SQL Server service and the SQL agent service. Specific agent operations include:

- Online—Brings the SQL Server 2000 service online.
- Offline—Takes the SQL Server 2000 service offline.
- Monitor—Queries the Service Control Manager (SCM) for the status of SQL Server 2000 services. See “[Monitoring options](#)” on page 21 for more information.
- Clean—Forcibly stops the SQL Server service.

Agent for MSSearch service

The agent brings the full-text search for a particular SQL instance online, monitors the status, and takes it offline. The agent provides high availability for full-text search index in a clustered environment. Specific agent operations include:

- Online—Brings the full-text search service for a particular instance online.
- Offline—Takes the full-text search service for a particular instance offline.
- Monitor—Monitors the full-text search service for a particular instance.
- Clean—Forcibly stops the MSSearch service for a particular instance.

Note: The MSSearch agent requires the SQL Server agent to be configured. Hence, you must configure MSSearch agent only on those cluster nodes that have SQL Server agent configured.

About the SQL Server 2005 agents

The agents for SQL Server 2005 monitor specific resources within an enterprise application, determines the status of these resources, brings them online, and takes them offline.

Agent for SQL Server 2005

The agent brings the SQL Server 2005 service online, monitors the status, and takes it offline. The agent provides high availability for SQL Server 2005 in a clustered environment. The SQL Server 2005 agent monitors the SQL Server service. Specific agent operations include:

- Online—Brings the SQL Server 2005 service online.
- Offline—Takes the SQL Server 2005 service offline.
- Monitor—Queries the Service Control Manager (SCM) for the status of SQL Server 2005 services. See “[Monitoring options](#)” on page 21 for more information.
- Clean—Forcibly stops the SQL Server service.

Agent for SQL Server 2005 agent service

The agent brings the SQL Server 2005 agent service online, monitors the status, and takes it offline. The agent provides high availability for SQL Server 2005 agent in a clustered environment. Specific agent operations include:

- Online—Brings the SQL Server 2005 agent service online.
- Offline—Takes the SQL Server 2005 agent service offline.
- Monitor—Monitors the SQL Server 2005 agent service.
- Clean—Forcibly stops the SQL Server 2005 agent service.

Agent for SQL Server 2005 analysis service

The agent brings SQL Server 2005 analysis service online, monitors the status, and takes it offline. The agent provides high availability for SQL Server 2005 analysis service in a clustered environment. Specific agent operations include:

- Online—Brings the SQL Server 2005 agent service online.
- Offline—Takes the SQL Server 2005 agent service offline.
- Monitor—Monitors the SQL Server 2005 agent service.
- Clean—Forcibly stops the SQL Server 2005 agent service.

Note: The VCS database agent for Microsoft SQL Server does not provide an agent for monitoring SQL Server 2005 MSSearch service. Instead, the SQL Server 2005 MSSearch service is monitored using a GenericService resource. Refer to *Veritas Cluster Server Bundled Agents Reference Guide* for more information about GenericService resources.

About the agent for MSDTC service

The MSDTC agent comprises two parts; MSDTC client and MSDTC server. The MSDTC client and the MSDTC server must not be configured on the same cluster node.

The MSDTC agent brings the MSDTC service online, monitors its status, and takes it offline. The agent provides high availability for an MSDTC service in a clustered environment. Specific agent operations include:

- Online—Brings the MSDTC service online.
- Offline—Takes the MSDTC service offline.
- Monitor—Monitors the MSDTC service.
- Clean—Forcibly stops the MSDTC service.

About SQL Server services

The VCS database agent for Microsoft SQL Server provides high availability for following services, MSSearch and MSDTC.

MSSearch service

SQL Server 2000 provides a full-text search index using the Microsoft Search (MSSearch) service. MSSearch is an indexing service that creates and manages full-text indices on specified columns within a database table. Full-text search creates an index entry for each word in the textual data for a specified column. Because each word in the textual data is indexed, queries looking for particular words become extremely fast.

Microsoft Data Transaction Coordinator (MSDTC) service

Microsoft Data Transaction Coordinator or the MSDTC service enables you to perform distributed transactions. A distributed transaction updates data on more than one computer in a network. The MSDTC service ensures that a transaction is successfully committed on each computer. A failure to commit on a single system aborts the transaction on all systems in the network.

If a transaction spans across more than one computer in the network, you must ensure that the MSDTC service is running on all the computers. Also, all the computers must be able to communicate with each other.

Monitoring options

The VCS database agent for Microsoft SQL Server provides two levels of application monitoring: basic and detail. Basic monitoring queries the SCM to verify whether the SQL Server services are continuously active. Detail monitoring updates a temporary table in the SQL Server database to verify the availability of the database instance. Refer to *VCS Database Agent for Microsoft SQL, Configuration Guide* for instructions on configuring detail monitoring for a SQL Server database.

If you start the SQL server services from outside VCS, then the SQL resource will go in an UNKNOWN state, because the VCS database agent for Microsoft SQL monitors the computer context of the services. If the SQL service is not started in the virtual server context the resource goes in an UNKNOWN state. You must ensure that you start all the SQL related services from within VCS.

Note: Detail monitoring can be configured only for the SQL Server 2000 and SQL Server 2005 agents.

How the agents make SQL Server highly available

The VCS database agent for Microsoft SQL Server detects an application failure if a configured virtual server becomes unavailable. The Network Appliance agents ensure consistent data access to the node on which SQL Server instances are running.

This section describes how the agents migrate SQL Server to another node in local clusters and in global disaster recovery configurations.

Local cluster configuration

When the VCS database agent for Microsoft SQL Server detects an application failure, the SQL Server service group is failed over to the next available system in the service group's system list. The configured SQL services and virtual server are started on the new system. The Network Appliance agents connect the virtual disks (LUNs) that contain the SQL Server data to the new node; thus ensuring continuous availability to SQL data.

Disaster recovery configuration

In a disaster recovery configuration, VCS first attempts to fail over the application to a node in the local cluster. If all nodes in the local cluster are unavailable, or if a disaster strikes the site, VCS attempts to fail over the application to the remote site. This involves the following steps:

- Connecting the virtual disks (LUNs) to the target hosts (using the NetAppSnapDrive agent)
- Performing a mirror break, which enables write access to the target (using the NetAppSnapMirror agent)
- Reversing the direction of replication by demoting the original source to a target, and begin replicating from the new source (using the NetAppSnapMirror agent)
- Starting the SQL services on the remote node (using the VCS database agent for SQL Server)

Running SQL Server in an active/active clustered environment

SQL Server allows up to 16 independent instances of SQL Server to run on a single machine. Using this feature, the VCS database agent for Microsoft SQL Server supports SQL Server in an active/active environment by allowing a node to run up to 16 instances. A SQL Server instance can fail over to any of the other configured nodes on its system list.

You can choose an active/active SQL Server configuration where several instances are intended to run on a single node. However, remember that you must configure failover nodes such that a single node can never host more than 16 instances.


Typical SQL Server 2000 configuration in a VCS cluster

A typical SQL Server 2000 configuration in a VCS cluster involves two cluster nodes accessing a shared storage. The SQL Server binaries are installed on the cluster nodes. The shared storage is used to store SQL Server data files and the MSDTC log files. The cluster nodes access the shared storage. The shared storage can be managed using Network Appliance suite of products.

The cluster nodes are configured to host the SQL Server 2000 and SQL Server 2000 MSSearch resource. The MSDTC resource can be configured on the same cluster nodes. You need not configure an MSDTC client if the MSDTC resource is configured on the same nodes that have SQL Server 2000 resource configured. However, if the MSDTC resource is configured on other nodes, you must configure an MSDTC client to point to the virtual server name of the MSDTC resource.

The following illustration depicts a two node cluster hosting a SQL Server service group with MSSearch service configured. MSDTC resource is also configured on the same nodes.

Figure 1-2 Typical SQL 2000 configuration in a VCS cluster




Typical SQL Server 2005 configuration in a VCS cluster

A typical SQL Server 2005 configuration in a VCS cluster involves two cluster nodes accessing a shared storage. The SQL Server binaries are installed on the cluster nodes. The shared storage is used to store SQL Server data files and the MSDTC log files. The cluster nodes access the shared storage. The shared storage can be managed using Network Appliance suite of products.

The cluster nodes are configured to host the SQL Server 2005 resource, the SQL Server 2005 agent service resource, the SQL Server 2005 MSSearch service resource, and the SQL Server 2005 analysis service resource. The MSDTC resource can be configured on the same cluster nodes. You need not configure an MSDTC client if the MSDTC resource is configured on the same nodes that have SQL Server 2005 resource configured. However, if the MSDTC resource is configured on other nodes, you must configure an MSDTC client to point to the virtual server name of the MSDTC resource.

The following illustration depicts a two node cluster hosting a SQL Server service group with the different services configured. MSDTC resource is also configured on the same nodes.


Figure 1-3 Typical SQL Server 2005 configuration in a VCS cluster



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 1-4 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.

Filer1 in the cluster on the primary site replicates to Filer2 in the cluster on the secondary site. Replication between the filers is set up using NetApp SnapMirror for SQL. Refer to NetApp documentation for more information on replication using NetApp filers.

If the Microsoft SQL Server server on System1 fails, SQL Server 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.

Installing VCS for Network Appliance SnapMirror

This chapter presents key considerations for software installation and describes how to install Veritas Cluster Server for Network Appliance SnapMirror.

Licensing

Licensing is based on the Microsoft Windows 2000 Server or Windows Server 2003 operating systems in use on a specific server. For each system that runs any of the Symantec products, you need a license.

Note: License keys for previous releases (4.2 or 4.3) are not supported in release 5.0 of VCS for NetApp SnapMirror.

Evaluation license key

An evaluation license key is embedded in the product. To use this key click Next at the license key entry screen of the product installer. This license key is valid for a limited evaluation period only.

You must purchase the product to obtain a permanent license key.

Virtual Server license policy

Each copy of the VCS for NetApp SnapMirror including all options and agents, whether used on a physical server or within a virtual machine must be separately licensed. Each Licensed Software license specifies the number of instances of the Licensed Software you may run on a particular server at one time.

License management

The product installer allows you to add and remove specific licenses. Adding a license for an option does not install the option. Use the Add/Remove function to install an option. License keys support installation on multiple systems.

Symantec License Inventory Manager

The Symantec License Inventory Manager is an enterprise asset management tracking tool. It determines all the Symantec software products and licenses being used in your network. The Symantec License Inventory Manager is available separately. To order a Symantec License Inventory Manager license and media kit, contact your Symantec sales representative.

See “[Configuring the Symantec License Inventory Agent](#)” on page 205.

Vxlicrep

Vxlicrep is a command line tool that generates a report of the licenses in use on your system.

To use Vxlicrep to display a license report:

Open the command prompt.

Enter **vxlicrep** without any options to generate a default report.

Optionally, use one of the following options to produce the type of report needed:

- g default report
- s short report
- e enhanced/detailed report
- h display this help

Prerequisites

Before you install VCS for NetApp SnapMirror, verify that your configuration meets the following criteria. Refer to the Hardware Compatibility List on the Symantec Support Web site at <http://entsupport.symantec.com> to determine the approved hardware for SFW HA.

Supported software

- Veritas Cluster Server (VCS) for Windows with the Veritas Cluster Server Database Agent for Microsoft SQL
- Microsoft SQL servers and their operating systems:
 - Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (SP4 required)
with
Windows Server 2003 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition (SP1 required for all editions)
or
Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition
or
Windows 2000 Server, Windows 2000 Advanced Server, or Windows 2000 Datacenter Server (all require Service Pack 4 with Update Rollup1)
 - Microsoft SQL Server 2005, (32-bit) Standard Edition or Enterprise Edition (SP1 required)
with
Windows Server 2003 (32-bit) Enterprise Edition or Datacenter Edition (SP1 required)
or
Windows Server 2003 R2 (32-bit) Standard Edition, Enterprise Edition, or Datacenter Edition

Note: Microsoft SQL Server 2000 and Microsoft SQL Server 2005 can exist in the same cluster and operate on or fail over to the same systems. However, only one default instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

- Network Appliance Data ONTAP 7.0.4

- Network Appliance SnapDrive 3.2
- Network Appliance SnapManager for SQL 2.0
- Microsoft iSCSI software Initiator 2.0

Software requirements

- The following software must be installed on each node:
 - Microsoft iSCSI Initiator
 - Network Appliance SnapDrive
When installing SnapDrive, you must specify a user account in the **SnapDrive Service Credentials** dialog box. The user account must be a domain user and part of the Administrators group of the local system and the filer.
 - Network Appliance SnapManager for SQL
- You must have licenses for Network Appliance SnapMirror and SnapRestore (for Disaster Recovery configurations).

Configuration requirements

- Make sure the filers reside in the same domain as the systems.
- Make sure the filers can replicate in both directions.
- The user account for the VCS Helper service must be a member of the administrators group on the filers.
- If you plan to set up a disaster recovery configuration, make sure the volumes at both sites are of the same size.
- Set up a private network for the shared storage. Assign an IP address to each adapter in the private network.
- Configure the Common Internet File System (CIFS) shares for each volume. See the Network Appliance documentation for instructions.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.

System requirements

- Shared disks to support applications that migrate between nodes in the cluster. Disaster Recovery configuration requires one array for each site.
- iSCSI Initiator supported NICs to access shared storage.
- Two NICs: one shared public and private, and one exclusively for the private network; Symantec recommends three NICs, see “[Best practices](#)” on page 34 for more information
- 1 GB of RAM for each system
- All servers running the same operating system.

Network requirements

- Install VCS for NetApp on servers in a Windows 2000 or Windows Server 2003 domain.
- Disable the Windows Firewall on systems running Windows Server 2003 SP1.
- Static IP addresses for the following purposes:
 - One IP address for each physical node in the cluster
 - One IP address for each cluster
 - One IP address for each virtual server
- Configure name resolution for each node
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.

Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in VCS because the Lanman agent uses DDNS to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This will enable the Lanman agent to refresh the resource records on the DNS servers. See the Lanman agent description in the *Veritas Cluster Server Bundled Agents Reference Guide* for more information.
- If the domain controller and the computer running the installation program are on different subnets, the installer may be unable to locate the computers selected for installation. In this situation, after the installation program displays an error message, type the host names and the IP addresses of the missing computers manually.

Permission requirements

- You must have administrative access to all systems selected for cluster operations. For this, add the domain user to the local Administrators group of each system.
- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.
- Domain Administrator or Account Operator privileges: if you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.

Additional requirements

- Installation media for all products and third-party applications
- Licenses for all products and third-party applications
- You must install the operating system in the same path on all systems. For example, if you install Windows 2003 on C : \WINDOWS of one node, installations on all other nodes must be on C : \WINDOWS. Make sure that the same drive letter is available on all nodes and that the system drive has adequate space for the installation.
- When installing, install only in a single domain.

Best practices

- Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Symantec recommends three network adapters (two NICs exclusively for the private network and one for the public network).
When using only two NICs, 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.
- Verify that you have set the Dynamic Update option for the DNS server to Secure Only.

Setting Windows driver signing options

Depending on the installation options you select, some Symantec drivers may not be signed. When installing on systems running Windows Server 2003, you must set the Windows driver signing options to allow installation.

[Table 2-1](#) describes the product installer behavior on local and remote systems when installing options with unsigned drivers.

Table 2-1 Installation behavior with unsigned drivers

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

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

To change the driver signing options on each system


- 1 Log on locally to the system.
 - 2 Open the Control Panel and click **System**.
 - 3 Click the **Hardware** tab and click **Driver Signing**.
 - 4 In the Driver Signing Options dialog box, note the current setting, and select **Ignore** or another option from the table that will allow installation to proceed.
 - 5 Click **OK**.
 - 6 Repeat for each computer.
- If you do not change the driver signing option, the installation may fail on that computer during validation. After you complete the installation, reset the driver signing option to its previous state.

Installing the software using the product installer

The product installer enables you to install VCS for NetApp SnapMirror. Ensure that you also select the option to install the Veritas Cluster Server Database Agent for SQL.

To install VCS for NetApp SnapMirror

- 1 From Windows Explorer, navigate to the root directory of the VCS for Network Appliance CD, and double-click **Setup.exe**.
- 2 From the Select a language drop-down list, select **English** and click **OK**.
- 3 On the product installation panel, click **VCS 5.0 for Network Appliance SnapMirror**.
- 4 Click Product Installation in the left pane, and then click **Complete/Custom** to install server components and optional client components.
Click **Administrative Console** if you want to install the client components.
- 5 Review the welcome message and click **Next**.
- 6 Read the License Agreement by using the scroll arrows in the view window. If you agree to the license terms, click **I accept the terms of the license agreement** and then click **Next**.
- 7 In the Enter a license key and click Add field, type the license key, click **Add**, and then click **Next**.
The License Key Details section displays more information about the selected key. To delete a license key, select the key from the key list and click **Remove**.
- 8 On the Option Selection panel, select the following product options and click **Next**:
 - **VCS 5.0 for NetApp Snap Mirror (Server)**
 - **Global Clustering Option** - if you plan to configure a disaster recovery environment.
 - **High Availability Hardware Replication Agents**
 - **High Availability Database Agents**
 - **VCS 5.0 for NetApp Snap Mirror (Client)** - This will install the VCS Java Console on the same nodes where the server components are installed.
- 9 On the Computer Selection panel, specify the domain and nodes for the installation and click **Next**.



- | | |
|-------------------------------------|---|
| Domain | Select the domain from the drop-down list. |
| Computer | <p>Select the computer name from the dropdown list.</p> <p>This may take some time to populate depending on the domain and network size, speed, and activity.</p> <p>Alternatively, if you know the computer name, type the name in the Computer field.</p> |
| Add | Click Add to add the computer for installation. |
| Selected computers for installation | The Selected computers for installation section displays the list of computers that you have added for installation. To remove a computer, click a computer name and click Remove . |
| Install Path | <p>The Install Path field displays the default path for installation. To specify a different installation path for a computer:</p> <ol style="list-style-type: none"> 1 Select the computer name from the Selected computers for installation list. 2 Type the path in the Install Path field. 3 Click Change. 4 Click Yes to confirm the changes. <p>To restore the default installation path for a computer:</p> <ol style="list-style-type: none"> 1 Select the computer name from the Selected computers for installation list. 2 Click Default. |

- 10 The installer validates the prerequisites for the selected computers and displays the results. Review the information and click **Next**.
If a computer fails validation, address the issue, and click **Validate Again**.
- 11 The Summary panel displays the installation settings summary. Review the information, click **Back** if you wish to make any changes, or click **Install** to begin the installation.
- 12 The Install/Uninstall panel displays the installation progress for the selected computers. After the installation is complete, click **Next**.
If the installation fails on any computer, the status screen shows a failed installation.
- 13 The Finish panel displays the installation report. Review the installation report and click **Next**.
If an installation fails, address the reason for failure. You may have to either repair the installation, or remove the software and re-install it.
- 14 Click **Reboot** to restart the remote computers.
- 15 After the remote computers have restarted, click **Next** and then click **Finish**.
- 16 Click **Yes** to restart the local computer.

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

To reset the driver signing options

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

Installing the software from the command line

You can perform a silent installation from the command prompt using the `Setup.exe` command. See the examples at the end of this section for reference.

To start the installation from the command window

- 1 Insert the software CD into a drive connected to the system.
- 2 Click **Start > Run**.
- 3 Enter **cmd** in the Open field and click **OK**.
- 4 From the command window, navigate to the root directory of the product CD.
- 5 Use the following command to install the software:

```
Setup.exe /s [INSTALL_MODE=InstallMode] [SOLUTION=Solution]  
LICENSEKEY="LicenseKey" [OPTIONS="a,b,c,..."]  
[INSTALLDIR="InstallDirPath"] [REBOOT=RebootMode]
```

Where the total length of the argument string is 512 characters.

Parameters for setup.exe

Information about the possible parameter values follows:

InstallMode

Set to indicate an installation or uninstallation.

1 = To install

5 = To uninstall

The default setting is **1** to install.

Example: **INSTALL_MODE=1**

Solution

Set to the type of installation.

6 = VCS for Network Appliance (Server Components)

7 = VCS for Network Appliance (Client Components)

The default setting is **6** for server components.

Example: **SOLUTION=6**

Note: To install the server and matching client components, run two `Setup.exe /s` commands, one with the `SOLUTION` parameter set for the server component and the other set for the client component. If you use a script to install the server and client, consider first installing the client and then the server, so that the script can reboot the system after server installation.

LicenseKey

Set the license key for the installation. Enter multiple keys by separating them with a comma (e.g.

123-345-567-789-123,321-543-765-789-321, etc.) The license key must start and end with a quotation mark (").

LicenseKey has no default setting.

Example: **LICENSEKEY="123-234-123-234-345"**

Options

Set the desired options, if any. The option must start and end with a quotation mark (").

There are no default settings.

The options for VCS for Network Appliance are:

- ***ntap***
- ***GCO***
- ***exchange***
- ***sql***

Example: **OPTIONS="GCO"**

InstallDirPath

Set the path of the installation directory. The path must start and end with a quotation mark (").

The default setting is *SystemDrive:\Program files\Veritas*

Example: **INSTALLDIR="C:\InstallationDirectory"**

RebootMode

Setting for the automatic reboot of the system at the completion of the installation.

0 = No reboot

1 = Reboot

The default setting is 0 for no system reboot.

Example: **REBOOT=1**

Note: Reboot the system at the end of installation to ensure that the software drivers for the server component are installed correctly. You do not have to reboot after installing the client.

Example for setup.exe: client installation

This command installs the client components at the specified installation path and tells the system not to reboot at the end of the installation.

```
Setup.exe /s INSTALL_MODE=1 SOLUTION=7  
INSTALLDIR="C:\InstallationDirectory"
```

Example for setup.exe: server installation

This command installs the server components with a license key of 123-234-123-234-345, and with the GCO option. Determines where the installation path is C:\InstallationDirectory and tells the system to reboot at the end of the installation.

```
Setup.exe /s INSTALL_MODE=1 SOLUTION=6  
LICENSEKEY="123-234-123-234-345" OPTIONS="GCO"  
INSTALLDIR="C:\InstallationDirectory" REBOOT=1
```

Configuring the cluster

After installing VCS for NetApp SnapMirror using the product installer, set up the components required to run a cluster. The VCS Configuration Wizard sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console, notification, and global clusters.


- If you plan to set up a disaster recovery environment, configure the wide-area connector process for global clusters.
- When configuring a user account for the VCS Helper service, choose the administrative account set up on the filer.

For instructions, see the chapter on Getting Started with VCS in the *Veritas Cluster Server Administrator's Guide*.

To configure a VCS cluster

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

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



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


- Clear the **Specify systems and users manually** check box.
- Click **Next**.

Proceed to [step 7](#) on page 44.

To specify systems and user names manually (recommended for large domains):


- Check the **Specify systems and users manually** check box.
Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
If you chose to retrieve the list of systems, proceed to [step 6](#) on page 44.
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 7](#) on page 44.


- 6
- On the System Selection panel, specify the systems to form a 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
- On the Cluster Configuration Options panel, click **Create New Cluster** and click **Next**.

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



Cluster Name Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.


Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster.

Warning: If you chose to specify systems and users manually in [step 4](#) or if you share a private network between more than one domain, make sure that the cluster ID is unique.

Operating System From the drop-down list, select the operating system that the systems are running.

Available Systems Select the systems that will be part of the cluster. 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.


- 9 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 8](#) on page 45, proceed to the next step. Otherwise, proceed to [step 11](#) on page 47.

- 10 On the Private Network Configuration panel, configure the VCS private network and click **Next**.



- Select the check boxes next to the two NICs to be assigned to the private network. Symantec recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
 - If you have only two NICs on a selected system, make sure you lower the priority of at least one NIC for that system. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - If your configuration contains teamed NICs, the wizard groups them as NIC Group #N where, N is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Symantec recommends that you do not select teamed NICs for the private network.
- 11 On the VCS Helper Service User Account panel, specify the name of a domain user context for the VCS Helper service. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper

Service user context to access the network. Do not use the Administrator account for the VCS Helper service.



- To specify an existing user, , do one of the following:
 - Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 43, 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**.

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

Do one of the following:

- To use the single sign-on feature



- Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

- If you know the name of the system that will serve as the root broker, click **Specify root broker system**, type the system name, and then click **Next**.

If you specify a cluster node, the wizard configures the node as the root broker and other nodes as authentication brokers. Authentication brokers reside one level below the root broker and serve as intermediate registration and certification authorities. These brokers can authenticate clients, such as users or services, but cannot authenticate other brokers. Authentication brokers have certificates signed by the root.

If you specify a system outside of the cluster, make sure that the system is configured as a root broker; the wizard configures all nodes in the cluster as authentication brokers.

- If you want to discover the system that will serve as root broker, click **Discover the root broker systems in the domain** and click **Next**. The wizard will discover root brokers in the entire domain, by default.

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition. For example, to search for systems managed by Administrator, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**.
- Click **Next**. The wizard discovers and displays a list of all the root brokers. Click to select a system that will serve as the root broker and then click **Next**.

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

- To use VCS user privilege

Veritas Cluster Server Configuration Wizard

Configure Security Service Option
Select security options for the cluster

☒ Domain Selection
☒ VCS Options
☒ Cluster Details
☒ Cluster Selection
☒ Validate Systems
☒ Edit Options
☒ NIC Selection
☒ Service Account
☒ Security
☐ Summary
☐ Finish

☐ Use Single Sign-on
 Configure the Symantec Product Authentication Service for secure communication in the cluster.

☒ Use VCS User Privileges
 Do not configure Symantec Product Authentication Service. You must configure a user with administrative privileges to the cluster.
 Provide the cluster administrator user information below.

User Name:
 Password:
 Re-enter Password:

Back Next Cancel

- Click **Use VCS User Privileges**. Accept the default user name and password for the VCS administrator account or type a new name and password.

The default user name for the VCS administrator is admin and the default password is password. Both are case-sensitive. Use this account

to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.

- Click **Next**.

- 13 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.

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


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

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

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

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

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




- Check the **Web Console** check box to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
- Check the **Notifier Option** check box to configure notification of important events to designated recipients.

Configuring Web Console

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

To configure Web Console

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




- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
 - If you choose to configure a new IP address, type the IP address and associated subnet mask.
 - Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.
- 2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started and click **Configure**.
 - 3 If you are finished with the configuration, click **Finish** to exit the wizard.

Configuring Notification

This section describes steps to configure notification.


To configure notification

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



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

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




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

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

Recipients	Severity
admin@example.com	Information

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

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




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

Configuring Wide-Area Connector process for global clusters

Configure the wide-area connector process only if you are configuring a disaster recovery environment.

To configure the wide-area connector process for global clusters

- 1 On the GCO Network Selection panel, specify the network information and click **Next**.



- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
 - If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
 - 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.
- 2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.
 - 3 Click **Finish** to exit the wizard.

Installing and configuring SQL Server

This chapter provides information for installing and configuring SQL Server. This environment uses an active/passive configuration with one to one failover capabilities.

To configure MSDTC service groups, see “[Configuring an MSDTC service group for High Availability](#)” on page 97. If you have existing SQL Server 2000/2005 deployment on a standalone server, see “[Making a standalone SQL Server highly available](#)” on page 105. If you are planning to deploy active/active configuration with multiple SQL instances, see “[Active/active configuration](#)” on page 115.

Configuring Microsoft iSCSI Initiator

The Microsoft iSCSI initiator enables communication between Windows systems and Network Appliance Filers. The initiator uses the iSCSI protocol to present the filer volume as a local block device to the system.

To configure Microsoft iSCSI initiator

- 1 Make sure the Microsoft iSCSI Initiator software version 2.0 is installed on all cluster nodes. Refer to Microsoft documentation for further information.
- 2 Start the Microsoft iSCSI initiator. Double-click the Microsoft iSCSI Initiator icon from the desktop.
- 3 Click the **Target Portals** tab, if not already selected.
- 4 Click **Add...**
- 5 In the Add Target Portals dialog box, specify the IP address or the DNS name for the Network Appliance Filer and click **OK**.
- 6 Click the **Available Targets** tab and click **Log On...**


- 7 In the Log On to Target dialog box, verify the target portal name and select the **Automatically restore this connection when the system reboots** check box.
- 8 Click **OK**.
- 9 Click the **Persistent Target** tab to verify that the newly added target portal is listed under the **Select a target** box.
- 10 Click **OK**.

Managing storage using Network Appliance Filer

Network Appliance manages data by creating volumes on physical disks. These volumes can further be divided into Logical Unit Numbers (LUNs). The LUNs are accessible from the cluster nodes, provided the nodes have Microsoft iSCSI Initiator and Network Appliance SnapDrive installed.

Note: Symantec does not support volumes created using qtree.

Figure 3-1 VCS cluster in a NetApp storage environment



The VCS enterprise agent for Microsoft SQL requires two LUNs to be created on the Network Appliance Filer, one for SQL Server data and the other for the registry replication information.

If you will be configuring an MSDTC service group, create additional volumes for MSDTC log, and MSDTC registry replication. These LUNs must be accessible from all cluster nodes.

Create the following volumes to manage your SQL Server database and logs:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB (minimum recommended size) volume for this purpose.

- INST1_DB1_VOL: contains the user database files
- INST1_DB1_LOG: contains the user database log files

Perform the following tasks to create LUNs on the Network Appliance Filer and to make them accessible from cluster nodes:

- Create volumes on the Network Appliance Filer.
- Share the volumes.
- Create LUNs on the shared volumes.

Refer to Network Appliance documentation for instructions on performing these tasks.

Connecting virtual disks to the cluster node

Once the virtual disks are created on the Network Appliance Filer, they must be connected (if not connected already) to the cluster nodes using Network Appliance SnapDrive.

To connect virtual disks to the cluster node

- 1 Start the Computer Management MMC on the cluster node where you want to connect the LUN. Click **Start > All Programs > Administrative Tools > Computer Management**.
- 2 From the left pane, expand **Storage** and double-click **SnapDrive**.
- 3 Double-click **Disks** to see the LUNs that are connected to the node.
- 4 Right-click the LUN you want to connect and then click **Connect Disk....**
- 5 In the Connect Disk alert box, click **OK**.

Disconnecting virtual disks from the cluster nodes

To disconnect virtual disks

- 1 Start the Computer Management MMC on the cluster node where you want to disconnect the LUN. Click **Start > All Programs > Administrative Tools > Computer Management**.
- 2 From the left pane, expand **Storage** and double-click **SnapDrive**.
- 3 Double-click **Disks** to see the LUNs that are connected to the node.
- 4 Right-click the LUN you want to disconnect and then click **Disconnect Disk....**
- 5 In the Disconnect Disk alert box, click **OK**.

Installing and configuring SQL Server 2000 on the first node

The following section describes how to install and configure SQL 2000 on the first cluster node.

Prerequisites

- Verify that VCS for NetApp SnapMirror is installed on all the nodes in the cluster.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW). See the *Veritas Cluster Server Administrator's Guide* for instructions.
- Verify that all nodes have adequate resources to run SQL Server 2000 or SQL Server 2005, and VCS for NetApp SnapMirror.
- Verify that the Microsoft iSCSI Initiator is configured to establish a persistent connection between the Network Appliance filer and the cluster nodes. See the Microsoft documentation for instructions.
- Make sure that you have an external, basic disk to create LUNs (or virtual disk). See “[Managing storage using Network Appliance Filer](#)” on page 61 for instructions.
- If your cluster has an Exchange service group configured, make sure to install SQL Server on a node, which is not in the SystemList attribute for the Exchange service group.

Privileges requirements

- You must be a domain user. The domain user must have administrator privileges.
- You must be a member of the Local Administrators group on all nodes on which Microsoft SQL will be installed.
- You must have write permissions for the Active Directory objects corresponding to these nodes.
- You must have write permissions on the DNS server to perform DNS updates.
- Domain Administrator or Account Operator privileges: if you plan to create a new user account for 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.

Installing Microsoft SQL Server 2000 on the first node

Before installing Microsoft SQL Server 2000 on the first cluster node, verify that the virtual disks on the filer are connected to the first cluster node. See [“Connecting virtual disks to the cluster node”](#) on page 62.

Install Microsoft SQL Server 2000 on the first node using the installation wizard provided with the product.


Install the SQL program files to a local disk and the SQL data files to the shared storage on the NetApp Filer. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2000 documentation for detailed installation information. Also refer to the Microsoft SQL Server 2000 documentation on the use of /PAE and /AWE switches if you are installing multiple instances of SQL and anticipate intensive memory use.


To install Microsoft SQL Server 2000

- 1 From the browser menu, select **SQL Server 2000 Components > Install Database Server**. Proceed with the installation steps.
- 2 If you are running Windows Server 2003, click **Continue** at the message that says SQL Server 2000 SP2 and below is not supported by this version of Windows. You will install SQL Server 2000 SP4 after installing SQL Server.
- 3 On the Welcome screen, click **Next**.
- 4 In the Computer Name panel, select **Local Computer** and click **Next**.
- 5 Proceed through the installation to the Installation Definition panel.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example INST1), and click **Next**.
Use the same instance name when installing this instance of SQL Server 2000 on failover nodes. If you are installing multiple instances of SQL in the cluster, each must have its own unique instance name.

- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for SQL Server program and data files.




- 9 In the Choose Folder dialog box, make the following selections and click **Next**:



- For Program Files, select a volume on the local disk.
- For Data Files, select the volume created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (\Program Files\Microsoft SQL Server) to remain. You must set the same path on all nodes.

- 10 In the Service Accounts panel, make the following selections:



- a Choose the **Customize the settings for each service** option.
 - b In the Services box, select the **SQL Server** option.
 - c In the Service Settings box, select **Use a Domain User account**.
 - d Specify the information for the user.
 - e Clear the **Auto Start Service** option.
 - f Repeat [step a](#) through [step e](#) for the SQL Server Agent option.
 - g Click **Next**.
- 11 Follow the wizard instructions to complete the installation.
- 12 When you have completed the initial installation of SQL Server 2000, install SQL Server 2000 SP4.
- 13 If you are running the Windows 2000 operating system, reboot when prompted. No reboot is necessary for Windows Server 2003.

Warning: Multiple instances of SQL Server 2000 must be installed in the same order on every node in the cluster.

Preparing to install SQL Server 2000 on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:


- [“Stopping the SQL Server 2000 Service”](#) on page 68
- [“Disconnecting virtual disks from the cluster nodes”](#) on page 62
- [“Connecting virtual disks to the cluster node”](#) on page 62
- [“Renaming shared SQL Server 2000 files”](#) on page 68

Stopping the SQL Server 2000 Service

Stop the SQL server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Click **Start > All Programs > Microsoft SQL Server > Service Manager** to open the SQL Server Service Manager.



- 2 Select the server to stop from the **Server** list.
- 3 Click **Stop**.
- 4 Click **Yes** in the SQL Service Manager dialog box to confirm that you do want to stop the service.

Disconnecting virtual disks from the first node

See [“Disconnecting virtual disks from the cluster nodes”](#) on page 62.

Connecting virtual disks to the second node

See [“Connecting virtual disks to the cluster node”](#) on page 62.

Renaming shared SQL Server 2000 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL

Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2000 on the second node

Follow the procedures provided in this section to install and configure SQL Server on additional nodes:

- [“Installing SQL Server”](#) on page 69
- [“Setting the internal name of the clustered instance”](#) on page 70
- [“Configure VCS SQL service group.”](#) on page 82

Installing SQL Server

Before installing Microsoft SQL Server 2000 on the second node, verify that the virtual disks on the filer are connected to the second cluster node. See [“Connecting virtual disks to the cluster node”](#) on page 62.

Install Microsoft SQL Server 2000 on additional nodes using the installation wizard provided with the product. See [“Installing Microsoft SQL Server 2000 on the first node”](#) on page 64 for the installation steps.

Removing shared SQL Server 2000 files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the Query Analyzer to set the internal name of the clustered instance to be the virtual server name.


Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do it from the last node, assuming that it is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.


To set the internal name of the clustered instance

- 1 Click **Start > All Programs > Microsoft SQL Server > Query Analyzer** to start the SQL Query Analyzer.
- 2 In the **Connect to SQL Server** window, provide connection information:



- In **SQL Server**, enter the SQL Server machine name in the format *System_Name\Instance_Name*. For example `SYSTEM2\INST1`.
- Select the **Start SQL server if it is stopped checkbox**.
- Enter valid user credentials and click **OK**.

3 Find the SQL Server name:



- In the upper pane of the query analyzer, enter the text “sp_helpserver”
 - Press F5.
 - Make note of the name listed in the lower pane, for example SYSTEM2\INST1. For a named instance, the name will be *System_Name\Instance_Name*. For a default instance, the name will be *System_Name*.
- 4 Delete the contents in the upper pane.
- 5 Disconnect the database:
- In the upper pane, enter the following:
“sp_dropserver ‘*System_Name\Instance_Name*.’”
where *System_Name\Instance_Name* is the name noted in step 3.
For example, for named instance:
“sp_dropserver ‘SYSTEM2\INST1.’”
For example, for a default instance:
“sp_dropserver ‘SYSTEM1.’”
 - Press F5.
- 6 Delete the contents in the upper pane.
- 7 Reconnect the database using the name of the virtual server:
- In the top pane, enter
“sp_addserver ‘*Virtual_Server_Name\Instance_Name*’,
local”
For example ‘INST1-VS\INST1’, local for a named instance, or
‘INST1-VS’, local for a default instance.

- Press F5.

Configure VCS SQL service group.

See “[Configuring the VCS SQL Server service group](#)” on page 83 for configuration steps.

Installing and configuring SQL Server 2005 on the first node

In preparation for installing Microsoft SQL Server 2005 ensure that the virtual disks on the filer are connected to the first cluster node. See “[Connecting virtual disks to the cluster node](#)” on page 62.

Complete the following procedures to install and configure Microsoft SQL Server 2005:

- Meet the prerequisites for installing SQL Server. See “[Prerequisites](#)” on page 63.
- Install SQL Server 2005 on the first node. See “[Installing Microsoft SQL Server 2005 on the first node](#)” on page 73.

Installing Microsoft SQL Server 2005 on the first node

Install Microsoft SQL Server 2005 on the first node using the installation wizard provided with the product.

Before you begin installing SQL Server 2005, verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

Install the SQL program files to a local disk and the SQL data files to the shared storage managed by the cluster disk group. As you progress through the installation, use the following guidelines to create an installation that will function properly in your environment.

Note: Only the portions of the SQL Server installation procedure relevant to the SFW environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install Microsoft SQL Server 2005


- 1 Navigate to the installation directory and launch **splash.hta**.
- 2 Under the Install section, click **Server components, tools, Books Online, and samples**.
- 3 Continue with the installation, following the instructions in the Microsoft SQL Server 2005 Setup Wizard.
Complete the SQL Server Component Update, System Configuration Check, and Registration Information panels.

- 4 In the Components to Install panel, select SQL Server Database Services and Workstation Components and optionally select any of the other components to install.

If you install optional components on one node, install the same components in the same order on other nodes.


- 5 Click **Advanced** and in the Feature Selection panel, specify the path for SQL Server data files and other selected services. Set the data files to the shared storage managed by the cluster disk group, as follows:

- Expand **Database Services**, click **Data Files**, and click **Browse**.




- In the Change Folder panel, set the installation path to the drive letter and location of the volume that was created for the SQL Server system data files (INST1_DATA_FILES). You can allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.

You must set the same path on all nodes.



- If you selected the Analysis Services option in [step 4](#), expand **Analysis Services**, click **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files. Click **OK** when done.
- 6 In the Feature Selection panel, expand **Client Components** and ensure that Connectivity Components and Management tools are selected. Click **Next**.
- 7 In the Instance Name panel, enter an instance name (for example, INST1) and click **Next**.
 Use the same instance name when you install this instance of SQL Server 2005 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name.
- 8 In the Service Accounts panel, make the following selections and click **Next**:

- Select **Use a domain user account**.



- Specify the user, password, and domain.
 - Clear all the check boxes that start services except for the SQL Browser, so that the SQL Server is not brought online.
- 9 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 10 Install any SQL service packs or hotfixes if required.
 - 11 Set SQL services to manual start.

Preparing to install SQL Server 2005 on the second node

Follow the procedures provided in this section before installing SQL Server on additional nodes:

- [“Stopping the SQL Server 2005 service”](#) on page 77
- [“Disconnecting virtual disks from the cluster nodes”](#) on page 62
- [“Connecting virtual disks to the cluster node”](#) on page 62
- [“Renaming Shared SQL Server 2005 files”](#) on page 78

Note: These procedures must be performed for every node that is intended to be a part of the cluster.

Stopping the SQL Server 2005 service

Stop a running SQL Server service on the configured node so the databases on the shared disk can be manipulated by the installation on the second node.

To stop the SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Stop**.
- 4 Repeat for all other SQL Server services that are running on the server.
- 5 Exit the SQL Server Configuration Manager.

Disconnecting virtual disks from the first node

See [“Disconnecting virtual disks from the cluster nodes”](#) on page 62.

Connecting virtual disks to the second node

See [“Connecting virtual disks to the cluster node”](#) on page 62.

Renaming Shared SQL Server 2005 files

Before installing SQL on the second node, open the SQL Server system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the installation of SQL Server on the additional nodes.

If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second node SQL Server installation. Once the installation completes successfully, you can then delete the renamed folder and its contents.

To rename shared SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the SQL Server data files.
- 2 Rename the folder that contains the SQL Server data files.

Installing SQL Server 2005 on the second node

Before installing Microsoft SQL Server 2005 on the second node, verify that the virtual disks on the filer are connected to the second cluster node. See [“Connecting virtual disks to the cluster node”](#) on page 62.

Install Microsoft SQL Server 2005 on additional nodes using the installation wizard provided with the product. See [“Installing Microsoft SQL Server 2005 on the first node”](#) on page 73 for the installation steps.

Repeat the procedures described in [“Preparing to install SQL Server 2005 on the second node”](#) on page 77 and [“Installing SQL Server 2005 on the second node”](#) on page 79 on any additional nodes.

Removing shared SQL Server 2005 files

If you renamed the shared SQL Server folder from the SQL Server system data files volume prior to installing SQL on the second node, you can delete the renamed folder and files now.

To delete the renamed SQL Server data files

- 1 On the computer on which the cluster disk group is imported, open the volume that holds the renamed SQL Server data files.
- 2 Select the renamed SQL Server data folder and files contained in the volume, and press Delete.

Setting the internal name of the clustered instance

Use the SQL Server Management Studio to set the internal name of the clustered instance to be the virtual server name\instance name (for example, `INST1-VS\INST1`).

Note: Do this procedure after you install and configure SQL Server on the last node for this instance, so that you need to do the procedure only once. Do the procedure from the last node, assuming that the node is still connected to the shared volumes.

The virtual server name you assign must be unique within the cluster. The virtual server name\instance name is used by the SQL Server clients to access the database. You specify the virtual server name again when configuring the VCS SQL service group for this instance.

Warning: For a disaster recovery configuration, the virtual server name on the secondary site cluster must match the one on the primary site cluster.


Before you set the internal name of the instance, start the SQL Server services on the node that is currently connected to the shared volumes.

To start a SQL Server service

- 1 Start the SQL Server Configuration Manager (**Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager**).
- 2 In the left pane, click **SQL Server 2005 Services**.
- 3 In the right pane, right-click the SQL Server instance and select **Start**.
- 4 Repeat for all other SQL Server services that are not running on the server.
- 5 Exit the SQL Server Configuration Manager.

To set the internal name of the clustered instance


- 1 Start the SQL Server Management Studio (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 In the **Connect to Server** window, provide connection information:



- Select the Database Engine from the server type drop down list.
 - Enter the name in the format *System_Name\Instance_Name*.
 - Select the appropriate authentication method.
 - Enter valid user credentials if using Domain authentication and click **Connect**.
- 3 Find the SQL Server name as follows:
 - Right-click the instance in the Object Explorer and click **New Query**.
 - In the right pane of the SQL Server Management Studio, enter the query text:

sp_helpserver

- Press **F5**. The right pane divides into an upper and lower pane.



- Make note of the name listed in the lower pane, which is in the format *System_Name\Instance_Name*, for example, SYSTEM1 \ INST1. (For a default instance, you see only *System_Name*.)
- 4 Delete the contents in the upper pane.
 - 5 Disconnect the database as follows:
 - In the upper pane, enter the following:
sp_dropserver "System_Name\Instance_Name"
 where **System_Name\Instance_Name** is the name noted in [step 3](#).
 For example, for a named instance:
 sp_dropserver "SYSTEM1 \ INST1"
 For example, for a default instance:
 sp_dropserver "SYSTEM1"
 - Press **F5**.
 - 6 Delete the contents in the upper pane.
 - 7 Reconnect the database using the name of the virtual server:
 - In the top pane, enter the following:
sp_addserver "Virtual_Server_Name\Instance_Name", local
 For example, for a named instance:
 sp_addserver "INST1-VS\INST1", local
 For example, for a default instance:
 sp_addserver "INST1-VS", local
 - Press **F5**.

- 8 Exit the SQL Server Management Studio.

Configure VCS SQL service group.

See “[Configuring the VCS SQL Server service group](#)” on page 83 for configuration steps.

Configuring the VCS SQL Server service group

Configuring the SQL Server service group involves creating resources for the NetApp and SQL resources. VCS provides several ways of configuring a service group, which include the service group configuration wizard, Cluster Manager (Java Console), and the command line. This chapter provides instructions on configuring a SQL service group using the SQL Server Configuration Wizard. The SQL Server Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

Prerequisites

- Verify that VCS for Network Appliance SnapMirror, along with the VCS database agent for SQL Server, is installed on all the cluster nodes.
- Verify that SQL Server is installed and configured identically on all the cluster nodes.
- Verify that you have configured a VCS cluster using VCS Configuration Wizard (VCW).
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be an Administrator for the NetApp Filer containing the LUNs created to store SQL Server components.
- Verify that the Veritas High Availability Daemon (HAD) is running on the system from where you run the wizard.
- Verify the virtual disks (LUNs) created to store the following data components are connected to the node where you run the wizard and disconnected from other nodes in the cluster.
 - SQL Server system data files
 - Registry replication information.
 - User database files
 - User database log files
- Verify the drive containing the SQL Server 2000 or SQL Server 2005 system data files and registry replication information, is mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- Verify the SQL Server 2000 or SQL Server 2005 instance is installed identically on all nodes that will participate in the service group.
- Assign a unique virtual server name to SQL Server. While it must be unique within the cluster, it must be the same as the virtual name of the SQL Server in the primary site.


Configuring a SQL Server service group

This section describes how to create a new SQL service group. To modify an existing service group, see “[Modifying a SQL service group configuration](#)” on page 95.

Symantec recommends reviewing the resource types and the attribute definitions of the agents before configuring the agents. You can find this information in the appendix. See “[Resource type definitions](#)” on page 173. For sample configuration files and resource dependency graphs of the SQL service groups, see “[Sample configurations](#)” on page 195.

To create a SQL Server service group on the cluster


- 1 Ensure that you have stopped the SQL Server service for the instance.
- 2 Start the SQL Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 3 In the Select Configuration Option panel, choose **MS SQL Server Service Group Configuration** and **Create**, and click **Next**.
- 4 Verify that you have met the prerequisites listed and click **Next**.
- 5 Specify the service group name and system list:



- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.

- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - Click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that need to be configured for high availability in your environment. Clear the **Configure NetApp SnapMirror Resource(s)** check box. This option is applicable only in case of a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. If you are setting up a disaster recovery environment, check this check box to configure the SnapMirror resource at the primary site. Note that you must configure the SnapMirror resource only after you have configured the cluster at the secondary site. Click **Next**.
 - 7 The User Databases List panel summarizes the databases on this instance of SQL. Click **Next**.

- 8 In the Detail Monitoring Configuration panel, optionally enable a monitoring script as follows:



- Select the check box for the SQL Server instance for which detail monitoring will be configured. Only the instances selected in [step 6](#) on page 86 are available for selection.
 - Specify the fully qualified user name and password for connecting to SQL Server database. Make sure the specified user has SQL Server log on permissions.
 - If the path of the script is same on all nodes, choose the **Global** option, click the **SQL Monitor Script** text box, and specify the path to the script on the first system displayed in the **System Name** list. If the path of the script is different on all nodes, choose the **Per System** option, and specify the path for the script on each node. Make sure the specified path exists on all the systems in the cluster.
 - Select the **Fail over service group if detail monitoring script execution fails** checkbox, if not already selected. This will enable the SQL agent to fail over the service group if the detail monitoring script execution fails.
 - Click **Apply**.
- 9 If you want to configure detail monitoring for additional instances, repeat [step 8](#) for all the instances for which detail monitoring will be configured.
 - 10 Click **Next**.

- 11 In the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**. Symantec recommends that RegRep resources and SQL data be in separate volumes.
- 12 Configure the virtual server as follows:

SQL Server Agent Configuration Wizard

Virtual Server Configuration
Enter a virtual server name for the application and specify the virtual IP information.

✓ Prerequisites
✓ System List
✓ Instance Selection
✓ Databases
✓ Detail Monitoring
✓ Registry Replication
▶ Virtual Computer

Summary
Configuration Status

Virtual Server Name:

Virtual IP Address:

Subnet Mask:

Specify the adapter to be used on each system.

System Name	Adapter Display Name
PC177VM-3	Public
PC177VM-4	Public

Advanced Settings...

< Back Next > Cancel

- Enter the virtual name for the server, for example INST1-VS. Ensure that the virtual server name you enter is unique in the cluster. It is the same as the virtual server name specified when setting the internal name of the clustered instance.
- Enter a unique virtual IP address that is currently not being used on your network, but is in the same subnet as the current computer.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Select the **Adapter Display Name** field to view the adapters associated with a system.
The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** checkbox, and select the Organizational Unit from the drop-down list.
- Click **Next**.

- 13 In the Service Group Summary, review the service group configuration. The **Resources** box lists the configured resources. Click a resource to view its attributes and their configured values in the **Attributes** box.
- 14 The wizard assigns unique names to resources based on their respective name rules. Optionally, change the names of the resources, if desired.
 - To edit a resource name, click the resource name or press the F2 key. Press Enter after editing each resource name.
 - To cancel editing a resource name, press Esc.
- 15 Click **Next** and when prompted to confirm creating the service group, click **Yes**. Messages indicate the status of the commands.
- 16 Complete the SQL Server service group configuration:
 - In the **Bring the service group online** check box, if you want to bring the service group online later, clear the check box.
You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.
 - Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.
The wizard marks all the resources in the service group as **CRITICAL**. If desired, use Cluster Manager (Java Console) or the command line to change the state.

To configure an MSDTC service group, see [“Configuring an MSDTC service group for High Availability”](#) on page 97.

Running SnapManager for SQL

After configuring the service group, you may want to run the SnapManager Configuration Wizard on the node on which the service group is online, to schedule backups of SQL Server database. You must adhere to the following requirements while running SnapManager for SQL:

- Make sure the SQL service group is online.
- Do not move the SQL Server database components.

If you are scheduling backups in a VCS cluster, schedule them on the node on which the service group is online. If the SQL service group fails over to another node, you must set up the backup schedule again on the new node.

See the Network Appliance documentation for more information about running SnapManager for SQL.

Creating a SQL Server user-defined database

The following tasks enable you to use VCS for NetApp to create and manage a SQL Server user-defined database.

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Server Configuration Wizard to add the NetAppFiler and NetAppSnapDrive resources for the user databases.

Creating new volumes

If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them now. In the sample deployment these volumes are named:

- INST1_DB1_VOL: contains a user-defined database file
- INST1_DB1_LOG: contains a user-defined database log file

Creating a new SQL Server database

Create a new SQL Server database and point the database files and transaction log to the new volumes created for them.

To create a new SQL Server 2000 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server > Enterprise Manager**).
- 2 Right-click on **Databases** and then click **New Database**.
- 3 In the New Database page, enter a name for the new database.
- 4 Click the browse button (...) in the Location column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 5 Choose other file properties as desired.
- 6 Click the **Transaction Log** tab.
- 7 Click ... (ellipsis button) in the Location column and browse to the location of the volume you created for the transaction log, and click **OK**.

To create a new SQL Server 2005 database

- 1 Open SQL Server Database Manager (**Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio**).
- 2 Expand the icon associated with your server.
- 3 Right-click on **Databases** and then click **New Database**.
- 4 In the New Database page, type a name for the new database.
- 5 Click ... (ellipsis button) in the Path column, browse to the location of the volume where you want to create your user database, and click **OK**.
- 6 Select and edit other file properties as desired.
- 7 Click ... (ellipsis button) in the Path column for the Transaction Log row and browse to the location of the volume you want to create for the transaction log, and click **OK**.
- 8 To add more data files if required:
 - Select **Add**.
 - Edit the properties in the new data file rows as required.
- 9 Click **OK**.

Adding NetAppSnapDrive resource

Before running the SQL Server Agent wizard to add the NetAppSnapDrive resource:

- Make sure the SQL service group is online.
- Make sure the volumes for the user database and transaction logs are mounted.

Note: NetAppSnapDrive resource is required only if the database is created on a new volume.

To add NetAppFile and NetAppSnapDrive resources using the SQL Server Configuration Wizard

- 1 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Select the **MS-SQL Server Service Group Configuration**, select the **Edit** option, and click **Next**.

- 3 Review the Prerequisites page and click **Next**.
- 4 In the Service Group Selection page, select the service group and click **Next**.
- 5 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 6 In the Service Group Configuration page, click **Next**.
- 7 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 8 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**.
- 9 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 10 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 11 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
- 12 To complete the user database configuration:
 - Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.
 - Click **Next** to configure another SQL service group or an MSDTC service group.

To configure an MSDTC service group, see “[Configuring an MSDTC service group for High Availability](#)” on page 97.

Verifying the service group configuration

Failover simulation is an important part of configuration testing. This section provides steps to verify the SQL Server service group configuration by bringing the service group online, taking the service group offline, and switching the service group to another cluster node.

Bringing the service group online

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

To bring a service group online from the Java Console

- 1 In the Cluster Explorer configuration tree, select the SQL service group to be taken online.
- 2 Right-click the service group name, and select **Enable Resources**. This enables all resources in the service group.
- 3 Right-click the service group name, and select the system on which to enable the service group. (Right-click > Enable > *system_name* or Right-click > Enable > All)
- 4 Save your configuration (**File > Close Configuration**).
- 5 Right-click the service group and select to online the service group on the system. (Right-click > Online > *system_name*)

Taking the service group offline

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

To take a service group offline from the Java Console

- 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

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

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

To switch service groups

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

If there is more than one service group, you must repeat this step until all the service groups are switched.
- 2 Verify that the service group is online on the node you selected to switch to in [step 1](#).
- 3 To move all the resources back to the original node, repeat [step 1](#) for each of the service groups.

To shut down an active cluster node

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

The service group you selected is taken offline and brought online on the node that you selected.

Administering a SQL Server service group

You can dynamically modify the SQL service group configuration in several ways, including the SQL Server Configuration Wizard, Cluster Manager (Java Console), Cluster Manager (Web Console), and the command line. The following steps describe how to modify the service group using the SQL Server Configuration Wizard.

Modifying a SQL service group configuration

Prerequisites

- If the SQL Server service group is online, you must run the wizard from a node on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV and VMDg resources for the service group should be online on the node where you run the wizard and offline on all other nodes.
- If you are running the wizard to remove a node from the service group's system list, do not run the wizard on the node being removed.
- If you are running the wizard to add or remove MountV resources for user defined databases, make sure the service group is online.

To modify a SQL Server service group

- 1 Start the SQL Server Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 In the Select Configuration Option panel, click **MS SQL Server - Service Group Configuration**, click **Edit**, and then click **Next**.
- 3 Review the prerequisites and click **Next**.
- 4 In the Service Group Selection panel, select the service group to modify and click **Next**.
- 5 In the Service Group Configuration panel, add or remove systems from the service group's SystemList and click **Next**.
- 6 In the SQL Server Instance Selection panel, select the SQL Server instance to be made highly available and click **Next**.

- 7 In the User Databases List panel, verify the master and user defined databases configured for the SQL instance. The wizard will create MountV resource for each database. Click **Next**.
- 8 Follow the wizard instructions and make desired modifications to the service group configuration.

Deleting a SQL service group

The following steps describe how to delete a SQL Server service group using the configuration wizard.

To delete a SQL Server service group

- 1 Start the SQL Server Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 In the Select Configuration Option dialog box, choose the **MS-SQL Server - Service Group Configuration** option, choose the **Delete** option, and click **Next**.
- 3 Review the prerequisites and click **Next**.
- 4 In the Service Group Selection dialog box, select the service group to delete and click **Next**.
- 5 In the Service Group Summary dialog box, click **Next**.
- 6 A message appears informing you that the wizard will run commands to delete the service group. Click **Yes** to delete the service group.
- 7 Click **Finish**.

Configuring an MSDTC service group for High Availability

This chapter describes the steps required to create a new MSDTC service group, configure the MSDTC Server, and configure the MSDTC client using the configuration wizard. The MSDTC agent comprises two parts; MSDTC client and MSDTC server. The MSDTC client and the MSDTC server must not run on the same cluster node.

Prerequisites

- You must be a cluster administrator. This user classification is required to create and configure a service group.
- You must be a local administrator on the node where you run the wizard.
- Verify that the SQL Server agent is installed on all cluster nodes.
- Verify that the VCS cluster is configured using the VCS Configuration Wizard (VCW).
- Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC service group.
- Assign a unique virtual server name and virtual IP address to the MSDTC server.
- Verify that the Distributed Transaction Coordinator service is stopped.
- Verify that you have created the LUNs (or virtual disks) and volumes for MSDTC log, and MSDTC registry replication, on the NetApp filer. See [“Managing storage using Network Appliance Filer”](#) on page 61.

- Verify that the volumes containing the MSDTC logs and registry replication directory are mounted on one node (the node on which you are configuring the service group) and unmounted on all other nodes.
- Mount the volumes that will be used by MSDTC service group. See [“Connecting virtual disks to the cluster node”](#) on page 62 for more information.

Reviewing the configuration

MSDTC servers can coexist with SQL servers on the same cluster nodes. If the MSDTC server and the SQL server are running on the same node, the MSDTC client is left in the default configuration. If the MSDTC server is not configured on the same node as the SQL server, then the MSDTC client must be configured on that node using the SQL Server Configuration Wizard.

For instance, a SQL Server configuration in a VCS cluster might span four nodes and two sets of shared storage. Two configurations are possible:

- SQL server is configured on different nodes than the MSDTC server
- SQL server is configured on the same node as the MSDTC server

Figure 5-1 MSDTC Server configured on different nodes than SQL Server





Figure 5-2 MSDTC Server configured on the same node as SQL Server



Creating an MSDTC service group

MSDTC is a global resource and is accessed by more than one SQL Server service group. Symantec recommends you to configure one MSDTC service group in a VCS cluster. VCS provides an SQL Server Configuration Wizard that guides you through the process of configuring an MSDTC service group. You can also use this wizard to modify an MSDTC service group configuration.

This section describes the steps required to create a new MSDTC service group using the wizard. For instructions on how to modify an existing service group, see the *Veritas Cluster Server Database Agent for Microsoft SQL Configuration Guide*.

To configure an MSDTC service group

- 1 Start the SQL Server Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
If you have just configured a SQL service group and you are in the Configuration Wizard, this step is not necessary.
- 2 In the Select Configuration Option panel, click **MSDTC Server - Service Group Configuration**, click **Create**, and then click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 Specify the service group name and system list.
 - Enter a name for MSDTC service group.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the service group's system list. Make sure to select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
To remove a system from the service group's system list, select the Systems in Priority Order list and click the left arrow.
To change a system's priority in the service group's system list, select the system from the Systems in Priority Order and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
 - Click **Next**. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 Specify the information related to the virtual server.
 - Enter the virtual name for the node. This is the server name on which DTC service is running. Ensure that the virtual server name you enter is unique in the cluster.

- Enter a unique virtual IP address for the MSDTC server.
- Enter the subnet mask to which the virtual IP address belongs.
- Click **Advanced...** to specify additional details for the Lanman resource.
 - Select the **AD Update required** check box to enable the Lanman resource to update the Active Directory with the virtual name.
 - Select the distinguished name of the Organizational Unit for the virtual server. By default, the Lanman resource adds the virtual server to the default container Computers.

The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

- Click **OK**.
 - For each system in the cluster, select the public network adapter name. Select the **Adapter Name** field to view the adapters associated with a system.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
 - If you require a computer object to be created in the Active Directory (AD), click **Advanced**, check **AD Update**, and select the Organizational Unit from the drop down list.
 - Click **Next**.
- 6 Specify the drive letter for the MSDTC log and replication directory and click **Next**. Symantec recommends using different paths for these directories. Clear the **Configure NetApp SnapMirror Resource(s)** check box. This option is applicable only in case of a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. If you are setting up a disaster recovery environment, check this check box to configure the SnapMirror resource at the primary site. The SnapMirror resource must be configured only after you have configured the cluster at the secondary site.
- 7 Review the service group configuration and change the resource name if desired.
- The **Resources** box lists the configured resources. Click on a resource to view its attributes and their configured values in the **Attributes** box.
- The wizard assigns unique names to resources. Change names of the resources, if desired.

To edit a resource name, select the resource name and either click it or press the F2 key. Press enter after editing each resource name. To cancel editing a resource name, press **Esc**.

- Click **Next** to create the service group.
 - A message appears informing you that the wizard will run commands to create the service group. Click **Yes** to create the service group. Various messages indicate the status of these commands. After the commands are executed, the completion dialog box appears.
- 8 In the Completing the MSDTC Configuration Wizard panel, check **Bring the service group online** to bring the configured service group online. To bring the service group online later, leave this option unchecked.
 - 9 Click **Finish**.
 - 10 Click **Next** to create an MSDTC client.

Creating an MSDTC client

Set the MSDTC client to run on nodes where:

- A SQL instance is configured to run
- The MSDTC server is not configured to run.

Before configuring the MSDTC configuration wizard to configure an MSDTC client:

- Verify the MSDTC service group is online in the cluster.
- Verify the node on which you run the wizard is not a part of an MSDTC service group system list.

To configure an MSDTC client

- 1 Start the SQL Server Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
- 2 Click **MSDTC Client Configuration** and click **Next**.
- 3 Verify that you have met the prerequisites and click **Next**.
- 4 Specify the nodes on which the MSDTC client will be configured.
 - Select the nodes in the Available Cluster Systems list. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
 - Click the right arrow to add them to the Selected systems list.

- The Unavailable Cluster Systems lists the nodes that have an MSDTC service group configured and are not available for setting up an MSDTC client.
MSDTC client must not be configured on a node that has MSDTC server configured.
 - Click **Next**.
- 5 If the MSDTC service group is not online in the cluster, an informational message appears informing you that the wizard will bring the MSDTC service group online. Click **Yes**.
 - 6 Specify the virtual DTC name and click **Next**.
 - 7 Click Next to return to the **Select Configuration Option** choices. See “[Configuring a SQL Server service group](#)” on page 85 if you have not already configured the SQL Server service group.
 - 8 Click **Finish** to exit the wizard.

Verifying the installation

Verify your installation by switching online nodes or by shutting down the computer that is currently online. Either process will test that the service group can be smoothly transferred between nodes.

Shutting down a node creates an actual failure, stressing your system, but more truly testing its high availability than by switching nodes. If you do shut down the online computer in your cluster, remember to bring it back up after you have confirmed that the service group successfully failed over to another node.

Complete the procedure “[Verifying the service group configuration](#)” on page 93 to verify the service group configuration.

Making a standalone SQL Server highly available

This chapter describes the procedure to convert a standalone SQL Server into a clustered SQL Server in a new Veritas Cluster Server environment. This environment involves an active/passive configuration with one to one failover capabilities.

If you are planning a new SQL Server deployment, refer to “[Installing and configuring SQL Server](#)” on page 59. If you are planning to deploy an active/active configuration with multiple SQL instances, refer to “[Active/active configuration](#)” on page 115.

Note: In addition to the information contained in this chapter, the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach are required. Refer to: <http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>

Prerequisites

Verify the requirements for your configuration are met before starting the VCS for NetApp installation.


See “[Prerequisites](#)” on page 31.

Reviewing the configuration

This chapter describes the tasks needed to incorporate an existing standalone SQL Server into a high available environment in order to ensure that the mission critical SQL resource is always available.

This chapter describes the tasks necessary to create a virtual server in an active/passive SQL configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node. At the end of this process, their environment will look like this:

Figure 6-1 Active/Passive configuration



The virtual SQL Server is online on SYSTEM1, serving client requests. The shared virtual disks/LUNs provide storage for the SQL Server databases. SYSTEM2 waits in a warm standby state as a backup node, prepared to begin handling client requests if SYSTEM1 becomes unavailable. 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 configuration

A sample setup is used through this guide to illustrate the installation and configuration tasks.

During the configuration process you will create virtual IP addresses for the following:

- SQL virtual server: the IP address should be the same on all nodes
- Cluster IP address: used by Veritas Cluster Manager (Web Console)

You should have these IP addresses available before you start deploying your environment. The following names describe the objects created and used during the installation and configuration:

Name	Object
SYSTEM1 & SYSTEM2	server names; SYSTEM1 is the existing standalone SQL server
INST1_SG	Microsoft SQL Server service group
SQL_CLUS1	virtual SQL server cluster
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_DB1_VOL	volume for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL server
INST1	SQL Instance Name
INST1-VS	name of the SQL Virtual Server

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Symantec recommends removing Internet Protocol 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. GAB 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.
- 5 Reboot each system. Verify that each system recognizes the attached shared disk.
- 6 Use Windows Disk Management on each system to verify that the attached shared virtual disks/LUNs are visible.

To verify the DNS settings for all systems on which SQL will be installed

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Open **Network and Dial-up Connections**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the **Advanced** menu, click **Advanced Settings**.
 - On the **Adapters and Bindings** tab, verify the public adapter is the first adapter in the **Connections** list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 From the status window, click **Properties**.

- 6 On the **General** tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify that the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 On the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Preparing the standalone SQL Server

Complete the following tasks before you begin the process of installing VCS for NetApp and creating a clustered environment:

- Backing up existing SQL data
- Setting SQL Server services to manual start

Backing up existing SQL data


Create a backup of the data on the existing standalone SQL server.

Setting SQL Server services to manual start

Set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).



- a Select the standalone server that you plan to incorporate into the cluster from the **Server** list.
 - b Select a service from the **Services** list.
 - c Clear the **Auto-start service when OS starts** check box.
- 2 Repeat [step b](#) and [step c](#) for all other SQL Server services that are running on the server.

Installing and configuring VCS for NetApp

Before converting the existing standalone SQL into a clustered server, complete the following procedures:

- “[Installing the software using the product installer](#)” on page 36 to install VCS for NetApp on all nodes of the cluster.
- “[Configuring the cluster](#)” on page 42 to configure VCS within the cluster.
- See “[Managing storage using Network Appliance Filer](#)” on page 61 to create the volumes/LUNs necessary to manage the SQL Server storage.

Installing and configuring SQL Server on additional nodes

Ensure the shared volume/LUN is imported on the node where you are installing SQL Server. Before installing Microsoft SQL Server on additional nodes, verify that the virtual disk on the filer is connected to the second cluster node. See “[Connecting virtual disks to the cluster node](#)” on page 62.

SQL Server must have the same configuration on all nodes in the cluster. You will need the following information:

- Instance name (if applicable)
- Destination Folder for Program Files and Data Files
- Authentication Mode

Complete the procedures specified in the following topics:

For SQL Server 2000, see “[Preparing to install SQL Server 2000 on the second node](#)” on page 67.

For SQL Server 2005, see “[Preparing to install SQL Server 2005 on the second node](#)” on page 77.

Moving the existing SQL Server data files and user databases

After completing the SQL installation and configuration on the additional nodes, move the existing standalone SQL Server data files and user databases from the local drive to the shared drives to ensure proper failover operations in the cluster. Complete the following tasks to move the databases:

- 1 Stop the SQL Server service. For SQL Server 2000, see “[Stopping the SQL Server 2000 Service](#)” on page 68. For SQL Server 2005, see “[Stopping the SQL Server 2005 service](#)” on page 77.
- 2 Verify that you have backed up your existing data.
- 3 Import the volume/LUNs to the node where the original files are located on the local drives and mount the volumes (add drive letters). See “[Connecting virtual disks to the cluster node](#)” on page 62 for instructions.
- 4 In case of SQL Server 2000, modify the SQL Server 2000 Data File and User Database locations.
Follow the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach.
<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>
- 5 Restart SQL Server.

To start the SQL service

- 1 Open the SQL Server Service Manager (**Start > All Programs > Microsoft SQL Server > Service Manager**).
- 2 Click **Start**.

Configuring the VCS SQL Server service group

The SQL Server Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

Complete the procedure “[Configuring a SQL Server service group](#)” on page 85 to configure the SQL Server service group.

Creating a SQL Server user-defined database

Complete the procedure “[Creating a SQL Server user-defined database](#)” on page 90 to create and manage a SQL Server user-defined database. The procedure includes the following tasks:

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Server Configuration Wizard to add the VMDg and MountV resources for the user databases.

Verifying the installation

Verify your installation by switching online nodes or by shutting down the computer that is currently online. Either process will test that the service group can be smoothly transferred between nodes.

Shutting down a node creates an actual failure, stressing your system, but more truly testing its high availability than by switching nodes. If you do shut down the online computer in your cluster, remember to bring it back up after you have confirmed that the service group successfully failed over to another node.

Complete the procedure “[Verifying the service group configuration](#)” on page 93 to verify the service group configuration.

Active/active configuration

This chapter describes how to install and configure VCS for NetApp for SQL Server that includes Active/Active clustering.

Prerequisites

Verify the requirements for your configuration are met before starting the VCS for NetApp installation. See [“Prerequisites”](#) on page 31.


Reviewing the configuration

A SQL Server instance is a completely independent SQL Server installation, with its own services, master database, storage, and memory resources. Each instance is defined uniquely by a separate SQL Server virtual server and service group.

SQL Server supports up to 16 independent instances of SQL Server to run on a single machine. A SQL Server instance can fail over to any of the other nodes configured nodes on its system list.

The following figure illustrates a two node active/active configuration. The SQL Server databases are configured on the shared storage on volumes/LUNs on the NetApp Filer. Each SQL Server virtual server is configured in a separate SQL Server service group. Each service group can fail over to the other node in the cluster.

Figure 7-1 Active/Active Configuration



For example, consider a two-node cluster hosting two SQL Server Virtual Servers, BILLING_VS and PAYROLL_VS. The table below and the sample configuration illustrate that the virtual servers are configured in two separate service groups with BILLING_VS online on SYSTEM1 but able to fail over to SYSTEM2, and PAYROLL_VS online on SYSTEM2 but able to fail over to SYSTEM1.

Table 7-1 Active/Active Configuration

SQL Virtual Server	Service Group	System List
BILLING_VS	BILLING_SG	SYSTEM1, SYSTEM2
PAYROLL_VS	PAYROLL_SG	SYSTEM2, SYSTEM1

Sample configuration

A sample setup is used to illustrate the installation and configuration tasks for two instances of SQL server, Billing and Payroll. During normal operation, one instance will be online on each of the two servers. If a failure occurs, the instance on the failing node will be brought online on the other server, resulting in two instances running on one server.

During the configuration process you will create virtual IP addresses for the following:

- Billing virtual server: virtual IP address is the same on all nodes
- Payroll virtual server: virtual IP address is the same on all nodes

■ Cluster IP address: used by Web Console

You should have these IP addresses available before you begin to deploy your environment.

The following names describe the objects created and used during the installation and configuration:

Name	Description
SYSTEM1 & SYSTEM2	server names
SQL_CLUS1	virtual SQL server cluster
BILLING_VS_SYS_FILES	volume for the SQL Server system data files for the billing instance
PAYROLL_VS_SYS_FILES	volume for the SQL Server system data files for the payroll instance
BILLING_DATA	volume for a SQL Server user-defined database for the billing instance
PAYROLL_DATA	volume for a SQL Server user-defined database for the payroll instance
BILLING_LOG	volume for a SQL Server user-defined database log file for the billing instance
PAYROLL_LOG	volume for a SQL Server user-defined database log file for the payroll instance
BILLING_REGREP	volume for the list of registry keys replicated among the nodes for the billing instance
PAYROLL_REGREP	volume for the list of registry keys replicated among the nodes for the payroll instance
BILLING_INST	instance name for the billing instance
PAYROLL_INST	instance name for the payroll instance
BILLING_VS	virtual SQL server name for the billing instance
PAYROLL_VS	virtual SQL server name for the payroll instance
BILLING_SG	SQL Server service group for the billing instance
PAYROLL_SG	SQL Server service group for the payroll instance

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers.
- 2 Connect the network adapters on each system.
To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Symantec recommends removing Internet Protocol 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. GAB 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.
- 5 Reboot each system. Verify that each system recognizes the attached shared disk.
- 6 Use Windows Disk Management on each system to verify that the attached shared disks are visible.

To verify the DNS settings for all systems on which SQL will be installed

- 1 Open the Control Panel (**Start > Control Panel**).
- 2 Open **Network and Dial-up Connections**.
- 3 Ensure the public network adapter is the first bound adapter:
 - From the **Advanced** menu, click **Advanced Settings**.
 - On the **Adapters and Bindings** tab, verify the public adapter is the first adapter in the **Connections** list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.
- 4 In the Network and Dial-up Connections window, double-click the adapter for the public network.
When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 5 From the status window, click **Properties**.

- 6 On the **General** tab:
 - Select the **Internet Protocol (TCP/IP)** check box.
 - Click **Properties**.
- 7 Select the **Use the following DNS server addresses** option.
- 8 Verify that the correct value for the IP address of the DNS server.
- 9 Click **Advanced**.
- 10 On the **DNS** tab, make sure the **Register this connection's address in DNS** check box is selected.
- 11 Make sure the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 12 Click **OK**.

Installing VCS for NetApp SnapMirror and configuring the cluster

Use the following procedures to install VCS for NetApp and then establish the cluster:

- [“Installing the software using the product installer”](#) on page 36 to install VCS for NetApp on all cluster nodes.
- [“Configuring the cluster”](#) on page 42 to configure VCS within the cluster.

Configuring virtual disks and volumes for SQL Server

See [“Managing storage using Network Appliance Filer”](#) on page 61 for instructions.

For each instance of SQL Server create a virtual disk on shared storage and a set of volumes.

For the Billing instance, create the following:

- BILLING_VS_SYS_FILES: volume for the SQL Server system data files
- BILLING-REGREP: volume for the list of registry keys replicated among cluster nodes for the Billing instance

For the Payroll Instance, create the following:

- PAYROLL_VS_SYS_FILES: volume for the SQL Server system data files
- PAYROLL_REGREP: volume for the list of registry keys replicated among cluster nodes for the Payroll instance

Installing and configuring the first instance of SQL Server

As you follow the procedures to install and configure SQL Server on the cluster, make the following changes to the process:

- Do not accept the default instance name. Specify an instance name for each SQL Server installation.
- Each SQL Server instance must be assigned a different port. The default port is 1433; ports for subsequent instances are generally assigned in descending order (1432, 1431, 1430, etc.).
- Set a unique internal name for each instance, for example BILLING_VS.

Complete the following tasks to install and configure the first SQL Server instance on all nodes of the cluster:

For SQL Server 2000

- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 63
- [“Preparing to install SQL Server 2000 on the second node”](#) on page 67
- [“Installing SQL Server 2000 on the second node”](#) on page 69
- [“Setting the internal name of the clustered instance”](#) on page 70

For SQL Server 2005

- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 73
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 77
- [“Installing SQL Server 2005 on the second node”](#) on page 79
- [“Setting the internal name of the clustered instance”](#) on page 79

Configuring the VCS SQL Server service group for the first instance

Consider the following points as you configure the SQL Server service group for the first instance:

- Assign a unique name to the SQL Server service group, for example BILLING_SG.
- Pay close attention to the priority order of the systems. For example, if the system priority for the first instance is SYSTEM1 then SYSTEM2; reverse the priority order for the second instance, so that SYSTEM2 has a higher priority.

See “[Configuring a SQL Server service group](#)” on page 85 for instructions.

Creating a SQL Server user-defined database

Complete the procedure “[Creating a SQL Server user-defined database](#)” on page 90 to create and manage a SQL Server user-defined database. The procedure includes the following tasks:

- Create volumes for a user-defined SQL Server database and its transaction log.
- Add new NetAppSnapDrive resources and link them to the existing NetAppFiler resource in the SQL Service Group.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.

Repeating SQL Server installation for additional instances

Repeat the previous steps to install and configure one or more additional SQL Server instances, making the same changes to the process:

- Do not accept the default instance name. Specify a unique instance name for each SQL Server installation.
- Each SQL Server instance must be assigned a different port. The default port is 1433; ports for subsequent instances are generally assigned in descending order (1432, 1431, 1430, etc.).
- Set a unique internal name for each instance, for example PAYROLL_VS.

Complete the following tasks to install and configure additional SQL Server instances on all nodes of the cluster:

For SQL Server 2000

- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 63
- [“Preparing to install SQL Server 2000 on the second node”](#) on page 67
- [“Installing SQL Server 2000 on the second node”](#) on page 69
- [“Setting the internal name of the clustered instance”](#) on page 70

For SQL Server 2005

- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 73
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 77
- [“Installing SQL Server 2005 on the second node”](#) on page 79
- [“Setting the internal name of the clustered instance”](#) on page 79

Consider the following points as you configure the SQL Server service groups for the additional instances:

- Assign a unique name to the SQL Server service group, for example PAYROLL_SG.
- Pay close attention to the priority order of the systems. For example, if the system priority for the first instance is SYSTEM1 then SYSTEM2; reverse the priority order for the second instance, so that SYSTEM2 has a higher priority.

See [“Configuring the VCS SQL Server service group”](#) on page 83 for instructions.

Create any SQL Server user-defined databases required for your environment.

See [“Creating a SQL Server user-defined database”](#) on page 90 for instructions.

Verifying the installation

See “[Verifying the service group configuration](#)” on page 93 for instructions.

Disaster recovery configuration

What is a disaster recovery solution?

A disaster recovery (DR) solution is a series of procedures you can use to safely and efficiently restore application data and services in the event of a catastrophic failure. A typical DR solution requires clusters on primary and secondary sites with replication between those sites. The cluster on the primary site provides data and services during normal operation; the cluster on the secondary site provides data and services if the cluster on the primary site fails.

This chapter describes how to set up a disaster recovery solution for SQL Server using the VCS agents for Network Appliance SnapMirror and Microsoft SQL.

Symantec recommends that you configure the secondary site only after you have established a local cluster with the GCO Option at the primary site.

Why implement 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.

Note: A DR solution requires a well-defined backup strategy. Refer to your backup product documentation for information on configuring backup.

Understanding replication

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

In the event that the primary site data center is destroyed, the application data is readily available at the remote site, and the application can be restarted at the remote site. Refer to the NetApp documentation for more information on replication in a NetApp storage environment.

What needs to be protected in a SQL Server environment?

The following components of a SQL server environment must be protected in the event of a disaster:


- **User Databases:** The most critical component in any SQL Server implementation is the user data that is stored in user-defined databases.
- **Logins:** Logins allow clients to connect to SQL Server and execute queries on user data. Logins are stored in the master database and each of the user-defined databases.
- **Jobs:** Jobs are a set of scheduled tasks that maintain SQL Server databases. The job configuration is stored in the msdb system database.
- **Alerts:** Alerts are actions that are taken when a specific event occurs. They are used to respond to and correct errors that occur in SQL Server. The alert configuration is stored in the msdb system database.
- **Operators:** Operators are contacts that address problems occurring in SQL Server. They are notified in the event of errors. The operator configuration is stored in the msdb system database.
- **Extended Stored Procedures:** Extended stored procedures are external routines that are called from within SQL Server. They are typically stored in DLL files on the file system.

- **Other Server Extensions:** SQL Server is a very flexible database engine and it is possible to extend its functionality in several ways. These extensions are also important to the operation of the SQL Server.

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 clusters at the primary and secondary sites are a part of the global cluster. 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. VCS continuously monitors and communicates events between clusters. Inter-cluster communication ensures that the global cluster is aware of the state of the global service group at all times.

Figure 8-1 Typical Disaster Recovery configuration



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.

Filer1 in the cluster on the primary site replicates to Filer2 in the cluster on the secondary site. Replication between the filers is set up using NetApp SnapMirror for SQL. Refer to NetApp documentation for more information on replication using NetApp filers.

If the Microsoft SQL Server server on System1 fails, SQL Server 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. VCS fails over the entire service group to the cluster at the secondary site. 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.

Disaster recovery: New SQL Server 2000 installation

This section provides information on how to install and configure the high availability and SQL 2000 components on the primary and secondary sites, with the intent of creating a parallel setup for the SQL service group on both sites. The configuration process is the same for both sites.

Reviewing the prerequisites

Verify that the hardware and software requirements for your configuration are met before starting the VCS for Netapp installation.

See “[Prerequisites](#)” on page 31.

Reviewing the configuration

During the configuration process you will create virtual IP addresses for the following:

- SQL virtual server: the IP address should be the same on all nodes
- Cluster IP address for the primary site--used for VCS Cluster Management Console (Single Cluster Mode), also known as Web Console.
- Cluster IP address for the secondary site--used for VCS Cluster Management Console (Single Cluster Mode), also known as Web Console.

You should have these IP addresses available before you start deploying your environment.

Installing VCS for NetApp

Perform the following tasks:

- Verify the driver signing option for the system
- Install VCS for NetApp
- Select the Global Cluster Option for VCS to enable wide-area failover
- Select the option to install Veritas Cluster Server Agent for Microsoft SQL Server

See “[Installing the software using the product installer](#)” on page 36 for instructions.

Configuring the cluster

Perform the following task:

- Configure cluster components including the Wide-Area Connector (WAC) resource for global clusters, using the Veritas Cluster Server Configuration Wizard (VCW).

See “[Configuring the cluster](#)” on page 42 for instructions.

Configuring storage on the NetApp Filer

Perform the following tasks:

- Create volumes on the NetApp Filer, share the volumes, and create virtual disks (LUNs) on the shared volumes
- Ensure that the virtual disks are connected to the first cluster node.

See “[Managing storage using Network Appliance Filer](#)” on page 61 for instructions.

Installing and configuring SQL Server 2000 at the primary site

Perform the following tasks:

- Install and configure SQL Server 2000 with SQL Server 2000 SP4 or higher
- Configure SQL services
- Stop the SQL services
- Disconnect the virtual disks from the first node
- Connect the virtual disks to the additional node
- Rename shared SQL 2000 files from the virtual disks.

- Install SQL Server 2000 on the additional node.
- Set the internal name of the clustered instance.

See “[Installing and configuring SQL Server 2000 on the first node](#)” on page 63 for instructions.

Configuring the VCS SQL service group

Perform the following tasks:

- Create a SQL Server service group using the VCS SQL Server Configuration Wizard.
- If required, create volumes for a user-defined database and transaction log, and then run the VCS SQL Server Configuration Wizard in the modify mode to bring these volumes under VCS control. This step is optional.
- If required, create a new user-defined database in SQL Server.
- If required, add NetApp resources for a user-defined database in VCS.
- Verify the installation.

See “[Configuring the VCS SQL Server service group](#)” on page 83 for instructions.

Creating a parallel environment on the secondary site

After setting up a high availability environment on the primary site, use the following guidelines to complete the same tasks on the secondary site:

Caution: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, the MSSearch resource (if present), and the SQL virtual IP resource on the primary site. Before you begin creating the SQL Server service group for the cluster at the secondary site, make sure that the SQL Server service group at the primary site is offline.

- [“Prerequisites”](#) on page 31
- [“Reviewing the configuration”](#) on page 129
- [“Configuring the cluster”](#) on page 42
- [“Managing storage using Network Appliance Filer”](#) on page 61
During the creation of virtual disks and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - NetApp Filer LUNs
 - Volume sizes
 - Volume names
 - Drive letters
- [“Installing and configuring SQL Server 2000 on the first node”](#) on page 63
Select the same options at the secondary site as you did at the primary site.
- [“Preparing to install SQL Server 2000 on the second node”](#) on page 67
The instance name must be the same on the primary site and secondary site.
- [“Installing SQL Server 2000 on the second node”](#) on page 69
- [“Configuring the VCS SQL Server service group”](#) on page 83
The service group name and virtual computer name must be the same on both the primary site and secondary site.

Configuring DR components on the primary and secondary sites

To complete the process of creating a DR solution, proceed to “[Configuring DR components](#)” on page 138 after performing the tasks outlined in this chapter.

Disaster recovery: New SQL Server 2005 installation

This section provides information on how to install and configure the high availability and SQL 2005 components on the primary and secondary sites, with the intent of creating a parallel setup for the SQL service group on both sites. The configuration process is the same for both sites.

Perform the following tasks to set up a disaster recovery environment for SQL Server 2005.

Reviewing the prerequisites

Verify that the hardware and software requirements for your configuration are met before starting the VCS for Netapp installation.

See “[Prerequisites](#)” on page 31.

Reviewing the configuration

During the configuration process you will create virtual IP addresses for the following:

- SQL virtual server: the IP address should be the same on all nodes at the primary and secondary sites
- Cluster IP address for the primary site--used for VCS Cluster Management Console (Single Cluster Mode), also known as Web Console.
- Cluster IP address for the secondary site--used for VCS Cluster Management Console (Single Cluster Mode), also known as Web Console.

You should have these IP addresses available before you start deploying your environment.

Installing VCS for NetApp

Perform the following tasks:

- Verify the driver signing option for the system
- Install VCS for NetApp
- Select the Global Cluster Option for VCS to enable wide-area failover
- Select the option to install Veritas Cluster Server Agent for Microsoft SQL Server

See “[Installing the software using the product installer](#)” on page 36 for instructions.

Configuring the cluster

Perform the following task:

- Configure cluster components including the Wide-Area Connector (WAC) resource for global clusters, using the Veritas Cluster Server Configuration Wizard (VCW).

See “[Configuring the cluster](#)” on page 42 for instructions.

Configuring storage on the NetApp Filer

Perform the following tasks:

- Create volumes on the NetApp Filer, share the volumes, and create virtual disks (LUNs) on the shared volumes
- Ensure that the virtual disks are connected to the first cluster node.

See “[Managing storage using Network Appliance Filer](#)” on page 61 for instructions.

Installing and configuring SQL Server 2005 at the primary site

Perform the following tasks:

- Create virtual disks on the NetApp Filer and ensure that they are connected to the first cluster node
- Install and configure SQL Server 2005
- Configure SQL services
- Stop the SQL services
- Disconnect the virtual disks from the first node
- Connect the virtual disks to the additional node

- Rename shared SQL 2005 files from the virtual disks
- Install SQL Server 2005 on the additional node
- Set the internal name of the clustered instance

See “[Installing and configuring SQL Server 2005 on the first node](#)” on page 73 for instructions.

Configuring the VCS SQL service group

Perform the following tasks:

- Create a SQL Server service group using the VCS SQL Server Configuration Wizard
- If required, create volumes for a user-defined database and transaction log
- If required, create a new user-defined database in SQL Server
- If required, add NetApp resources for a user-defined database in VCS
- Verify the installation

See “[Configuring the VCS SQL Server service group](#)” on page 83 for instructions.

Creating a parallel environment on the secondary site

After setting up a high availability environment on the primary site, use the following guidelines to complete the same tasks on the secondary site:

Caution: Before you begin to configure the secondary site, offline the SQL Server resource, the SQL virtual server name resource, the MSSearch resource (if present), and the SQL virtual IP resource on the primary site. Before you begin creating the SQL Server service group for the cluster at the secondary site, make sure that the SQL Server service group at the primary site is offline.

- [“Prerequisites”](#) on page 31
- [“Reviewing the configuration”](#) on page 134
- [“Configuring the cluster”](#) on page 42
- [“Managing storage using Network Appliance Filer”](#) on page 61
During the creation of virtual disks and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:
 - Volume sizes
 - Volume names
 - Drive letters
- [“Installing and configuring SQL Server 2005 on the first node”](#) on page 73
Select the same options at the secondary site as you did at the primary site.
- [“Preparing to install SQL Server 2005 on the second node”](#) on page 77
The instance name must be the same on the primary site and secondary site.
- [“Installing SQL Server 2005 on the second node”](#) on page 79
- [“Configuring the VCS SQL Server service group”](#) on page 83
The service group name and virtual computer name must be the same on both the primary site and secondary site.

Installing DR components on the primary and secondary sites

To complete the process of creating a DR solution, proceed to [“Configuring DR components”](#) on page 138 after performing the tasks outlined in this chapter.

Configuring DR components

After configuring the high availability and SQL components on the primary and secondary sites, complete the disaster recovery solution by configuring the disaster recovery components for both sites. Perform the following tasks:

- Set up volume replication using Network Appliance SnapMirror. See [“Configuring replication using Network Appliance SnapMirror”](#) on page 139.
- Configuring SnapMirror resources in the SQL service group at the primary site. See [“Configuring SnapMirror resources at the primary site”](#) on page 140.
- Configure GCO option for wide-area failover. See [“Configuring the Global Cluster Option for wide-area failover”](#) on page 142 .

Configuring replication using Network Appliance SnapMirror

You can replicate SQL Server data by establishing a SnapMirror relationship between the filers at the primary and secondary sites. Before configuring replication, make sure the service group is offline at the secondary site.

SnapMirror replicates snapshots taken on a filer and applies them to a remote filer over a wide area network; these snapshots can be used by the target host to provide rapid failover in case of a disaster.

If required, you can transfer the initial base snapshot image from the primary to secondary via tape, and then set up incremental SnapMirror updates to the destination filer. After you set up a SnapMirror relationship, ensure that the state of the volumes (that are to be replicated) at the primary site shows as SnapMirrored.

Refer to Network Appliance documentation for more information.

Configuring SnapMirror resources at the primary site

Configure NetAppSnapMirror resources at the primary site to monitor data replication from the primary to the secondary site. Creating a resource at the primary site will enable the filer to replicate from the primary to the secondary site.

You may want to repeat this procedure and create a NetAppSnapMirror resource at the secondary site. This is required in cases where:

- the service group is online at the secondary site (either it is failed over or switched to the secondary site) and the filer should replicate from secondary to primary site
- if you want to fail over or switch the service group from the secondary to the primary site

The following steps describe how to add the resources using the SQL Server Configuration Wizard.

Verify that the LUNs created to store the registry replication information and the SQL Server database are connected to the node on which you run the wizard, and disconnected from other nodes in the cluster.

To configure SnapMirror resource using the SQL Server Configuration Wizard

- 1 Start the SQL Server Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Database Agent for SQL Server > SQL Server Configuration Wizard**.
If you are running this procedure on the secondary site, ensure that the service group is online up to the snapdrive resources.
- 2 In the Select Configuration Option panel, click **MS SQL Server Service Group Configuration**, click **Edit**, and then click **Next**.
- 3 Review the prerequisites and then click **Next**.
- 4 In the Service Group Selection panel, select the service group to modify and click **Next**.
- 5 In the Service Group Configuration panel, verify the list of systems in the service group and click **Next**.
- 6 On the SQL Server Instance Selection panel, check **Configure the NetApp SnapMirror resource(s)** and click **Next**.

An informational message appears if you chose to configure a SnapMirror resource without configuring replication between Network Appliance filers

at primary and secondary sites. Review the message and click **Yes** to continue.

- 7 Accept default values in the subsequent wizard panels and click **Next** till you reach the wizard completion panel.
- 8 On the Configuration Status panel, do not check the **Bring the service group online** check box. Click **Finish**.

Configuring the Global Cluster Option for wide-area failover

The Global Cluster option is required to manage global clustering for wide-area disaster recovery. Creating a global cluster environment involves:

- Connecting standalone clusters by adding a remote cluster to a local cluster.
- Converting the local service group that is common to all the clusters to a global service group.

You need to create a wide-area connector resource for global clusters.

You can use the VCS Java Console or Web Console to perform global cluster operations; this guide only provides procedures for the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more information on GCO operations available from the Java and Web consoles.

Prerequisites

Creating a global cluster environment requires:

- Wide-area Connector process is configured and the ClusterService group is online at both sites.
See “[Configuring Wide-Area Connector process for global clusters](#)” on page 57 for instructions.
- All service groups properly configured and able to come online.
- The service group serving as the global group has the same unique name across all applicable clusters.
- The clusters use the same version of VCS.
- The clusters use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.
- The names of the clusters at the primary and secondary sites and the virtual IP addresses associated with them are registered in the DNS with reverse lookup.

Linking clusters: Adding a remote cluster to a local cluster

The VCS Java Console provides a wizard to create global clusters by linking standalone clusters or bringing a standalone cluster into an existing global environment.

- If you are creating a global cluster environment for the first time with two standalone clusters, run the wizard from either the cluster on the primary site or the cluster on the secondary site.
- If you are adding a standalone cluster to an existing global cluster environment, run the wizard from a cluster already in the global cluster environment.

The following information is required for the Remote Cluster Configuration Wizard in Cluster Explorer:

- The active host name or IP address of each cluster in the global configuration and of the cluster being added to the configuration.
- The user name and password of the administrator for each cluster in the configuration.
- The user name and password of the administrator for the cluster being added to the configuration.

Note: Symantec does not support adding a cluster that is already part of a global cluster environment. To merge the clusters of one global cluster environment (for example, cluster A and cluster B) with the clusters of another global environment (for example, cluster C and cluster D), separate cluster C and cluster D into standalone clusters and add them one by one to the environment containing cluster A and cluster B.

To add a remote cluster in Cluster Explorer

- 1 From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.
or
From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.
- 2 Review the required information for the **Remote Cluster Configuration Wizard** and click **Next**.
- 3 In the **Wizard Options** dialog box:
 - a Click **Add Cluster**.
 - b Click **Next**.
- 4 Enter the details of the new cluster:

If the cluster is not running in secure mode:

- a Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- b If necessary, change the default port number.
- c Enter the user name.
- d Enter the password.
- e Click **Next**.

If the cluster is running in secure mode:

- a Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- b Verify the port number.
- c Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.

If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.

- d Click **Next**.
- 5 Click **Finish**. After running the wizard, the configurations on all the relevant clusters are in read-write mode; the wizard does not close the configurations.
- 6 Verify that the heartbeat connection between clusters is alive. From the command window enter `hahb -display`. The state attribute in the output should show **alive**.

If the state is **unknown**, then offline and online the ClusterService group.

Converting a local service group to a global service group

After linking the clusters, use the Global Group Configuration wizard to convert a local service group that is common to the global clusters to a global group. This wizard also enables you to convert global groups into local groups.

To convert a local service group to a global group

- 1 From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
or

From the Cluster Explorer configuration tree, right-click the cluster, and click **Configure Global Groups**.

or

From the Cluster Explorer configuration tree, right-click the service group, click **Configure As Global**, and proceed to step 3b.

- 2 Review the information required for the Global Group Configuration wizard and click **Next**.
- 3 Enter the details of the service group to modify:
 - a Click the name of the service group that will be converted from a local group to a global group, or vice versa.
 - b From the **Available Clusters** box, click the clusters on which the group can come online. Click the right arrow to move the cluster name to the **Clusters for Service Group** box; for global to local cluster conversion, click the left arrow to move the cluster name back to the **Available Clusters** box. A priority number (starting with 0) indicates the cluster on which the group will attempt to come online. If necessary, double-click the entry in the **Priority** column and enter the new value.
 - c Select the policy for cluster failover:
 - **Manual** prevents a group from automatically failing over to another cluster.
 - **Auto** enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster fails.
 - **Connected** enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster.
 - d Click **Next**.
- 4 Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:

Cluster not in secure mode:

 - a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
 - b Verify the port number.
 - c Enter the user name.
 - d Enter the password.
 - e Click **OK**.
 - f Repeat these steps for each cluster in the global environment.

Cluster in secure mode:

- a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
- b Verify the port number.
- c Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and domain.
If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
- d Click **OK**.
- e Repeat these steps for each cluster in the global environment.

5 Click **Next**.

6 Click **Finish**.

At this point, you must bring the global service group online from Cluster Explorer.

Bringing a global service group online

After converting the local service group that is common to the global clusters to a global group, use the Cluster Explorer to bring the global service group online.

To bring a remote global service group online from Cluster Explorer

- 1 In the **Service Groups** tab of the configuration tree, right-click the service group.
or
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Online**, and click **Remote online**.
- 3 In the Online global group dialog box:
 - a Click the remote cluster to bring the group online.
 - b Click the specific system, or click **Any System**, to bring the group online.
 - c Click **OK**.

Administering global service groups

Administering global groups requires the following conditions:

- A group that will serve as the global group must have the same name across all applicable clusters.
- You must know the user name and password for the administrator to each cluster in the configuration.

Use the VCS Java Console or Web Console to bring a global group online, take a global group offline, or switch a global group on a remote cluster. The section below provides additional procedures for administering global groups from the Java Console. Refer to the *Veritas Cluster Server Administrator's Guide* for more information on global cluster operations from the Java Console and Web Console.

Note: For remote cluster operations, the user must have the same name and privilege as the user logged on to the local cluster.

Taking a remote global service group offline

Use Cluster Explorer to take a remote global service group offline.

To take a remote global service group offline from Cluster Explorer

- 1 In the **Service Groups** tab of the configuration tree, right-click the service group.
or
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Offline**, and click **Remote offline**.
- 3 In the Offline global group dialog box:
 - a Click the remote cluster to take the group offline.
 - b Click the specific system, or click **All Systems**, to take the group offline.
 - c Click **OK**.

Switching a remote service group

Use Cluster Explorer to switch a remote service group.

To switch a remote service group from Cluster Explorer

- 1 In the **Service Groups** tab of the configuration tree, right-click the service group.
or
Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Click **Switch To**, and click **Remote switch**.
- 3 In the Switch global group dialog box:
 - a Click the cluster to switch the group.
 - b Click the specific system, or click **Any System**, to take the group offline.
 - c Click **OK**.

Deleting a remote cluster

If necessary, use the Remote Cluster Configuration wizard to delete a remote cluster. This operation involves the following tasks:

- Taking the wide area cluster (wac) resource in the ClusterService group offline on the cluster that will be removed from the global environment. For example, to delete cluster C2 from a global environment containing C1 and C2, log on to C2 and take the wac resource offline.
- Removing the name of the specified cluster (C2) from the cluster lists of the other global groups using the Global Group Configuration wizard. Note that the Remote Cluster Configuration wizard in Cluster Explorer automatically updates the cluster lists for heartbeats. Log on to the local cluster (C1) to complete this task before using the Global Group Configuration wizard.
- Deleting the cluster (C2) from the local cluster (C1) through the Remote Cluster Configuration wizard.

Note: You cannot delete a remote cluster if the cluster is part of a cluster list for global service groups or global heartbeats, or if the cluster is in the **RUNNING**, **BUILD**, **INQUIRY**, **EXITING**, or **TRANSITIONING** states.

Use Cluster Explorer to take the wide area cluster resource offline, remove a cluster from the cluster list for a global group, and delete a remote cluster from the local cluster.

To take the wide area cluster (wac) resource offline

- 1 From Cluster Monitor, log on to the cluster that will be deleted from the global cluster environment.


- 2 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **wac** resource under the **Application** type in the **ClusterService** group.
or
Click a service group in the configuration tree, click the **Resources** tab, and right-click the **wac** resource in the view panel.
- 3 Click **Offline**, and click the appropriate system from the menu.

To remove a cluster from a cluster list for a global group


- 1 From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
- 2 Click **Next**.
- 3 Enter the details of the service group to modify:
 - a Click the name of the service group.
 - b For global to local cluster conversion, click the left arrow to move the cluster name from the cluster list back to the **Available Clusters** box.
 - c Click **Next**.
- 4 Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:
If the cluster is not running in secure mode
 - a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
 - b Verify the port number.
 - c Enter the user name.
 - d Enter the password.
 - e Click **OK**.
 If the cluster is running in secure mode
 - a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
 - b Verify the port number.
 - c Choose to connect to the remote cluster using the connected cluster's credentials, or enter new credentials, including the user name, password, and domain.
 - d Click **OK**.
- 5 Click **Next**.
- 6 Click **Finish**.

To delete a remote cluster from the local cluster

- 1 From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.
or
From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Clusters**.
- 2 Review the required information for the **Remote Cluster Configuration** wizard and click **Next**.
- 3 On the **Wizard Options** dialog box:




- a Click **Delete Cluster**.
 - b Click **Next**.
- 4 In the Delete Cluster dialog box:



- a Click the name of the remote cluster to delete.
 - b Click **Next**.

- 5 Review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:



If the cluster is not running in secure mode

- a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
- b Verify the port number.
- c Enter the user name.
- d Enter the password.
- e Click **OK**.

If the cluster is running in secure mode

- a Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
- b Verify the port number.
- c Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.

If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.

- d Click **OK**.

- 6 Click **Finish**.

Removing VCS for NetApp SnapMirror

This chapter describes steps to remove VCS for Network Appliance SnapMirror.

Prerequisites

- You must have administrative access to all systems selected for cluster operations. For this, add the domain user to the local Administrators group of each system.
- You must be a domain user.
- You must be a member of the Local Administrators group for all nodes where you are installing.

Removing the VCS configuration

Before you remove the VCS for Network Appliance SnapMirror, use the VCS Configuration Wizard to delete the VCS configuration from all nodes.

For instructions, see the chapter on Modifying the Cluster Configuration in the *Veritas Cluster Server Administrator's Guide*.

Uninstalling the agents

This section describes steps for uninstalling VCS for Network Appliance SnapMirror.

Prerequisites

- Verify that you have local administrative privileges on the node where you are uninstalling the agent.
- Verify that all SQL service groups are offline on all nodes in the cluster.

Instructions

Follow these steps to remove VCS for Network Appliance SnapMirror.

To remove VCS for Network Appliance SnapMirror

- 1 In the **Add/Remove Programs** applet, click VCS for Network Appliance (Server Components) and click **Remove**.
- 2 Review the Welcome page and click **Next**.
- 3 Select the check box if you want to remove the client components and click **Next**.
- 4 Specify the nodes from which you want to remove the agents.
 - Select the node. This may take some time depending on the size of the domain and network conditions.
 - Click **Add**. To remove a node, highlight the node in the computer list and click **Remove**.
 - Click **Next**.
- 5 The installer validates the system for uninstallation. After the node is accepted, click **Next**.

If a node is rejected, the **Comments** column displays the cause for rejecting the node. Highlight the node to view a detailed information about the failure in the **Details** box. Resolve the error, highlight the node in the selected systems list, and click **Validate Again**. Once all the nodes are accepted, click **Next**.
- 6 Review the summary of your selections and click **Uninstall**.
- 7 The installer displays the status of uninstallation.
- 8 After the uninstallation is complete, review the report and click **Next**.
- 9 Click **Finish**.

License management

The product installer also allows you to add or remove license keys for options in your installation of VCS for NetApp SnapMirror components.

To add or remove license keys

- 1 Open the Windows Control Panel and click **Add or Remove Programs**.
- 2 Select **VCS for Network Appliance (Server Components)** and click **Change**.
- 3 The Symantec Product Installer screen appears. Select **License Management**. Click **Next**.
- 4 The license key screen appears. Enter the license key you want to add and click **Update**. If you want to remove a license key, select the license key in the Licenses field and click **Remove**.

Troubleshooting the agents

This chapter describes how to troubleshoot common problems in the VCS enterprise agents for Network Appliance and Microsoft SQL Server. The chapter lists the error messages, and describes the problem associated with the agent. Recommended solution is included, where applicable.

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

A typical agent log resembles:

```
2004/01/12 11:22:47 VCS NOTICE V-16-20020-10
SQLServer2000:SQL-Group-SQLServer2000:monitor:SQL Server
Instance name is not specified. Agent will operate on default
instance.
```

Here,

- Timestamp denotes the date and time when the message was logged.
- Mnemonic denotes which Veritas product logs the message. For VCS enterprise agent for Microsoft SQL, mnemonic is 'VCS'.
- Severity denotes the seriousness of the message. Severity of the VCS error messages is classified into the following types:

- 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 that VCS has initiated an action.
- INFO informs about various state messages or comments.
Of these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.
- UMI or Unique Message ID is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the SQLServer agent would resemble: V-16-20020-13
Originator ID for all VCS products is 'V-16.' Category ID for SQLServer agent is 20020, for MSDTC is 20021, and that for MSSearch agent is 20022. Message ID is a unique number assigned to the message text.
- 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 also written to the Windows event log.

VCW logs

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

Here, %allusersprofile% is the file system directory containing application data for all users. A typical path is C:\Documents and Settings\All Users\.

The format of the VCW log is

ThreadID | Message Text

- *ThreadID*: the ID of the thread initiated by VCW.
- *Message Text*: the actual message generated by VCW.

A typical VCW log resembles:

```
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
```

VCWsilent logs

The VCWsilent log is located at `<currentdirectory>\vcwsilent.log`.

Here, `<currentdirectory>` is the directory from where the VCWsilent.exe is run.

A typical VCWsilent log resembles:

```
00005540-00000064: 5540: STARTING - Discovering NICs on the
selected machines...
00009956-00000064: 9956: STARTING - Generating private network
related files...
00009956-00000048: 9956: COMPLETED - Generating LLT host
files...
00009956-00000048: 9956: COMPLETED - Generating GAB tab files...
00009956-00000048: 9956: COMPLETED - Generating main.cf file...
00009956-00000064: 9956: STARTING - Configuring LLT on all the
nodes.
00009956-00000048: 9956: COMPLETED - Configuring LLT on all the
nodes.
```

Network Appliance agents error messages

This section describes the error messages for the NetApp agents.

Table 10-1 NetApp agents error messages

Message	Description
Failed to open connection to filer %s.	Make sure that the HAD Helper account is part of the administrator's group on the local host and the filer. Make sure the private network is functioning properly. Verify you can ping the IP used for the private storage network. This is the IP defined the StorageIP attribute of the NetAppFiler resource.
Failed to initialize ONTAPI on system	The agent could not find the file NTAPADMIN.DLL on the system. Verify the file exists in the %VCS_HOME%\bin directory
Invalid attributes exist in the configuration	Some agent attributes have not been defined or have been defined incorrectly. Verify the configuration definition for the agent.
ONTAP API called failed for <i>object_name</i> on <i>filer_name</i> .	The specified API failed on the specified object. See the NetApp ONTAP API documentation for information about the associated error message
Volume %s on filer %s is not a SnapMirror replicated volume	Verify replication is set up on the specified volume.
Multiple snapmirror destinations for a volume is not supported by this agent. 'snapmirror status' for volume %s on filer %s returned multiple status entries. Administrative intervention required	There should be only one destination per source volume.

Table 10-1 NetApp agents error messages (continued)

Message	Description
Initialize VLibNetAppHost::Initialize() failed. (error_type: %s, error_code: 0x%s)	The agent could not detect the iSCSI Initiator on the host. Make sure that you have installed and configured Microsoft iSCSI Initiator on each node.
Failed to connect/disconnect virtual disk. (error_type: %s, error_code: 0x%s. error_message: %s)	This could occur because one or more of the following parameters are defined incorrectly in the VCS configuration: ■ Filer name ■ Volume name/LUN name ■ Share name ■ Storage IP Verify the configuration definition of the resource. Make sure each attribute is defined correctly.
Unable to create/delete online lock file %s. Error code %s,	Make sure you have write permissions on the specified directory.

Error messages and descriptions

The following table lists the messages of type ERROR and WARNING. Each message includes a description and a recommended solution, if applicable.

Agent for SQL Server 2000

This section describes the error messages for the SQL Server 2000 agent.

Table 10-2 SQL Server 2000 agents error messages

Message	Description
Invalid value specified for attribute <i>attribute name</i> .	No value provided for the specified attribute. Solution: Provide a value for the attribute.
Failed to lookup the account name ' <i>account name</i> '. Error = <i>Error code</i> .	Invalid user information specified for detail monitoring. Solution: Provide the correct user information
Failed to open the SCM handle. Error = <i>Error code</i> .	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
The service ' <i>service name</i> ' is not in stopped state.	The service is in an invalid state. Solution: Check the state of the service.
Failed to start the service ' <i>service name</i> '. Error = <i>Error code</i> .	The agent failed to start the service. Solution: See the associated Windows error code for more information.
SQL Server Instance name is not specified. Agent will operate on the default instance.	No value spaced-out for SQL Server instance name. Agent would operate on the default SQL Server instance.
The service ' <i>service name</i> ' did not start within the specified time limit.	The agent failed to start the service within the time limit as specified in the SQLOnlineTimeout attribute.
The ' <i>service name</i> ' service is not in stopped or running state. State= <i>state name</i> .	The service is in an invalid state.

Table 10-2 SQL Server 2000 agents error messages (continued)

Message	Description
The ' <i>service name</i> ' service is in <i>state name</i> state, while the ' <i>service name</i> ' service is in <i>state name</i> state.	The SQL Server and the SQL agent service are not in the same state. Both the services must be in ONLINE state to bring the service group online.
Failed to get the password attribute. Error = <i>Error code</i> .	Incorrect encrypted password specified for detail monitoring. Solution: Provide a password that is encrypted using the 'VCSencrypt' utility.
Failed to open the service ' <i>service name</i> '. Error = <i>Error code</i> .	The agent failed to open the service from the Service Control Manager. Solution: Check whether the service is present in the Service Control Manager.
The service ' <i>service name</i> ' is not in running state. Attempt to stop it might be unsuccessful.	The SQL Server service could be in PAUSE, PAUSE PENDING, or START PENDING state. Solution: Resume the service and then attempt to stop it.
The service ' <i>service name</i> ' did not stop. Error = <i>Error code</i> .	The agent failed to stop the service. See the associated Windows error code for more information.
The service ' <i>service name</i> ' did not stop within the specified timeout. Error = <i>Error code</i> .	The agent failed to stop the service within the time limit as specified in the SQLOfflineTimeout attribute.
Sql script has failed with error <i>error code</i> .	The SQL script for detail monitoring failed. See the associated Windows error code for more information.
Could not find the database path for instance <i>instance name</i> . Please correct the instance name specified.	The instance name specified for the SQL Server instance is incorrect. Solution: Specify the correct SQL Server instance name.
Mount specified does not contain the database for instance <i>instance name</i> . Please correct the mount path.	The specified mount does not contain the data pertaining to a given instance name. Solution: Specify the correct mount path.
The password attribute has not been configured.	The password attribute used for detail monitoring is not configured.

Table 10-2 SQL Server 2000 agents error messages (continued)

Message	Description
Failed to start the Sql script. (User = <i>user name</i> , Domain = <i>domain name</i>) Error : <i>Error code</i> .	The agent failed to execute the script for detail monitoring. See the associated Windows error code for more information.
Sql script failed. Script output : <i>script output</i>	The SQL script failed to monitor the SQL Server instance. See the script output for more information.

Agent for MSSearch

This section describes the error messages for the MSSearch agent.

Table 10-3 MSSearch agent error messages

Message	Description
The AppName attribute has not been configured.	No value specified for the AppName attribute. Solution: Specify a valid value for the AppName attribute.
Failed to instantiate the Microsoft Search COM object. Error = <i>Error code</i> .	MSSearch is not installed. See the associated Windows error code for more information.
Failed to open the Microsoft Search service. Error = <i>Error code</i> .	The agent failed to open the MSSearch service from the Service Control Manager (SCM). Solution: Check whether the service is present in the Service Control Manager.

Agent for MSDTC

This section describes the error messages for the MSDTC agent.

Table 10-4 MSDTC agent error messages

Message	Description
Lanman attribute has not been configured.	No value specified for the LanmanResName attribute. Solution: Specify a valid value for the LanmanResName attribute.
MountResName attribute has not been configured.	No value specified for MountResName attribute. Solution: Specify a valid value for the MountResName attribute.
LogPath attribute has not been configured.	No value specified for LogPath attribute. Solution: Specify a valid value for the MountResName attribute.
Failed to open the SCM handle. Error = <i>Error code</i> .	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
Failed to open the MSDTC service. Error = <i>Error code</i> .	The agent failed to open the MSDTC service from the Service Control Manager (SCM). Solution: Check whether the service is present in the Service Control Manager.
Failed to start the MSDTC service. Error = <i>Error code</i> .	The agent failed to start the MSDTC service. See the associated Windows error code for more information.
The MSDTC log path is ' <i>path name</i> '. Configured one is ' <i>path name</i> '.	The specified path for the MSDTC logs is different from the actual path. Solution: Specify the correct MSDTC log path.
The MSDTC service is not in running state. Offline might be unsuccessful.	The MSDTC service could be in PAUSE, PAUSE PENDING, or START PENDING state. Solution: Resume the service and then attempt to stop it.

Table 10-4 MSDTC agent error messages (continued)

Message	Description
Failed to stop the MSDTC service. Error = <i>Error code</i> .	The MSDTC service could not be stopped. See the associated Windows error code for more information.
Failed to wait for the MSDTC service to stop. Error = <i>Error code</i> .	The agent could not stop the service within the specified time limit of 20 seconds. See the associated Windows error code for more information.

Agent for SQL Server 2005

This section describes the error messages for the SQL Server 2005 agent.

Table 10-5 SQL Server 2005 agents error messages

Message	Description
Failed to convert the argument list. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Invalid value specified for attribute <i>attribute name</i> .	No value provided for the specified attribute. Solution: Provide a value for the attribute.
Failed to initialize the SQLServer2005 agent.	The agent failed to initialize the SQLServer2005 agent for SQL Server 2005.
Failed to open the SCM handle. Error = <i>Error code</i> .	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
The service <i>service name</i> is not in stopped state.	The service is in an invalid state. Solution: Check the state of the service.
Failed to start the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to start the service. Solution: See the associated Windows error code for more information.
SQL Server Instance name is not specified. Agent will operate on the default instance.	No value spaced-out for SQL Server instance name. Agent would operate on the default SQL Server instance.

Table 10-5 SQL Server 2005 agents error messages (continued)

Message	Description
Failed to set the virtual computer name in the environment of the service <i>service name</i> . Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The service <i>service name</i> did not start within the specified time limit.	The agent failed to start the service within the time limit as specified in the SQLOnlineTimeout attribute.
Failed to wait for the service <i>service name</i> to start. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The <i>service name</i> service is not in stopped or running state. State= <i>state name</i> .	The service is in an invalid state.
Failed to get the password attribute. Error = <i>Error code</i> .	Incorrect encrypted password specified for detail monitoring. Solution: Provide a password that is encrypted using the 'VCSencrypt' utility.
Failed to convert the password attribute. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to open the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to open the service from the Service Control Manager. Solution: Check whether the service is present in the Service Control Manager.
Failed to query the status of the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to query the state of the service. Solution: Check whether the service is present in the Service Control Manager.
The service <i>service name</i> is not in running state. Attempt to stop it might be unsuccessful.	The SQL Server service could be in PAUSE, PAUSE PENDING, or START PENDING state. Solution: Resume the service and then attempt to stop it.
The service <i>service name</i> did not stop. Error = <i>Error code</i> .	The agent failed to stop the service. See the associated Windows error code for more information.
The service <i>service name</i> did not stop within the specified timeout. Error = <i>Error code</i> .	The agent failed to stop the service within the time limit as specified in the SQLOfflineTimeout attribute.

Table 10-5 SQL Server 2005 agents error messages (continued)

Message	Description
Sql script has failed with error <i>error code</i> .	The SQL script for detail monitoring failed. See the associated Windows error code for more information.
Error occurred while getting the process exit code. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
WaitForSingleObject failed with error <i>error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The password attribute has not been configured.	The password attribute used for detail monitoring is not configured.
Failed to start the Sql script. (User = <i>user name</i> , Domain = <i>domain name</i>) Error = <i>Error code</i> .	The agent failed to execute the script for detail monitoring. See the associated Windows error code for more information.
Unable to convert the buffer to UNICODE. Error = <i>Error code</i>	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Sql script failed. Script output : <i>output</i>	The SQL script failed to monitor the SQL Server instance. See the script output for more information.
Failed to get the temporary file path. Error : <i>Error code</i>	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to create the temporary file. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed read the temporary file. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to remove the virtual name environment for the service <i>service name</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.

Agent for SQL Server 2005 Agent service

This section describes the error messages for the SQL Server 2005 Agent service agent.

Table 10-6 SQL Server 2005 Agent service agent error messages

Message	Description
Failed to convert the argument list. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Invalid value specified for attribute <i>attribute name</i> .	No value provided for the specified attribute. Solution: Provide a value for the attribute.
Failed to initialize the SQLAgService2005 agent.	The agent failed to initialize the SQLAgService2005 agent for the SQL Server 2005 Agent Service. Solution: Verify whether the SQLAgService2005 agent exists.
Failed to open the SCM handle. Error = <i>Error code</i> .	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
The service <i>service name</i> is not in stopped state.	The service is in an invalid state. Solution: Check the state of the service.
Failed to start the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to start the service. Solution: See the associated Windows error code for more information.
SQL Server Instance name is not specified. Agent will operate on the default instance.	No value spaced-out for SQL Server instance name. Agent would operate on the default SQL Server instance.
Failed to set the virtual computer name in the environment of the service <i>service name</i> . Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The service <i>service name</i> did not start within the specified time limit.	The agent failed to start the service within the time limit as specified in the SQLOnlineTimeout attribute.

Table 10-6 SQL Server 2005 Agent service agent error messages (continued)

Message	Description
Failed to wait for the service <i>service name</i> to start. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The <i>service name</i> service is not in stopped or running state. State= <i>state name</i> .	The service is in an invalid state.
Failed to open the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to open the service from the Service Control Manager. Solution: Check whether the service is present in the Service Control Manager.
Failed to query the status of the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to query the state of the service. Solution: Check whether the service is present in the Service Control Manager.
The service <i>service name</i> is not in running state. Attempt to stop it might be unsuccessful.	The SQL Server service could be in PAUSE, PAUSE PENDING, or START PENDING state. Solution: Resume the service and then attempt to stop it.
The service <i>service name</i> did not stop. Error = <i>Error code</i> .	The agent failed to stop the service. See the associated Windows error code for more information.
The service <i>service name</i> did not stop within the specified timeout. Error = <i>Error code</i> .	The agent failed to stop the service within the time limit as specified in the SQLOfflineTimeout attribute.
Failed to remove the virtual name environment for the service <i>service name</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.

Agent for SQL Server 2005 Analysis service

This section describes the error messages for the SQL Server 2005 Analysis service agent.

Table 10-7 SQL Server 2005 Analysis service agent error messages

Message	Description
Failed to convert the argument list. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Invalid value specified for attribute <i>attribute name</i> .	No value provided for the specified attribute. Solution: Provide a value for the attribute.
Failed to initialize the SQLOLapService2005 agent.	The agent failed to initialize the SQLOLapService2005 agent for the SQL Server 2005 Analysis Service. Solution: Verify whether the SQLOLapService2005 agent exists.
Failed to open the SCM handle. Error = <i>Error code</i> .	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
The service <i>service name</i> is not in stopped state.	The service is in an invalid state. Solution: Check the state of the service.
Failed to start the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to start the service. Solution: See the associated Windows error code for more information.
SQL Server Instance name is not specified. Agent will operate on the default instance.	No value spaced-out for SQL Server instance name. Agent would operate on the default SQL Server instance.
Failed to set the virtual computer name in the environment of the service <i>service name</i> . Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The service <i>service name</i> did not start within the specified time limit.	The agent failed to start the service within the time limit as specified in the SQLOnlineTimeout attribute.

Table 10-7 SQL Server 2005 Analysis service agent error messages (continued)

Message	Description
Failed to wait for the service <i>service name</i> to start. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The <i>service name</i> service is not in stopped or running state. State= <i>state name</i> .	The service is in an invalid state.
Failed to open the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to open the service from the Service Control Manager. Solution: Check whether the service is present in the Service Control Manager.
Failed to query the status of the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to query the state of the service. Solution: Check whether the service is present in the Service Control Manager.
The service <i>service name</i> is not in running state. Attempt to stop it might be unsuccessful.	The SQL Server service could be in PAUSE, PAUSE PENDING, or START PENDING state. Solution: Resume the service and then attempt to stop it.
The service <i>service name</i> did not stop. Error = <i>Error code</i> .	The agent failed to stop the service. See the associated Windows error code for more information.
The service <i>service name</i> did not stop within the specified timeout. Error = <i>Error code</i> .	The agent failed to stop the service within the time limit as specified in the SQLOfflineTimeout attribute.
Failed to remove the virtual name environment for the service <i>service name</i> .	This is a VCS internal error. Solution: Contact Symantec Technical Support.

Resource type definitions

This appendix lists the resource type definitions and attribute definitions of the agents. The resource type represents the VCS configuration definition of the agent, and specifies how the agent is defined in the configuration file `main.cf`. The Attribute Definitions lists the attributes associated with the agent. The Required attributes table lists the attributes that must be configured for the agent to function properly.

The sample dependency graph depicts the resource types, resources, and resource dependencies required to create the service groups.

VCS enterprise agent for Network Appliance

NetApp Filer agent

Resource type definition

```
type NetAppFiler (  
    static int MonitorInterval = 30  
    static i18nstr ArgList[] = { FilerName, StorageIP }  
    static str Operations = None  
    str FilerName  
    str StorageIP  
)
```

Attribute definitions

Table A-1 NetApp Filer agent attributes

Attribute	Type - Dimension	Description
FilerName	string-scalar	DNS-resolvable name or IP address of the locally attached filer.
StorageIP	string -scalar	The private storage IP address of the filer.

NetApp SnapDrive agent

Resource type definition

```
type NetAppSnapDrive (  
    static int MonitorInterval = 30  
    static int NumThreads = 1  
    static i18nstr ArgList[] = { FilerResName,  
        "FilerResName:FilerName",  
        "FilerResName:StorageIP",  
        VolumeName, ShareName, LUN, MountPath, Initiator  
    }  
  
    str FilerResName  
    str VolumeName  
    str ShareName  
    str LUN  
    str MountPath  
    str Initiator  
)
```

Attribute definitions

Table A-2 NetApp SnapDrive agent attributes

Attribute	Type - Dimension	Description
FilerResName	string-scalar	Name of the VCS NetAppFiler-type resource in the service group.
VolumeName	string-scalar	Name of the volume containing the virtual disk. Define the volume name in the same case as on the filer.
ShareName	string-scalar	Name of the CIFS share containing the virtual disk.
LUN	string-scalar	Name of the LUN (virtual disk) on the filer that is presented to the host for mounting. Define the LUN name in the same case as on the filer.
MountPath	string-scalar	Drive letter to be assigned to the virtual disk.
Initiator	string-scalar	<i>Name of iSCSI initiator the host uses to connect virtual disks. You can retrieve this value from the Disk Management console.</i>

NetApp SnapMirror agent

Resource type definition

```
type NetAppSnapMirror (
    static keylist SupportedActions = { fbsync }
    static int MonitorInterval = 300
    static int NumThreads = 1
    static i18nstr ArgList[] = { FilerResName,
        "FilerResName:FilerName",
        "FilerResName:StorageIP", VolumeName,
    SnapMirrorArguments,
        SnapMirrorSchedule, AppResName }
    str FilerResName
    str VolumeName
    str SnapMirrorArguments
    str SnapMirrorSchedule
    str AppResName
)
```

Attribute definitions

Table A-3 NetApp SnapMirror agent attributes

Attribute	Type - Dimension	Description
FilerResName	string-scalar	Name of the VCS NetAppFiler-type resource in the group.
VolumeName	string-scalar	Name of the filer volume that is to be mounted via iSCSI. Define the volume name in the same case as on the filer.
SnapMirrorArguments	string-scalar	Specifies the SnapMirror arguments such as maximum transfer speed and restart mode.
SnapMirrorSchedule	string-scalar	Specifies the schedule the destination uses for updating data. Do not assign a value for this attribute if you use SnapManager. Default = - - - -

Table A-3 NetApp SnapMirror agent attributes (continued)

Attribute	Type - Dimension	Description
AppResName	string-scalar	Name of the resource configured to monitor the application being made highly available. When used with Exchange Server, set this value to the name of the ExchService resource configured to monitor the Exchange Information Store.

SQL Server 2000 agents

SQL Server 2000 agents include agents for SQL Server 2000 and MSSearch service.

Agent for SQL Server 2000

The agent for SQL Server 2000 is represented by the `SQLServer2000` resource type.

Resource type definition

```
type SQLServer2000 (  
    static il8nstr ArgList[] = {Instance,  
    "LanmanResName:VirtualName",  
  
    "MountResName:MountPath",  
  
    SQLOnlineTimeout, SQLOfflineTimeout,  
  
    DetailMonitor, SQLDetailMonitorTimeout,  
  
    Username, Domain, Password, SQLFile }  
    str Instance  
    str LanmanResName  
    int SQLOnlineTimeout = 90  
    int SQLOfflineTimeout = 90  
    boolean DetailMonitor = 0  
    boolean FaultOnDMScriptFailure = 1  
    int SQLDetailMonitorTimeout = 30  
    il8nstr Username  
    il8nstr Domain  
    str Password  
    il8nstr SQLFile  
    str MountResName  
)
```

Attribute definitions

Review the following information to familiarize yourself with the required agent attributes for a SQLServer2000 resource type. This information will assist you during the agent configuration.

Table A-4 SQL Server 2000 agent required attributes

Required Attributes	Type and Dimension	Definition
Instance	string-scalar	Name of instance to monitor. If the attribute is blank, the agent monitors the default instance.
LanmanResName	string-scalar	The Lanman resource name on which the SQL Server 2000 resource depends.
MountResName	string-scalar	The mount resource name on which the SQL Server 2000 resource depends.
SQLOnlineTimeout	integer-scalar	Number of seconds that can elapse before online entry point aborts. Default is 90.
SQLOfflineTimeout	integer-scalar	Number of seconds that can elapse before offline entry point aborts. Default is 90.

Table A-5 SQL Server 2000 agent optional attributes

Optional Attributes	Type and Dimension	Definition
DetailMonitor	boolean	Defines whether the agent performs detail monitoring of SQL Server 2000 database. The value 0 indicates the agent will not monitor the database in detail. The value 1 indicates the agent will. Default = 0. Note: If the attribute is set to 1, the attributes Username, Password, Domain, SQLDetailMonitorTimeOut, and SQLFile must be assigned appropriate values.

Table A-5 SQL Server 2000 agent optional attributes (continued)


Optional Attributes	Type and Dimension	Definition
FaultOnDMScriptFailure	boolean	Defines whether the agent fails over the service group if the detail monitoring script execution fails. Default = 1 The value 1 indicates that the agent fails over the service group if detail monitoring script fails to execute. The value 0 indicates that it does not.
SQLDetailMonitor Timeout	integer-scalar	Number of seconds that can elapse before the detail monitor routine aborts. Default is 30.
Username	string-scalar	The Microsoft Windows authentication name when logging in to a database for detail monitoring. This attribute must not be null if “DetailMonitor” attribute is set to 1. Note: This attribute can take localized values.
Domain	string-scalar	Domain for the user account. This attribute is used to create a trusted connection to the SQL Server 2000 instance if “DetailMonitor” attribute is set to 1. Note: This attribute can take localized values.
Password	string-scalar	Password for logging in to a database for detail monitoring. This attribute must not be null if “DetailMonitor” attribute is set to 1.
SQLFile	string-scalar	The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the “DetailMonitor” attribute is set to 1. Note: This attribute can take localized values.

Sample dependency graph

The following example describes typical service groups configured to monitor the state of SQL Server 2000 in a VCS cluster. A sample configuration graphically depicts the resources and their dependencies within the service group. For more information about these resources and resource types, see the chapter VCS Resource Types and Agents in the *Veritas Cluster Server Administrator's Guide*.

In the sample dependency graph below, the shared virtual disk is configured using the NetApp Filer (NetAppFiler) resource. The virtual name for the SQL Server is created using the Lanman resource. The service group IP address for the SQL Server is configured using the IP and NIC resources. The Netapp SnapDrive mount point is created using the NetappSnapDrive resource. SQL Server 2000 registry is replicated using the RegRep and RegRepNetAppSnapDrive resources. The SQL Server 2000 resource comes online after each of these resources are brought online.

Figure A-1 SQL Server 2000 service group dependency graph



Agent for MSSearch service

The agent for MSSearch service is represented by the MSSearch resource type.

Resource type definition

```
type MSSearch (  
    static i18nstr ArgList[] = { AppName }  
    str AppName  
)
```

Attribute definitions

Review the following information to familiarize yourself with the required agent attributes for an MSSearch resource type. This information will assist you during the agent configuration.

Table A-6 MSSearch service agent required attribute


Required Attribute	Type and Dimension	Definition
AppName	string-scalar	The name of MSSearch instance to be monitored. See “ MSSearch service ” on page 20 for more information.

Sample dependency graph

The following examples describe typical service groups configured to monitor the state of MSSearch service in a VCS cluster. A sample configuration graphically depicts the resources and their dependencies within the service group. For more information about these resources and resource types, see the chapter VCS Resource Types and Agents in the *Veritas Cluster Server Administrator's Guide*.

In the sample configuration shown in the dependency graph below, the shared virtual disk is configured using the NetApp File (NetAppFile) resource. The MSSearch service depends upon the SQL Server 2000 service. The virtual name for the SQL Server is created using the Lanman resource. The service group IP address for the SQL Server is configured using the IP and NIC resources. The Netapp SnapDrive mount point is created using the NetappSnapDrive resource. SQL Server 2000 and MSSearch registry is replicated using the RegRep and RegRepNetAppSnapDrive resources. Different mounts are used for SQL Server 2000 and MSSearch registry replication. The MSSearch resource comes online after each of these resources are brought online.

Figure A-2 Typical MSSearch service service group dependency graph



SQL Server 2005 agents

SQL Server 2005 agent includes agents for SQL Server 2005, SQL Agent service, and SQL Server 2005 Analysis service.

Agent for SQL Server 2005

The agent for SQL Server 2005 is represented by the `SQLServer2005` resource type.

Resource type definition

```
type SQLServer2005 (  
    static il8nstr ArgList[] = {  
Instance, "LanmanResName:VirtualName",  
SQLOnlineTimeout, SQLOfflineTimeout,  
  
DetailMonitor, SQLDetailMonitorTimeout,  
  
Username, Domain, Password, SQLFile }  
    str Instance  
    str LanmanResName  
    int SQLOnlineTimeout = 90  
    int SQLOfflineTimeout = 90  
    boolean DetailMonitor = 0  
    boolean FaultOnDMScriptFailure = 1  
    int SQLDetailMonitorTimeout = 30  
    il8nstr Username  
    il8nstr Domain  
    str Password  
    il8nstr SQLFile  
)
```

Attribute definitions

Review the following information to familiarize yourself with the required agent attributes for a SQLServer2005 resource type. This information will assist you during the agent configuration.

Table A-7 SQL Server 2005 agent required attributes

Required Attributes	Type and Dimension	Definition
Instance	string-scalar	Name of instance to monitor. If the attribute is blank, the agent monitors the default instance.
LanmanResName	string-scalar	The Lanman resource name on which the SQL Server 2005 resource depends.
SQLOnlineTimeout	integer-scalar	Number of seconds that can elapse before online entry point aborts. Default is 90.
SQLOfflineTimeout	integer-scalar	Number of seconds that can elapse before offline entry point aborts. Default is 90.

Table A-8 SQL Server 2005 agent optional attributes

Optional Attributes	Type and Dimension	Definition
DetailMonitor	boolean	Defines whether the agent performs detail monitoring of SQL Server 2005 database. The value 0 indicates the agent will not monitor the database in detail. The value 1 indicates the agent will. Default = 0. Note: If the attribute is set to 1, the attributes Username, Password, Domain, SQLDetailMonitorTimeOut, and SQLFile must be assigned appropriate values.
FaultOnDMScriptFailure	boolean	Defines whether the agent fails over the service group if the detail monitoring script execution fails. Default = 1 The value 1 indicates that the agent fails over the service group if detail monitoring script fails to execute. The value 0 indicates that it does not.

Table A-8 SQL Server 2005 agent optional attributes (continued)


Optional Attributes	Type and Dimension	Definition
SQLDetailMonitor Timeout	integer-scalar	Number of seconds that can elapse before the detail monitor routine aborts. Default is 30.
Username	string-scalar	The Microsoft Windows authentication name when logging in to a database for detail monitoring. This attribute must not be null if “DetailMonitor” attribute is set to 1. Note: This attribute can take localized values.
Domain	string-scalar	Domain for the user account. This attribute is used to create a trusted connection to the SQL Server 2005 instance if “DetailMonitor” attribute is set to 1. Note: This attribute can take localized values.
Password	string-scalar	Password for logging in to a database for in-depth monitoring. This attribute must not be null if “DetailMonitor” attribute is set to 1.
SQLFile	string-scalar	The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the “DetailMonitor” attribute is set to 1. Note: This attribute can take localized values.

Sample dependency graph

The following example describes typical service groups configured to monitor the state of SQL Server 2005 in a VCS cluster. A sample configuration graphically depicts the resources and their dependencies within the service group. For more information about these resources and resource types, see the chapter VCS Resource Types and Agents in the *Veritas Cluster Server Administrator's Guide*.

In the sample dependency graph below, the shared virtual disk is configured using the NetApp Filer (NetAppFiler) resource. The virtual name for the SQL Server is created using the Lanman resource. The service group IP address for the SQL Server is configured using the IP and NIC resources. The Netapp SnapDrive mount point is created using the NetappSnapDrive resource. SQL Server 2005 registry is replicated using the RegRep and RegRepNetAppSnapDrive resources. The SQL Server 2005 resource comes online after each of these resources are brought online.

Figure A-3 SQL Server 2005 service group dependency graph



Agent for SQL Server 2005 agent service

The agent for SQL Server 2005 Agent service is represented by the SQLAgService2005 resource type.

Resource type definition

```
type SQLAgService2005 (
    static i18nstr ArgList[] = {
"SQLServer2005ResName:Instance",

"LanmanResName:VirtualName" }
    str SQLServer2005ResName
    str LanmanResName
)
```

Attribute definitions

Review the following information to familiarize yourself with the required agent attributes for a SQLAgService2005 resource type. This information will assist you during the agent configuration.

Table A-9 SQL Server 2005 Agent service agent required attributes


Required Attributes	Type and Dimension	Definition
SQLServer2005ResName	string-scalar	The name of the SQLServer2005 resource on which the SQL Server 2005 Agent service resource depends.
LanmanResName	string-scalar	The Lanman resource name on which the SQL Server 2005 resource depends.

Sample dependency graph

The following example describes typical service groups configured to monitor the state of SQL Server 2005 Agent service in a VCS cluster. A sample configuration graphically depicts the resources and their dependencies within the service group. For more information about these resources and resource types, see the chapter VCS Resource Types and Agents in the *Veritas Cluster Server Administrator's Guide*.

In the sample dependency graph below, the shared virtual disk is configured using the NetApp Filer (NetAppFiler) resource. The virtual name for the SQL Server is created using the Lanman resource. The service group IP address for the SQL Server is configured using the IP and NIC resources. The Netapp SnapDrive mount point is created using the NetappSnapDrive resource. SQL Server 2005 registry is replicated using the RegRep and RegRepNetAppSnapDrive resources. The SQL Server 2005 Agent service resource depends on the SQL Server 2005 resource. The SQL Server 2005 Agent service resource comes online after each of these resources are brought online.

Figure A-4 SQL Server 2005 Agent Service service group dependency graph



Agent for SQL Server 2005 analysis service

The agent for SQL Server 2005 Analysis service is represented by the `SQLOLapService2005` resource type.

Resource type definition

```
type SQLOLapService2005 (
    static i18nstr ArgList[] = {
"SQLServer2005ResName:Instance",

"LanmanResName:VirtualName" }
    str SQLServer2005ResName
    str LanmanResName
)
```

Attribute definitions

Review the following information to familiarize yourself with the required agent attributes for a `SQLOLapService2005` resource type. This information will assist you during the agent configuration.

Table A-10 SQL Server 2005 Analysis service agent required attributes


Required Attributes	Type and Dimension	Definition
SQLServer2005ResName	string-scalar	The name of the SQLServer2005 resource on which the SQL Server 2005 Analysis service resource depends.
LanmanResName	string-scalar	The Lanman resource name on which the SQL Server 2005 resource depends.

Sample dependency graph

The following example describes typical service groups configured to monitor the state of SQL Server 2005 Analysis service in a VCS cluster. A sample configuration graphically depicts the resources and their dependencies within the service group. For more information about these resources and resource types, see the chapter VCS Resource Types and Agents in the *Veritas Cluster Server Administrator's Guide*.

In the sample dependency graph below, the shared virtual disk is configured using the NetApp Filer (NetAppFiler) resource. The virtual name for the SQL Server is created using the Lanman resource. The service group IP address for the SQL Server is configured using the IP and NIC resources. The Netapp SnapDrive mount point is created using the NetappSnapDrive resource. SQL Server 2005 registry is replicated using the RegRep and RegRepNetAppSnapDrive resources. The SQL Server 2005 Analysis service resource depends on the SQL Server 2005 resource. The SQL Server 2005 Analysis service resource comes online after each of these resources are brought online.

Figure A-5 SQL Server 2005 Analysis Service service group dependency graph



MSDTC agent

The MSDTC agent is represented by the MSDTC resource type.

Resource type definition

```
type MSDTC (
    static i18nstr ArgList[] = {"LanmanResName:VirtualName",

    "MountResName:MountPath", LogPath }
    str LanmanResName
    str MountResName
    i18nstr LogPath
)
```

Attribute definitions

Review the following information to familiarize yourself with the required agent attributes for an MSDTC resource type. This information will assist you during the agent configuration

Table A-11 MSDTC agent required attributes


Required Attributes	Type and Dimension	Definition
LanmanResName	string-scalar	Name of the Lanman resource on which the MSDTC resource depends.
MountResName	string-scalar	The mount resource name on which the MSDTC resource depends.
LogPath	string-scalar	The path for MSDTC logs. This attribute can take localized values.

Sample dependency graph

The following examples describe typical service groups configured to monitor the state of an MSDTC service group in a VCS cluster. A sample configuration graphically depicts the resources and their dependencies within the service group. For more information about these resources and resource types, see the chapter VCS Resource Types and Agents in the *Veritas Cluster Server Administrator's Guide*.

In the sample configuration shown in the dependency graph below, the shared virtual disk is configured using the NetApp Filer (NetAppFiler) resource. The virtual name for the MSDTC Server is created using the Lanman resource. The service group IP address for the MSDTC Server is configured using the IP and NIC resources. The Netapp SnapDrive mount point is created using the NetappSnapDrive resource. MSDTC registry is replicated using the RegRep and RegRepNetAppSnapDrive resources. The MSDTC resource comes online after each of these resources are brought online.

Figure A-6 MSDTC service group dependency graph



Sample configurations

The sample configurations in this appendix describe typical service groups configured to monitor the state of SQL Server 2000 and SQL Server 2005 in a VCS cluster using Network Appliance Filers to manage shared storage.

Sample configuration for a SQL Server 2000 cluster

The following sample configuration depicts a two node cluster. A SQL Server service group with MSSearch instance is configured in the cluster. The MSDTC service group is also configured on the same nodes.

```
include "types.cf"

cluster clsec (
    UserNames = { a = bIIc }
    ClusterAddress = "10.217.60.172"
    Administrators = { a }
)
remotecluster cl48_49 (
    ClusterAddress = "10.217.60.169"
)
heartbeat Icmp (
    ClusterList = { cl48_49 }
    AYATimeout = 30
    Arguments @cl48_49 = { "10.217.60.169" }
)
system VCSW2K244 (
)
group ClusterService (
    SystemList = { VCSW2K244 = 0 }
    AutoStartList = { VCSW2K244 }
)
IP csg_ip (
    Address = "10.217.60.172"
    SubNetMask = "255.255.252.0"
    MACAddress @VCSW2K244 = "00:0B:DB:95:9C:A0"
)
```

Sample configuration for a SQL Server 2000 cluster

```

NIC csg_nic (
    MACAddress @VCSW2K244 = "00:0B:DB:95:9C:A0"
)
Process wac (
    StartProgram @VCSW2K244 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wac.exe"
    StopProgram @VCSW2K244 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacstop.exe"
    MonitorProgram @VCSW2K244 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacmonitor.exe"
)
wac requires csg_ip
csg_ip requires csg_nic
// resource dependency tree
//
// group ClusterService
// {
//     Process wac
//     {
//         IP csg_ip
//         {
//             NIC csg_nic
//         }
//     }
// }

group sq (
    SystemList = { VCSW2K244 = 0 }
    ClusterList = { cl48_49 = 1, clsec = 0 }
    Authority = 1
    ClusterFailOverPolicy = Auto
)
IP sq-IP (
    Address = "10.217.60.173"
    SubNetMask = "255.255.252.0"
    MACAddress @VCSW2K244 = "00-0B-DB-95-9C-A0"
)
Lanman sq-Lanman (
    VirtualName = VTEST1
    IPResName = sq-IP
)
MSSearch sq-MSSearch (
    AppName = "SQLServer$TEST1"
)
NIC sq-NIC (
    MACAddress @VCSW2K244 = "00-0B-DB-95-9C-A0"
)
NetAppFiler sq-NetAppFiler (
    FilerName = netapp-vcs
    StorageIP = "10.212.88.100"
)
NetAppSnapDrive sq-NetAppSnapDrive (

```

```

        FilerResName = sq-NetAppFiler
        VolumeName = vol9
        ShareName = data09
        LUN = "d90.lun"
        MountPath = I
        Initiator @VCSW2K244 =
"iqn.1991-05.com.microsoft:vcs2k244.vcs.dev"
    )
    NetAppSnapDrive sq-NetAppSnapDrive-1 (
        FilerResName = sq-NetAppFiler
        VolumeName = vol9
        ShareName = data09
        LUN = "d91.lun"
        MountPath = J
        Initiator @VCSW2K244 =
"iqn.1991-05.com.microsoft:vcs2k244.vcs.dev"
    )
    NetAppSnapMirror sq-NetAppSnapMirror (
        FilerResName = sq-NetAppFiler
        VolumeName = vol9
    )
    RegRep sq-RegRep-MSSQL (
        MountResName = sq-NetAppSnapDrive-1
        ReplicationDirectory = "\\RegRep\\sq-RegRep-MSSQL"
        Keys = {
            "HKLM\\SOFTWARE\\Microsoft\\MSSQLServer\\Client" = "",
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\TEST1" = "",
            "HKLM\\SOFTWARE\\Network Appliance\\SnapManager for SQL
Server\\Server" = "" }
        ExcludeKeys = {
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\TEST1\\Setup",
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\TEST1\\SQLServerAgent\\Subsystems",
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\TEST1\\Tracking" }
    )
    RegRep sq-RegRep-MSSearch (
        MountResName = sq-NetAppSnapDrive-1
        ReplicationDirectory = "\\RegRep\\sq-RegRep-MSSearch"
        Keys = {

"HKLM\\Software\\Microsoft\\Search\\1.0\\Applications\\SQLServer$TE
ST1" = "",

"HKLM\\Software\\Microsoft\\Search\\1.0\\CatalogNames\\SQLServer$TE
ST1" = "",

"HKLM\\Software\\Microsoft\\Search\\1.0\\Databases\\SQLServer$TEST1
" = "",

```

Sample configuration for a SQL Server 2000 cluster

```

"HKLM\\Software\\Microsoft\\Search\\1.0\\Gather\\SQLServer$TEST1" =
"",
    "HKLM\\Software\\Microsoft\\Search\\1.0\\Gathering
Manager\\Applications\\SQLServer$TEST1" = "",

"HKLM\\Software\\Microsoft\\Search\\1.0\\Indexer\\SQLServer$TEST1"
= "" }
)
SQLServer2000 sq-SQLServer2000 (
    Critical = 0
    Instance = TEST1
    LanmanResName = sq-Lanman
    MountResName = sq-NetAppSnapDrive
    DetailMonitor = 1
    Username = administrator
    Domain = "vcs.dev"
    Password = AKEmBMiMPmFEgBgCGd
    SQLFile = "c:\\script.sql"
)
sq-SQLServer2000 requires sq-RegRep-MSSQL
sq-SQLServer2000 requires sq-Lanman
sq-SQLServer2000 requires sq-NetAppSnapDrive
sq-NetAppSnapDrive requires sq-NetAppSnapMirror
sq-RegRep-MSSQL requires sq-NetAppSnapDrive-1
sq-NetAppSnapDrive-1 requires sq-NetAppSnapMirror
sq-IP requires sq-NIC
sq-Lanman requires sq-IP
sq-NetAppSnapMirror requires sq-NetAppFiler
sq-MSSearch requires sq-SQLServer2000
sq-MSSearch requires sq-RegRep-MSSearch
sq-RegRep-MSSearch requires sq-NetAppSnapDrive-1

// resource dependency tree
//
// group sq
// {
//     MSSearch sq-MSSearch
//     {
//         SQLServer2000 sq-SQLServer2000
//         {
//             RegRep sq-RegRep-MSSQL
//             {
//                 NetAppSnapDrive sq-NetAppSnapDrive-1
//                 {
//                     NetAppSnapMirror sq-NetAppSnapMirror
//                     {
//                         NetAppFiler sq-NetAppFiler
//                     }
//                 }
//             }
//         }
//     }
//     Lanman sq-Lanman

```

```
//      {
//      IP sq-IP
//      {
//      NIC sq-NIC
//      }
//      }
//      NetAppSnapDrive sq-NetAppSnapDrive
//      {
//      NetAppSnapMirror sq-NetAppSnapMirror
//      {
//      NetAppFiler sq-NetAppFiler
//      }
//      }
//      }
//      RegRep sq-RegRep-MSSearch
//      {
//      NetAppSnapDrive sq-NetAppSnapDrive-1
//      {
//      NetAppSnapMirror sq-NetAppSnapMirror
//      {
//      NetAppFiler sq-NetAppFiler
//      }
//      }
//      }
//      }
//      }
```

Sample configuration for a SQL Server 2005 cluster

The following sample configuration depicts a two node SQL Server cluster with MSDTC server configured. A SQL Server service group with MSSearch service, SQL Server Agent service, and SQL Server Analysis service is configured in the cluster. The MSDTC service group is also configured on the same nodes.

```
include "types.cf"

cluster clnew4533 (
    UserNames = { a = eHHb }
    ClusterAddress = "10.217.60.166"
    Administrators = { a }
)

remoteclass clremote (
    ClusterAddress = "10.217.119.60"
)

heartbeat Icmp (
    ClusterList = { clremote }
    AYATimeout = 30
    Arguments @clremote = { "10.217.119.60" }
)

system VCSW2K246 (
)

system VCSW2K247 (
)

group ClusterService (
    SystemList = { VCSW2K246 = 0, VCSW2K247 = 1 }
    AutoStartList = { VCSW2K246, VCSW2K247 }
)

IP csg_ip (
    Address = "10.217.60.166"
    SubNetMask = "255.255.252.0"
    MACAddress @VCSW2K246 = "00:0B:DB:95:9E:6C"
    MACAddress @VCSW2K247 = "00:0B:DB:95:9D:F4"
)

NIC csg_nic (
    MACAddress @VCSW2K246 = "00:0B:DB:95:9E:6C"
    MACAddress @VCSW2K247 = "00:0B:DB:95:9D:F4"
)

Process wac (
```

```
        StartProgram @VCSW2K246 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wac.exe"
        StartProgram @VCSW2K247 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wac.exe"
        StopProgram @VCSW2K246 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacstop.exe"
        StopProgram @VCSW2K247 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacstop.exe"
        MonitorProgram @VCSW2K246 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacmonitor.exe"
        MonitorProgram @VCSW2K247 = "C:\\Program
Files\\Veritas\\Cluster Server\\bin\\wacmonitor.exe"
    )
```

```
csg_ip requires csg_nic
wac requires csg_ip
```

```
// resource dependency tree
//
// group ClusterService
// {
//   Process wac
//   {
//     IP csg_ip
//     {
//       NIC csg_nic
//     }
//   }
// }
// }
```

```
group sqlg (
    SystemList = { VCSW2K246 = 0, VCSW2K247 = 1 }
    ClusterList = { clnew4533 = 1, clremote = 0 }
    Authority = 1
    ClusterFailOverPolicy = Auto
)

IP sqlg-IP (
    Address = "10.217.60.167"
    SubNetMask = "255.255.252.0"
    MACAddress @VCSW2K246 = "00-0B-DB-95-9E-6C"
    MACAddress @VCSW2K247 = "00-0B-DB-95-9D-F4"
)

Lanman sqlg-Lanman (
    VirtualName = VHD
    IPResName = sqlg-IP
)

NIC sqlg-NIC (
```

Sample configuration for a SQL Server 2005 cluster

```

        MACAddress @VCSW2K246 = "00-0B-DB-95-9E-6C"
        MACAddress @VCSW2K247 = "00-0B-DB-95-9D-F4"
    )

    NetAppFiler sqlg-NetAppFiler (
        FileName = vcsfiler01
        StorageIP = "10.217.60.153"
    )

    NetAppSnapDrive sqlg-NetAppSnapDrive (
        Critical = 0
        FilerResName = sqlg-NetAppFiler
        VolumeName = vol3
        ShareName = data03
        LUN = "d30.lun"
        MountPath = I
        Initiator @VCSW2K246 =
"iqn.1991-05.com.microsoft:vcs2k246.vcs.dev"
        Initiator @VCSW2K247 =
"iqn.1991-05.com.microsoft:vcs2k247.vcs.dev"
    )

    NetAppSnapDrive sqlg-NetAppSnapDrive-1 (
        Critical = 0
        FilerResName = sqlg-NetAppFiler
        VolumeName = vol3
        ShareName = data03
        LUN = "d31.lun"
        MountPath = J
        Initiator @VCSW2K246 =
"iqn.1991-05.com.microsoft:vcs2k246.vcs.dev"
        Initiator @VCSW2K247 =
"iqn.1991-05.com.microsoft:vcs2k247.vcs.dev"
    )

    NetAppSnapMirror sqlg-NetAppSnapMirror (
        FilerResName = sqlg-NetAppFiler
        VolumeName = vol3
    )

    RegRep sqlg-RegRep-MSSQL (
        MountResName = sqlg-NetAppSnapDrive-1
        ReplicationDirectory = "\\RegRep\\sqlg-RegRep-MSSQL"
        Keys @VCSW2K246 = {
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\MSSQLServer" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_MSSQLServer.reg",
            "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\PROVIDERS" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_PROVIDERS.reg",

```

```

        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\Replication" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_Replication.reg",
        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\SQLServerAgent" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_SQLServerAgent.reg",
        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\SQLServerSCP" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_SQLServerSCP.reg" }
        Keys @VCSW2K247 = {
        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\MSSQLServer" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_MSSQLServer.reg",
        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\PROVIDERS" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_PROVIDERS.reg",
        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\Replication" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_Replication.reg",
        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\SQLServerAgent" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_SQLServerAgent.reg",
        "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
Server\\MSSQL.1\\SQLServerSCP" =
"SaveRestoreFile:sqlg-RegRep-MSSQL_SQLServerSCP.reg" }
        ExcludeKeys = { "" }
    )

    SQLServer2005 sqlg-SQLServer2005 (
        Critical = 0
        Instance = HD
        LanmanResName = sqlg-Lanman
    )

    sqlg-SQLServer2005 requires sqlg-RegRep-MSSQL
    sqlg-SQLServer2005 requires sqlg-Lanman
    sqlg-SQLServer2005 requires sqlg-NetAppSnapDrive
    sqlg-NetAppSnapDrive requires sqlg-NetAppSnapMirror
    sqlg-RegRep-MSSQL requires sqlg-NetAppSnapDrive-1
    sqlg-NetAppSnapDrive-1 requires sqlg-NetAppSnapMirror
    sqlg-IP requires sqlg-NIC
    sqlg-Lanman requires sqlg-IP
    sqlg-NetAppSnapMirror requires sqlg-NetAppFiler

    // resource dependency tree
    //
    // group sqlg
    // {
    //     SQLServer2005 sqlg-SQLServer2005
    //     {
    //         RegRep sqlg-RegRep-MSSQL

```

Sample configuration for a SQL Server 2005 cluster

```

//      {
//      NetAppSnapDrive sqlg-NetAppSnapDrive-1
//      {
//      NetAppSnapMirror sqlg-NetAppSnapMirror
//      {
//      NetAppFiler sqlg-NetAppFiler
//      }
//      }
//      }
//      Lanman sqlg-Lanman
//      {
//      IP sqlg-IP
//      {
//      NIC sqlg-NIC
//      }
//      }
//      NetAppSnapDrive sqlg-NetAppSnapDrive
//      {
//      NetAppSnapMirror sqlg-NetAppSnapMirror
//      {
//      NetAppFiler sqlg-NetAppFiler
//      }
//      }
//      }
//      }
//      }

```

Configuring the Symantec License Inventory Agent

This appendix includes the following topics:

- [About the Symantec License Inventory Manager](#)
- [When the Symantec License Inventory Agent is installed](#)
- [When the server and access points are installed](#)
- [What you can do with the agent](#)
- [How to remove the agent](#)

About the Symantec License Inventory Manager

The Symantec License Inventory Manager (license inventory manager) is an enterprise asset management tracking tool. It inventories Symantec Information Availability products in your network and consolidates critical information on the deployment of these products. You can use the information to:

- Determine all the Symantec software products and licenses being used in your enterprise
- Achieve easier license self-compliance management
- Know your Enterprise License Agreement deployment status
- Reduce administrative overhead for managing license compliance
- Renew the support and maintenance agreements that are based on the licenses you have deployed
- Gain more control over your Symantec software usage
- Manage the department chargebacks based on actual software usage

- Use more flexible licensing and pricing models
- Exploit detailed deployment data to perform return on investment analyses for purchased software

The license inventory manager is a three-tiered system that consists of a server tier, access point tier, and an agent tier. The server tier is the Symantec License Inventory Server, which consolidates and stores the information that it gathers from the agents and access points.

The optional access point tier includes Symantec License Inventory Access Points and serves as a consolidation layer between the agents and server.

The agent tier includes Symantec License Inventory Agents, which are deployed on individual hosts in a network. Each agent gathers product information on the supported Symantec products that are installed on the agent's host. The agents then send the information they have gathered to an access point or the server.

When the Symantec License Inventory Agent is installed

The Symantec License Inventory Manager is installed using installation media available separately. To order a Symantec License Inventory Manager license and installation media kit, contact your Symantec sales representative.

The installation media provides online documentation for the Symantec License Inventory Manager. You can order printed copies of the documentation from your sales representative. The documents you can order include:

- *Symantec License Inventory Manager Installation and Configuration Guide*
- *Symantec License Inventory Manager Administrator's Guide*
- *Symantec License Inventory Manager User's Guide*

The installation media provides online documentation with details on all of the topics that are discussed in this appendix.

For the latest information on updates, patches, and software issues regarding this product, read the release notes located at:

<http://entsupport.symantec.com/docs/285602>

The Symantec product installer installs or upgrades the agent on the host with the Symantec product. The agent is installed in the following folder:

C:\Program Files\Veritas\License Inventory Manager\Agent

The agent is installed with a default configuration that minimizes its impact on a running system. The minimum configuration prevents remote communication with the agent to keep its data and interfaces secure.

When the server and access points are installed

The server and access points are not installed automatically. If you want to use the Symantec License Inventory Manager, you must manually install the server and, optionally, the access points. After you install the server and access points, the agents can gather information and you can create inventory reports.

You can install the server and access points from the Symantec License Inventory Manager installation media.

What you can do with the agent

After the agent is installed, you can use it to track Symantec products on the system on which it was installed. You can remove the agent if you choose not to use it.

To use the agent, however, you must manually configure it to enable remote communication between the agent and its server or access point. Complete instructions for reconfiguring the agent are provided in the *Symantec License Inventory Manager 4.2 Release Notes*. You can download this document from the following URL: <http://entsupport.symantec.com/docs/285602>

How to remove the agent

If you do not want to use the Symantec License Inventory Manager, you can remove the agent using Add/Remove Programs. It is listed as the Symantec License Inventory Agent program in Add/Remove Programs.

Later, you can reinstall the agent with the Symantec License Inventory Manager installation disc. This disc is available in the Symantec License Inventory Manager kit.

Index

A

- Active/Active
 - configuration diagram 116
 - configuration overview 115
 - configuring network and storage 118
 - configuring VCS SQL Server service group 122
 - creating a user-defined database 122
- Active/Active environment
 - definition of 23
- Administrative Console 36
- application failure, detecting 21
- attribute definitions
 - MSDTC agent 192
 - MSSearch agent 182
 - SQL Server 2000 agent 179
 - SQL Server 2005 agent 185
 - SQL Server Agent Service agent 188
 - SQL Server Analysis Service agent 190

C

- clusters
 - configuring (DR) 130, 135
 - configuring (HA) 42
 - setting the internal name of the clustered instance (2000) 70
 - setting the internal name of the clustered instance (2005) 79
- command line
 - installing 39
- Complete/Custom 36
- configurations
 - Disaster Recovery 28, 127
- configuring
 - cluster (DR) 130, 135
 - cluster (HA) 42
 - SQL Server 2000 DR overview 129
 - SQL Server 2005 DR overview 134
 - SQL Server service group 2000 131
 - SQL Server service group 2005 136
- converting
 - moving data files and databases (HA) 112

standalone SQL conversion to HA 106

D

- database
 - new database for SQL Server 2000 90
 - user-defined for SQL Server 2000 90
- dependency graph
 - MSDTC agent 193
 - MSSearch agent 183
 - SQL Server 2000 agent 181
 - SQL Server 2005 agent 187
 - SQL Server Agent Service agent 189
 - SQL Server Analysis Service agent 191
- disaster recovery (DR)
 - deploying for SQL Server 2000 129
 - deploying for SQL Server 2005 134
 - illustrated 28, 127
 - protection in a SQL Server environment 126
 - rationale 125
 - setting up a secondary SQL Server 2000 site 132
 - setting up a secondary SQL Server 2005 site 137
 - typical configuration 127
- Disaster Recovery configuration 28, 127
- driver signing options 35
- resetting 38

G

- Global Cluster Option (GCO)
 - prerequisites 142

I

- installation
 - installing the base product 36
 - license management 30
 - repairing 155
 - using the command line 39
- installing
 - before installing SQL Server 2000 on a second

- node 67
- before installing SQL Server 2005 on a second node 77
- command line 39
- SQL Server 2000 on a second node (VCS) 69
- SQL Server 2000 on the first node 63
- SQL Server 2000 tasks (DR) 129
- SQL Server 2000, secondary site tasks 132
- SQL Server 2005 on a second node (VCS) 79
- SQL Server 2005 on the first node 73, 135
- SQL Server 2005 tasks (DR) 134
- SQL Server 2005, secondary site tasks 137
- VCS for NetApp (DR) 130, 135
- verifying the installation 104
- internal name
 - setting for the clustered instance (SQL 2000) 70
 - setting for the clustered instance (SQL 2005) 79

L

- licensing 29
 - adding or removing 30
 - management 30
- logging
 - VCW logs 158
 - VCWsilent logs 159
- Logs
 - VCS 157
 - VCW 158
 - VCWsilent 159

M

- MSDTC
 - configuring the client (VCS) 103
 - creating an MSDTC service group 101
 - HA configuration overview 98
 - service group configuration (VCS) 98
- MSDTC agent
 - attributes 192
 - operations 19
- MSDTC agent attributes
 - LanmanResName 192
 - LogPath 192
 - MountResName 192
- MSDTC service 20
- MSSearch agent
 - attributes 182

- operations 17
- MSSearch agent attribute, AppName 182
- MSSearch service 20

O

- option
 - driver signing 35, 38

P

- parameters
 - setup.exe 39
- prerequisites
 - SQL Server 2000 (DR) 129
 - SQL Server 2005 (DR) 134
 - SQL Server standalone conversion (HA) 106

Q

- Query Analyzer 70

R

- repairing the installation 155
- replication
 - defined 126
- resetting
 - driver signing options 38
- resource type definition
 - MSDTC agent 192
 - MSSearch agent 182
 - SQL Server 2000 agent 178
 - SQL Server 2005 agent 184
 - SQL Server Agent service agent 188
 - SQL Server Analysis Service agent 190

S

- sample configurations
 - SQL Server (HA) 116
 - standalone SQL Server conversion (HA) 107
- Security Services
 - configuring 48
- service groups
 - administering global groups 147
- setting
 - signing options 35
- setup.exe
 - parameters 39
- setup.exe parameters

- InstallDirPath 40
 - licenseKey 40
 - options 40
 - RebootMode 40
 - solution 39
 - SQL agent
 - removing 154
 - uninstalling 154
 - SQL cluster
 - Disaster Recovery setup 28, 127
 - SQL cluster configuration
 - Disaster Recovery 28, 127
 - SQL Server 2000
 - before installing on a second node 67
 - configuring the SQL Server service group 131
 - creating a new database 90
 - DR configuration overview 129
 - installing on a second node (HA) 69
 - installing on the first node 63
 - new installation tasks (DR) 129
 - overview of user-defined database 90
 - removing shared files 68
 - removing shared files on the second node 79
 - stopping the SQL Server service 68
 - tasks for setting up a secondary site 132
 - SQL Server 2000 database agent
 - MSSearch agent 17
 - SQL Server agent 17
 - SQL Server 2005
 - before installing on a second node 77
 - configuring the SQL Server service group 136
 - DR configuration overview 134
 - HA sample configuration 116
 - installing on a second node 79
 - installing on the first node 73, 135
 - Management Studio 79
 - new installation tasks (DR) 134
 - removing shared files 78
 - start a SQL Server service 80
 - stopping the SQL Server service 77
 - tasks for setting up a secondary site 137
 - SQL Server 2005 agent
 - attributes 185
 - SQL Server agent
 - attributes 179
 - operations 17, 18
 - SQL Server Agent Service agent
 - attributes 188
 - SQL Server Analysis Service agent
 - attributes 190
 - SQL Server database agent
 - monitoring options 21
 - MSDTC agent 19
 - SQL Server services
 - MSDTC 20
 - MSSearch 20
 - standalone SQL Server
 - configuring storage hardware and network 108
 - installing and configuring SFW (HA) 111
 - overview of HA configuration 106
 - preparing for conversion (HA) 110
 - sample HA configuration 107
- ## U
- uninstalling
 - SQL agent 154
- ## V
- ### VCS
- configuring the cluster for DR 130, 135
 - configuring the cluster for HA 42
 - configuring the SQL Server service group 131, 136
 - configuring the SQL Server service group (Active-Active) 122
- ### VCS for NetApp
- uninstalling 153, 154
- ### VCS logs
- 158

