

Veritas™ Cluster Server Database Agent for Microsoft SQL Configuration Guide

Windows Server 2003
Windows Server 2008

5.1 Service Pack 2



Veritas Cluster Server Database Agent for Microsoft SQL Configuration Guide

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

Product version: 5.1. Service Pack 2

Document version: 5.1.SP2.0

Legal Notice

Copyright © 2010 Symantec Corporation. All rights reserved.

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

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

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

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

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

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

Technical Support

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

Symantec's support offerings include the following:

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

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

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

All support services will be delivered in accordance with your support agreement

and the then-current enterprise technical support policy.

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level

- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Contents

Chapter 1	Introducing the VCS database agent for SQL	
	About the VCS database agent for SQL	12
	Agents for SQL Server 2000	12
	Agents for SQL Server 2005	12
	About SQL Server services	13
	MSSearch service	13
	Agent service	13
	Analysis service	13
	Microsoft Data Transaction Coordinator (MSDTC) service	13
	About the VCS agents for SQL Server 2000	14
	Agent for SQL Server 2000	14
	Agent for MSearch service	14
	About the VCS agents for SQL Server 2005	15
	Agent for SQL Server 2005	15
	Agent for SQL Server 2005 agent service	15
	Agent for SQL Server 2005 analysis service	16
	About the VCS agent for MSDTC	16
	About application monitoring options	17
	How the agent detects application failure	17
	SQL Server in an Active-Active clustered environment	18
	Installing SQL Server 2000 in an Active-Active environment	19
	Typical SQL Server 2000 configuration in a VCS cluster	20
	Typical SQL Server 2005 configuration in a VCS cluster	21
	Setting up a SQL Server cluster	23
Chapter 2	Installing the VCS database agent for SQL	
	About installing the VCS database agent for SQL	26
	Before you install the agent	26
	Installing the VCS database agent for SQL	26
	Configuring the cluster	28
	Configuring notification	39
Chapter 3	Installing and configuring SQL Server	
	About installing and configuring SQL	44
	Prerequisites installing SQL Server	44

About managing storage using SFW	44
About installing and configuring SQL Server 2000	53
About installing SQL Server 2000 on the first node	53
About installing SQL Server 2000 on failover nodes	56
Setting the internal name to clustered instance (SQL 2000)	57
About installing and configuring SQL Server 2005	59
About installing SQL Server 2005 on the first node	59
Installing SQL Server 2005	60
Setting SQL Server 2005 services to manual start	62
Stopping the SQL Server 2005 service	62
Installing and configuring SQL Server 2005 on failover nodes	63
Assigning ports for multiple SQL Server instances	65

Chapter 4 Configuring the SQL Server service group

About configuring the SQL service group	68
Prerequisites for configuring the SQL service group	68
Configuring SQL Server in a VCS cluster	69
About creating a SQL Server user-defined database	78
Create new volumes	78
Creating a SQL Server database	78
Adding the VCS storage resources	80

Chapter 5 Configuring the MSDTC service group

About configuring the MSDTC service group	84
Reviewing the configuration	84
Prerequisites for configuring the MSDTC service group	86
Configure MSDTC in a VCS cluster	88
Configuring an MSDTC service group	88
Configuring an MSDTC client	93

Chapter 6 Verifying the service group configuration

About verifying the service group	96
Bringing the service group online	96
From the Cluster Manager (Java Console)	96
From the Web Console	97
Switching the service group	97
From the Cluster Manager (Java Console)	97
From the Web Console	98
Taking the service group offline	98
From the Cluster Manager (Java Console)	98
From the Web Console	99
Disabling the service group	99

	From the Cluster Manager (Java Console)	99
	From the Web Console	99
Chapter 7	Removing the VCS database agent for SQL	
	About removing the VCS database agent for SQL	102
	Removing the VCS database agent for SQL	102
Chapter 8	Administering the service group	
	About administering the SQL service group	106
	Administer a SQL service group	106
	Modifying a SQL Server service group	106
	Deleting a SQL Server service group	107
	Administer an MSDTC service group	108
	Modifying an MSDTC service group	108
	Deleting an MSDTC service group	109
	Detail monitoring of a SQL Server database instance	110
	Enabling detail monitoring	110
	Disabling detail monitoring	113
Chapter 9	Troubleshooting VCS agents	
	About troubleshooting the VCS database agent for SQL	116
	VCS logging	116
	Error messages and descriptions	118
	Agent for SQL Server 2000	118
	Agent for MSSearch	120
	Agent for MSDTC	121
	Agent for SQL Server 2005	122
	Agent for SQL Server 2005 agent service	126
	Agent for SQL Server 2005 analysis service	128
Appendix A	Resource type definitions	
	About resource type definitions	132
	Agent for SQL Server 2000	132
	Resource type definition	132
	Attribute definitions	133
	Sample dependency graph	134
	Agent for MSSearch service	135
	Resource type definition	135
	Attribute definitions	136
	Sample dependency graph	136
	Agent for SQL Server 2005	137

Resource type definition	137
Attribute definitions	138
Sample dependency graph	140
Agent for SQL Server 2005 Agent service	140
Resource type definition	140
Attribute definitions	141
Sample dependency graph	141
Agent for SQL Server 2005 Analysis service	142
Resource type definition	142
Attribute definitions	143
Sample dependency graph	143
MSDTC agent	144
Resource type definition	144
Attribute definitions	145
Sample dependency graph	146

Appendix B Using the virtual MMC viewer

About using the virtual MMC viewer	148
Viewing DTC transaction information	148

Appendix C Sample configurations

About sample configurations	152
Sample configuration for a SQL Server 2000 cluster	152
Sample configuration for a SQL Server 2005 cluster	156

Index	161
-------------	-----

Introducing the VCS database agent for SQL

This chapter contains the following topics:

- [“About the VCS database agent for SQL”](#) on page 12
- [“About SQL Server services”](#) on page 13
- [“About the VCS agents for SQL Server 2000”](#) on page 14
- [“About the VCS agents for SQL Server 2005”](#) on page 15
- [“About the VCS agent for MSDTC”](#) on page 16
- [“About application monitoring options”](#) on page 17
- [“How the agent detects application failure”](#) on page 17
- [“SQL Server in an Active-Active clustered environment”](#) on page 18
- [“Typical SQL Server 2000 configuration in a VCS cluster”](#) on page 20
- [“Typical SQL Server 2005 configuration in a VCS cluster”](#) on page 21
- [“Setting up a SQL Server cluster”](#) on page 23

About the VCS database agent for SQL

Microsoft SQL Server is a relational database management system (RDBMS) used for building, managing, and deploying business applications. The SQL Server infrastructure provides services such as jobs, notification, and in-built replication. SQL Server also uses MSSearch service for full-text search indexing and Microsoft Distributed Transaction Coordinator (MSDTC) to co-ordinate transactions.

The VCS database agent for Microsoft SQL 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 VCS database agent for Microsoft SQL provides “Active-Active” support for SQL Server. VCS provides separate agents for SQL Server 2000 and SQL Server 2005.

Agents for SQL Server 2000

Agents for SQL Server 2000 are as follows:

- 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

Agents for SQL Server 2005 are as follows:

- 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.
- Agent for SQL Server MSDTC service
The VCS database agent for Microsoft SQL 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 SQL Server services

The VCS database agent for Microsoft SQL provides high availability for the MSSearch and MSDTC services.

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.

Agent service

SQL Server Agent is a Microsoft SQL Server 2005 component that lets you automate some administrative tasks. SQL Server Agent runs jobs, monitors SQL Server, and processes alerts.

Analysis service

Microsoft SQL Server 2005 Analysis Services (SSAS) uses both server and client components to supply online analytical processing (OLAP) and data mining functionality for business intelligence applications. Analysis Services allows you to design, create, and manage multidimensional structures that contain detail and aggregated data from multiple data sources, such as relational databases, in a single unified logical model supported by built-in calculations. Analysis Services provides fast, intuitive, top-down analysis of large quantities of data built on this unified data model, which can be delivered to users in multiple languages and currencies.

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.

About the VCS agents for SQL Server 2000

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 functions include the following:

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 “ About application monitoring options ” on page 17.
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.

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.

Specific agent operations include the following:

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

About the VCS agents for SQL Server 2005

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 the following:

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 “ About application monitoring options ” on page 17.
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 the following:

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.

Note: VCS database agent for Microsoft SQL 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 agent.

Specific agent operations include the following:

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

About the VCS agent for MSDTC

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 the following:

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

About application monitoring options

The VCS database agent for Microsoft SQL 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.

See “[Detail monitoring of a SQL Server database instance](#)” on page 110.

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 virtual computer context of the services. If the SQL service is not started in the virtual computer 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 agent detects application failure

Detecting failure of SQL Server

The SQL Server 2000 agent monitors the SQL Server 2000 service and the SQL agent service. If the SQL Server 2000 service is not running, the agent returns a failure status. If the SQL agent service is not running, the agent declares the state as UNKNOWN.

[Table 1-1](#) on page 17 lists the SQL Server 2000 agent behavior.

Table 1-1 SQL agents behavior

SQL Server 2000 Service	SQL 2000 Agent Service	Action taken by the Monitor function
STARTED	STARTED	Returns the state as ONLINE.
STARTED	STOPPED	Attempts to start the SQL agent service. If the monitor routine succeeds, it returns ONLINE. Otherwise, it returns UNKNOWN.

The SQL Server 2005 agent monitors only the SQL Server 2005 service. If the SQL Server 2005 service is not running, the agent returns a failure status and declares the state as OFFLINE.

The VCS database agent for Microsoft SQL provides a sample script for detail monitoring of the SQL database. The script contains SQL statements and is available under the directory %VCS_HOME%\Samples.

If detail monitoring is enabled, the agent launches this script. If the SQL statements are executed successfully, the agent declares the service group as online.

Other than the default sample script, you can provide a custom script file for detail monitoring of the SQL Server database.

Detecting failure of SQL Server 2000 configured with MSSearch

If SQL Server is configured with MSSearch, the MSSearch instance is attached to a particular SQL instance. In such a scenario, the agent also monitors the state of the MSSearch service for the particular SQL instance. If the MSSearch instance is not detected, the agent declares only the MSSearch service as OFFLINE.

Detecting failure of SQL Server 2005 agent service

The agent for SQL Server 2005 agent service monitors the SQL Server 2005 agent service. If the agent does not detect the SQL Server 2005 agent service, the agent declares the SQL Server 2005 agent service as OFFLINE.

Detecting failure of SQL Server 2005 analysis service

The agent for SQL Server 2005 analysis service monitors the SQL Server 2005 analysis service. If the analysis service is not detected, the agent declares the SQL Server 2005 analysis service as offline.

Detecting failure of MSDTC

The MSDTC agent monitors the MSDTC service to detect failure. The agent detects an MSDTC failure if the MSDTC service is not running.

SQL Server in an Active-Active clustered environment

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 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 so that a single node can never host more than 16 instances.

Installing SQL Server 2000 in an Active-Active environment

To install SQL Server 2000 in an Active-Active environment, you must repeat the following procedures for each additional instance in an Active-Active configuration:

- See [“Installing the VCS database agent for SQL”](#) on page 25.
- See [“Installing and configuring SQL Server”](#) on page 43.
- See [“Configuring the SQL Server service group”](#) on page 67.

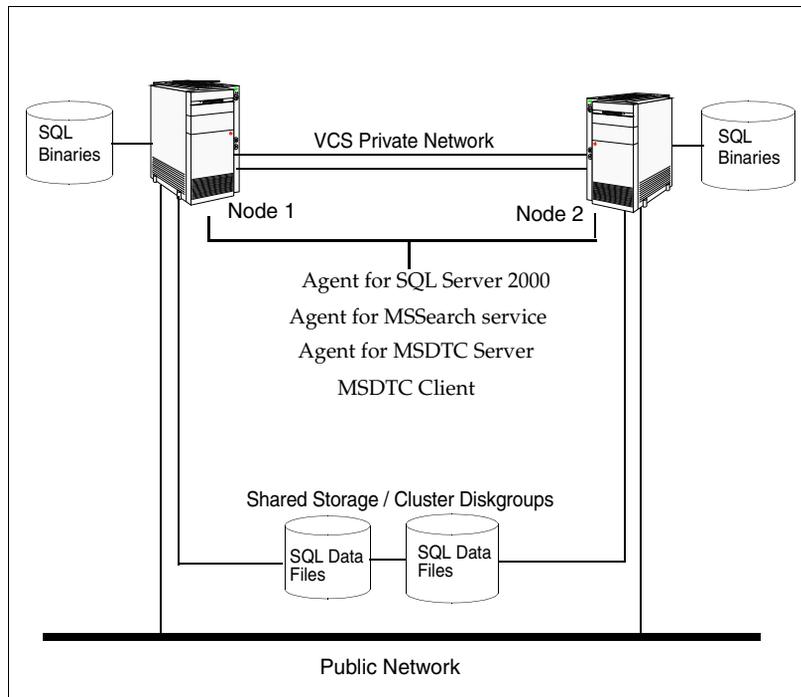
Typical SQL Server 2000 configuration in a VCS cluster

A typical SQL Server 2000 configuration involves two cluster nodes accessing a shared storage. The SQL Server binaries are installed on the nodes. The shared storage is used to store SQL Server data files and the MSDTC log files. The nodes access the shared storage.

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 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.

[Figure 1-1](#) on page 20 depicts a two node cluster hosting a SQL Server 2000 service group with MSSearch service configured. MSDTC resource is also configured on the same nodes.

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



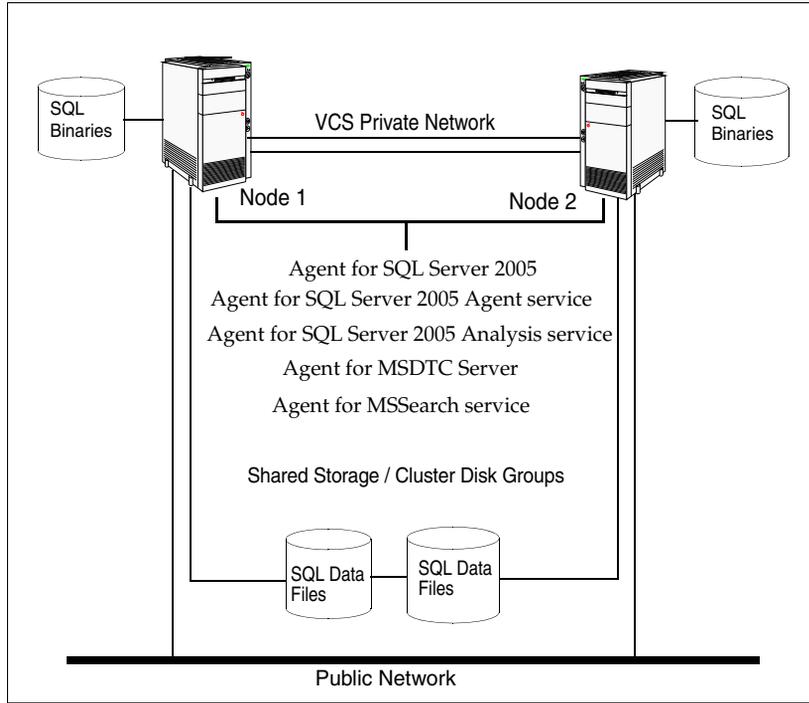
Typical SQL Server 2005 configuration in a VCS cluster

A typical SQL Server 2005 configuration involves two cluster nodes accessing a shared storage. The SQL Server binaries are installed on the nodes. The shared storage is used to store SQL Server data files and the MSDTC log files. The nodes access the shared storage.

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 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.

[Figure 1-2](#) on page 22 depicts a two node cluster hosting a SQL Server 2005 service group with the different services configured. MSDTC resource is also configured on the same nodes.

Figure 1-2 Typical SQL Server 2005 configuration in a VCS cluster



Setting up a SQL Server cluster

Setting up a SQL Server cluster in a VCS environment involves the following tasks:

- Installing the VCS database agent for Microsoft SQL.
See “[Installing the VCS database agent for SQL](#)” on page 25 for more information.
- Setting up a VCS cluster.
See “[Configuring the cluster](#)” on page 28.
- Installing and configuring SQL Server.
See “[Installing and configuring SQL Server](#)” on page 43 for more information.
- Configuring a SQL Server service group.
See “[Configuring the SQL Server service group](#)” on page 67 for more information.

Each task is described in detail in subsequent chapters.

Installing the VCS database agent for SQL

This chapter contains the following topics:

- [“About installing the VCS database agent for SQL”](#) on page 26
- [“Before you install the agent”](#) on page 26
- [“Installing the VCS database agent for SQL”](#) on page 26
- [“Configuring the cluster”](#) on page 28

About installing the VCS database agent for SQL

This chapter describes how to install the VCS database agent for Microsoft SQL in a VCS cluster. The agent is installed using the installer for Veritas Storage Foundation for Windows. The installer installs the agent on selected nodes in the cluster and adds the SQL Server resource types to the cluster configuration. Perform these steps only if you have not installed the VCS database agent for SQL while installing VCS.

Before you install the agent

This section lists the prerequisites for installing VCS database agent for Microsoft SQL Server in a VCS cluster.

- Verify that Storage Foundation HA for Windows (SFW HA) is installed on all nodes in the cluster.
Refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for installation steps.
- Verify you have local Administrator privileges on the node where you are installing the agent.

Installing the VCS database agent for SQL

Perform the following steps to install the VCS database agent for Microsoft SQL. Repeat these steps on all systems on which you wish to configure SQL with VCS.

To install the agent

- 1 Start Windows Add/Remove Programs, click **Storage Foundation HA 5.1 SP2 for Windows (Server Components)** and click **Change**.
- 2 In the Storage Foundation 5.1 SP2 for Windows panel, click **Add or Remove** and click **Next**.
- 3 In the Storage Foundation HA options panel, click **Next**.
- 4 Check **Veritas Cluster Server Database Agent for SQL** and click **Next**.
The disk space required for the installation is displayed towards the bottom of the screen. When you add or remove an option, the total space changes.
- 5 The installer validates the system for prerequisites. After the system is accepted, click **Next**.
If a system is rejected, the Comments column displays the cause for rejecting the system. Highlight the system to view detailed information

about the failure in the Details box. Resolve the error, highlight the system from the list, and click **Validate Again**.

- 6 An informational message appears if you selected the DMP option. Review the information and click **OK** to continue.
- 7 Review the summary of your selections and click **Update** to start the installation. The installer displays the status of installation.
- 8 After the installation is complete, review the installation report and click **Next** and then click **Finish**.

Configuring the cluster

After installing the agent, set up the components required to run a cluster. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also 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.

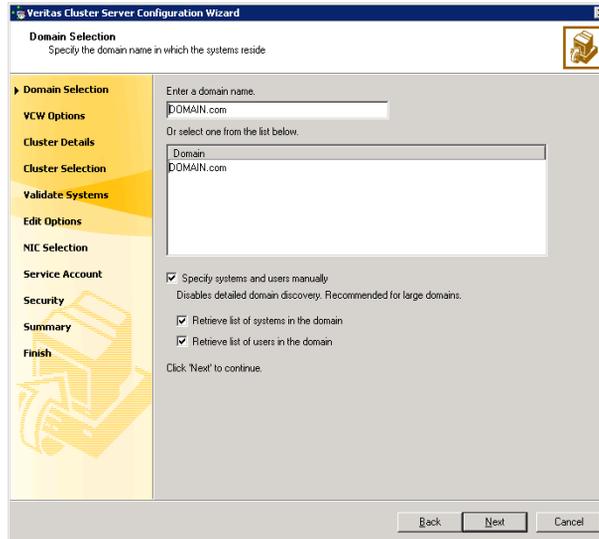
Complete the following tasks before configuring a cluster:

- You must have administrator privileges on the system where you run the wizard. The user account must be a domain account.
- You must have administrative access to all systems selected for cluster operations. Add the domain user to the Local Administrators group of each system.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group.
- When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. 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.
This account does not require Domain Administrator privileges.
- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.

To configure a VCS cluster

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

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

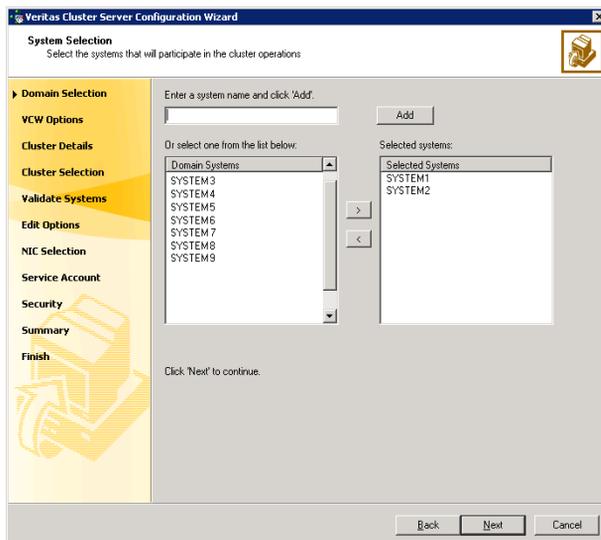


Do one of the following:

- To discover information about all systems and users in the domain:
 - Clear the **Specify systems and users manually** check box.
 - Click **Next**.Proceed to [step 8](#) on page 30.
- 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 30. Otherwise, proceed to the next step.

- 5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster. Proceed to [step 8](#) on page 30.

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



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

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

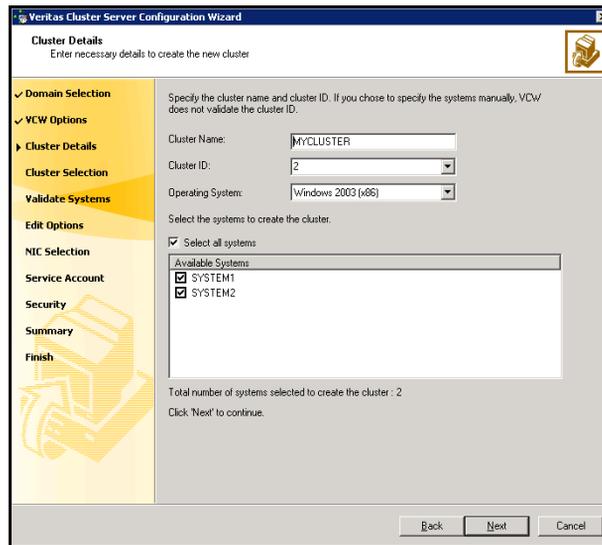
A system can be rejected for any of the following reasons:

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

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

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

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



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

Cluster ID Select a cluster ID from the suggested cluster IDs in the drop-down list or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255.

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

Operating System From the drop-down list select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a 32-bit and a 64-bit system in the same cluster.

Available Systems Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously.

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

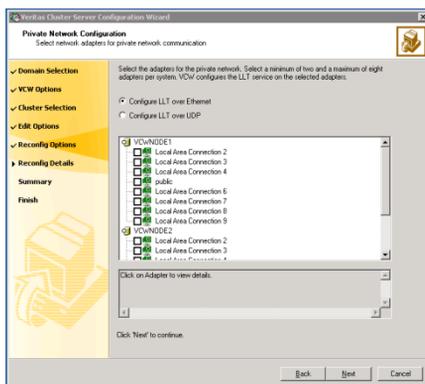
10 The wizard validates the selected systems for cluster membership. After the systems are validated, click **Next**.

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

If you chose to configure a private link heartbeat in the earlier step, proceed to the next step. Otherwise, proceed to [step 12](#) on page 34.

11 On the Private Network Configuration panel, configure the VCS private network and then click **Next**. You can configure the VCS private network either over the ethernet or over the User Datagram Protocol (UDP) layer. Do one of the following:

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



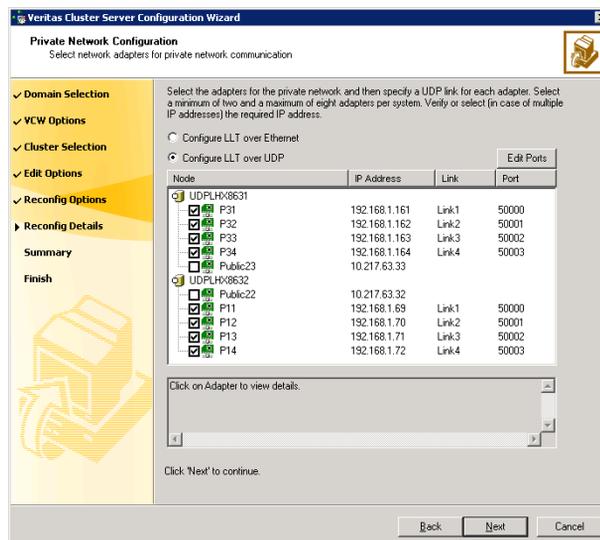
- Select the check boxes next to the two NICs to be assigned to the private network.

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

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

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

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

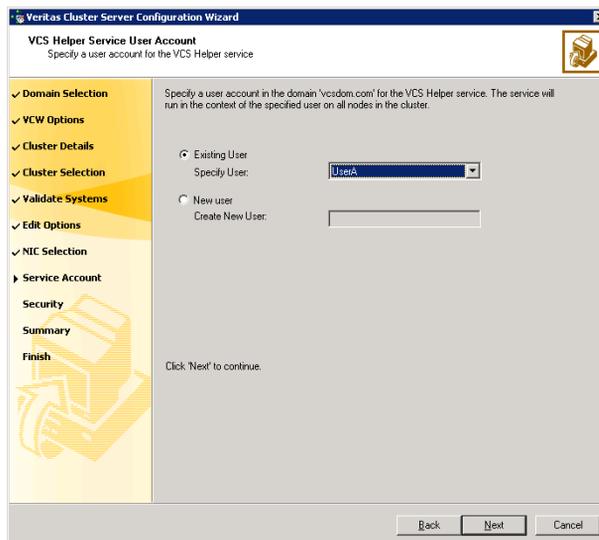


- Select **Configure LLT over UDP**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Symantec recommends reserving at least two NICs exclusively for the VCS private network.

- Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. Each IP address can be in a different subnet.
The IP address is used for the VCS private communication over the specified UDP port.
- For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

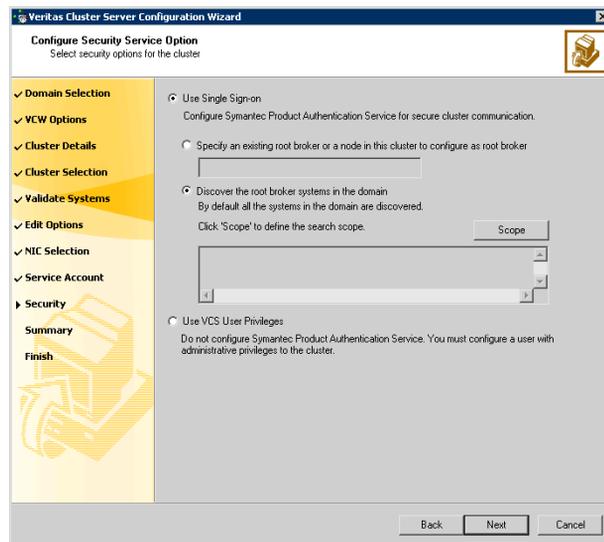
- 12 On the VCS Helper Service User Account panel, specify a domain user account for the VCS Helper service. The VCS high availability engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require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



Specify a domain user as follows:

- To specify an existing user, do one of the following:

- Click **Existing user** and select a user name from the drop-down list
 - If you chose not to retrieve the list of users in [step 4](#) on page 29, type the user name in the **Specify User** field, and then click **Next**.
 - To specify a new user, click **New user** and type a valid user name in the Create New User field, and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
 - In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.
- 13 On the Configure Security Service Option panel, specify the security options for the cluster and then click **Next**.
Do one of the following:
- To use the single sign-on feature, complete the following steps:



- Click **Use Single Sign-on**. In this mode, the Symantec Product Authentication Service is used to secure communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode. For more information about secure communications in a cluster, see the *Veritas Storage Foundation and High Availability Solutions Quick Start Guide for Symantec Product Authentication Service*.

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

[Table 2-1](#) contains some more examples of search criteria.

Table 2-1 Search criteria examples

1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Name	is (exactly)	*system	Displays all systems with names that end with <i>system</i> .
Name	is (exactly)	*vcsnode*	Displays all systems with names that contain <i>vcsnode</i> .

Table 2-1 Search criteria examples

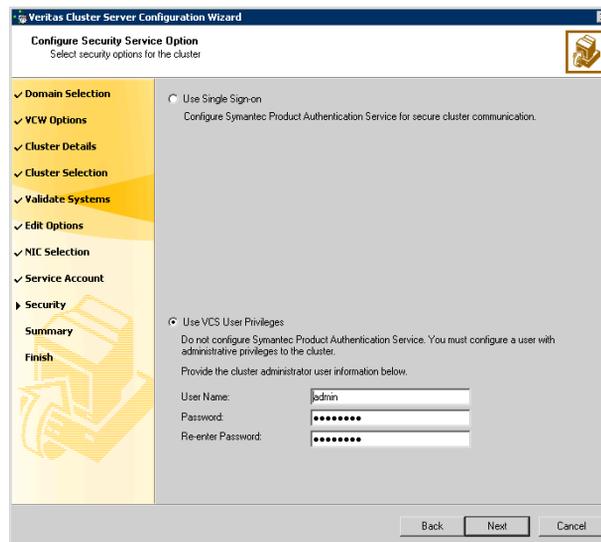
1st drop-down list value	2nd drop-down list value	Adjacent field entry	Search result
Operating System	is (exactly)	*2003*	Displays all Windows Server 2003 systems.
Operating System	is (exactly)	*Enterprise*	Displays all Windows Server Enterprise Edition systems.
Operating System Version	is (exactly)	5.*	Displays all systems whose operating system version is 5.x, where x could be 0, 1, 2, etc.

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

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

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

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



- Click **Use VCS User Privileges** and then type a user name and password. The wizard configures this user as a VCS cluster administrator. In this mode, communication between cluster nodes and clients, including Java console, occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCS`Encrypt` utility to encrypt the user password. The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.
 - Click **Next**.
- 14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network. If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard. The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.
- 15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for notification, and for global clusters. To configure the ClusterService group later, click **Finish**. At this stage, the wizard has collected the information required to set up the cluster configuration. After the wizard completes its operations, with or without the ClusterService group components, the cluster is ready to host application service groups. The wizard also starts the VCS engine (HAD) and the Veritas Command Server at this stage.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for complete details on the Notification resource.

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

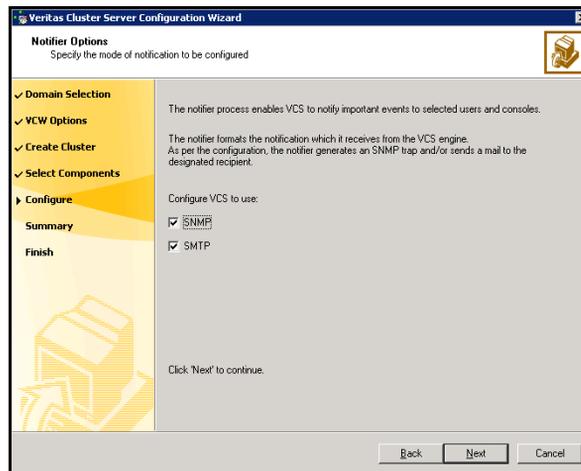
- 16 On the Cluster Service Components panel, select the components to be configured in the ClusterService service group and click **Next**.
 - Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
See “[Configuring notification](#)” on page 39.

Configuring notification

This section describes steps to configure notification.

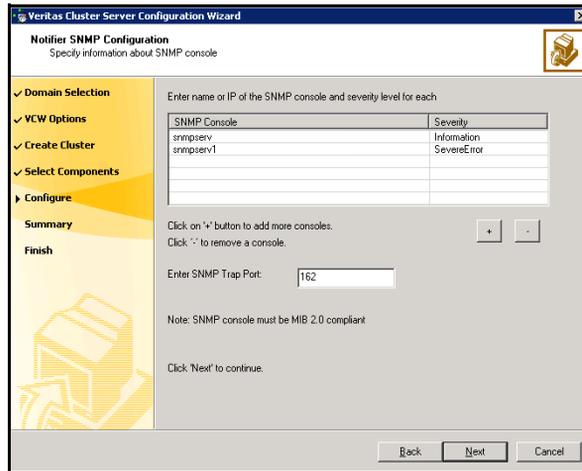
To configure notification

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



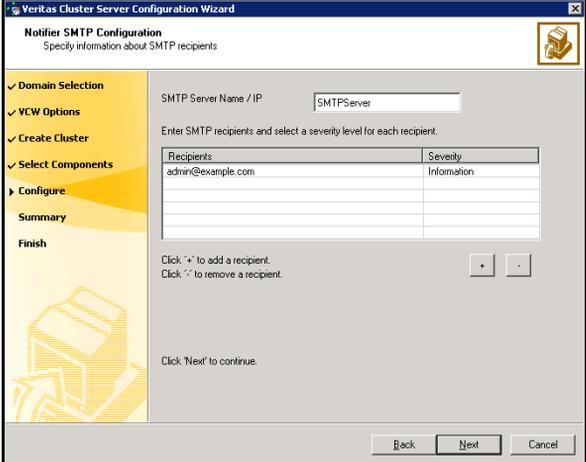
You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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



- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '+' to add a field; click '-' to remove a field.
- Enter an SNMP trap port. The default value is "162".

- 3 If you chose to configure SMTP, specify information about SMTP recipients and click **Next**.

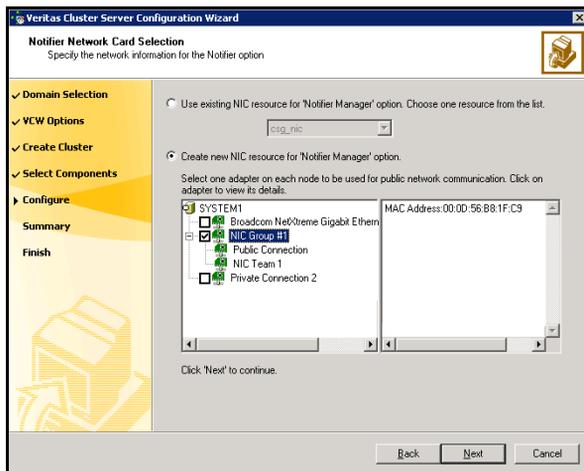


The screenshot shows the 'Notifier SMTP Configuration' window of the Veritas Cluster Server Configuration Wizard. The window title is 'Veritas Cluster Server Configuration Wizard' and the subtitle is 'Notifier SMTP Configuration'. Below the subtitle, it says 'Specify information about SMTP recipients'. On the left side, there is a navigation pane with the following items: 'Domain Selection' (checked), 'VCW Options' (checked), 'Create Cluster' (checked), 'Select Components' (checked), 'Configure' (selected), 'Summary', and 'Finish'. The main area contains the following fields and controls:

- 'SMTP Server Name / IP' text box with the value 'SMTPServer'.
- 'Enter SMTP recipients and select a severity level for each recipient.' instruction.
- A table with two columns: 'Recipients' and 'Severity'. The first row contains 'admin@example.com' and 'Information'. There are three empty rows below.
- 'Click '+' to add a recipient.' and 'Click '-' to remove a recipient.' instructions with corresponding '+' and '-' buttons.
- 'Click 'Next' to continue.' instruction.
- 'Back', 'Next', and 'Cancel' buttons at the bottom.

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

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



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

Installing and configuring SQL Server

This chapter contains the following topics:

- [“About installing and configuring SQL”](#) on page 44
- [“Prerequisites installing SQL Server”](#) on page 44
- [“About installing and configuring SQL Server 2000”](#) on page 53
- [“About installing and configuring SQL Server 2005”](#) on page 59

About installing and configuring SQL

A basic configuration of SQL Server has one or more instances of SQL Server in the cluster. Each instance is managed by a service group that is configured with a distinct set of nodes in the cluster. Also, SQL Server 2000 and SQL Server 2005 instances can co-exist on a cluster node.

This chapter describes steps to install SQL Server 2000 and SQL Server 2005 in a VCS cluster. The procedures ensure installations of SQL Server are uniform throughout the cluster. When you complete the installation, each node that will participate in a SQL Server service group will have a copy of the SQL program files located on the local disk, and access to a single copy of the SQL data files on a shared storage.

For more information on installing SQL Server, refer to the Microsoft SQL Server documentation.

Prerequisites installing SQL Server

Note the following prerequisites before installing SQL Server:

- Verify that SFW HA is installed on all nodes in the cluster. Refer to the *Veritas Storage Foundation and High Availability Solutions Installation and Upgrade Guide* for instructions.
- Verify that you have run VCS Cluster Configuration Wizard (VCW) to configure VCS on all nodes. See “[Configuring the cluster](#)” on page 28.
- Create the required volumes on the shared disks before installing SQL Server. See “[About managing storage using SFW](#)” on page 44.
- If your cluster has an Exchange service group configured, make sure to install and configure SQL Server on separate failover nodes in the cluster.

About managing storage using SFW

Veritas Storage Foundation for Windows (SFW) manages data by creating disk groups from physical disks. These disk groups are further divided into volumes, which can be mounted from the cluster nodes.

Symantec recommends that you have a minimum of two volumes on shared disks, one each for SQL Server data and registry replication information. If you wish to configure an MSDTC Server service group, create additional volumes for MSDTC log and MSDTC registry replication. Ensure that these volumes are accessible from all cluster nodes.

If your storage devices are SCSI-3 compliant and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA > Control Panel > Storage Agent > System Settings > SCSI Support).

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

Perform the following tasks to create volumes and make them accessible from cluster nodes:

- Create disk groups.
See "[Creating disk groups](#)" on page 45.
- Import disk groups.
See "[Importing disk groups](#)" on page 47.
- Create volumes on disk groups.
See "[Creating volumes](#)" on page 47.
- Mount volumes.
See "[Mounting volumes](#)" on page 51.

Creating disk groups

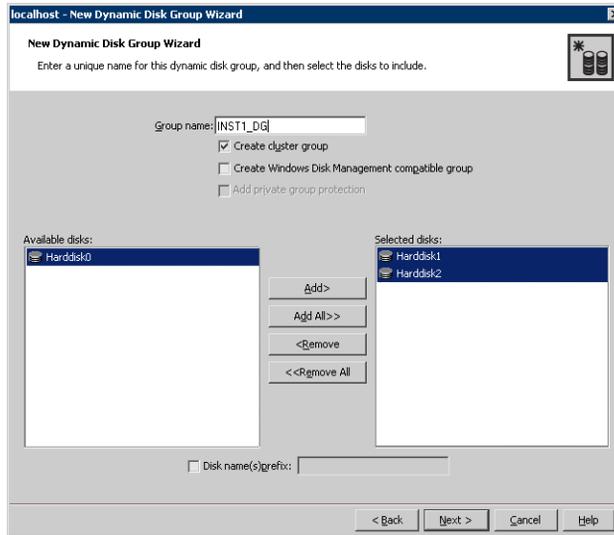
This section describes how to create disk groups.

To create a dynamic (cluster) disk group

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

- 1 Open the VEA console by clicking **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
- 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
- 6 Provide information about the cluster disk group:

- In the **Group Name** field, enter the name of the disk group, for example, SMPDG.
 - Enter the name of the disk group (for example, EVS1_SG1_DG).
- 7 Provide information about the cluster disk group.



- In the **Group name** field, enter a name for the disk group (for example, INST1_DG).
- Check the **Create cluster group** check box.
- Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier. For example, entering TestGroup as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.

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

Note: For Windows Server 2008, Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).

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

Importing disk groups

These steps describe how to import a disk group.

To import a disk group

- 1 From the VEA console, right-click a disk name in the dynamic group or the dynamic group name in the **General** tab or tree view.
- 2 Click **Import Dynamic Disk Group** from the menu.
- 3 In the Import Dynamic Disk Group dialog box, click **OK**.

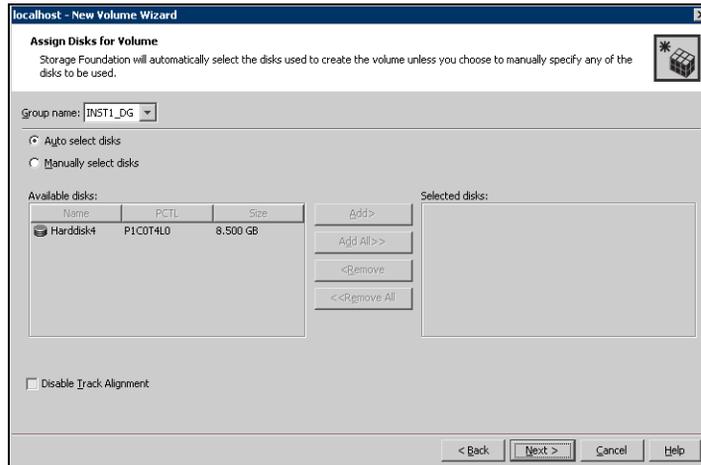
Creating volumes

This section describes how to create volumes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created, for example INST1_DG.
- 5 At the New Volume wizard opening screen, click **Next**.

- 6 Select the disks for the volume. Make sure the appropriate disk group name appears in the Group name drop-down list.

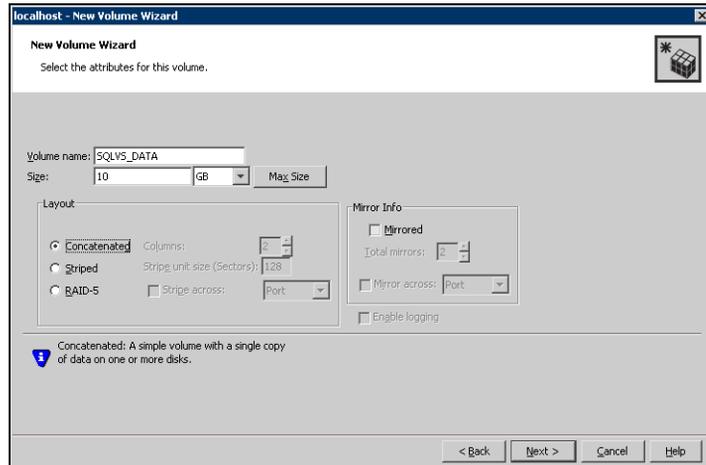


- 7 Automatic disk selection is the default setting. To manually select the disks, click the **Manually select disks** radio button and use the **Add** and **Remove** buttons to move the appropriate disks to the “Selected disks” list. Manual selection of disks is recommended.

You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling Track Alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.

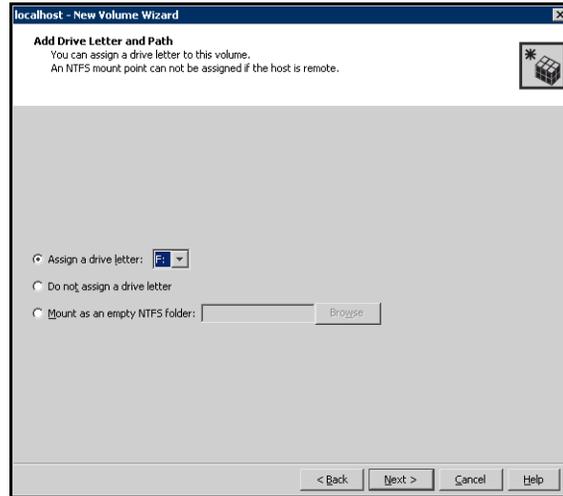
- 8 Click **Next**.

9 Specify the parameters of the volume.



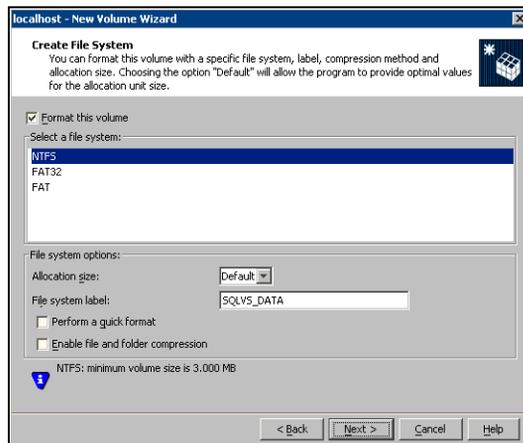
- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
 - Select a volume layout type. To select mirrored striped, click both the **Mirrored** checkbox and the **Striped** radio button.
 - If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.
 - Provide a size for the volume.
 - If you click on the **Max Size** button, a size appears in the Size box that represents the maximum possible volume size for that layout in the dynamic disk group.
 - In the Mirror Info area, select the appropriate mirroring options.
- 10 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.

- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.



11 Click **Next**.

12 Create an NTFS file system.



- Make sure the **Format this volume** checkbox is checked and click **NTFS**.
- Select an allocation size or accept the **Default**.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.

- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 13 Click **Finish** to create the new volume.
 - 14 Repeat these steps to create additional volumes.
Create the cluster disk group and volumes on the first node of the cluster only.

Mounting volumes

Mounting a volume involves importing the disk group and assigning a drive letter or the folder path. Unmounting a volume involves removing the drive letter or the folder path and deporting the disk group.

To mount a volume

- 1 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
- 2 On the Veritas Enterprise Administrator (VEA) screen, click **File > Connect**.
- 3 On the Connection dialog box, type the host name, and click **OK**.
- 4 In the left pane of the VEA console, select a system name and click the plus sign (+) next to the system name to expand the tree view.
- 5 If the disk group is not imported, import it. Right-click the disk group and then click **Import Dynamic Group**.
- 6 Right-click the volume and then click **File System > Change Drive Letter and Path**.
- 7 In the Drive Letter and Paths dialog box, click **Add**.
- 8 In the Assign Drive Letter panel, choose one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder, and click **OK**.
 - To assign a drive letter
Select **Assign a Drive Letter** and select a drive letter from the drop-down list.
 - To mount the volume as a folder
Select **Mount as an empty NTFS folder** and click **Browse** to locate an empty folder on the shared disk.

You must ensure that you use the same disk on all the nodes. To identify a disk, run `vmgetdrive -details` from the command prompt on one

node. Note the disk number and signature of the disk. Rerun the command on the second node and identify the disk that has the same signature as the one on the first node. Once you have identified two identical disks, assign the same drive letter to the volumes on each disk.

- 9 Click **OK**.
- 10 Repeat [step 6](#) through [step 9](#) for all the volumes to be mounted.

About installing and configuring SQL Server 2000

This section describes steps to install and configure SQL Server 2000 on the first cluster node and failover cluster nodes.

About installing SQL Server 2000 on the first node

Perform the installation and configuration tasks on the first node in the cluster. Then repeat the steps on each failover node in the cluster, as instructed under [“About installing SQL Server 2000 on failover nodes”](#) on page 56.

Overview of tasks

Complete the following tasks:

- Create volumes on the shared storage.
See [“Creating volumes”](#) on page 47.
- Mount volumes for the SQL Server 2000 installation.
See [“Mounting volumes”](#) on page 51.
- Install SQL Server 2000 on the first node.
See [“Installing SQL Server 2000”](#) on page 54.
- Unmount the volumes.
See [“Unmounting the volumes”](#) on page 55.

Installing SQL Server 2000

Before installing Microsoft SQL Server 2000, verify that the cluster disk group is imported on the first node and the volumes are mounted (are assigned drive letters).

Install Microsoft SQL Server 2000 on the first node using the installation wizard from the Microsoft software disk. Install the SQL program files on a local disk and the SQL data files on a shared storage managed using SFW. Use the following guidelines to install SQL Server to function properly in a VCS environment.

This section describes steps required for a SQL Server installation relevant in a VCS environment.

Refer to 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.

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

To install SQL Server 2000

- 1 Insert the SQL Server 2000 software disk in a disk drive connected to the node and run the Setup program.
- 2 From the browser menu, select **SQL Server 2000 Components > Install Database Server**.
- 3 If you are installing SQL Server 2000 on a node running a flavor of Windows Server 2003, a dialog box appears informing that SQL Server 2000 SP2 and below is not supported on this version of Windows. Click **Continue**. You will install SQL Server 2000 SP4 after installing SQL Server.
- 4 Read the information in the Welcome panel and click **Next**.
- 5 In the Computer Name panel, choose **Local Computer** option and click **Next**. Proceed through the installation to the Installation Definition dialog box.
- 6 In the Installation Definition panel, choose the **Server and Client Tools** option and click **Next**.
- 7 In the Instance Name panel, select the default instance or type an instance name and click **Next**. Use the same instance name when installing SQL Server 2000 on failover nodes.

- 8 In the Setup Type panel, select the type of installation and click **Browse** to specify the destination folder for installing 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 mounted using SFW. 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 dialog box, make the following selections and click **Next**:
 - Choose the **Customize the settings for each service** option.
 - In the Services box, select **SQL Server** option.
 - In the Service Settings box, select **Use a Domain User account** and specify the username, password, and domain.
 - Clear the **Auto Start Service** option.
 - Repeat these steps for the SQL Server Agent option.
- 11 Follow the wizard instructions to complete the installation. Reboot the node if prompted to do so.
- 12 Install Service Pack 4 or higher for SQL Server 2000. Refer to Microsoft documentation for further information.

Unmounting the volumes

These steps describe how to unmount volumes.

To unmount a volume

- 1 Stop the SQL Server service from the Service Control Manager (SCM).
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** to open the Veritas Enterprise Administrator.
- 3 From the VEA console's tree view, right-click on the volume and then click **File System > Change Drive Letter and Path**.
- 4 On the Drive Letter and Paths dialog box, select **Remove** and then click **OK**.
- 5 Repeat [step 3](#) through [step 4](#) for all the volumes to be unmounted.
- 6 From the VEA console's tree view, right-click the disk and then click **Deport Dynamic Group**.

About installing SQL Server 2000 on failover nodes

To ensure application installations are identical, the values you specify while installing and configuring SQL Server 2000 on failover nodes must match the values specified during the installation on the first node. Use the following steps to install and configure SQL Server 2000 on each failover node.

Overview of tasks for the failover node

Before proceeding, ensure that the SQL Server service is stopped on the first node.

Complete the tasks as per the following sequence:

- Import the disk group.
See “[Importing disk groups](#)” on page 47.
- Mount the volumes (using the same drive letter as assigned on the first node).
See “[Mounting volumes](#)” on page 51.
- Rename the previous SQL Server 2000 installation directory on the shared disk.
- Install SQL Server 2000 on the failover node.
See “[Installing SQL Server 2000](#)” on page 54.
- Unmount the volumes.
See “[Unmounting the volumes](#)” on page 55.
- After you have installed and configured SQL Server 2000 on all the failover nodes, you must set the internal name of the SQL server to a clustered instance name.
See “[Setting the internal name to clustered instance \(SQL 2000\)](#)” on page 57.

Setting the internal name to clustered instance (SQL 2000)

Use the SQL Query Analyzer to set the internal name of the SQL instance to the clustered instance name. Before you proceed ensure that the volume containing the SQL data files is mounted, and the SQL Server service is started.

Note: Perform these steps only once per instance and must be done after SQL Server is installed and configured on the last failover node.

To set the internal name to a 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 dialog box, select the system on which SQL Server was installed last.
- 3 Specify valid user credentials.
 - If you select **Windows authentication**, the user credential with which you are logged in is applicable by default.
 - If you select **SQL Server authentication**, enter a valid username and password.
- 4 Click **OK**.
- 5 In the Query dialog box, perform the following steps:
 - Retrieve the SQL Server name
 - Type **sp_helpserver** in the upper pane.
 - Press **F5**.Make a note of the name listed in the lower pane. For a default instance, the name will be System Name. For a named instance, the name will be System Name\Instance Name.
You must provide this name in the next step.
 - Disconnect the database
 - Delete the contents in the upper pane.
 - Type **sp_dropserver 'name of the system displayed in previous step'**
 - Press **F5**.
 - Reconnect the database using the clustered instance name
 - Delete the contents in the upper pane
 - Type **sp_addserver 'name of clustered instance', local**

- Press **F5**.
- 6 Exit the SQL Query Analyzer.
- 7 Stop the SQL Server instance.
- 8 Unmount the volumes.
See “[Unmounting the volumes](#)” on page 55.

About installing and configuring SQL Server 2005

This section describes steps to install and configure SQL Server 2005 on the first cluster node and failover cluster nodes.

About installing SQL Server 2005 on the first node

Perform the installation and configuration tasks on the first node in the cluster. Then repeat the steps on each failover node in the cluster, as instructed under [“Overview of tasks for the failover node”](#) on page 63.

Overview of tasks

Complete the tasks as per the following sequence:

- Create volumes on the shared storage.
See [“Creating volumes”](#) on page 47.
- Mount volumes for the SQL Server 2005 installation.
See [“Mounting volumes”](#) on page 51.
- Install SQL Server 2005 on the first node.
See [“Installing SQL Server 2005”](#) on page 60.
- Unmount the volumes.
See [“Unmounting the volumes”](#) on page 63.

Installing SQL Server 2005

Before installing SQL Server 2005, verify the shared drives created to store SQL data are mounted on the node where you run the wizard.

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 VCS environment are documented. Refer to the Microsoft SQL Server 2005 documentation for detailed installation information.

To install SQL Server 2005

- 1 To begin the SQL Server 2005 installation, navigate to the installation directory, launch **splash.hta**, and review the hardware and software requirements for SQL 2005.
- 2 Under the **Install** section, select **Server components, tools, Books Online, and samples**. Continue with the installation, following the instructions from Microsoft. Complete the SQL Server Component Update, System Configuration Check, and Registration Information pages.
- 3 In the Components to Install dialog box:
 - Select the **SQL Server Database Services**.
 - Select the optional components:
 - Analysis Service
 - Notification Services
 - Integration Services
 - Workstation Components
 - Click the **Advanced** option.
- 4 In the **Feature Selection** dialog box, specify the path for SQL Server data files and other services.
 - Expand **SQL Server Database Services** and select **Data Files**.
 - Select **Browse** to reset the installation path.
 - Set the installation path in the Change Folders dialog box to the drive letter and location of the volume created for the SQL Server system data files (INST1_DATA_FILES). Allow the rest of the path (Program Files\Microsoft SQL Server) to remain and click **OK**.
This must be the same as the path on all nodes.

- If you selected the **Analysis Services** option in [step 3](#), expand **Analysis Services**, select **Data Files**, and click **Browse** to specify the same location as for the SQL Server data files.
This must be the same as the path on all nodes.
 - Click **Next**.
- 5 In the **Instance Name** dialog box, type an instance name or accept the default and click **Next**. Only one default instance is allowed per cluster. Use the same instance name when installing SQL Server 2005 on the first node and on all failover nodes.
 - 6 In the **Service Accounts** dialog box, make the following selections and click **Next**:
 - Select **Use a domain user account** and then specify the user name, password, and domain.
If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group.
See Technote <http://support.veritas.com/docs/281828>.
 - Clear all the check boxes that start services except the SQL Browser, so that SQL Server is not brought online.
 - 7 Follow the wizard instructions to complete the installation. Reboot if prompted.
 - 8 Install any SQL service packs or hotfixes, if required.
 - 9 Set SQL services to manual start.

Setting SQL Server 2005 services to manual start

If you did not clear the Auto-start Service option during installation, set all SQL Server services to manual start.

To set SQL Server services to manual start

- 1 Click **Start > All Programs > Microsoft SQL Server 2005 > Configuration Tools > SQL Server Configuration Manager** to access the SQL Server Configuration Manager.
- 2 Access SQL Server properties:
 - In the left pane, select **SQL server 2005 Services**.
 - Right-click **SQL Server** instance name in the right pane.
 - Select **Properties**.
- 3 To set **Start Mode**.
 - Click the **Service** tab.
 - Select **Start Mode**.
 - Select **Manual** in the drop down list.
 - Click **OK**.
- 4 Repeat for all other SQL Server services that are running on the server for this specific instance.

Stopping the SQL Server 2005 service

Stop a running SQL Server service on the configured node so that 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 2005 > Configuration Tools > SQL Configuration Manager** to Access the SQL Server Configuration Manager.
- 2 Access SQL Server 2005 Services:
 - In the left pane of the SQL Server Configuration Manager, select **SQL Server 2005 Services**.
 - Right-click the SQL Server instance in the right pane.
 - Select **Stop**.
- 3 Repeat for all other SQL Server services that are running on the server.

Unmounting the volumes

These steps describe how to unmount volumes.

To unmount a volume

- 1 Stop the SQL Server service from the Service Control Manager.
- 2 Click **Start > All Programs > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator** to open the Veritas Enterprise Administrator.
- 3 From the VEA console's tree view, right-click on the volume and then click **File System > Change Drive Letter and Path**.
- 4 On the Drive Letter and Paths dialog box, select **Remove** and then click **OK**.
- 5 Repeat [step 3](#) through [step 4](#) for all the volumes to be unmounted.
- 6 From the VEA console's tree view, right-click the disk and then click **Deport Dynamic Group**.

Installing and configuring SQL Server 2005 on failover nodes

To ensure that application installations are identical, the values you specify while installing and configuring SQL Server 2005 on failover nodes must match the values specified during the installation on the first node. Use the following steps to install and configure SQL Server 2005 on each failover node.

Note: If the failover node is already a member of the system list of a functional SQL Server 2005 service group in the cluster, verify that SQL Server 2005 is installed on a volume with the same drive letter as the first node and proceed to [“Configuring SQL Server in a VCS cluster”](#) on page 69. If the failover node is *not* a member of the system list of an existing functional SQL Server 2005 service group in the cluster, proceed to the steps below.

Overview of tasks for the failover node

Before proceeding, make sure the SQL server service is stopped on the first node.

Complete the tasks as per the following sequence:

- Import the disk group.
See [“Importing disk groups”](#) on page 47.
- Mount the volumes (using the same drive letter as assigned on the first node).
See [“Mounting volumes”](#) on page 51.

- Rename the previous SQL Server 2005 installation directory on the shared disk.
- Install SQL Server 2005 on the failover node.
See “[Installing SQL Server 2005](#)” on page 60.
- Unmount the volumes.
See “[Unmounting the volumes](#)” on page 55.

Assigning ports for multiple SQL Server instances

If you are running multiple SQL Server instances, you must assign a different port to each named instance. You can assign static or dynamic ports.

Refer to the Microsoft Knowledge Base for the instructions on assigning ports. At the time of this release, this information is in the following article:

Microsoft Knowledge Base Article - 823938: How to configure an instance of SQL Server to listen on a specific TCP port or a dynamic port

Refer to:

<http://support.microsoft.com/kb/823938/en-us>

If you wish to change the port after configuring the SQL service group, you must perform the steps in the following order:

- Bring the SQL service group online or partially online (upto the registry replication resource) on a cluster node.
- On the node on which the SQL service group is online or partially online, change the port assigned to the SQL instance.
Refer to the instructions mentioned in the Microsoft Knowledge Base article specified earlier.
- Take the SQL service group offline on the node, and then bring it online again. The configuration changes will be replicated to the remaining nodes.

Configuring the SQL Server service group

This chapter contains the following topics:

- [“About configuring the SQL service group”](#) on page 68
- [“Prerequisites for configuring the SQL service group”](#) on page 68
- [“Configuring SQL Server in a VCS cluster”](#) on page 69
- [“About creating a SQL Server user-defined database”](#) on page 78

About configuring the SQL service group

Configuring a SQL Server service group involves two major tasks, configuring a SQL Server service and configuring an MSDTC service group. This chapter provides instructions on configuring the SQL Server service.

To configure an MSDTC service group, refer to the respective chapter.

See “[Configuring the MSDTC service group](#)” on page 83.

VCS provides several ways to configure a service group, including the service group configuration wizard, Cluster Manager (Java Console), Cluster Manager (Web Console), and the command line. This chapter provides instructions on configuring the service group using the configuration wizard.

Prerequisites for configuring the SQL service group

Note the following prerequisites before configuring the SQL service group:

- If you have configured Windows Firewall, add the following to the Firewall Exceptions list:
 - Port 14150 or the VCS Command Server service
`%vcs_home%\bin\CmdServer.exe`
Here, `%vcs_home%` is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
 - Port 14141
For a detailed list of services and ports used by SFW HA, refer to the *Veritas Storage Foundation and High Availability Solutions for Windows Installation and Upgrade Guide*.
- 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. If you wish to configure detail monitoring, you must have the permission to log on to the respective SQL instance.
- Verify that Storage Foundation HA for Windows (SFW HA), along with the VCS database agent for Microsoft SQL, is installed and configured on all cluster nodes.
See “[Installing the VCS database agent for SQL](#)” on page 25 for instructions.
- Verify that a VCS cluster is configured using VCS Cluster Configuration Wizard (VCW).
See “[Configuring the cluster](#)” on page 28.

- Verify the SQL Server instance to be made highly available is installed identically on all nodes that will participate in the service group.
- Verify the shared volumes created to store SQL Server data files and registry replication information are mounted on the node where you run the wizard and unmounted from other nodes in the cluster.
See “[About managing storage using SFW](#)” on page 44.
- Assign a unique virtual server name and IP address to SQL Server instance.
- Make sure the SQL Server service is stopped for the SQL instance that you wish to configure in the service group.
- This is applicable in case of SFW only. If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA > Control Panel > Storage Agent > System Settings > SCSI Support). See the Veritas Storage Foundation Administrator’s Guide for more information.

Configuring SQL Server in a VCS cluster

VCS provides the SQL Configuration Wizard that guides you through the process of configuring a SQL Server service group. This section describes the steps required to create a SQL Server service group using the wizard.

You can also use this wizard to modify and delete a SQL Server service group.

See “[Modifying a SQL Server service group](#)” on page 106.

Before configuring the service group, review the resource types and the attribute definitions of the SQL Server agents, and the sample configuration and resource dependency graphs of the SQL Server service group.

See [Appendix A, “Resource type definitions”](#) on page 131.

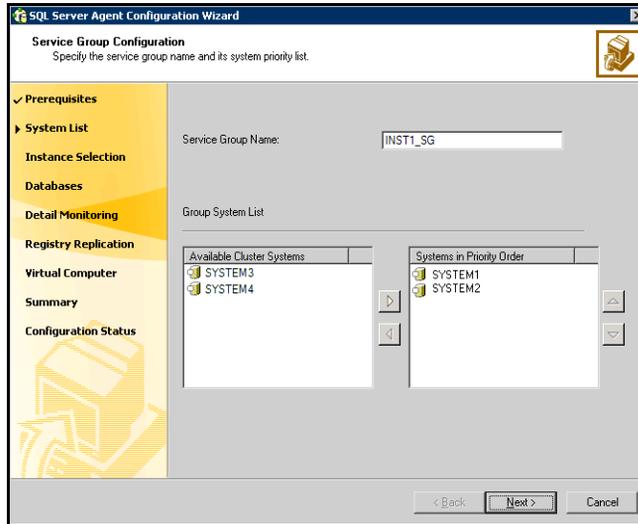
See [Appendix C, “Sample configurations”](#) on page 151.

Note: Symantec recommends that you configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.

To create an SQL Server service group

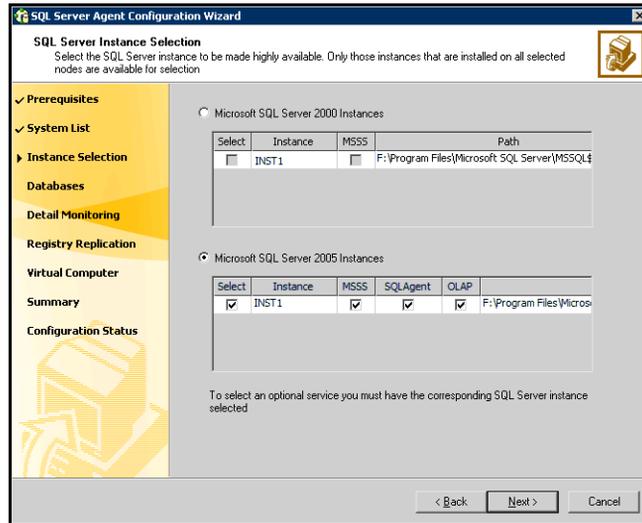
- 1 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL Configuration Wizard.
- 2 In the Select Configuration Option panel, click **MS-SQL Server - Service Group Configuration**, click **Create**, and then click **Next**.

- 3 Review the prerequisites for configuring a SQL Server service group and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and system priority list and then click **Next**.
The wizard starts validating your configuration. Various messages indicate the validation status.



- | | |
|---------------------------|---|
| Service Group Name | Type a name for the SQL Server service group. |
| Available Cluster Systems | 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 box. |
| Systems in Priority Order | This list represents the service group's system list. To remove a system from the service group's system list, select the a system and click the left arrow. To change a system's priority in the service group's system list, select the system 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. |

- 5 On the SQL Server Instance Selection panel, select the SQL Server instance to be made highly available and then click **Next**.



The wizard shows all the SQL Server instances, for SQL Server 2000 and SQL Server 2005, installed on the cluster nodes.

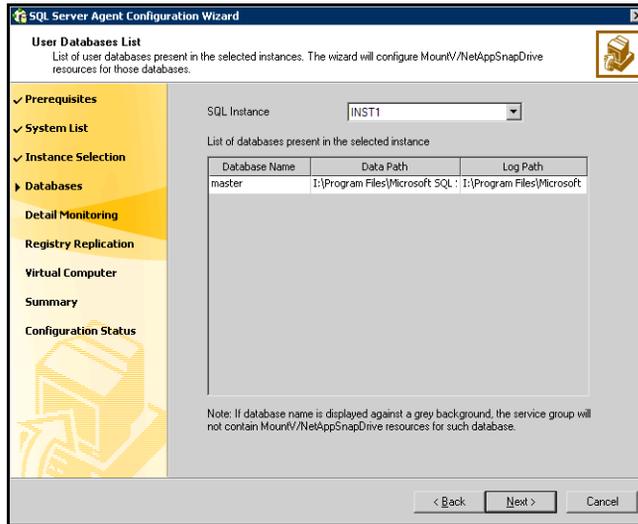
Microsoft SQL Server 2000 Instances

To configure a SQL Server 2000 instance, click **Microsoft SQL Server 2000 Instances** and then check the box adjacent to the instance name. Check the box in the **MSSS** column to make the associated MSSearch service highly available. To select the dependent services, you must select the corresponding SQL instance.

Microsoft SQL Server 2005 Instances

To configure a SQL Server 2005 instance, click **Microsoft SQL Server 2005 Instances** and then check the box adjacent to the instance name. Check the box in the **MSSS** column to make the associated MSSearch service highly available. Check the box in the **SQLAgent** column to make the associated SQL agent service highly available. Check the box in the **OLAP** column to make the associated Analysis service highly available. To select the dependent services, you must select the corresponding SQL instance.

- 6 On the User Databases List panel, verify the databases configured for the SQL instances and then click **Next**.

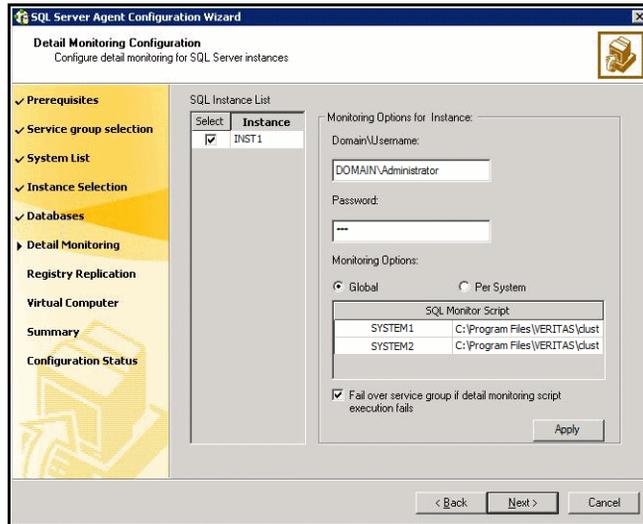


The **SQL Instance** drop-down list box contains all the instances selected on the SQL Server Instance Selection panel, in [step 5](#). Select a SQL instance to see the master database configured for the instance. The wizard will create MountV resource for the master database.

If you add user defined databases for a SQL instance, make sure you run the SQL configuration wizard in modify mode. This ensures the user defined databases are included in the VCS configuration.

See “[Modifying a SQL Server service group](#)” on page 106 for instructions.

- 7 On the Detail Monitoring Configuration panel, configure detail monitoring of SQL Server database, if required, and then click **Next**.



SQL Instance List

Check the box adjacent to the SQL Server instance for which you want to configure detail monitoring. Only the instances selected in [step 5](#) will be available for selection.

Domain\Username

Type the fully qualified user name for connecting to SQL Server database. Make sure the specified user has SQL Server login permissions.

The user name must be the Fully Qualified Domain Name (FQDN) in the format domain.com\user name.

Password

Type the password for the user name specified in the Domain\Username field.

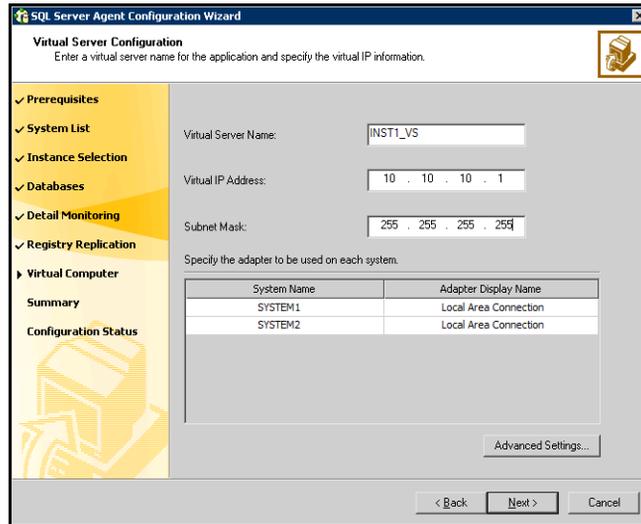
Global

Click **Global** if the path of the monitor script is the same on all nodes and then click SQL Monitor Script and type the path for the monitor script.

Per System	Click Per System if the path of the monitor script is different on all the nodes and then click SQL Monitor Script and type the path for the monitor script for each system. Make sure that the specified path exists on all the systems in the cluster.
Fail over service group if detail monitoring script execution fails	Check this option if you want the SQL service group to fail over when the detail monitor script execution fails.
Apply	Click Apply to configure detail monitoring for the selected SQL instance. Repeat this for all instances for which you want to configure detail monitoring.

- 8 On the Registry Replication Path panel, specify the location where the registry replication information will be stored and click **Next**. Symantec recommends that the mount path for RegRep resources be different from the mount path SQL data files.

- 9 On the Virtual Server Configuration panel, specify the information related to the virtual server and then click **Next**.



- Virtual Server Name** Type a virtual name for the node. This is the server\instance name that is used by the SQL Server 2000 clients to access the database. Ensure that the virtual computer name you enter is unique in the cluster.
- Virtual IP Address** Type a unique virtual IP address for the SQL Server.
- Subnet Mask** Type the subnet mask to which the virtual IP address belongs.

Advanced

Click **Advanced** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.

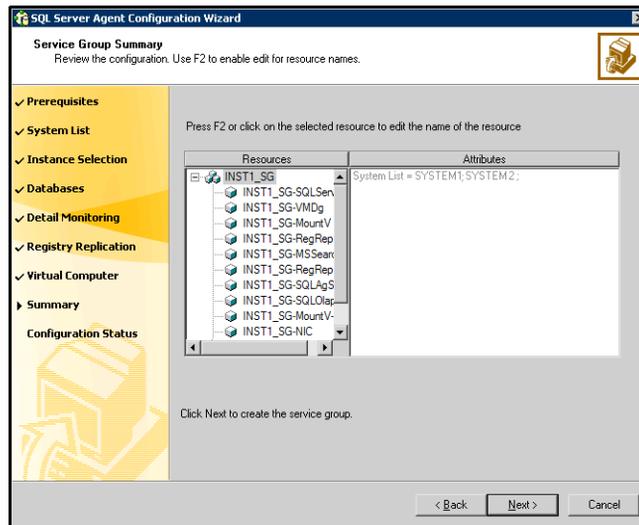
On the Lanman Advanced Configuration dialog box, complete the following:

- Check **Active Directory Update required** check box to enable the Lanman agent to update the Active Directory with the virtual name. This sets the Lanman agent attributes `ADUpdateRequired` and `ADCriticalForOnline` to true.
- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format `CN=containername,DC=domainname,DC=com`. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
- Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

Adapter Display Name

This field displays the adapters associated with a system. For each system in the cluster, select the public network adapter name. The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Ensure that you select the adapters to be assigned to the public network, and not those assigned to the private network.

- 10 On the Service Group Summary panel, review the service group configuration and change the resource names, if desired, and then click **Next**.



The **Resources** box lists the configured resources. Click 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. Edit the resource name, and press the **Enter** key to confirm the changes. To cancel editing a resource name, press the **Esc** key.

- 11 Click **Yes** on the message that prompts you that the wizard will run commands to modify the service group configuration. Various messages indicate the status of these commands. The wizard marks all the resources in the service group as **CRITICAL**.
 Use Cluster Manager (Java Console) or the command-line to change their state, if required.
- 12 On the Configuration Complete panel, check **Bring the service group online** to bring the configured service group online on the local system, and click **Next** to configure additional SQL Server service groups or an MSDTC Server service group or click **Finish** to exit the wizard.
 See “[Configure MSDTC in a VCS cluster](#)” on page 88.

About creating a SQL Server user-defined database

The following tasks describe how to use VCS to create and manage a SQL Server user-defined database.

You must complete the tasks in the following order:

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a 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 appropriate VCS storage resources for the user databases.

Create 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 the following volumes are used:

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

Creating a SQL Server database

Create a 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 Click **Start > All Programs > Microsoft SQL Server > Enterprise Manager** to open SQL Server Database Manager.
- 2 Right-click **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 On the Transaction Log tab, 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 Click **Start > All Programs > Microsoft SQL Server 2005 > SQL Server Management Studio** to open SQL Server Database Manager.
- 2 Expand the icon associated with your server.
- 3 Right-click **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 the VCS storage resources

Ensure the following before running the SQL Server Configuration Wizard to add the VMDg and MountV resources:

- Make sure the SQL Server resources are online.
- Make sure the volumes for the user database and transaction logs are mounted on the node.

To add VMDg and MountV 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 Tools > 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 message 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**. Databases that are highlighted will not contain MountV resources.
- 9 If a database is not configured correctly, a 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 message indicates the configuration will be modified.
- 12 To complete the user database configuration, choose one of the following:
 - 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.

Configuring the MSDTC service group

This chapter contains the following topics:

- [“About configuring the MSDTC service group”](#) on page 84
- [“Prerequisites for configuring the MSDTC service group”](#) on page 86
- [“Configure MSDTC in a VCS cluster”](#) on page 88

About configuring the MSDTC service group

Microsoft Distributed Transaction Coordinator (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.

To configure high availability for MSDTC, you first configure the MSDTC Server service group using the wizard and then manually configure the MSDTC client.

VCS provides several ways to configure a service group, including the SQL Server Configuration Wizard, Cluster Manager (Java Console), Cluster Management Console (Web Console), and the command line.

This chapter describes the steps involved in configuring an MSDTC Server service group using the wizard.

Reviewing the configuration

MSDTC servers can co-exist 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 configured 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. In general, you must configure the MSDTC client on all nodes except the node on which the MSDTC Server is configured to fail over.

For instance, a SQL Server configuration in a Veritas Cluster Server cluster might span four nodes and two sets of shared storage.

The following configurations are possible:

- SQL Server and MSDTC Server configured on different nodes
- SQL Server is configured on the same node as the MSDTC Server
- SQL Server and MSDTC Servers configured on nodes in different clusters

[Figure 5-1](#) illustrates a configuration where MSDTC Server and SQL Server are configured on different nodes in a cluster.

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

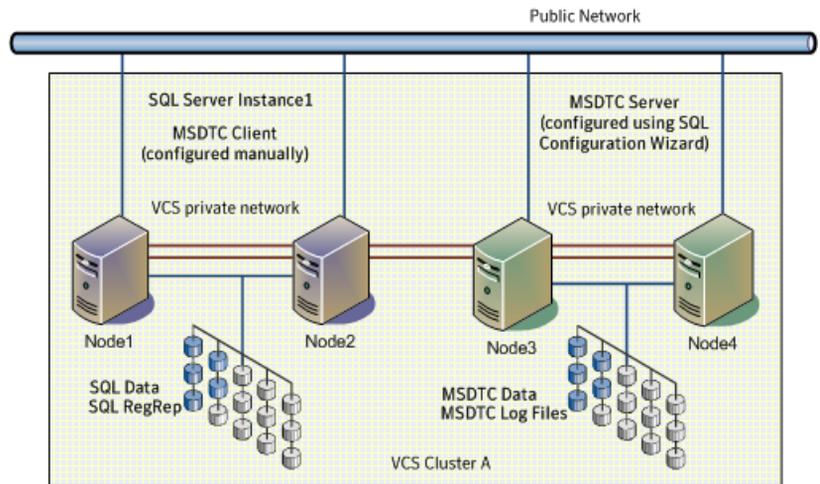


Figure 5-2 illustrates a configuration where MSDTC Server and SQL Server are configured on the same node in a cluster.

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

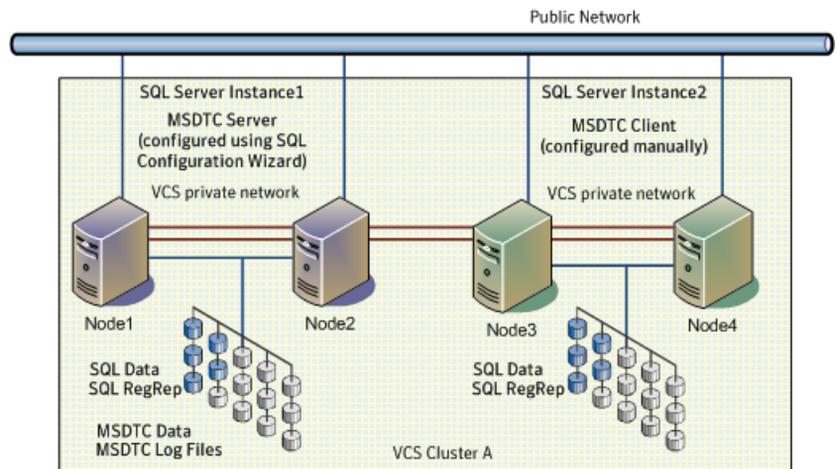
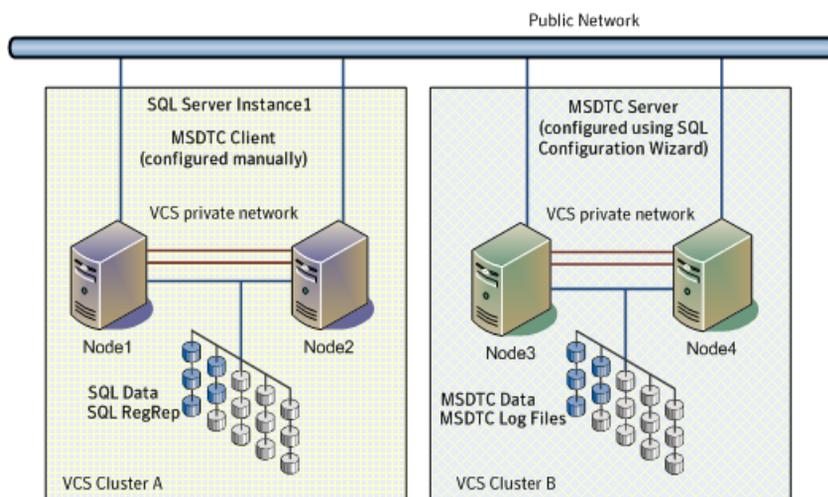


Figure 5-3 illustrates a configuration where the MSDTC Server and SQL Server are configured on nodes in different clusters.

Figure 5-3 MSDTC Server and SQL Server configured in different clusters



Prerequisites for configuring the MSDTC service group

Note the following prerequisites before configuring the MSDTC service group:

- 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 Storage Foundation HA for Windows (SFW HA), along with the VCS database agent for Microsoft SQL, is installed and configured on all cluster nodes.
- Verify that the VCS cluster is configured using the VCS Cluster Configuration Wizard (VCW).
See “[Configuring the cluster](#)” on page 28.
- Verify the MSDTC service is installed on all nodes that will participate in the MSDTC service group.
- This is applicable in case of SFW only. If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (**VEA > Control Panel > Storage Agent > System Settings > SCSI Support**). See the Veritas Storage Foundation Administrator’s Guide for more information.

- Verify the shared volumes created to store MSDTC log files and registry replication information are mounted on the node where you run the wizard and unmounted from other nodes in the cluster.
See “[About managing storage using SFW](#)” on page 44.
- Make sure the Distributed Transaction Coordinator (DTC) service is stopped.
- If you have configured a firewall, add the following to the firewall exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141For a detailed list of services and ports used by VCS, refer to the *Veritas Cluster Server Installation and Upgrade Guide*.
- Keep the following information ready with you; the wizard prompts you for these details:
 - A unique virtual server name for the MSDTC Server. This is the name that is used by MSDTC clients to connect to the MSDTC Server. The DTC service runs under this name.
 - A unique virtual IP address for the MSDTC Server.

Configure MSDTC in a VCS cluster

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 a SQL Server Configuration Wizard that guides you through the process of configuring an MSDTC Server service group. You can also use this wizard to modify an MSDTC service group configuration. After configuring the service group proceed to configuring the MSDTC client.

Note: You can configure only one MSDTC Server service group in a cluster.

Configuring MSDTC in a VCS cluster involves the following tasks:

- Configuring an MSDTC Server service group.
See “[Configuring an MSDTC service group](#)” on page 88.
- Configuring an MSDTC client.
See “[Configuring an MSDTC client](#)” on page 93.

Configuring an MSDTC service group

VCS provides the SQL Configuration Wizard that guides you through the process of configuring an MSDTC service group. This section describes the steps required to create a new MSDTC service group using the wizard.

Before configuring the service group, review the resource types and the attribute definitions of the MSDTC agent, and the sample configuration files and resource dependency graphs of the SQL Server service group.

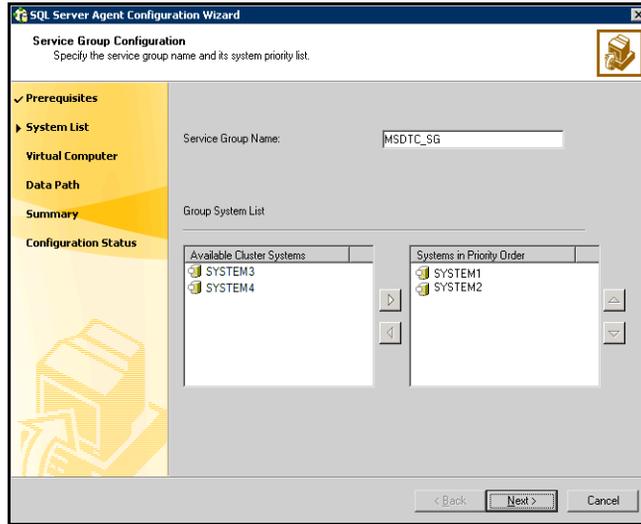
See [Appendix A, “Resource type definitions”](#) on page 131.

See [Appendix C, “Sample configurations”](#) on page 151.

To create an MSDTC service group

- 1 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL Configuration Wizard.
- 2 In the Select Configuration Option panel, click **MSDTC Server - Service Group Configuration**, click **Create**, and click **Next**.
- 3 Review the prerequisites for configuring an MSDTC service group and click **Next**.
- 4 On the Service Group Configuration panel, specify the service group name and system priority list and then click **Next**.

The wizard starts validating your configuration. Various messages indicate the validation status.



Service Group Name

Type a name for the MSDTC service group.

Available Cluster Systems

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 box.

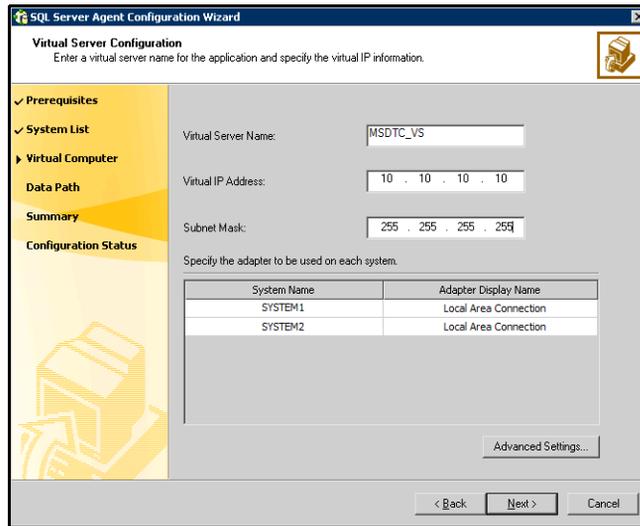
Make sure you select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.

Systems in Priority Order

This list represents the service group's system list. To remove a system from the service group's system list, select the a system and click the left arrow.

To change a system's priority in the service group's system list, select the system 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.

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



Virtual Server Name

Type a virtual name for the node. This is the server name on which the DTC service is running. Ensure that the virtual computer name you enter is unique in the cluster.

Virtual IP Address

Type a unique virtual IP address for the MSDTC Server.

Subnet Mask

Type the subnet mask to which the virtual IP address belongs.

Advanced Settings

Click **Advanced Settings** if you wish to configure the Lanman agent to perform Windows AD update. These settings are applicable to the Lanman resource in the service group.

On the Lanman Advanced Configuration dialog box, complete the following:

- Check **Active Directory Update required** check box to enable the Lanman agent to update the Active Directory with the virtual name. This sets the Lanman agent attribute ADUpdateRequired and ADCriticalForOnline to true.
- In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format *CN=containername,DC=domainname,DC=com*. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
- Click **OK**. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

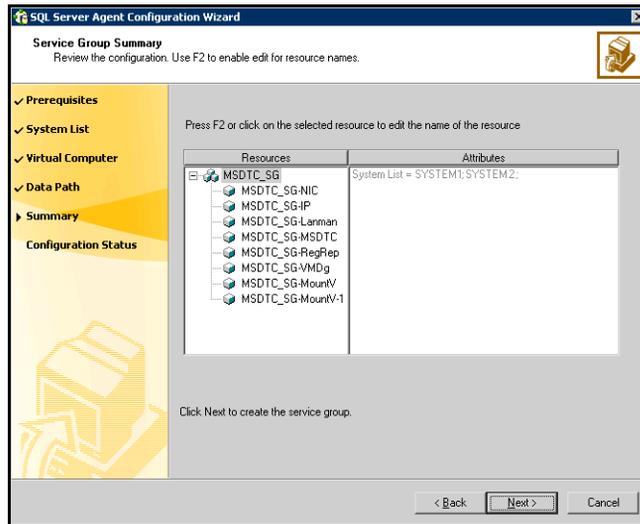
Adapter Display Name

This field displays the adapters associated with a system. For each system in the cluster, select the public network adapter name.

The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Ensure that you select the adapters to be assigned to the public network, and not those assigned to the private network.

- 6 Specify the MSDTC log and replication directory and click **Next**. Make sure the paths for both the directories are different. If the directory does not exist, the wizard creates it.

- 7 On the Service Group Summary panel, review the service group configuration and change the resource names, if desired, and click **Next**.



The **Resources** box lists the configured resources. Click 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. Edit the resource name, and press the Enter key to confirm the changes. To cancel editing a resource name, press the Esc key.

- 8 Click **Yes** on the message that prompts you that the wizard will run commands to modify the service group configuration. Various messages indicate the status of these commands.
- 9 On the Completing the MSDTC Configuration Wizard panel, check **Bring the service group online** check box to bring the configured service group online on the local system and then click **Finish**.

To bring the service group online later, clear the **Bring the service group online** check box.

This completes the MSDTC Server service group configuration. You can now proceed to configuring the MSDTC client.

See “[Configuring an MSDTC client](#)” on page 93.

Configuring an MSDTC client

Configure the MSDTC client after configuring the service group for the MSDTC Server. Set the MSDTC client to run on nodes where a SQL instance is configured to run and the MSDTC server is not configured to run. In general, you must configure the MSDTC client on all nodes except the nodes on which the MSDTC Server is configured. You do not need to configure the MSDTC client on the nodes that are part of the MSDTC Server service group.

The MSDTC client and the MSDTC Server must not run on the same cluster nodes.

Note: You have to configure the MSDTC client manually. You cannot use the SQL Server Configuration Wizard to configure the MSDTC client.

Before configuring the MSDTC client:

- Verify that the MSDTC Server service group is online in the cluster.
- Verify that the node on which you configure the client is not a part of an MSDTC service group's SystemList.

Procedures for Windows 2003 and Windows 2008 are different. Follow the appropriate procedure depending on the operating system.

Configuring MSDTC client on Windows 2003

Complete the MSDTC client and security configuration on Windows 2003 systems as described below.

To configure an MSDTC client on Windows 2003

- 1 Ensure that the MSDTC service group is online.
- 2 Launch the Windows Component Services Administrative tool.
Click **Start > Programs > Administrative Tools > Component Services**
or
Click **Start > Run**, type **dcomcnfg** and click **OK**.
- 3 In the console tree of the Component Services administrative tool, expand **Component Services > Computers**, right-click **My Computer** and then click **Properties**.
- 4 On the MSDTC tab perform the following steps:
 - Clear the **Use local coordinator** check box.
 - In the Remote Host field, specify the virtual server name that you specified while creating the MSDTC Server service group.

If you are unsure of the exact name, click **Select** to search from a list of all computers on the network and select the virtual computer name from the list.

- Click **Apply** and then click **OK**.

Configuring MSDTC client on Windows 2008

Complete the MSDTC client and security configuration on Windows 2008 systems as described below.

To configure an MSDTC client on Windows 2008

- 1 Ensure that the MSDTC service group is online.
- 2 Launch the Windows Component Services Administrative tool.
Click **Start > All Programs > Administrative Tools > Component Services**
or
Click **Start > Run**, type **dcomcnfg** and click **OK**.
- 3 In the console tree of the Component Services administrative tool, expand **Component Services > Computers**, right-click **My Computer** and then click **Properties**.
- 4 On the MSDTC tab, perform the following steps:
 - Clear the **Use local coordinator** check box.
 - In the Remote Host field, specify the virtual server name that you specified while creating the MSDTC Server service group.
If you are unsure of the exact name, click **Select** to search from a list of all computers on the network and select the virtual computer name from the list.
 - Click **Apply** and then click **OK**.

Verifying the service group configuration

This chapter contains the following topics:

- [“About verifying the service group”](#) on page 96
- [“Bringing the service group online”](#) on page 96
- [“Switching the service group”](#) on page 97
- [“Taking the service group offline”](#) on page 98
- [“Disabling the service group”](#) on page 99

About verifying the service group

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

Bringing the service group online

This section describes steps required to bring the service group online.

From the Cluster Manager (Java Console)

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

To bring the service group online

- 1 In the Cluster Explorer configuration tree, select the newly created service group.
- 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 systems 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*)

From the Web Console

Perform the following steps to bring the service group online from the Veritas Cluster Management Console (Single Cluster Mode) also referred to as Web Console.

To bring the service group online

- 1 On the Service Group page (**Cluster Summary > All Groups > Service Group**), click **Online**.
- 2 In the Online Group dialog box, select the system on which to bring the service group online. Click **OK**.

Switching the service group

This section describes steps required to switch the service group. The process of switching a service group involves taking it offline on its current system and bringing it online on another system.

From the Cluster Manager (Java Console)

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

To switch the service group online

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

From the Web Console

Perform the following steps to switch the service group from the Veritas Cluster Management Console (Single Cluster Mode) also referred to as Web Console.

To switch the service group online

- 1 From the Service Group page (**Cluster Summary>All Groups>Service Group**), click **Switch**.
- 2 On the Switch Group dialog box, select the system to switch the service group to.
- 3 Click **OK**.

Taking the service group offline

This section describes steps required to take the service group offline.

From the Cluster Manager (Java Console)

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

To take the service group offline

- 1 In 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*).

From the Web Console

Perform the following steps to take the service group offline from the Veritas Cluster Management Console (Single Cluster Mode) also referred to as Web Console.

To take the service group offline

- 1 On the **Service Group** page (**Cluster Summary > All Groups > Service Group**), click **Offline**. This opens the Offline Group dialog box.
- 2 Select the system on which to take the service group offline.
- 3 Click **OK**.

Disabling the service group

To disable the agent, you must change the SQL Server service group to an OFFLINE state. You can stop the application completely or switch the agent to another system.

From the Cluster Manager (Java Console)

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

To disable the service group

- 1 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.
or
Select the cluster in the Cluster Explorer configuration tree, select the **Service Groups** tab, and right-click the service group icon in the view panel.
- 2 Choose **Disable**, and choose the appropriate system from the pop-up menu.

From the Web Console

Perform the following steps to disable the service group from the Veritas Cluster Management Console (Single Cluster Mode) also referred to as Web Console.

To disable the service group

- 1 In the **Service Group** page (**Cluster Summary>All Groups>Service Group**), click **Disable**. This opens the Disable Group dialog box.
- 2 Select the system on which to disable the service group. To disable the service group on all systems, click **All Systems**.

Removing the VCS database agent for SQL

This chapter contains the following topics:

- [“About removing the VCS database agent for SQL”](#) on page 102
- [“Removing the VCS database agent for SQL”](#) on page 102

About removing the VCS database agent for SQL

This chapter describes steps for uninstalling the VCS database agent for Microsoft SQL using the installer for Storage Foundation for Windows.

Removing the VCS database agent for SQL

Before uninstalling the agent:

- Verify you have local administrator privileges on the node where you are uninstalling the agent.
- Verify the agent is not running when you attempt to remove it. If the agent is removed while running, information about the resource type and the service groups remains in the configuration files and causes VCS to fail.
- Verify all SQL Server service groups are offline on all nodes in the cluster and all SQL Server resource types are deleted.

To uninstall the agent

- 1 In the Windows Add/Remove Programs applet, click **Storage Foundation HA 5.1 SP2 for Windows (Server Components)** and click **Change**.
- 2 In the Storage Foundation 5.1 for Windows dialog box, click **Add or Remove** and click **Next**.
- 3 In the SFW product options panel, click **Next**.
- 4 Uncheck the check box for **Veritas Cluster Server Database Agent for SQL** and click **Next**.
- 5 The installer validates the system for prerequisites. After the system is accepted, click **Next**.
If a system is rejected, the Comments column displays the cause for rejecting the system. Highlight the system to view a detailed information about the failure in the Details box. Resolve the error, highlight the system from the list, and click **Validate Again**.
- 6 An informational message appears if you selected the DMP option. Review the information and click **OK** to continue.
- 7 Review the summary of your selections and click **Update** to start the uninstallation. The installer displays the status of uninstallation.
- 8 After the uninstallation is complete, review the report and click **Next** and click **Finish**.
Repeat these steps on all systems where VCS database agent for Microsoft SQL is to be uninstalled.

Administering the service group

This chapter contains the following topics:

- [“About administering the SQL service group”](#) on page 106
- [“Administer a SQL service group”](#) on page 106
- [“Administer an MSDTC service group”](#) on page 108
- [“Detail monitoring of a SQL Server database instance”](#) on page 110

About administering the SQL service group

This chapter describes the administrative tasks you can perform on SQL Server and MSDTC service groups such as modifying the service group configuration and deleting a service group. The chapter also describes steps to configure detail monitoring for a SQL Server database.

Administer a SQL service group

The SQL Configuration Wizard enables you to modify a SQL Server service group or delete the service group.

Modifying a SQL Server service group

You can modify the configuration of the service groups using the configuration wizards, the Java Console, and the Cluster Management Console (Single Cluster Mode) also referred to as the Web Console.

This section describes the steps to modify a SQL Server service group using the SQL Configuration Wizard.

Note the following prerequisites before modifying a SQL Server service group:

- 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 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL 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 and click **Next**. The wizard will create MountV resource for each database.
- 8 Follow the wizard instructions and make desired modifications to the service group configuration.
See "[Configuring SQL Server in a VCS cluster](#)" on page 69.

Deleting a SQL Server 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 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL Configuration Wizard.
- 2 In the Select Configuration Option panel, click **MS-SQL Server - Service Group Configuration**, click **Delete**, and then click **Next**.
- 3 Review the prerequisites and click **Next**.
- 4 In the Service Group Selection panel, select the service group to delete and click **Next**.
- 5 In the Service Group Summary panel, 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**.

Administer an MSDTC service group

The SQL configuration wizard enables you to modify an MSDTC service group or delete the service group.

Modifying an MSDTC service group

You can modify the configuration of the service groups using the configuration wizards, the Java Console, and the Cluster Management Console (Single Cluster Mode) also referred to as the Web Console. This section describes the steps to modify an MSDTC service group using the SQL configuration wizard.

Note the following prerequisites before modifying an MSDTC service group:

- If the MSDTC 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 SystemList, do not run the wizard on the node being removed.

To modify an MSDTC service group

- 1 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL Configuration Wizard.
- 2 In the Select Configuration Option panel, click **MSDTC 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 Follow the wizard instructions and make desired modifications to the service group configuration.
See “[Configuring an MSDTC service group](#)” on page 88.

Deleting an MSDTC service group

If you delete the service group using the wizard, the wizard may not be able to revert the cluster settings for the MSDTC service, on the node. In such cases, after the service group is deleted, you will have to uninstall and re-install the MSDTC service, and then restart the node.

To uninstall the MSDTC service, type the following on the command line:

```
C: /> msdtc -uninstall
```

To install the MSDTC service, type the following on the command line:

```
C: /> msdtc -install
```

The following steps describe how to delete an MSDTC service group using the SQL Configuration Wizard.

To delete an MSDTC service group

- 1 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL Configuration Wizard.
- 2 In the Select Configuration Option panel, click **MSDTC Server - Service Group Configuration**, click **Delete**, and then click **Next**.
- 3 Review the prerequisites and click **Next**.
- 4 In the Service Group Selection panel, select the service group to delete and click **Next**.
- 5 In the Service Group Summary panel, 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**.

Detail monitoring of a SQL Server database instance

Use detail monitoring capability of VCS database agent for Microsoft SQL to monitor the status of a database. Before setting up detail monitoring, you must have the SQL Server agent running at the basic level of monitoring, that is, the DetailMonitor attribute must be set to 0.

The SQL Server agent uses a script to monitor the status of the database. A sample SQL script, located at %VCS_HOME%\bin\SQLServer<2000/2005>\sample_script.sql, is provided with the agent for the purpose. If the script is successfully executed during monitoring, the agent considers the database instance available. If the execution fails, the database instance is considered not available and the service group faults and fails over to the failover nodes. You can customize the script to meet your configuration requirements.

Note: You should use a separate script for each SQL Server service group that exists in the cluster. The script should exist on all the nodes in the service group.

Enabling detail monitoring

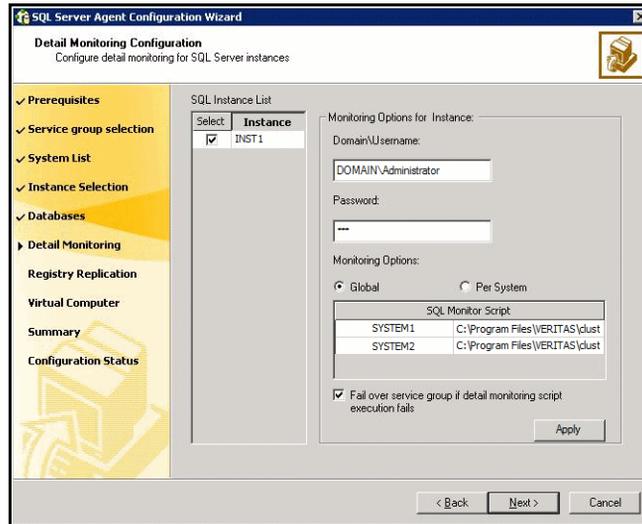
The following steps describe how to enable detail monitoring using the configuration wizard.

Note: You must be logged on as a Domain Administrator on the node where you run the wizard.

To enable detail monitoring

- 1 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL 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 configured for the instance to be monitored in detail and click **Next**.
- 5 In the SQL Server Instance Selection panel, click **Next**.

- 6 In the Detail Monitoring Configuration panel, specify information to enable detail monitoring and then click **Next**.



SQL Instance List

Check the box adjacent to the SQL Server instance for which you want to configure detail monitoring. Only the instances selected in [step 5](#) on page 110 will be available for selection.

Domain\Username

Type the fully qualified user name for connecting to SQL Server database. Make sure the specified user has SQL Server login permissions

Password

Type the password for the user name specified in the Domain\Username field.

Global

Click **Global** if the path of the monitor script is the same on all nodes and then click SQL Monitor Script and type the path for the monitor script.

Per System

Click **Per System** if the path of the monitor script is different on all the nodes and then click SQL Monitor Script and type the path for the monitor script. Make sure that the specified path exists on all the systems in the cluster.

Detail monitoring of a SQL Server database instance

Fail over service group if detail monitoring script execution fails

Check this option if you want the SQL service group to fail over when the detail monitor script execution fails.

Apply

Click **Apply** to configure detail monitoring for the selected SQL instance.

Repeat this for all instances that you want to configure detail monitoring for.

- 7 Follow the wizard instructions and accept the default values in the subsequent wizard panels.
See “[Configuring SQL Server in a VCS cluster](#)” on page 69.

Disabling detail monitoring

The following steps describe how to disable detail monitoring using the configuration wizard.

Note: You must be logged on as a Domain Administrator on the node where you run the wizard.

To disable detail monitoring

- 1 Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the SQL 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 configured for the instance for which detail monitoring is being disabled, and then click **Next**.
- 5 In the SQL Server Instance Selection panel, click **Next**.
- 6 In the Detail Monitoring Configuration panel, uncheck the check box corresponding to the SQL Server instance for which detail monitoring is being disabled, click **Apply** and then click **Next**.
- 7 Follow the wizard instructions and accept the default values in the subsequent wizard panels.
See [“Configuring SQL Server in a VCS cluster”](#) on page 69.

Troubleshooting VCS agents

This chapter contains the following topics:

- [“About troubleshooting the VCS database agent for SQL”](#) on page 116
- [“VCS logging”](#) on page 116
- [“Error messages and descriptions”](#) on page 118

About troubleshooting the VCS database agent for SQL

This chapter describes how to troubleshoot common problems in the VCS database agent for Microsoft SQL. 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

The agent log message components are defined as follows:

- Timestamp denotes the date and time when the message was logged.
- Mnemonic denotes which Symantec product logs the message. For VCS database 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 the user that VCS has initiated an action.
 - INFO informs the user of various state messages or comments. Of these, CIRITICAL, 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-16-20024-13
Originator ID for all VCS products is 'V-16.<number>' Category ID for SQLServer agent is 20020, for MSDTC agent is 20021, while 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 written to the Windows event log.

A typical agent log resembles:

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

Error messages and descriptions

Occasionally you may encounter problems when using the VCS agents. The following sections list the most common messages of type ERROR and WARNING. Each message includes a description and a recommended solution, if applicable.

Agent for SQL Server 2000

[Table 9-1](#) on page 118 describes the error messages for the VCS agent for SQL Server 2000.

Table 9-1 SQL Server 2000 agent 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.

Table 9-1 SQL Server 2000 agent error messages (Continued)

Message	Description
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.
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.

Table 9-1 SQL Server 2000 agent error messages (Continued)

Message	Description
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.
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

[Table 9-2](#) on page 120 describes the error messages for the VCS agent for MSSearch.

Table 9-2 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

[Table 9-3](#) on page 121 describes the error messages for the VCS agent for MSDTC.

Table 9-3 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.

Table 9-3 MSDTC agent error messages (Continued)

Message	Description
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.
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

[Table 9-4](#) on page 122 describes the error messages for the VCS agent for SQL Server 2005.

Table 9-4 SQL Server 2005 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 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.

Table 9-4 SQL Server 2005 agent error messages (Continued)

Message	Description
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.
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.

Table 9-4 SQL Server 2005 agent error messages (Continued)

Message	Description
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.
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.

Table 9-4 SQL Server 2005 agent error messages (Continued)

Message	Description
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

Table 9-5 on page 126 describes the error messages for the VCS agent for SQL Agent service.

Table 9-5 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 9-5 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

Table 9-6 on page 128 describes the error messages for the VCS agent for SQL Analysis service.

Table 9-6 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 9-6 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 chapter contains the following topics:

- [“About resource type definitions”](#) on page 132
- [“Agent for SQL Server 2000”](#) on page 132
- [“Agent for MSSearch service”](#) on page 135
- [“Agent for SQL Server 2005”](#) on page 137
- [“Agent for SQL Server 2005 Agent service”](#) on page 140
- [“Agent for SQL Server 2005 Analysis service”](#) on page 142
- [“MSDTC agent”](#) on page 144

About resource type definitions

This appendix lists resource type definitions and attribute definitions of the SQL Server 2000 agents, SQL Server 2005 agents, and MSDTC agent.

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 Definition table explains 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.

Agent for SQL Server 2000

The resource type definitions and attribute definitions for the VCS agent for SQL Server 2000 are as follows. This information will assist you during the agent configuration.

Resource type definition

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

```
type SQLServer2000 (  
    static i18nstr ArgList[] = {Instance,  
        "LanmanResName:VirtualName", "MountResName:MountPath",  
        SQLOnlineTimeout, SQLOfflineTimeout, DetailMonitor,  
        FaultOnDMScriptFailure, SQLDetailMonitorTimeout, Username,  
        Domain, Password, SQLFile }  
    str Instance  
    str LanmanResName  
    int SQLOnlineTimeout = 90  
    int SQLOfflineTimeout = 90  
    boolean DetailMonitor = 0  
    boolean FaultOnDMScriptFailure = 1  
    int SQLDetailMonitorTimeout = 30  
    i18nstr Username  
    i18nstr Domain  
    str Password  
    i18nstr SQLFile  
    str MountResName  
)
```

Attribute definitions

[Table A-1](#) on page 133 describes the attributes associated with the VCS agent for SQL Server 2000. The Required attributes table lists the attributes that must be configured for the agent to function properly.

Table A-1 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-2](#) on page 133 describes the optional attributes associated with the VCS agent for SQL Server 2000.

Table A-2 SQL Server 2000 agent optional attributes

Optional Attributes	Type and Dimension	Definition
DetailMonitor	boolean	<p>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.</p> <p>Default = 0</p> <p>Note: If the attribute is set to 1, the attributes Username, Password, Domain, SQLDetailMonitorTimeOut, and SQLFile must be assigned appropriate values.</p>

Table A-2 SQL Server 2000 agent optional attributes (Continued)

Optional Attributes	Type and Dimension	Definition
FaultOnDMScriptFailure	boolean	<p>Defines whether the agent fails over the service group if the detail monitoring script execution fails.</p> <p>Default = 1</p> <p>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.</p>
SQLDetailMonitorTimeout	integer-scalar	<p>Number of seconds that can elapse before the detail monitor routine aborts. Default is 30.</p>
Username	string-scalar	<p>Note: 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.</p>
Domain	string-scalar	<p>Note: 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.</p>
Password	string-scalar	<p>Password for logging in to a database for detail monitoring. This attribute must not be null if DetailMonitor attribute is set to 1.</p>
SQLFile	string-scalar	<p>Note: The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the DetailMonitor attribute is set to 1.</p>

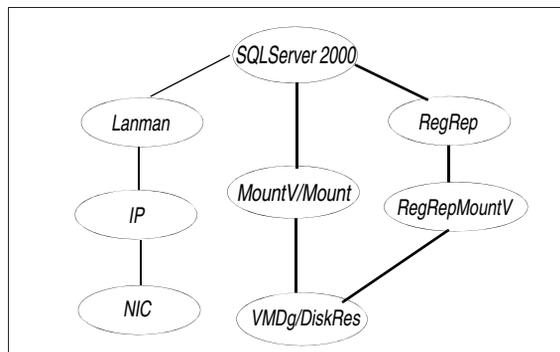
Sample dependency graph

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*.

Figure A-1 on page 135 illustrates a sample dependency graph for a SQL Server 2000 service group. The shared disk group is configured using Volume Manager Diskgroup (VMDg) 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 MountV mount point is created using the MountVresource. SQL Server 2000 registry is replicated using the RegRep and RegRepMountV 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 resource type definitions and attribute definitions for the VCS agent for MSSearch service are as follows.

Resource type definition

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

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

Attribute definitions

[Table A-3](#) on page 136 describes the attribute associated with the VCS agent for MSSearch service. The Required attribute must be configured for the agent to function properly.

Table A-3 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 13 for more information.

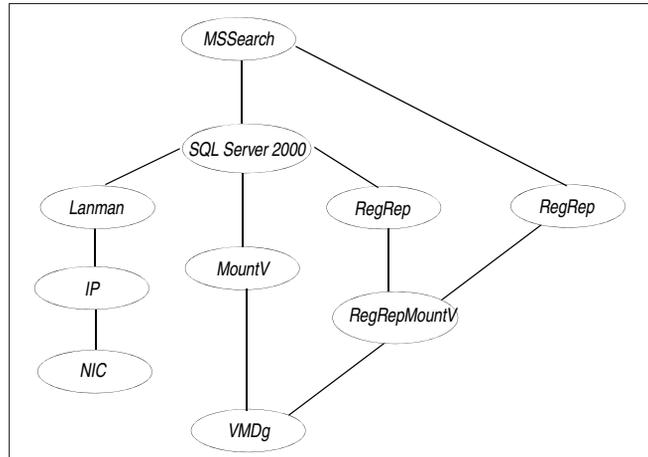
Sample dependency graph

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*.

[Figure A-2](#) on page 137 illustrates the dependency graph for an MSSearch resource configured in a SQL Server 2000 service group. The shared disk group is configured using Volume Manager Diskgroup (VMDg) 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 MountV mount point is created using the MountV resource. SQL Server 2000 and MSSearch registry is replicated using the RegRep and RegRepMountV 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 MSSearch service group dependency graph



Agent for SQL Server 2005

The resource type definitions and attribute definitions for the VCS agent for SQL Server 2005 are as follows.

Resource type definition

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

```

type SQLServer2005 (
  static i18nstr ArgList[] =
  {Instance, "LanmanResName:VirtualName", SQLOnlineTimeout,
  SQLOfflineTimeout, DetailMonitor, FaultOnDMScriptFailure,
  SQLDetailMonitorTimeout, Username, Domain, Password,
  SQLFile }
  str Instance
  str LanmanResName
  int SQLOnlineTimeout = 90
  int SQLOfflineTimeout = 90
  boolean DetailMonitor = 0
  boolean FaultOnDMScriptFailure = 1
  int SQLDetailMonitorTimeout = 30
  i18nstr Username
  i18nstr Domain
  str Password
  i18nstr SQLFile
)

```

Attribute definitions

[Table A-4](#) on page 138 describes the attributes associated with the VCS agent for SQL Server 2005. The Required attributes table lists the attributes that must be configured for the agent to function properly.

Table A-4 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-5](#) on page 138 describes the optional attributes associated with the VCS agent for SQL Server 2005.

Table A-5 SQL Server 2005 agent optional attributes

Optional Attributes	Type and Dimension	Definition
DetailMonitor	boolean	<p>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.</p> <p>Default = 0.</p> <p>Note: If the attribute is set to 1, the attributes Username, Password, Domain, SQLDetailMonitorTimeOut, and SQLFile must be assigned appropriate values.</p>

Table A-5 SQL Server 2005 agent optional attributes (Continued)

Optional Attributes	Type and Dimension	Definition
FaultOnDMScriptFailure	boolean	<p>Defines whether the agent fails over the service group if the detail monitoring script execution fails.</p> <p>Default = 1</p> <p>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.</p>
SQLDetailMonitorTimeout	integer-scalar	<p>Number of seconds that can elapse before the detail monitor routine aborts. Default is 30.</p>
Username	string-scalar	<p>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.</p> <p>Note: This attribute can take localized values.</p>
Domain	string-scalar	<p>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.</p> <p>Note: This attribute can take localized values.</p>
Password	string-scalar	<p>Password for logging in to a database for in-depth monitoring. This attribute must not be null if DetailMonitor attribute is set to 1.</p>
SQLFile	string-scalar	<p>The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the DetailMonitor attribute is set to 1.</p> <p>Note: This attribute can take localized values.</p>

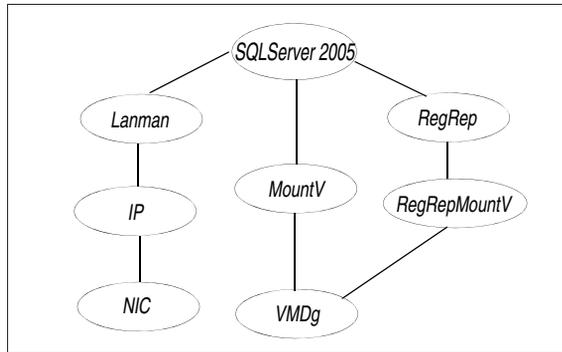
Sample dependency graph

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*.

Figure A-3 on page 140 illustrates the dependency graph for a SQL Server 2005 service group. The shared disk group is configured using Volume Manager Diskgroup (VMDg) 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 MountV mount point is created using the MountV resource. SQL Server 2005 registry is replicated using the RegRep and RegRepMountV 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 resource type definitions and attribute definitions for the VCS agent for SQL Server 2005 Agent service are as follows.

Resource type definition

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

```

type SQLAgService2005 (
  static i18nstr ArgList[] =
  { "SQLServer2005ResName:Instance",
    "LanmanResName:VirtualName" }
  str SQLServer2005ResName
  str LanmanResName
)

```

Attribute definitions

[Table A-6](#) on page 141 describes the attributes associated with the VCS agent for SQL Server 2005 Agent service. The Required attributes must be configured for the agent to function properly.

Table A-6 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.

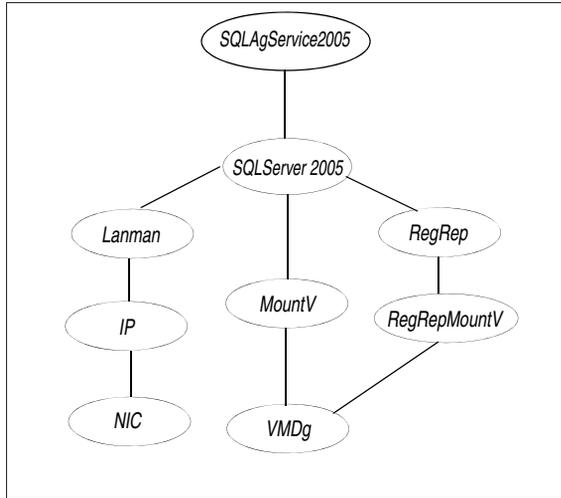
Sample dependency graph

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*.

[Figure A-4](#) on page 142 illustrates the dependency graph for a SQL Server 2005 Agent service resource in a SQL Server 2005 service group. The shared disk group is configured using Volume Manager Diskgroup (VMDg) 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 MountV mount point is created using the MountV resource. SQL Server 2005 registry is replicated using the RegRep and RegRepMountV 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 resource type definitions and attribute definitions for the VCS agent for SQL Server 2005 Analysis service are as follows.

Resource type definition

The agent for SQL Server 2005 Analysis service is represented by the SQL0lapService2005 resource type.

```
type SQL0lapService2005 (  
    static i18nstr ArgList[] =  
    { "SQLServer2005ResName:Instance",  
      "LanmanResName:VirtualName" }  
    str SQLServer2005ResName  
    str LanmanResName  
)
```

Attribute definitions

[Table A-7](#) on page 143 describes the attributes associated with the VCS agent for SQL Server 2005 Analysis service. The Required attributes must be configured for the agent to function properly.

Table A-7 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.

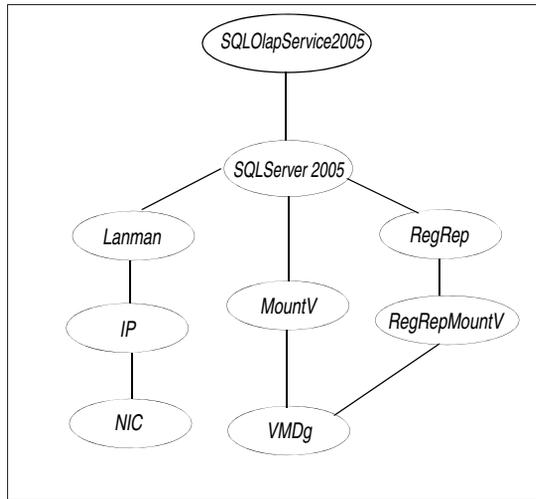
Sample dependency graph

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*.

[Figure A-5](#) on page 144 illustrates the dependency graph for a SQL Server 2005 Analysis service resource in a SQL Server 2005 service group. The shared disk group is configured using Volume Manager Diskgroup (VMDg) 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 MountV mount point is created using the MountV resource. SQL Server 2005 registry is replicated using the RegRep and RegRepMountV 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 resource type definitions and attribute definitions for the VCS agent for MSDTC are as follows.

Resource type definition

The MSDTC agent is represented by the MSDTC resource type.

```
type MSDTC (  
    static i18nstr ArgList[] = {"LanmanResName:VirtualName",  
    "MountResName:MountPath", LogPath }  
    str LanmanResName  
    str MountResName  
    i18nstr LogPath  
)
```

Attribute definitions

[Table A-8](#) on page 145 describes the attributes associated with the VCS agent for MSDTC. The Required attributes must be configured for the agent to function properly.

Table A-8 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.

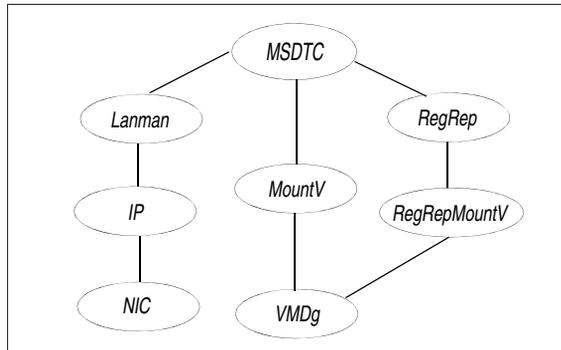
Sample dependency graph

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*.

Figure A-6 on page 146 illustrates the dependency graph for an MSDTC service group. The shared disk group is configured using Volume Manager Diskgroup (VMDg) 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 MountV mount point is created using the MountV agent. MSDTC registry is replicated using the RegRep and RegRepMountV resources. The MSDTC resource comes online after each of these resources are brought online.

Figure A-6 MSDTC service group dependency graph



Using the virtual MMC viewer

This chapter contains the following topics:

- [“About using the virtual MMC viewer”](#) on page 148
- [“Viewing DTC transaction information”](#) on page 148

About using the virtual MMC viewer

VCS starts the MSDTC service in a cluster under the context of the virtual server. Because the MMC snap-in is not aware of such a configuration, it is not possible to view the transactions on the DTC virtual server from the node where the MSDTC resource is online.

VCS provides a virtual MMC viewer, the VCS Application Manager (VAM) utility, that enables you to view the distributed transactions statistics on the DTC virtual server from a node where the MSDTC resource is online.

Note: If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

Viewing DTC transaction information

In cases where a communication line fails or a distributed transaction application leaves unresolved transactions, you might want to view transaction lists and statistics, control which transactions are displayed, set transaction time-out periods, and control how often transactions are updated.

The following steps describe how to view the DTC transactions information.

Prerequisites for viewing DTC transaction information are as follows:

- An MSDTC service group must be configured in the cluster.
- MSDTC client must be configured on the nodes where you wish to view the transactions.
- The MSDTC service group must be online on the node where you run the VCS Application Manager utility.

To view transactions from a node where MSDTC resource is online

- 1 Start the VCS Application Manager utility.
Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Application Manager**.

or

In the Solutions Configuration Center (SCC), under Tools, click **VCS Application Manager**.

The VCS Application Manager displays a list of supported application service groups configured in the cluster. For each service group it also displays the state of the service group, the name of the virtual server

resource (Lanman resource) and the corresponding management tools used for that application.

- 2 Select **MSDTC** from the Select the resource type drop-down list.
- 3 Select the MSDTC resource that is online and then click **Manage**, or double-click the MSDTC resource name.
VAM launches the Component Services snap-in in the virtual server context.
- 4 In the console tree of the Component Services administrative tool, expand **Component Services > Computers > My Computer > Distributed Transaction Coordinator > Local DTC**.
- 5 Click **Transaction List** to view all transactions, their status, and their identifiers. Right-click a transaction and click **View > Properties** to list the parent transaction and its children.
- 6 Click **Transaction Statistics** to view statistical information about the transactions in which a server participated. You can use transaction statistics to get an overview of DTC performance.
Refer to the Microsoft documentation for further information.

To view transactions from any node in the domain

- 1 Launch the Windows Component Services Administrative tool.
Click **Start > All Programs > Administrative Tools > Component Services**
- 2 In the console tree of the Component Services administrative tool, double-click **Component Services**, right-click **Computers**, click **New > Computer**.
- 3 In the Add Computer dialog box, specify the virtual server name that you specified while creating the MSDTC Server service group. If you are unsure of the exact name, click **Browse** to search from a list of all computers on the network and select the virtual computer name from the list.
- 4 Click **OK**. The virtual computer entry is added to the Computers container.
- 5 Expand the newly added virtual computer entry and double-click **Distributed Transaction Coordinator**.
- 6 Click **Transaction List** to view all transactions, their status, and their identifiers. Right-click a transaction and click **View > Properties** to list the parent transaction and its children.
- 7 Click **Transaction Statistics** to view statistical information about the transactions in which a server participated. You can use transaction statistics to get an overview of DTC performance.
Refer to the Microsoft documentation for further information.

Sample configurations

This chapter contains the following topics:

- [“About sample configurations”](#) on page 152
- [“Sample configuration for a SQL Server 2000 cluster”](#) on page 152
- [“Sample configuration for a SQL Server 2005 cluster”](#) on page 156

About 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.

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 MYCLUSTER (
    UserNames = { admin = fQRjQLqNRmRRpZRlQO }
    Administrators = { admin }
    CredRenewFrequency = 0
    CounterInterval = 5
)

system SYSTEM1 (
)

system SYSTEM2 (
)

group sql2k_sg (
    SystemList = { SYSTEM1 = 0, SYSTEM@ = 1 }
)

IP sql2k_sg-IP (
    Address = "10.182.147.149"
    SubNetMask = "255.255.252.0"
    MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
    MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

Lanman sql2k_sg-Lanman (
    VirtualName = VSQl2K
    IPResName = sql2k_sg-IP
)

MSSearch sql2k_sg-MSSearch (
    AppName = "SQLServer$SQL2K"
)

MountV sql2k_sg-MountV (
    MountPath = "O:"
    VolumeName = sysdata
)
```

```
VMDGResName = sql2k_sg-VMDg
)

MountV sql2k_sg-MountV-1 (
  MountPath = "P:"
  VolumeName = regrep
  VMDGResName = sql2k_sg-VMDg
)

NIC sql2k_sg-NIC (
  MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
  MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

RegRep sql2k_sg-RegRep-MSSQL (
  MountResName = sql2k_sg-MountV-1
  ReplicationDirectory = "\\RegRep\\sql2k_sg-RegRep-MSSQL"
  Keys =
  { "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL Server\\SQL2K",
    "HKLM\\SOFTWARE\\Microsoft\\MSSQLServer\\Client" }
  ExcludeKeys =
  { "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\SQL2K\\Setup",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\SQL2K\\SQLServerAgent\\Subsystems",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\SQL2K\\Tracking" }
)

RegRep sql2k_sg-RegRep-MSSearch (
  MountResName = sql2k_sg-MountV-1
  ReplicationDirectory = "\\RegRep\\sql2k_sg-RegRep-MSSearch"
  Keys =
  { "HKLM\\Software\\Microsoft\\Search\\1.0\\
    Gathering Manager\\Applications\\SQLServer$SQL2K",
    "HKLM\\Software\\Microsoft\\Search\\1.0\\
    Gather\\Applications\\SQLServer$SQL2K",
    "HKLM\\Software\\Microsoft\\Search\\1.0\\
    Applications\\Applications\\SQLServer$SQL2K",
    "HKLM\\Software\\Microsoft\\Search\\1.0\\
    CatalogNames\\Applications\\SQLServer$SQL2K",
    "HKLM\\Software\\Microsoft\\Search\\1.0\\
    Indexer\\Applications\\SQLServer$SQL2K",
    "HKLM\\Software\\Microsoft\\Search\\1.0\\
    Databases\\Applications\\SQLServer$SQL2K" }
)

SQLServer2000 sql2k_sg-SQLServer2000 (
  Instance = SQL2K
  LanmanResName = sql2k_sg-Lanman
  MountResName = sql2k_sg-MountV
)
```

Sample configuration for a SQL Server 2000 cluster

```

VMDg sql2k_sg-VMDg (
    DiskGroupName = SQL2K_DG
    DGGuid = bd9d1f73-aed9-4b36-bdbb-24e6e653f69e
)

sql2k_sg-SQLServer2000 requires sql2k_sg-MountV
sql2k_sg-SQLServer2000 requires sql2k_sg-Lanman
sql2k_sg-SQLServer2000 requires sql2k_sg-RegRep-MSSQL
sql2k_sg-MountV requires sql2k_sg-VMDg
sql2k_sg-RegRep-MSSQL requires sql2k_sg-MountV-1
sql2k_sg-MSSearch requires sql2k_sg-SQLServer2000
sql2k_sg-MSSearch requires sql2k_sg-RegRep-MSSearch
sql2k_sg-RegRep-MSSearch requires sql2k_sg-MountV-1
sql2k_sg-MountV-1 requires sql2k_sg-VMDg
sql2k_sg-IP requires sql2k_sg-NIC
sql2k_sg-Lanman requires sql2k_sg-IP

group MSDTC_SG (
    SystemList = { SYSTEM1 = 0, SYSTEM2 = 1 }
)
IP MSDTC_SG-IP (
    Address = "10.182.147.128"
    SubNetMask = "255.255.252.0"
    MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
    MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

Lanman MSDTC_SG-Lanman (
    VirtualName = VDTC
    IPResName = MSDTC_SG-IP
    DNSUpdateRequired = 1
    DNSCriticalForOnline = 1
)

MSDTC MSDTC_SG-MSDTC (
    LanmanResName = MSDTC_SG-Lanman
    MountResName = MSDTC_SG-MountV
    LogPath = "\\MSDTC"
)

MountV MSDTC_SG-MountV (
    MountPath = "L:"
    VolumeName = log
    VMDGResName = MSDTC_SG-VMDg
)

MountV MSDTC_SG-MountV-1 (
    MountPath = "Q:"
    VolumeName = regrep
    VMDGResName = MSDTC_SG-VMDg
)

```

```
NIC MSDTC_SG-NIC (
  MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
  MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

RegRep MSDTC_SG-RegRep (
  MountResName = MSDTC_SG-MountV-1
  ReplicationDirectory = "\\RegRep"
  Keys = { "HKCR\\CID", "HKCR\\SVCID",
           "HKLM\\Software\\Microsoft\\MSDTC\\Security" }
  ExcludeKeys =
  { "HKCR\\CID\\2d032e98-c348-4dc0-ba9b-a51c3f5330d0\\
    CustomProperties\\Service\\Path" }
)

VMDg MSDTC_SG-VMDg (
  DiskGroupName = MSDTC1
  DGGuid = "1eb1a20f-4737-4bee-bdfc-c828c614fcf1"
)

requires group msdtc_rvg_sg online local hard
MSDTC_SG-IP requires MSDTC_SG-NIC
MSDTC_SG-Lanman requires MSDTC_SG-IP
MSDTC_SG-MSDTC requires MSDTC_SG-Lanman
MSDTC_SG-MSDTC requires MSDTC_SG-RegRep
MSDTC_SG-MSDTC requires MSDTC_SG-MountV
MSDTC_SG-RegRep requires MSDTC_SG-MountV-1
MSDTC_SG-MountV requires MSDTC_SG-VMDg
MSDTC_SG-MountV-1 requires MSDTC_SG-VMDg
```

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 MYCLUSTER (
    UserNames = { admin = GPQiPKpMQlQQoYQkPN }
    ClusterAddress = "10.182.146.127"
    Administrators = { admin }
    CredRenewFrequency = 0
    CounterInterval = 5
)

system SYSTEM1 (
)

system SYSTEM2 (
)

group SQL_SG (
    SystemList = { SYSTEM1 = 0, SYSTEM2 = 1 }
)

GenericService SQL_SG-MSSearch (
    ServiceName = "msftesql$FIRST"
)

IP SQL_SG-IP (
    Address = "10.182.147.127"
    SubNetMask = "255.255.252.0"
    MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
    MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

Lanman SQL_SG-Lanman (
    VirtualName = VSQL2K5
    IPResName = SQL_SG-IP
)

MountV SQL_SG-MountV (
    MountPath = "S:"
    VolumeName = data
    VMDGResName = SQL_SG-VMDg
)

MountV SQL_SG-MountV-1 (
    MountPath = "R:"
```

```
VolumeName = regrep
VMDGResName = SQL_SG-VMDg
)

NIC SQL_SG-NIC (
  MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
  MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

RVGPrimary SQL_SG-RVGPrimary (
  RvgResourceName = sql_rvg_sg-VvrRvg
)

RegRep SQL_SG-RegRep-MSSQL (
  MountResName = SQL_SG-MountV-1
  ReplicationDirectory = "\\RegRep\\SQL_SG-RegRep-MSSQL"
  Keys @SYSTEM1 =
  { "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\MSSQLServer",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\Replication",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\SQLServerAgent",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\PROVIDERS",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\SQLServerSCP" }
  Keys @SYSTEM2 =
  { "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\MSSQLServer",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\Replication",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\SQLServerAgent",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\PROVIDERS",
    "HKLM\\SOFTWARE\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\SQLServerSCP" }
  ExcludeKeys = { "" }
)

RegRep SQL_SG-RegRep-MSSearch (
  MountResName = SQL_SG-MountV-1
  ReplicationDirectory = "\\RegRep\\SQL_SG-RegRep-MSSearch"
  Keys @SYSTEM1 =
  { "HKLM\\Software\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\MSSearch\\GlobalAdminBlob" }
  Keys @SYSTEM2 =
  { "HKLM\\Software\\Microsoft\\Microsoft SQL
    Server\\MSSQL.1\\MSSearch\\GlobalAdminBlob" }
  ExcludeKeys = { "" }
)
```

Sample configuration for a SQL Server 2005 cluster

```

SQLAgService2005 SQL_SG-SQLAgService (
    SQLServer2005ResName = SQL_SG-SQLServer2005
    LanmanResName = SQL_SG-Lanman
)

SQLOlapService2005 SQL_SG-SQLOlapService (
    SQLServer2005ResName = SQL_SG-SQLServer2005
    LanmanResName = SQL_SG-Lanman
)

SQLServer2005 SQL_SG-SQLServer2005 (
    Instance = FIRST
    LanmanResName = SQL_SG-Lanman
)

VMDg SQL_SG-VMDg (
    DiskGroupName = sql12k5_dg
    DGGuid = "3ecadf51-2ad5-48da-b93a-0a193718fb94"
)

SQL_SG-SQLServer2005 requires SQL_SG-MountV
SQL_SG-SQLServer2005 requires SQL_SG-RegRep-MSSQL
SQL_SG-SQLServer2005 requires SQL_SG-Lanman
SQL_SG-MountV requires SQL_SG-VMDg
SQL_SG-RegRep-MSSQL requires SQL_SG-MountV-1
SQL_SG-MountV-1 requires SQL_SG-VMDg
SQL_SG-IP requires SQL_SG-NIC
SQL_SG-Lanman requires SQL_SG-IP
SQL_SG-MSSearch requires SQL_SG-SQLServer2005
SQL_SG-MSSearch requires SQL_SG-RegRep-MSSearch
SQL_SG-RegRep-MSSearch requires SQL_SG-MountV-1
SQL_SG-SQLAgService requires SQL_SG-SQLServer2005
SQL_SG-SQLOlapService requires SQL_SG-SQLServer2005

group MSDTC_SG (
    SystemList = { SYSTEM1 = 0, SYSTEM2 = 1 }
)
IP MSDTC_SG-IP (
    Address = "10.182.147.128"
    SubNetMask = "255.255.252.0"
    MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
    MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

Lanman MSDTC_SG-Lanman (
    VirtualName = VDTC
    IPResName = MSDTC_SG-IP
    DNSUpdateRequired = 1
    DNSCriticalForOnline = 1
)

```

```

MSDTC MSDTC_SG-MSDTC (
    LanmanResName = MSDTC_SG-Lanman
    MountResName = MSDTC_SG-MountV
    LogPath = "\\MSDTC"
)

MountV MSDTC_SG-MountV (
    MountPath = "L:"
    VolumeName = log
    VMDGResName = MSDTC_SG-VMDg
)

MountV MSDTC_SG-MountV-1 (
    MountPath = "Q:"
    VolumeName = regrep
    VMDGResName = MSDTC_SG-VMDg
)

NIC MSDTC_SG-NIC (
    MACAddress @SYSTEM1 = "00-11-43-D1-9F-E3"
    MACAddress @SYSTEM2 = "00-11-43-D1-9E-74"
)

RegRep MSDTC_SG-RegRep (
    MountResName = MSDTC_SG-MountV-1
    ReplicationDirectory = "\\RegRep"
    Keys = { "HKCR\\CID", "HKCR\\SVCID",
            "HKLM\\Software\\Microsoft\\MSDTC\\Security" }
    ExcludeKeys =
    { "HKCR\\CID\\2d032e98-c348-4dc0-ba9b-a51c3f5330d0\\
      CustomProperties\\Service\\Path" }
)

VMDg MSDTC_SG-VMDg (
    DiskGroupName = MSDTC1
    DGGuid = "1eb1a20f-4737-4bee-bdfc-c828c614fcf1"
)

requires group msdtc_rvg_sg online local hard
MSDTC_SG-IP requires MSDTC_SG-NIC
MSDTC_SG-Lanman requires MSDTC_SG-IP
MSDTC_SG-MSDTC requires MSDTC_SG-Lanman
MSDTC_SG-MSDTC requires MSDTC_SG-RegRep
MSDTC_SG-MSDTC requires MSDTC_SG-MountV
MSDTC_SG-RegRep requires MSDTC_SG-MountV-1
MSDTC_SG-MountV requires MSDTC_SG-VMDg
MSDTC_SG-MountV-1 requires MSDTC_SG-VMDg

```


Index

A

- Active/Active environment
 - definition of 18
 - installing SQL 19
- application failure, detecting 17
- attribute definitions
 - MSDTC agent 145
 - MSSearch agent 136
 - SQL Server 2000 agent 133
 - SQL Server 2005 agent 138
 - SQL Server Agent Service agent 141
 - SQL Server Analysis Service agent 143

C

- cluster
 - configure LLT over ethernet 32
 - configure LLT over UDP 33
- clusters
 - configuring cluster 28
- configure
 - LLT over ethernet 32
 - LLT over UDP using VCW 33
- configure cluster
 - ethernet 32
 - UDP 33
- configuring
 - cluster (HA) 28

D

- database
 - new database for SQL Server 2000 78
 - user-defined for SQL Server 2000, SQL Server 2005 78
- dependency graph
 - MSDTC agent 146
 - MSSearch agent 136
 - SQL Server 2000 agent 134
 - SQL Server 2005 agent 140
 - SQL Server Agent Service agent 141
 - SQL Server Analysis Service agent 143

- detail monitoring
 - definition of 110
 - disabling 113
 - enabling 110
- disk groups, importing 47

E

- enterprise agent. See SQL Server database agent

F

- failover
 - MSDTC 18
 - SQL Server 2000 17
 - SQL Server 2000 with MSSearch 18
 - SQL Server 2005 agent 18

L

- LLT over ethernet
 - configuring using VCW 32
- LLT over UDP
 - configuring using VCW 33

M

- MSDTC
 - HA configuration overview 84
 - service group configuration (VCS) 84
- MSDTC agent
 - attributes 145
 - operations 16
- MSDTC agent attributes
 - LanmanResName 145
 - LogPath 145
 - MountResName 145
- MSDTC client, configuring using wizard 93
- MSDTC service 13
- MSDTC service group
 - configuring 83
 - modifying configuration 108
- MSSearch agent

- attributes 136
- operations 14
- MSSearch agent attribute, AppName 136
- MSSearch service 13
- multiple SQL Server instances
 - assigning the port 65

O

- offlining service groups 98
- onlining service groups 96

P

- port assignment for multiple instances 65

R

- removing agent 96
- resource type definition
 - MSDTC agent 144
 - MSSearch agent 135
 - SQL Server 2000 agent 132
 - SQL Server 2005 agent 137
 - SQL Server Agent service agent 140
 - SQL Server Analysis Service agent 142

S

- Security Services
 - configuring 35
- service groups
 - bringing online 96
 - switching 97
 - taking offline 98
- SFW
 - creating volumes 47
 - importing disk groups 47
- SQL Server 2000
 - creating a new database 78
 - installing in Active/Active environment 19
 - overview of user-defined database 78
- SQL Server 2000 database agent
 - MSSearch agent 14
 - SQL Server agent 14
- SQL Server 2005
 - setting services to manual start 62
 - stopping the SQL Server service 62
- SQL Server 2005 agent
 - attributes 138
- SQL Server agent

- attributes 133
- configuring 67
- operations 14, 15
- SQL Server Agent Service agent
 - attributes 141
- SQL Server Analysis Service agent
 - attributes 143
- SQL Server database agent
 - disabling 99
 - installing 25
 - monitoring options 17
 - MSDTC agent 16
- SQL Server service group
 - modifying configuration 106
- SQL Server services
 - MSDTC 13
 - MSSearch 13
- switching service groups 97

U

- utilities
 - Veritas Application Manager 148

V

- VCS
 - configuring the cluster for HA 28
- VCS utilities
 - Veritas Application Manager 148
- Veritas Application Manager utility 148