

Veritas[™] Cluster Server One Agent for Microsoft SQL Configuration Guide

Windows Server 2003
Windows Server 2008

5.0 Service Pack 1



Veritas Cluster Server One by Symantec 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: 5.0. Service Pack 1

Document version: 5.0.SP1.0

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 maintenance 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 Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

For information about Symantec's Maintenance Programs, you can visit our Web site at the following URL:

<http://www.symantec.com/business/support/index.jsp>

Contacting Technical Support

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

http://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:

<http://customercare.symantec.com>

Customer service

Customer service information is available at the following URL:

<http://customercare.symantec.com>

Customer Service is available to assist with 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 maintenance 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

Maintenance agreement resources

If you want to contact Symantec regarding an existing maintenance agreement, please contact the maintenance 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

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Chapter 1	Introducing the VCS One agents for SQL Server	
	About the VCS One agents for Microsoft SQL Server	12
	About the agent for SQL Server Database Engine	13
	Attribute definitions for SQL Server Database Engine agent	14
	About the GenericService agent for SQL Server services	17
	Attribute definitions for GenericService agent	18
	About the agent for FileStream	21
	Attribute definitions for FileStream agent	22
	About monitoring options	23
	How the agents make SQL Server highly available	24
	Local cluster configuration	24
	Disaster recovery configuration	24
	Running SQL Server in an Active-Active clustered environment	27
	SQL Server failover service group- resource dependency graph	27
	SQL Server standalone (start/stop) service group- resource dependency graph	29
	Typical SQL Server configuration in a VCS One cluster	31
Chapter 2	Installing, modifying and removing the VCS One database agent for SQL Server	
	Before installing the agent	35
	Installing the agent	36
	Adding the agent resource type to the Policy Master Server-Using CLI	38
	Importing the Binary Message Catalog (BMC) files	39
	Modifying or removing the agent installation	40
	Uninstalling the agent pack	42
Chapter 3	Installing and configuring SQL Server	
	About installing and configuring SQL Server	46
	Supported software	47
	Before installing SQL Server	48
	About managing storage	48
	About managing storage using Windows Logical Disk Management (LDM)	49
	About managing storage using SFW	52

Unmounting the volumes	56
Installing SQL Server	57
Installing SQL Server on the first cluster system	57
Installing SQL Server on the failover cluster systems	58
Changing the default datapath to the shared drive in case of SQL 2005	59
Changing the default datapath to the shared drive in case of SQL 2008	60

Chapter 4 Configuring the SQL service group

About configuring the failover SQL Server service group	64
About configuring the start/stop SQL Server service group	64
Before you configure a SQL Server service group	66
Creating a SQL Server service group using the wizard	68

Chapter 5 Making an existing SQL Server highly available

About making existing SQL server highly available	80
Reviewing the configuration	81
Before you begin to install VCS One	82
Configuring the storage, hardware and network	82
Converting the existing SQL server into a clustered server	83
Moving the existing SQL Server data files and user databases	84

Chapter 6 Administering the service group

Modifying the SQL Server service group	86
Prerequisites	86
Modifying a SQL Server service group	87
Creating a SQL Server user-defined database	89
Adding storage agent resources to the SQL service group	90

Chapter 7 Upgrading SQL Server

About upgrading SQL Server	94
Upgrading from Microsoft SQL Server 2005 to SQL Server 2008	94
Upgrading SQL on the first cluster system	95
Upgrading SQL Server on the failover systems	95
Upgrading the SQL Server service packs	97
Upgrading SQL Server 2005 to SQL Server 2005 SP1	97
Upgrading SQL Server 2005 to 2005 SP2 or later	98
Prerequisites	98
Preliminary installation information	99
Installing the Service Pack on “passive” cluster systems	100

	Installing the Service Pack on “active” cluster systems	101
	Prerequisites	103
	Upgrading the cluster	103
Appendix A	Resource type definitions	
	About resource type definitions	106
	Agent for SQL Server	106
	Resource type definition	106
	Agent for SQL Server FileStream	108
	Resource type definition	108
Appendix B	Troubleshooting VCS One agents	
	About troubleshooting the VCS One database agent for SQL	112
	Newly added SQL Server 2008 Databases are not accessible on failover	112
	SQL Server service group faults upon a failover with an error - Failed to find Disk Group	113
	Error messages and descriptions	113
	VCS One logging	119
Index		121

Introducing the VCS One agents for SQL Server

This chapter contains the following topics:

- [“About the VCS One agents for Microsoft SQL Server”](#) on page 12
- [“About monitoring options”](#) on page 23
- [“How the agents make SQL Server highly available”](#) on page 24
- [“Running SQL Server in an Active-Active clustered environment”](#) on page 27
- [“SQL Server failover service group- resource dependency graph”](#) on page 27
- [“SQL Server standalone \(start/stop\) service group- resource dependency graph”](#) on page 29
- [“Typical SQL Server configuration in a VCS One cluster”](#) on page 31

About the VCS One agents for Microsoft SQL Server

The VCS One agents for Microsoft SQL Server provides high availability and “Active-Active” support for Microsoft SQL Server 2005 and Microsoft SQL Server 2008 and its services, in a VCS One cluster. In an Active-Active configuration, several SQL server instances are intended to run on a single system when necessary.

The agents monitor the configured resources, determine the status of these resources, bring them online, and take them offline. It detects an application failure if the configured SQL Server instance or services become unavailable. Further, depending on the type of service group configuration selected, it either failover the configured service group to the next available system or brings the service group offline.

For more details refer to,

[“About configuring the failover SQL Server service group”](#) on page 64

[“About configuring the start/stop SQL Server service group”](#) on page 64

The VCS One agents for SQL Server services includes:

Agent for SQL Server Database Engine service	Provides high availability for SQL Server Database Engine service. The agent monitors the status of the service and if the service is not running in the VCS One context, the agent declares the service resource as OFFLINE.
--	---

For more details refer to,

[“About the agent for SQL Server Database Engine”](#) on page 13

You can install this agent using the agent pack installer.

Agent for SQL Server 2008 FILESTREAM	The agent provides high availability for SQL Server 2008 FILESTREAM feature. The agent monitors the Windows FILESTREAM configuration settings for the SQL Server instance.
--------------------------------------	--

For more details refer to,

[“About the agent for FileStream”](#) on page 21

GenericService agent for SQL Server Agent service, Analysis service, and Full-Text Indexing and Search service	<p>Provides high availability for the following SQL Server services:</p> <ul style="list-style-type: none"> ■ SQL Server Agent service ■ Analysis service ■ Full-Text Indexing and Search service (in case of SQL Server 2005 only)
--	--

The GenericService agent monitors these services. If the services are not running in the VCS One context, the agent declares the service resources as OFFLINE.

For more details refer to,

[“About the GenericService agent for SQL Server services”](#) on page 17

This agent is installed by default along with the VCS One client.

About the agent for SQL Server Database Engine

The VCS One agent for SQL Server Database Engine service monitors the database engine service and if the service is not running, the agent declares the state as OFFLINE. Depending on the detail monitoring configuration, the agent checks the health of critical SQL databases or executes a monitoring script. If detail monitoring is configured and is successful, the agent declares the service group as online.

Specific agent functions include the following:

Online	<p>Brings the SQL Server Database Engine service online.</p> <p>For SQL Server 2008, in addition to bringing the SQL Server DB Engine online, this function also validates if the SQL Service account is a member of the configured domain group account. If the validation succeeds then it adds the SQL Service account to the local SQL instance domain group account.</p>
Offline	Takes the SQL Server Database Engine service offline.
Monitor	Queries the status of SQL Server Database Engine service. Also, if detail monitoring is configured, the agent performs a database health check or executes a monitoring script depending on the configuration.
Clean	Forcibly stops the SQL Server service.

Attribute definitions for SQL Server Database Engine agent

Review the following information to familiarize yourself with the SQL Server database engine agent attributes. The agent is configured as a SQLServer resource type.

Table 1-1 describes the required attributes associated with the agent.

Table 1-1 SQL Server Database Engine agent required attributes

Required Attributes	Definition
Instance	Name of the SQL Server instance to monitor. If the attribute is blank the agent monitors the default instance (MSSQLSERVER). Type and dimension: string-scalar
LanmanResName	The Lanman resource name on which the SQL Server resource depends. It defines the virtual server name for the SQL Server. If the attribute value is blank the agent starts the service in physical server name context. This is in optional attribute for a standalone service group. Type and dimension: string-scalar
SQLOnlineTimeout	Number of seconds that can elapse before the agent online function aborts. Default = 90 Type and dimension: integer-scalar
SQLOfflineTimeout	Number of seconds that can elapse before the agent offline function aborts. Default = 90 Type and dimension: integer-scalar

[Table 1-2](#) describes the optional attributes associated with the agent.

Table 1-2 SQL Server Database Engine agent optional attributes

Required Attributes	Definition
DetailMonitorInterval	<p>Defines whether the agent performs detail monitoring of SQL Server. The default value 0 means that the agent will not monitor SQL Server in detail. A non-zero value indicates the number of online monitor cycles that the agent must wait before performing detail monitoring.</p> <p>Note: If this attribute is set to a non-zero value, then the following attributes must be assigned appropriate values:</p> <p>For script-based monitoring: SQLFile, Username, Password, Domain, and SQLDetailMonitorTimeOut</p> <p>For databases list based monitoring: DBList</p> <p>Type and dimension: integer-scalar</p>
FaultOnDMFailure	<p>Defines whether the agent fails over the service group if the detail monitoring fails.</p> <p>If the value is set to True, the agent fails over the service group if detail monitoring fails. In case of a standalone service group configuration however, the agent brings the service group offline.</p> <p>If set to False, the agent does not fail over the service group if detail monitoring fails. In such a case, the SQL Server resource goes in to the UNKNOWN state.</p> <p>Note: The value set gets operational only after the SQL Server resource under consideration comes online.</p> <p>Default = True</p> <p>Type and dimension: boolean</p>
SQLDetailMonitor Timeout	<p>Number of seconds that can elapse before script based detail monitor routine aborts.</p> <p>Default = 30</p> <p>Type and dimension: integer-scalar</p>

Table 1-2 SQL Server Database Engine agent optional attributes (continued)

Required Attributes	Definition
SQLClusterAccount	<p>Use this attribute if the SQL service account is not a member of the local Administrators group on the systems.</p> <p>Specify a domain group or the SQL Server service name. The agent assigns the account with Full Control privileges to the SQL Server databases and log files. If you specify a domain group, then the SQL service account must be part of this domain group.</p> <p>For a domain group specify in the format Domain.com\DomainGroup.</p> <p>For SQL Server service name specify in the format MSSQL\$InstanceName. For the default instance, the service name is MSSQLServer.</p>
Username	<p>The user account in whose context the script-based detail monitoring is performed. The agent uses this account context to execute the monitoring script.</p> <p>This attribute must not be null if DetailMonitorInterval attribute is set to a non-zero value and script-based detail monitoring is configured.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
Domain	<p>Domain of the user account.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
Password	<p>Password for the user account to launch the SQL script for Detail Monitoring.</p> <p>Type and dimension: string-scalar</p>

Table 1-2 SQL Server Database Engine agent optional attributes (continued)

Required Attributes	Definition
SQLFile	<p>The location of the SQL file executed during a monitor cycle. This attribute must not be null if script-based detail monitoring is configured and the DetailMonitorInterval attribute is set to a non-zero value.</p> <p>This attribute is related only to script-based Detail Monitoring. This attribute can be null if DBList based detail monitoring is configured.</p> <p>Note that if both script-based (SQLFile attribute) and databases list based (DBList attribute) detail monitoring are configured, then the DBList attribute takes precedence. The agent performs databases list based (DBList attribute) detail monitoring.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
DBList	<p>List of databases for which the agent will perform detail monitoring.</p> <p>Note that if both script-based (SQLFile attribute) and databases list based (DBList attribute) detail monitoring are configured, then the DBList attribute takes precedence. The agent performs databases list based (DBList attribute) detail monitoring.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-vector</p>

About the GenericService agent for SQL Server services

VCS One uses the GenericService agent to make the following SQL Server services highly available:

- SQL Server Agent service
- SQL Server Analysis service
- SQL Server Full-Text Indexing and Search service (in case of SQL Server 2005 only)

The GenericService agent brings these services online, monitors their status, and takes them offline. Note that a service is an application type supported by Windows that conforms to the interface rules of the Service Control Manager (SCM). The SQL Server services are defined as resources of type GenericService.

You can configure the GenericService agent to monitor multiple services by defining a resource for each service to be monitored. You can start and monitor services in the virtual computer context by enabling the UseVirtualName attribute.

Note: In SQL 2008, if the services are configured in the virtual computer context, only one analysis service can be configured per lanman resource.

Specific agent functions include the following:

Online	Brings the configured SQL Server service online.
Offline	Takes the configured SQL Server service offline.
Monitor	Queries the Service Control Manager (SCM) for the status of the configured SQL Server service.
Clean	Forcibly stops the configured SQL Server service.

Attribute definitions for GenericService agent

Review the following information to familiarize yourself with the GenericService agent attributes. The agent is configured as a GenericService resource type.

[Table 1-3](#) describes the required attributes associated with the VCS One GenericService agent.

Table 1-3 GenericService agent required attributes

Required Attributes	Definition
ServiceName	<p>Name of the service to be monitored. The service name must be the Service Key Name.</p> <p>The names for SQL Server services are configured as follows:</p> <p>Agent service: <i>SQLSERVERAGENT</i>, for default instance <i>SQLAgent\$<INSTANCENAME></i>, for named instance</p> <p>Analysis service: <i>MSOLAP\$<SQLinstancename></i>, for named instance and <i>MSSQLServerOLAPService</i> for default instance.</p> <p>Full-text Search service: ??</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>

[Table 1-4](#) describes the optional attributes associated with the VCS One GenericService agent.

Table 1-4 GenericService agent optional attributes

Optional Attributes	Definition
DelayAfterOffline	<p>Number of seconds the offline routine waits for the service to go offline.</p> <p>Default = 10</p> <p>Type and dimension: integer-scalar</p>
DelayAfterOnline	<p>Number of seconds the online routine waits for the service to come online.</p> <p>Default = 10</p> <p>Type and dimension: integer-scalar</p>

Table 1-4 GenericService agent optional attributes

Optional Attributes	Definition
UserAccount	<p>A valid user account in whose context the service will be monitored. Username can be of the form username@domain.com or domain.com\username.</p> <p>If you do not specify a value for this attribute, then the user account of the service in the SCM is ignored. To monitor service under built-in accounts, you must provide explicit values.</p> <p>For example:</p> <p>On Windows 2003: UserAccount='LocalSystem', 'Local Service', or 'Network Service'. Domain='NT Authority'.</p> <p>The 'NT Authority' domain is not applicable for the 'LocalSystem' account.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
Password	<p>The password of the user in whose context the service is started. If the UserAccount attribute is empty or contains a built-in service account, this attribute is ignored.</p> <p>Type and dimension: string-scalar</p>
Domain	<p>The domain name to which the user specified in the UserAccount attribute belongs.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
service_arg	<p>An array of arguments passed to the service.</p> <p>Type and dimension: string-vector</p>
UseVirtualName	<p>Set this value to enable the service specified by the ServiceName attribute to inherit the virtual server context of Lanman resource specified in LanmanResName attribute. If this attribute is set, the LanmanResName becomes a mandatory attribute.</p> <p>Note: It is recommended that you do not set this attribute for system services. If you configure system services using the GenericService agent, then this attribute is not supported.</p> <p>Type and dimension: boolean-scalar</p>

Table 1-4 GenericService agent optional attributes

Optional Attributes	Definition
LanmanResName	<p>The name of the Lanman resource on which the GenericService resource depends.</p> <p>Note: It is recommended that you do not set this attribute for system services. If you configure system services using the GenericService agent, then this attribute is not supported.</p> <p>Type and dimension: string-scalar</p>

About the agent for FileStream

The VCS One agent for SQL Server 2008 FILESTREAM enables FILESTREAM for the specified SQL Server 2008 instance, monitors its status, and disables it. It also makes FILESTREAM highly available in a clustered environment.

FILESTREAM in SQL Server 2008 enables SQL Server-based applications to store unstructured data, such as documents and images, on the file system. FILESTREAM integrates the SQL Server Database Engine with an NTFS file system by storing varbinary(max) binary large object (BLOB) data as files on the file system. Transact-SQL statements can insert, update, query, search, and back up FILESTREAM data. Win32 file system interfaces provide streaming access to the data.

Specific agent functions include the following:

Online	Enables FILESTREAM on the node on which the service group comes online.
Offline	Disables FILESTREAM on the node on which the service group goes offline.
Monitor	Monitors FILESTREAM status on the node on which the service group is online. If the agent is unable to query the status of FILESTREAM, the FILESTREAM resource in the service group faults.
Clean	Forcibly stops the configured FileStream services.

Attribute definitions for FileStream agent

Review the following information to familiarize yourself with the agent attributes for a SQLFilestream resource type.

[Table 1-5](#) describes the required attributes associated with the VCS One agent for SQL Server 2008 FILESTREAM.

Table 1-5 SQL Server 2008 Filestream agent required attributes

Required Attributes	Definition
InstanceName	<div>The name of the SQLServer2008 instance to which the FILESTREAM is bound. If this attribute is blank, the agent monitors the default SQL server instance (MSSQLSERVER).</div> <div>Type and dimension: string-scalar</div>

About monitoring options

The VCS One agents for Microsoft SQL Server provides two levels of application monitoring: *basic* and *detail*. In the basic monitoring option the agent continuously queries the Windows Service Control Manager (SCM) to verify if the configured SQL Server services are active. In detail monitoring option the agent queries the SQL Server databases to verify the availability of database instance.

You can configure database monitoring in the following ways:

- DBList detail monitoring

The SQL Server agent monitors only the list of databases specified in the SQL Server agent's DBList attribute. The agent uses Microsoft ActiveX Data Objects (ADO) to establish a connection with the selected databases to verify the health of those databases. If the connection is successful the agent considers the database as available. If the connection fails, the database instance is considered not available and, if the FaultOnDMFailure agent attribute value is set to True, the agent fails over the service group to the failover systems.

- Script-based detail monitoring

The SQL Server agent uses a script to monitor the status of the database. 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, if the FaultOnDMFailure attribute value is set to True, the agents fails over the service group to the failover systems. A sample script is provided as a reference. You can customize the script to meet your configuration requirements.

The script is located at:

`%VCSONE_HOME%\bin\SQLServer\sample_script.sql`

Here, the variable %VCSONE_HOME% is the default installation directory for VCS One, typically it is C:\Program Files\Veritas\Cluster Server One.

You should use a separate script for each SQL Server service group that exists in the farm. The script should exist at the location that is configured on all the systems in the service group.

You can enable and configure detail monitoring by running the SQL Server Configuration Wizard for VCS One. Refer to the instructions for configuring a SQL Server service group for more information.

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

How the agents make SQL Server highly available

The VCS One agents for Microsoft SQL Server detects an application failure if a configured virtual server becomes unavailable.

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

Local cluster configuration

When the VCS database agent for Microsoft SQL Server detects an application failure, the SQL Server service group is failed over to the next available system in the service group’s system list. The configured SQL services and virtual server are started on the new system.

Disaster recovery configuration

In a disaster recovery configuration, VCS One first attempts to fail over the application to a cluster system in the local cluster. If all systems in the local cluster are unavailable, or if a disaster strikes the site, VCS One provides the capability to fail over the applications in that cluster to a system in another cluster on the remote site.

Figure 1-1 depicts the high-level workflow to set up a global cluster for application high availability.

Figure 1-1 High-level workflow to set up VCS One global clusters




Figure 1-2 depicts the tasks to configure VCS One clusters.

Figure 1-2 Tasks to configure VCS One clusters

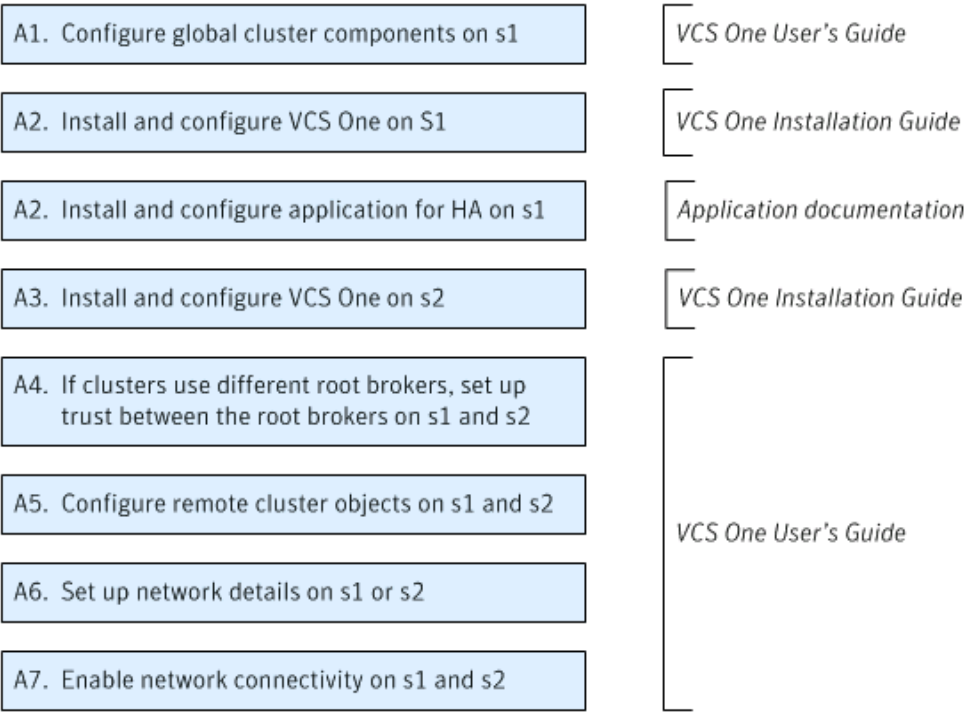


Figure 1-3 depicts the tasks to configure application on the secondary site

Figure 1-3 Tasks to configure application on secondary site

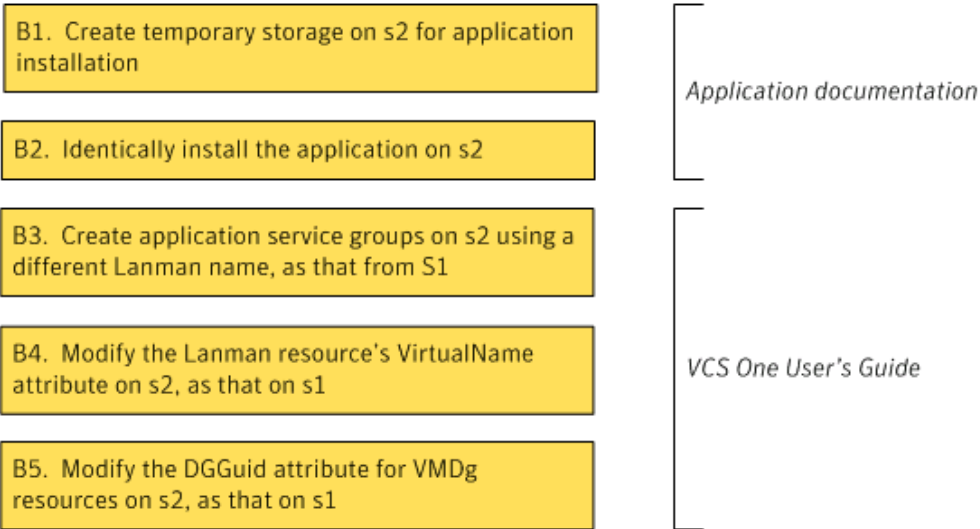


Figure 1-4 depicts the tasks to set up replication between the primary and secondary clusters

Figure 1-4 Tasks to set up replication between the sites





Figure 1-5 depicts the tasks to configure the composite service groups

Figure 1-5 Tasks to configure composite service groups



Running SQL Server in an Active-Active clustered environment

In an Active-Active SQL Server configuration, several instances are intended to run on a single system when necessary. A SQL Server instance is a completely independent SQL Server installation, with its own services, master database, storage, and memory resources. Each instance is defined uniquely by a separate SQL Server virtual server and service group.

SQL Server allows multiple independent instances of SQL Server to run on a single machine. Using this feature, the VCS One database agent for SQL Server supports SQL Server in an Active-Active environment by allowing a system to run up as many instances as supported by SQL. A SQL Server instance can fail over to any of the other systems that you specify when you configure the SQL Server service group.

You can choose an active-active SQL Server configuration where several instances are intended to run on a single system. However, remember that you must configure failover systems such that a single system can never host more instances than what is supported by SQL Server.

Refer to the Microsoft documentation for more information on multiple instance support.

SQL Server failover service group- resource dependency graph

The following SQL Server sample configuration depicts the resources and their dependencies within a SQL Server failover service group in a VCS One cluster.

A SQL Server failover service group has a shared storage configured using the MountV and VMDg resources. The virtual name for the SQL Server is created using the Lanman resource, which in turn is dependent on the IP and NIC

resource. The IP address for the service group is configured using the IP and NIC resources. The SQL Server data files and log files are stored in the volume specified in the mount resource and the required registry settings are replicated accross machine using the RegRep resource and its corresponding mount.

The SQL Server Analysis service (MSOlap) and SQL Server Agent service (SQLServerAgent) are configured as GenericService resources.

The SQL Server resource comes online after each of the underlying resources are brought online.

Figure 1-6 SQL Server 2005 failover service group- resource dependency

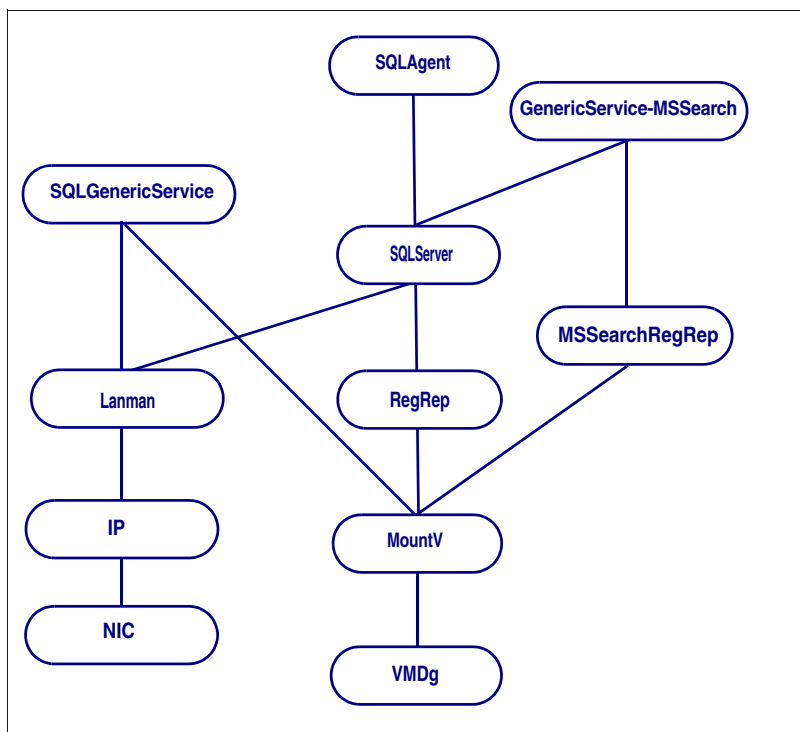
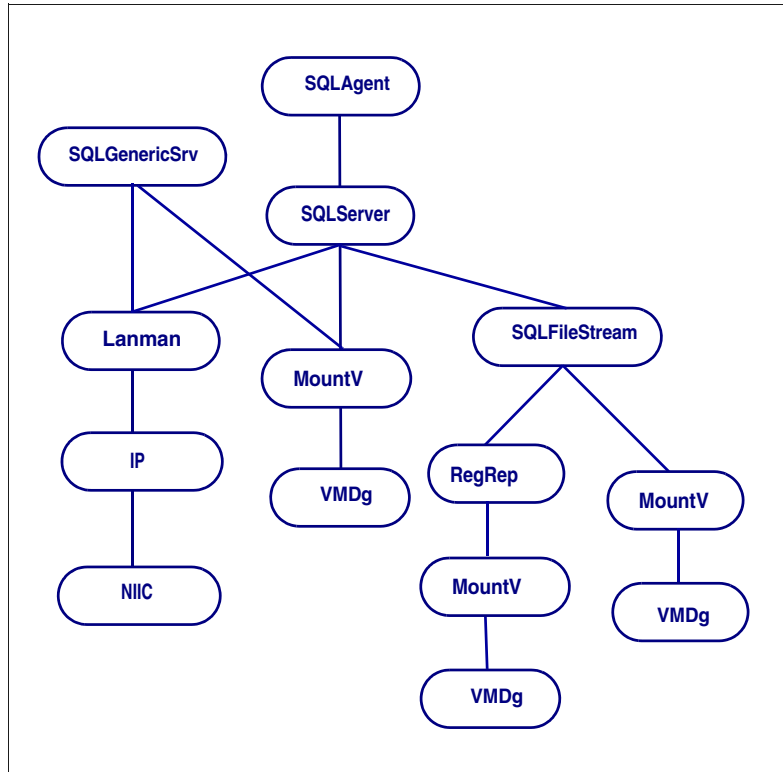


Figure 1-7 SQL Server 2008 failover service group- resource dependency



SQL Server standalone (start/stop) service group- resource dependency graph

The following SQL Server sample configuration depicts the resources and their dependencies within a SQL Server standalone service group in a VCS One cluster.

A SQL Server standalone service group has the local storage configured using the MountMonitor resource. If the Lanman resource is configured, the SQL Server Service runs under the virtual server name context. Else, it runs under the Local system name context.

The IP address for the service group is configured using the IP and NIC resources.

The SQL Server data files and log files are stored in the volume specified in the MountMonitor resource.

The SQL Server Analysis service (MSOlap) and SQL Server Agent service (SQLServerAgent) are configured as GenericService resources.

The SQL Server resource comes online after each of the underlying resources are brought online.

Figure 1-8 SQL Server 2008 start/stop service group- resource dependency

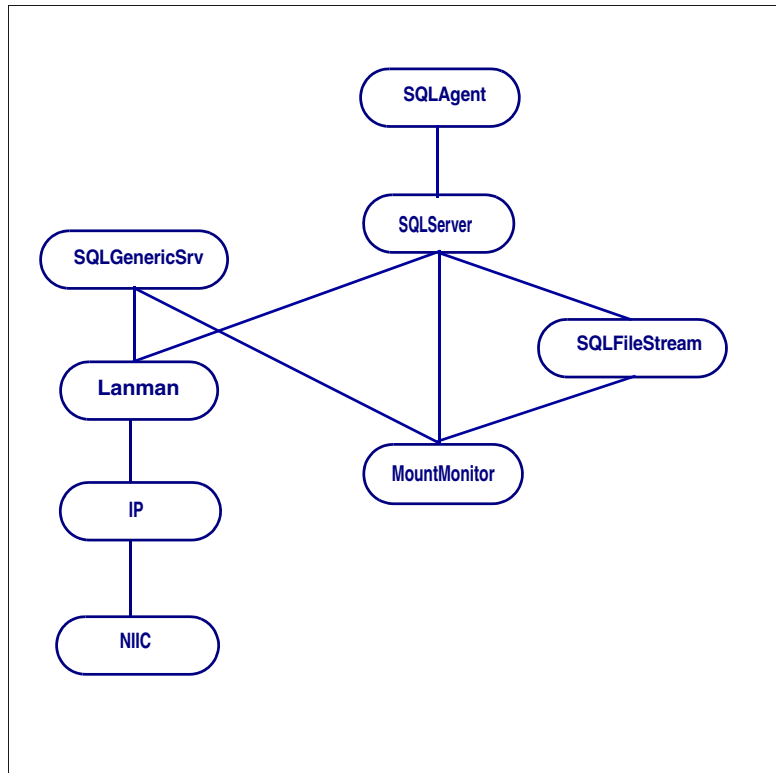
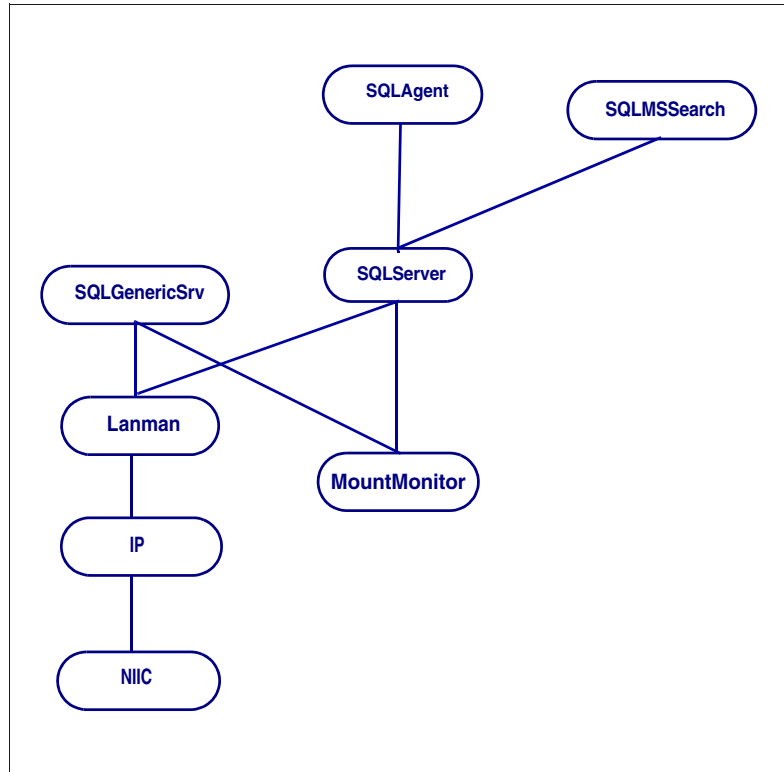


Figure 1-9 SQL Server 2005 start/stop service group- resource dependency



Typical SQL Server configuration in a VCS One cluster

VCS One supports the following types of SQL Server configuration.

- Failover
- Start/Stop

In a failover configuration, the VCS One cluster involves two systems accessing a shared storage. The SQL Server binaries are installed on both the systems. The shared storage is used to store SQL Server data files .

In a Start/Stop configuration, the VCS One cluster involves a single system. The SQL Server binaries and the data files are installed on the local system.

Note: The MSDTC resource can be configured on the same cluster systems. You need not configure an MSDTC client if the MSDTC resource is configured on the same systems that have SQL Server resource configured. However, if the MSDTC resource is configured on other systems, you must configure an MSDTC client to point to the virtual server name of the MSDTC resource.

Figure 1-10 depicts a typical SQL Server failover service group with the different services configured.

Figure 1-10 Typical failover SQL HA configuration in a VCS One cluster

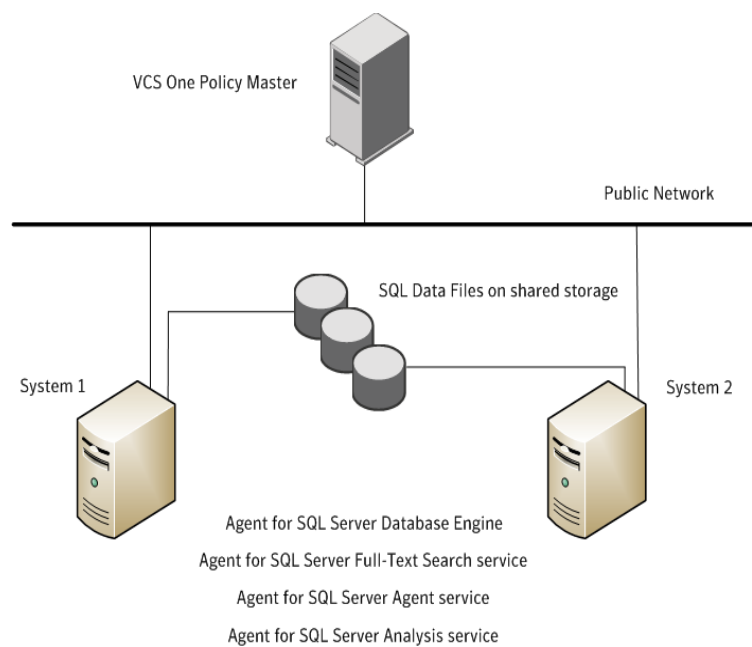
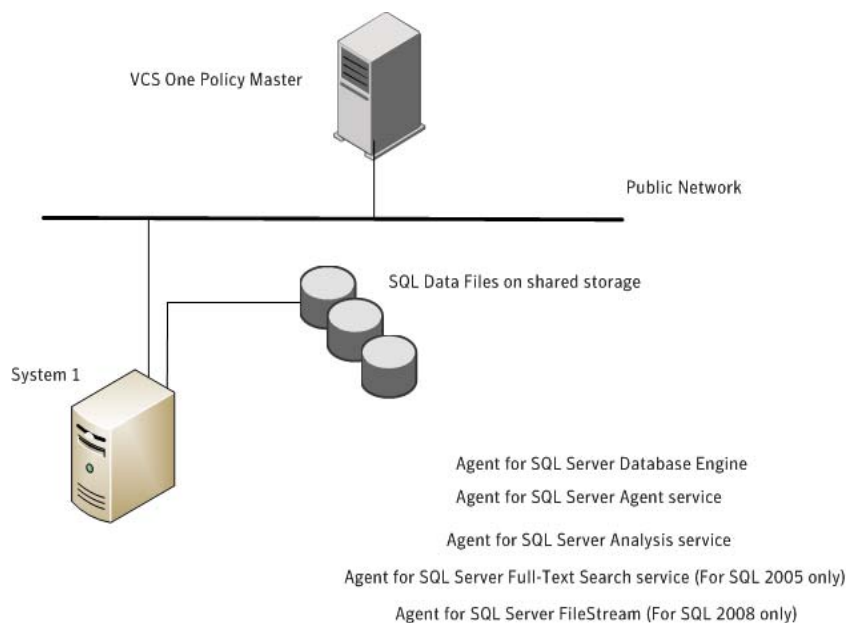


Figure 1-11 Typical start/stop SQL configuration in a VCS One cluster



Installing, modifying and removing the VCS One database agent for SQL Server

This chapter contains the following topics:

- [Before installing the agent](#)
- [Installing the agent](#)
- [Adding the agent resource type to the Policy Master Server- Using CLI](#)
- [Importing the Binary Message Catalog \(BMC\) files](#)
- [Modifying or removing the agent installation](#)
- [Uninstalling the agent pack](#)

Before installing the agent

Ensure that you have performed the following tasks before you proceed to install the VCS One agent for SQL Server.

- Verify that you have installed and configured VCS One client on all cluster systems. Refer to the *Veritas Cluster Server One Installation and Configuration Guide* for instructions.
- Verify that you have Local Administrator privileges on the system where you want to install the agent.

Installing the agent

Use the VCS One Agent Pack Installer, to install the agent.

You can launch the installer from the Policy Master Server and install the agent on the desired client systems or launch the installer from the client system and install the agent on a local as well as remote client systems.

Along with the agent installation, you can also update the agent resource types on the Policy Master Server.

Note: Updating the agent resource type is an one time activity. If you plan to update the resource type using the installer, you must run the installer from a client system.

To install the agent

- 1 Insert the agent pack software disc and run the setup.exe file.
- 2 On the CD Browser Welcome panel, under the Product Installation menu, click **VCS One Agent Pack**.
- 3 On the Agent Pack Installer Welcome panel, review the pre-requisites and click **Next**.
- 4 On the License Agreement panel, review the End User License Agreement and select **I accept the terms of License Agreement** and then click **Next**.
- 5 On the Option Selection panel, select the agents you want to install and click **Next**.
- 6 On the System Selection panel, add the systems on which you want to install the VCS One Agent Pack. You can perform this in one of the following ways:
 - In the System Name text box, manually type the system name and click **Add**.
 - Alternatively, browse to select the systems.
On the Select Systems panel, the systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**.

Once you add or select a system, the wizard performs the verification checks and notes the verification details. To review the details, click the corresponding information icon.

By default the wizard uses %ProgramFiles%\Veritas as the installation directory. However, if you have customized your installation directory

during the VCS One client installation, the agent pack uses the customized directory as the installation directory.

- 7
- On the System Selection panel, click **Next**.
Note that the installer fails to proceed with the installation, unless all the selected systems have passed the verification checks and are ready for installation. In case the verification checks have failed on any of the system, review the details and rectify the issue. Before you choose to proceed with the installation click **Re-verify** to re-initiate the verification checks for this system.
- 8
- On the Pre-install Summary panel, review the summary and click **Next**.
- 9
- On the Installation panel, review the progress of installation and click **Next** when the installation is complete.
- 10
- On the Post-install Summary panel, review the installation result and click **Next**.
If the installation has failed on any of the system, refer to the log file for details.
- 11
- On the Finish panel, click **Finish**, if you plan to only install the agent. You can now configure the agent on the systems where the installation was successful.
However, if you also plan to update the agent resource types to the Policy Master server on Windows, click **Next**.

Note: If you are updating the resource types to the Policy Master Server on UNIX, See [“Adding the agent resource type to the Policy Master Server-Using CLI”](#) on page 38.

For updating the agent resource type, specify the following details on the Policy Master details page, and then click **Next**.

Policy Master: IP Address: Port:	Type the virtual IP address of the Policy Master (PM). Do not specify the physical IP address of the PM server. Type the port number used by the PM. The default port is 14151.
Use Logged-on user credentials	This check box is selected by default. Continue with the selection if you want the wizard to use the logged on user account context to perform the client configuration tasks. The credentials values are populated in the respective fields below and the fields are disabled by default. If you do not want to use the logged on user account context, clear this check box and enter the relevant credentials in the fields below.

User Name	Type a user name. Note: This user account must be configured on the Policy Master and must have read access to the VCS One OU and the privileges to add and modify service groups. The user account must also have the privileges to attach service groups to the desired OUValue. If you choose to use the logged-on user, ensure that the logged-on user account has these privileges.
Password	Type the password for the specified user.
Domain Type	Select the appropriate domain type from the drop-down list. For a Windows PM, select nt.
Domain Name	Type the qualified domain name of the Policy Master (PM) server. For example, if your PM is in the domain <i>mydomain.com</i> , type mydomain .

- 12 On the Policy Master Type Update panel, review the progress and click **Finish** after the process is complete.

Adding the agent resource type to the Policy Master Server- Using CLI

During the agent installation, the Agent Pack Installer enables you to add the agent resource type to the Policy Master Server. However, if you had not added the types during the agent installation but now want to add them, you can perform the task using the CLI.

To add the agent resource type to the Policy Master on UNIX, perform the following steps from the Policy Master Server

- 1 Create a temporary folder on the Policy Master Server.
mkdir addtypes_tmp
- 2 Change your working directory to the temporary directory created.
chdir addtypes_tmp
- 3 Mount the VCS One software disc and navigate to the following directory:
DVD\Installer\PMTypeUpdate
- 4 Copy the resource type XMLs for the required agent type.
- 5 Convert the type XML file to type cmd file.
/opt/VRTSvcstone/bin/haconf -xmltocmd
resource-type-XML target-directory

- Run the newly generated .cmd file to add the resource type to the Policy Master database.

Importing the Binary Message Catalog (BMC) files

After you have successfully installed the agent and updated the agent types on the Policy Master Server, you must also import the BMC files to the Policy Master Server. This enables you to view the attribute descriptions when you manage the application service groups or its resources through VCS One Web Console.

The BMC files are copied to the client machine during the agent pack installation.

To import the BMC files to the Policy Master Server

- On the client system navigate to the following path.

Client system on Windows	%VCSONE_HOME%\messages\en
--------------------------	---------------------------

Client system on UNIX	/opt/VRTS/messages/en
-----------------------	-----------------------

- Copy the BMC files to the following on the Policy Master.

Policy Master Server on Windows	%VCSONE_HOME%\messages\en
---------------------------------	---------------------------

Policy Master Server on UNIX	/opt/VRTS/messages/en
------------------------------	-----------------------

- On the Policy Master Server, update the VCSOne.bmcmap file using the following command:

For PM on UNIX

```
bmcmap -create en VCSOne -path /opt/VRTS/messages/en
```

For PM on Windows

```
bmcmap -create en VCSOne -path
VCSONE_HOME%\messages\en
```

- Using the VCS One Web Console, refresh the Message Catalog files (**Administration > Global Settings > Refresh Message Catalog from Configuration**)

Modifying or removing the agent installation

To remove the agents installed or to install additional VCS One agents, for high availability of other applications, you must run the VCS One Agent Pack Installer in the modify mode.

Note that you can run the VCS One Agent Pack Installer in the modify mode, only on the local client system.

To modify the agent installation

- 1 Go to **Start > Settings > Control Panel > Add or Remove Programs**.
In case you are working on Windows Server 2008, go to Programs and Features. (**Start > Settings > Control Panel > Programs and Features**)
- 2 Scroll to VCS One Agent Pack and select it.
- 3 Click **Change**.
- 4 On the Mode Selection panel, select **Add or Remove**, to perform either of the following tasks:
 - Install another VCS One agent available in the agentpack
 - Remove the installed VCS One agent
- 5 On the Option Selection panel, the agents that are already installed are selected by default. Depending on the task you want to perform,
 - Clear the selection for the agents you want to remove
 - Select the agents you want to add
- 6 On the System Selection panel, click **Next**.
- 7 On the Pre-Install panel, review the pre-update report and click **Next**.
- 8 On the Installation panel, review the progress and click **Next** after the process is complete.
- 9 On the Post-Install panel, review the Post-update report and click **Next**.
- 10 On the Finish panel, click **Finish** if you have removed any agent. However, if you have added any agent click **Next** to update the agent resource type on the Policy Master Server.
- 11 On the Policy Master details page, specify the following details and then click **Next**.

Policy Master: IP	Type the virtual IP address of the Policy Master (PM).
Address: Port:	Do not specify the physical IP address of the PM server.
	Type the port number used by the PM. The default port is 14151.

Use Logged-on user credentials	<p>This check box is selected by default.</p> <p>Continue with the selection if you want the wizard to use the logged on user account context to perform the client configuration tasks.</p> <p>The credentials values are populated in the respective fields below and the fields are disabled by default.</p> <p>If you do not want to use the logged on user account context, clear this check box and enter the relevant credentials in the fields below.</p>
User Name	<p>Type a user name.</p> <p>Note: This user account must be configured on the Policy Master and must have read access to the VCS One OU and the privileges to add and modify service groups. The user account must also have the privileges to attach service groups to the desired OUValue. If you choose to use the logged-on user, ensure that the logged-on user account has these privileges.</p>
Password	Type the password for the specified user.
Domain Type	Select the appropriate domain type from the drop-down list. For a Windows PM, select nt.
Domain Name	<p>Type the qualified domain name of the Policy Master (PM) server.</p> <p>For example, if your PM is in the domain <i>mydomain.com</i>, type mydomain.</p>

- On the Policy Master Type Update panel, review the progress and click **Finish** after the process is complete.

Uninstalling the agent pack

You can uninstall the VCS One Agent Pack, using the VCS One Agent Pack Installer.

You can uninstall the agents from the local as well as multiple remote client systems.

To uninstall the agent

- 1 Go to **Start > Settings > Control Panel > Add or Remove Programs**.
In case you are working on Windows Server 2008, go to Programs and Features. (**Start > Settings > Control Panel > Programs and Features**)
- 2 Scroll to VCS One Agent Pack and select it.
- 3 Click **Remove**.
In case of Windows Server 2008, click **Uninstall**.
- 4 On the VCS One Agent Pack Installer welcome page, review the list of prerequisites and click **Next**.
- 5 On the System Selection panel, add the systems from which you want to uninstall the VCS One client. You can perform this in one of the following ways:
 - In the System Name text box, manually type the system name and click **Add**.
 - Alternatively, browse to select the systems.
The systems that belong to the domain in which you have logged in are listed in the Available Systems list. Select one or more systems and click the right arrow to move them to the Selected Systems list. Click **OK**.
Once you add or select a system, the wizard performs the verification checks and notes the verification details. To review the details, click the corresponding information icon.

Note: By default the local system is selected for un-installation. In case you are performing a remote un-installation and do not want to uninstall the AgentPack from the local system, click the corresponding <remove> icon to remove the system from the list.

- 6 Click **Next**.
Note that the wizard fails to proceed with the un-installation, unless all the selected systems have passed the validation checks and are ready for un-installation. In case the validation checks have failed on any of the system, review the details and rectify the issue. Before you choose to

proceed with the un-installation click **Re-verify** to re-initiate the verification checks for this system.

- 7 On the Pre-uninstall Summary panel, review the summary and click **Next**.
- 8 On the VCS One Agent Pack Un-installation panel, review the uninstallation progress and click **Next** when the uninstallation is complete.
- 9 On the Post-uninstall Summary panel, review the uninstallation results and click **Next**.
If the un-installation has failed on any of the system, review its summary report and check the log file for details.
- 10 On the Finish panel, click **Finish**.

Installing and configuring SQL Server

This chapter contains the following topics:

- [“About installing and configuring SQL Server”](#) on page 46
- [“Supported software”](#) on page 47
- [“Before installing SQL Server”](#) on page 48
- [“Installing SQL Server”](#) on page 57

About installing and configuring SQL Server

This chapter provides information for installing and configuring SQL Server and its services in a VCS One cluster...

When you complete the installation, each system that participates in a SQL Server service group has 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.

Also refer to the points to be followed while installing SQL Server on the cluster system

See [“Before installing SQL Server”](#) on page 48.

See [“Installing SQL Server”](#) on page 57.

A basic configuration of SQL Server has one or more instances of SQL Server in the Veritas Cluster Server One (VCS One) cluster. These instances are managed by a service group that is configured with a distinct set of systems in the VCS One cluster. .

For details refer to,

[“About configuring the failover SQL Server service group”](#) on page 64

[“About configuring the start/stop SQL Server service group”](#) on page 64

Supported software

This release of VCS One agent for SQL Server supports the following software versions.

Note that all systems must be running the same operating system.

Table 3-1 SQL Server and Windows Server support

SQL Server versions	Windows Server versions
<ul style="list-style-type: none">■ Microsoft SQL Server 2005, 32-bit Standard Edition or Enterprise Edition■ Microsoft SQL Server 2005, 64-bit Standard Edition or Enterprise Edition	<ul style="list-style-type: none">■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)■ Windows Server 2008 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2008 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)
<ul style="list-style-type: none">■ Microsoft SQL Server 2008, 32-bit Standard Edition or Enterprise Edition.■ Microsoft SQL Server 2008, 64-bit Standard Edition or Enterprise Edition.	<ul style="list-style-type: none">■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)■ Windows Server 2003 Standard x64 Edition, Enterprise x64 Edition, or Datacenter x64 Edition (SP2 required)■ Windows Server 2003 x64 Editions (for AMD64 or Intel EM64T): Standard x64 R2 Edition, Enterprise x64 R2 Edition, or Datacenter x64 R2 Edition (SP2 required)

Only one default instance can exist on a system at one time. Additional instances that operate on or fail over to that system must be named instances. The number of named instances is limited only by the normal limit of instances for that version of SQL Server.

Before installing SQL Server

Ensure that you meet the following prerequisites before you install the SQL Server.

- Verify that VCS One client is installed and configured on all systems in the VCS One cluster. The VCS One Client installer also installs the VCS One generic service agents for SQL Server.
Refer to the *Veritas Cluster Server One Installation Guide* for instructions.
- Verify that the VCS One agent for SQLServer is installed on all the systems.
Refer to “[Before installing the agent](#)” on page 35
- Create the required volumes on the shared disks and ensure that they are accessible from the first cluster systems.
See “[About managing storage](#)” on page 48.
- When configuring the VCS One Client Remoting service ensure that the user account for the service is a domain user with administrative privileges on the cluster systems.
- Ensure that the logged-on user is a domain user with local Administrator privileges.
- Ensure that the logged-on user is a member of the local Administrators group on all nodes on which you want to install Microsoft SQL Server.
- Ensure that the logged-on user or the VCS One Remoting Service user account have write permissions for the Active Directory objects corresponding to the cluster systems.
- The logged-on user or the VCS One Remoting Service user account have write permissions on the DNS server to perform DNS updates.

About managing storage

Symantec recommends that you have a minimum of two volumes on shared disks, one each for the following:

- SQL Server data
- Registry replication information

If required, configure additional volumes for the following:

- User defined databases
- User defined databases and their log files

For example, you can create the following volumes to manage your SQL Server database and logs:

- **INST1_DATA_FILES:** contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- **INST1_REGREP_VOL:** contains the list of registry keys that must be replicated among server farm systems for the SQL Service. Create a 100 MB (minimum recommended size) volume for this purpose.
- **INST1_DB1_VOL:** contains the user database files
- **INST1_DB1_LOG:** contains the user database log files

Refer to the appropriate instructions depending on your storage environment. See “[About managing storage using Windows Logical Disk Management \(LDM\)](#)” on page 49.

See “[About managing storage using SFW](#)” on page 52.

About managing storage using Windows Logical Disk Management (LDM)

Note the following limitations in case you are using Windows LDM as your storage manager.

- For using LDM, your storage devices must be configured to use SCSI-3 disk reservations. SCSI-2 is not supported.
- LDM support is not applicable for Disaster Recovery (DR) configurations. Currently only High Availability (HA) configurations are supported.

Perform the following tasks to create volumes and make them accessible from all the systems in the VCS One cluster:

- Unassign a drive letter
See “[Unassigning a drive letter](#)” on page 49.
- Reserve disks.
See “[Reserving disks \(if you use Windows LDM\)](#)” on page 50.
- Create volumes.
See “[Creating volumes \(if you use Windows LDM\)](#)” on page 51.
- Mount volumes.
See “[Mounting volumes \(if you use Windows LDM\)](#)” on page 51.

Unassigning a drive letter

While installing an application on multiple systems, you must first unassign drive letters and release the disks from one node, and then reserve the disks, and mount the volumes using the same drive letters and then install the application on the failover node.

Note: You must run Disk Management on all systems each time you add a shared disk. This ensures each disk has a valid signature written to it and that the device paths and symbolic links are updated.

Complete these steps to unassign the drive letters from a node.

To unassign drive letter

- 1 Log in as Administrator.
- 2 Open Disk Management. Type the following on the command prompt:
`C:\> diskmgmt.msc`
- 3 Right-click the partition or logical drive and click **Change Drive Letter and Path**.
- 4 In the **Change Drive Letter and Paths** dialog box, click the drive letter and click **Remove**.

Reserving disks (if you use Windows LDM)

Complete the following steps to reserve the disks on the system on which you plan to install SQL Server.

To reserve the disks

- 1 To display all the disks, type the following on the command line:
`C:\>haval -scsitest /l`
You will see a table that lists all the disks that are visible from the current system. Make a note of the disk numbers (Disk# column in the table). You will need it in the next step.
- 2 To reserve a disk, type the following on the command line:
`C:\>haval -scsitest /RES:<disk #>`
For example, to reserve disk 4, type:
`C:\>haval -scsitest /RES:4`
Make a note of the disk number and the corresponding signature. You will require these details to identify and reserve the disks while installing SQL Server, and configuring the SQL Server service group, on additional nodes in the server farm.

Creating volumes (if you use Windows LDM)

Perform the following steps to create volumes.

To create volumes

- 1 Use the Windows Disk Management tool to verify that the disks are visible on the nodes and then create volumes on the reserved disks.
- 2 After creating the required volumes on a node, release the reserved disks on that node.
See [“Releasing disks”](#) on page 51.
- 3 Rescan the disks on all the remaining nodes in the server farm.
Refer to Microsoft Windows documentation for more information about the Disk Management tool.

Mounting volumes (if you use Windows LDM)

Perform the following steps to mount volumes on a server farm node.

To mount a volume

- 1 Use the Windows Disk Management tool to mount the volumes.
- 2 After mounting the volumes, run the CHKDSK command and verify that there are no errors on the mounted volumes.
- 3 Make a note of the drive letters that you assign to the mounted volumes.
Use the same drive letters while mounting these volumes on the remaining nodes.
Refer to Microsoft Windows documentation for more information about the CHKDSK command and the Disk Management tool.

Releasing disks

If you are using Windows LDM as your storage, you must release the reserved disks.

To release a reserved disk, type the following on the command line:

```
C:\>havol -scsitest /REL:<disk #>
```

For example, to release disk 4, type:

```
C:\>havol -scsitest /REL:4
```

Make a note of the disk number and the corresponding signature. You will require these details to identify and reserve the disks while installing SQL Server, and configuring the SQL Server service group, on additional systems in the VCS One cluster.

About managing storage using SFW

Veritas Storage Foundation for Windows (SFW) manages data by creating clustered disk groups from physical disks. These disk groups are further divided into volumes that can be mounted on the VCS One cluster systems.

Perform the following tasks to create volumes and make them accessible from all the VCS One cluster systems:

- Create disk groups.
See “[Creating disk groups](#)” on page 52.
- Import disk groups.
See “[Importing disk groups](#)” on page 53.
- Create volumes on disk groups.
See “[Creating volumes](#)” on page 53.
- Mount volumes.
See “[Mounting volumes](#)” on page 55.

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** 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 a name for the disk group (for example, INST1_DG).

- Click the checkbox for **Create cluster group**.
- 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.

- Click **Next**.
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 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.

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 SQL Server

Installing SQL Server in a VCS One cluster involves installing SQL Server on the first system within the VCS One cluster and then installing on each of the failover system in the cluster..

If you wish to have a failover service group configuration, during the installation ensure that you install the SQL Server data files and the Analysis Service data files on the shared storage managed by the disk group. Additionally, verify that the shared drives created to store SQL data are mounted on the system where you would run the wizard to configure the service groups.

For a start/stop service group configuration, you can continue to install the data files on the local disk.

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

If the failover node is already a member of the system list of a functional SQL Server service group in the VCS One cluster, verify if SQL Server is installed on a volume with the same drive letter as the first system and then proceed to configure the SQL Server service group. If the failover system is *not* a member of the system list of an existing functional SQL Server service group in the VCS One cluster, proceed to the steps below.

Only the portions of the SQL Server installation procedure relevant to the VCS One environment are documented here. Refer to the Microsoft SQL Server documentation for detailed installation information.

Installing SQL Server on the first cluster system

Run the Microsoft SQL Server installer to install SQL Server on the first cluster node and choose to install SQL Server in the stand-alone installation mode.

While installing SQL, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure for high availability.

Refer to the Microsoft documentation for instructions.

Additionally, note the following points while installing SQL Server on the first cluster system:

- Ensure that the required volumes are mounted or connected to the system. See [“About managing storage”](#) on page 48.
- If you wish to have a failover service group configuration, during the installation ensure that you change the datapath for the selected

components (SQL Server, Analysis server etc) from the default local drive to the shared drive path.

For a standalone service group configuration, you can continue to install the data files on the local disk.

See “[Changing the default datapath to the shared drive in case of SQL 2005](#)” on page 59.

See “[Changing the default datapath to the shared drive in case of SQL 2008](#)” on page 60.

- Make a note of the SQL instance name and instance ID. You must use the same instance name and instance ID when you install the SQL Server instance on additional failover nodes.
- If you are installing multiple instances of SQL in the cluster, provide a unique instance name and instance ID for each instance.
- While specifying a user name for the SQL Server services account, specify a domain user account.

Note: 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. Install any SQL Service Packs or hotfixes, if required.

See Technote <http://support.veritas.com/docs/281828>

- Install the SQL Service Packs or hotfixes, if required.
- Set SQL services to manual start, using SCM.

Installing SQL Server on the failover cluster systems

Run the Microsoft SQL Server installer to install SQL Server on the first cluster node and choose to install SQL Server in the stand-alone installation mode. While installing SQL, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure for high availability.

Refer to the Microsoft documentation for instructions.

Additionally, note the following points while installing SQL Server on each failover cluster system:

- Install the SQL Server data files and analysis services files to a local disk. You do not have to install these files to the shared storage. However, if you choose to install these data files to a shared storage, make sure that the SQL server service is stopped on the first system and ensure

that you have renamed the previous SQL Server data directory on the shared disk. Provide the shared data directory path identical to the one provided on the first system.

- The instance name and instance ID is same as the one provided on the first cluster system.
- If you are installing multiple instances of SQL in the cluster, provide a unique instance name and instance ID for each instance.
- While specifying a user name for the SQL Server services account, specify a domain user account.
- Install any SQL Service Packs or hotfixes, if required.
- Set SQL services to manual start, using SCM.

Once the installation is complete on all the systems, run the SQL Server Service Group Configuration Wizard from the first system to create the service groups. While creating the service groups, the wizard synchronizes the data path on all the systems.

Changing the default datapath to the shared drive in case of SQL 2005

To change the datapath from the default local drive location to the shared drive location, perform the following steps during the SQL Server installation.

- On the Components to Install panel:
 - Select the **SQL Server Database Services**.
 - Select the optional components:
 - Analysis Service
 - Notification Services
 - Integration Services
 - Workstation Components

Note: VCS One provide high availability only to SQL Server Database Services and the Analysis Service.

- Click **Advanced**.
- 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 to the shared drive location.

This must be identical on all systems, if you choose to install SQL Server following the ABC method.

- 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. Also, make a note of the volume (drive letter). You will need this while configuring the service group.

- If you had selected the **Analysis Services** option, expand **Analysis Services**, select **Data Files**, and click **Browse** to specify the shared drive location.
- Click **Next**.

Changing the default datapath to the shared drive in case of SQL 2008

To change the datapath from the default local drive location to the shared drive location, perform the following steps during the SQL Server installation.

- On the Database Engine Configuration panel:
 - Select the **Data Directories** tab.
 - Browse to change the Data root directory path to the shared drive.
This must be identical on all systems, if you choose to install SQL Server following the ABC method.
 - Click **Next**.
- If you wish to configure the Analysis Services, on the Analysis Services Configuration dialog box:
 - Select the **Data Directories** tab.
 - Browse to change the data directory path to the shared drive.
 - Click **Next**.

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 VCS One cluster system.
- 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 service group

This chapter contains the following topics:

- [About configuring the failover SQL Server service group](#)
- [About configuring the start/stop SQL Server service group](#)
- [Before you configure a SQL Server service group](#)
- [Creating a SQL Server service group using the wizard](#)

About configuring the failover SQL Server service group

Veritas Cluster Server One provides high availability for all the SQL Server services that are configured in the VCS One cluster. If the active system fails, the VCS One SQL Server agent brings the configured SQL Server instance(s) online on another system.

Configuring the SQL Server failover service group involves creating resources for shared storage, SQL Server and related services.

A SQL Server failover service group is configured in a virtual server name context. Clients thus access the SQL databases using the virtual server name.

VCS One provides several ways of configuring a service group, which includes the service group configuration wizard, VCS One Web console, and the command line. This chapter provides instructions on configuring a SQL service group using the SQL Server Service Group Configuration Wizard.

The SQL Server Service Group Configuration Wizard enables you to create a SQL Server service group, define the attributes for its resources, and modify the SQL Server service group on all the systems within the VCS One cluster. Also, if you plan to have multiple systems as a part of the SQL service group configuration, you can specify the failover priority for the systems while configuring the service group.

About configuring the start/stop SQL Server service group

The SQL Server Service Group Configuration Wizard also provides an option to configure a start/stop service group for the SQL Server services, in addition to the failover service group.

A standalone service group is configured on a single system wherein VCS One,

- brings the service group online and takes it offline, on that system.
- monitors the SQL Server services and the underlying storage and network components used by SQL Server, on that system.

Configuring the SQL Server standalone service group involves creating resources for local storage, SQL Server and related services.

You can configure the start/stop service group in the local system context or the virtual server context. Depending on the configuration, clients can access SQL Server instances using the local system name or the virtual server name.

In case of an application failure, VCS One faults the service group and takes it offline. Since this service group does not contain multiple systems, it cannot be failed over. Administrative intervention is thus required to bring the service group online again.

VCS One provides several ways of configuring a service group, which includes the service group configuration wizard, VCS One Web console, and the command line. This chapter provides instructions on configuring a SQL service group using the SQL Server Service Group Configuration Wizard.

The SQL Server Service Group Configuration Wizard enables you to create a SQL Server service group, define the attributes for its resources, and modify the SQL Server service group on all the systems within the VCS One cluster.

Before you configure a SQL Server service group

Note the following prerequisites before you configure a SQL Server service group:

- Verify that VCS One 5.0 SP1 for Windows and the SQL agent pack, is installed on all the systems on which you want to configure the service group.
- Verify that you have configured the VCS One client on all the systems on which you want to configure the SQL Server service group.
- Using VEA, rescan the shared disks from all the cluster systems.
- Verify that you have the privileges to add, modify and delete a service group on the OU level under which you wish to configure the service group.
- Verify that you have a user account configured on the Policy Master Server. This account must have privileges to add, modify or delete service groups in the VCS One cluster.

If you choose to use the logged-on user credentials, ensure that the logged-on user account has these privileges.

- Ensure that you are a local administrator on the system where you run the wizard.
- Ensure that the domain user account in which the SQL Server service is installed or running is a part of the local administrator group.
- Verify that the VCS One Remoting and Client Daemon services are running on all the systems that are going to be part of the service group.
In case of a start/stop service group, these services must be running on the system where you run the wizard.
- Verify that the SQL Server instances that you wish to configure in the service group are installed on all the systems that are going to be a part of the service group.
- Verify that the drive containing the SQL Server data files is available and mounted on the system where you run the wizard. This is required to store the registry replication information.
- Ensure that you have a unique virtual IP address for the SQL Server. You need to specify this IP address while configuring the service group.
The virtual IP address is optional for a start/stop service group.
- Identify a virtual server name to set the internal server name for clustered SQL server instance. You need to specify this virtual server name while configuring the service group.

- If you wish to use a script for detail monitoring, either store the script file on a shared drive or ensure that the same file exists in the same location on all the systems. A sample script is supplied in
`%vcsone_home%\bin\SQLServer\sample_script.sql`
 Here `%vcsone_home%` is the default installation directory for VCS One, typically `C:\Program Files\VCS One`.
- Ensure that the SQL Server service is stopped for the SQL instances that you wish to configure in the service group.
- In case, you are configuring the SQL Server service group on a virtual machine, ensure that you have referred to the *VCS One User's Guide*, for details on the following tasks:
 - Configure communications between VCS One and the VirtualCenter Management console.
 - Configure VCS One to discover the VMware VirtualCenter environment.
 - Monitor the discovery of VirtualCenter objects using the VCS One console logs.
 - Place applications in a VMware environment under VCS One control.

Creating a SQL Server service group using the wizard

This section describes how to configure a SQL Server service group.

To modify an existing service group, see “[Modifying the SQL Server service group](#)” on page 86.

Symantec recommends reviewing the resource types and the attribute definitions of the agents before configuring the service group.

See “[About the VCS One agents for Microsoft SQL Server](#)” on page 12.

To create a SQL Server service group

- 1 Ensure that you have stopped the SQL Server service for the instances that you wish to configure in the service group.
- 2 Click **Start > All Programs > Symantec > VCS One > SQL Server Configuration Wizard** to start the SQL Server Configuration Wizard.
- 3 Review the information in the Welcome panel and click **Next**.
- 4 On the Policy Master details panel, specify the login credentials as follows and then click **Next**:

Policy Master: IP Address: Port:	The wizard populates the virtual IP address and port number of the Policy Master (PM) server.
Use Logged-on user credentials	<p>This check box is selected by default.</p> <p>Continue with the selection if you want the wizard to use the logged-on user account context to perform the SQL service group configuration tasks.</p> <p>The credentials values are populated in the respective fields below and the fields are disabled by default.</p> <p>If you do not want to use the logged-on user account context, clear this check box and enter the relevant credentials in the fields below.</p>
User Name	<p>Type a user name.</p> <p>Note: This user account must be configured on the Policy Master and must have read access to the VCS One OU and the privileges to add and modify service groups. The user account must also have the privileges to attach service groups to the desired OUValue. If you choose to use the logged-on user, ensure that the logged-on user account has these privileges.</p>
Password	Type the password for the specified user.

Domain Type	Select the appropriate domain type from the drop-down list. For a Windows PM, select nt. For a UNIX PM select unixpwd.
Domain Name	Type the domain name of the specified user. For example, if Tom is part of the domain <i>mydomain.com</i> , type mydomain . Note that <i>Tom@mydomain</i> must already be configured on the PM.

- 5 On the Service Group Options panel, select from the following options and then click **Next**:
 - Click **Create failover service group** if you wish to configure a multi-system failover service group.
 - Click **Create start/stop service group** if wish to configure a single system standalone service group.
- 6 On the OUValue Details panel, review the VCS One organization tree structure and specify the OUvalue under which you want to create the service group. Click **Next**.
 You can either type the OUValue or select the OUValue from the tree. Type the OUValue path in the format /OUName=OUValue.
 For example, /VCSOneCO=Retail or /VCSOneCO=Marketing/B2B=Dept
 You cannot attach service groups with OUName; you must specify a OUValue.
 Note that only the OUs on which you have access will be available.
 If you are configuring a start/stop service group, proceed to [step 9](#).
- 7 For a multi-system failover service group, on the System Selection panel, select the systems for the SQL Server service group and then click **Next**.
 This step is not applicable if you are configuring a start/stop service group. The Available Systems box displays the systems in the VCS One cluster that the specified user account has access to.

Note: If you are configuring the service group on a virtual machine that has virtual machine high availability (VMHA) policy enabled, the wizard does not allow you to add any other system to the service group. You cannot add more than one vframe to a service group if a vframe is enabled for VMHA.

You can also search systems based on specific system properties.

[Table](#) lists the various filter options that you can use to search systems.

Table 4-1 System selection panel: Filter criteria

Filter Criteria	Usage and Example
Organization Unit	<p>Search systems based on the VCS One OU objects that they belong to. You can either type the OU object details or choose the desired OU object from the Organization Tree displayed.</p> <p>Example: To list the systems belonging to the OUName <i>VCSOneCO</i> and OUValue <i>Retail</i>, either select Retail, or type /VCSOneCO=Retail.</p>
Attributes	<p>Search systems based on the number of CPUs or the amount of random access memory (RAM) available, or a combination of both.</p> <p>Example: To list systems with 4 CPUs and having 4 GB of RAM, type 2 and 4 in the CPU and Memory fields respectively.</p>
Expressions	<p>Search systems based on the VCS One Extended Attributes (EA) or OU expressions, or a combination of both.</p> <p>Specify the details in the following format:</p> <p>EA Expression: <i>EA</i>='EA value'</p> <p>OU Expression: Type <i>/OUName=OUValue</i></p> <p>For EA Expressions, specify the EA value in single quotes.</p> <p>Example: To list systems administered (EA=<i>Admin</i>) by Tom (EA value=<i>Tom</i>) and that belong to the OUValue <i>Retail</i>, specify the details as follows:</p> <p>EA Expression: Admin="Tom"</p> <p>OU Expression: /VCSOneCO=Retail</p> <p>To use EA expressions, the extended attributes must already be defined in the farm.</p>

Table 4-1 System selection panel: Filter criteria

Filter Criteria	Usage and Example
Extended Attributes	<p>Search systems based on the VCS One Extended Attributes. Select an EA from the Name box, then type or select an appropriate EA value in the Value box, and then click Add.</p> <p>Use the boolean AND or OR buttons to combine multiple Extended Attributes and achieve the desired filter criteria.</p> <p>Example: To list systems administered (EA=<i>Admin</i>) by Tom (EA value=<i>Tom</i>) and having operating system (EA=<i>OS</i>) as Windows Server 2008 (EA value=<i>Win2008</i>), do the following:</p> <p>Select Admin and Tom in the Name and Value box respectively, and then click Add.</p> <p>Click AND.</p> <p>Select OS and Win2008 in the Name and Value box respectively, and then click Add.</p> <p>Your search filter should appear as follows:</p> <p>Admin='Tom' AND OS='Win2008'</p> <p>To use the Extended Attributes, they must already be defined in the VCS One cluster.</p>
Keyword	<p>Search systems based on keywords. You can use * and % as wildcard characters to filter systems.</p> <p>Example: To list systems whose names start with the word Win, type Win* or Win%. To list systems whose names contain the word Win, type *Win* or %Win%.</p> <p>The keywords are not case sensitive. *win and *WIN will fetch the same results.</p> <p>Do not use the wildcard characters in between keywords. *SystemName* or %SystemName% is a valid input. System*Name or System%Name is an invalid input.</p>

To filter the systems, click **Edit Filter** and select the appropriate filter from the **Filter systems by** drop-down list and click **Update Filter**.

The Available Systems box now displays the systems as per your chosen search criterion.

You can now arrange the selected systems in the failover priority order. Use the up and down arrows to move the systems to arrange them in their failover priority order. The systems on which the service group needs to come online first must be at the top of the list.

For an active/active configuration, ensure that the priority for active and failover systems is set differently for each instance. For example, if the

system priority for the first instance is SYSTEM1 and then SYSTEM2, the system priority for the second instance should be SYSTEM2 and then SYSTEM1.

- 8
- For a multi-system failover service group, on the Network configuration panel, specify the network details and the virtual computer name for the SQL Server instance and then click **Next**.

This step is not applicable if you are configuring a start/stop service group.

Virtual Server Name	Type a unique virtual server name by which the SQL Server will be known to the clients.
Virtual IP Address	Type a unique virtual IP address for the SQL server that is currently not being used on your network, but is in the same subnet as the current system.
Subnet Mask	Type the subnet to which the virtual IP belongs.
Adapter Display Name	<p>For each system in the cluster, select the public network adapter name from the Adapter Display Name drop-down list.</p> <p>The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.</p>
Advanced Settings	<p>Click Advanced Settings and specify additional details for Active Directory (AD) and DNS updates. These settings apply to the Lanman resource in the service group.</p> <p>On the Advanced Settings dialog box, set the following as desired:</p> <ul style="list-style-type: none">■ Check Active Diretory Update Required check box to enable the Lanman resource to update the Active Directory with the virtual computer name.■ In the Organizational Unit field, type the distinguished name of the Organizational Unit (OU) in the Windows AD. Type the OU in the format CN=containername,DC=domainname,DC=com. To browse for an OU, click Browse 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."■ Check the AD Update Critical For Online check box to set the Lanman attribute ADCriticalForOnline to 1. The Lanman resource will fault if the AD update fails.

- Check the **DNS Update Required** check box to set the Lanman attribute DNSUpdateRequired attribute to 1. The Lanman agent updates the DNS with the specified virtual IP address.
 - Check the **DNS Update Critical For Online** check box to set the Lanman attribute DNSCriticalForOnline to 1. The Lanman resource faults if the agent fails to update the Active Directory.
 - Check the **Ignore Default** check box to set the value of the Lanman attribute DNSOptions to IgnoreDefault. The Lanman agent will update only the DNS servers specified in the AdditionalDNSServers attribute.
 - Check the **Purge Duplicate** check box to set the value of the Lanman attribute DNSOptions to PurgeDuplicate. The Lanman agent will remove duplicate entries from the DNS servers.
 - Click **OK**.
The user account for VCS One Client Helper service must have adequate privileges on the specified container to create and update computer accounts.
- 9 For a start/stop service group, on the Network Configuration panel, choose the desired option and then click **Next**.
This step is not applicable if you are configuring a multi-system failover service group.
Do one of the following:
- To configure the SQL Server in the local system context, click **Use physical network** and then select an IP address from the list.
 - To configure the SQL Server in the virtual server context, click **Use virtual network** and then specify the virtual server name and the network details for the file share server.
- | | |
|----------------------|---|
| Virtual Server Name | Type the unique virtual server name for the SQL Server. This is the same virtual server name that you specified when setting the internal name of the clustered instance. |
| Virtual IP Address | Type a unique virtual IP address for the SQL Server. |
| Subnet Mask | Type the subnet to which the virtual IP belongs. This field displays the subnet of the local system by default. |
| Adapter Display Name | Select a public network adapter from the drop-down list. This list displays all the TCP/IP enabled adapters on the system. |

Advanced Settings

Click **Advanced Settings** and specify additional details for Active Directory (AD) and DNS updates. These settings apply to the Lanman resource in the service group.

On the Advanced Settings dialog box, set the following options as desired:

- Check **Active Directory Update Required** check box to enable the Lanman resource to update the Active Directory with the virtual computer name.
- In the Organizational Unit field, specify the distinguished name of the Windows AD Organizational Unit (OU). To browse for an OU, click **Browse** and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server details to the default container "Computers."
- Check the **AD Update Critical For Online** check box to set the Lanman attribute ADCriticalForOnline to 1. The Lanman resource will fault if the AD update fails.
- Check the **DNS Update Required** check box to set the Lanman attribute DNSUpdateRequired to 1. The Lanman agent updates the DNS with the specified virtual server details.
- Check the **DNS Update Critical For Online** check box to set the Lanman attribute DNSCriticalForOnline to 1. The Lanman resource faults if the agent fails to update the DNS.
- Check the **Ignore Default** check box to set the Lanman attribute DNSOptions to IgnoreDefault. The Lanman agent will update only the DNS servers specified in the AdditionalDNSServers attribute.
- Check the **Purge Duplicate** check box to set the Lanman attribute DNSOptions to PurgeDuplicate. The Lanman agent will remove duplicate entries from the DNS servers.
- Click **OK**.

Note: The Lanman agent uses the VCS One Remoting Service user account context to perform AD and DNS updates. This account must have privileges to create and update computer accounts on the specified container.

10 On the SQL Instance selection panel, perform the following:

- Select the desired SQL Server Version from the **Select SQL Server Version** drop-down list.

The wizard displays the corresponding SQL Server instances that are installed identically on all the systems and which satisfy the following criteria:

- Commonly installed instances
- Instances that are not configured in other service groups

Select the instance(s) that you wish to configure in the service group.

Note: At a time you can create the service groups for only one of the SQL Server version.

- Select the FILESTREAM checkbox to configure FileStream in the service group. This checkbox is selected by default if FileStream is enabled in SQL Configuration Manager.

Note that, if you have not configured the FileStream option during the SQL Server 2008 installation, the selection of FILESTREAM checkbox will configure the FileStream option and create the resource with level 3 enabled.

Note: You can clear this checkbox only if the FileStream is disabled.

- 11 On the SQL Instance database details panel, review the list of databases for the selected SQL instances.

The panel displays the databases and the respective files for which the wizard configures storage resources. Click a database name to view its database files.

Databases that appear with a red cross indicate that the wizard does not configure the storage agent resources for those items. These databases either do not reside on shared storage or the wizard is unable to locate them. If you wish to configure resources for these databases, ensure that the database are located on shared storage and then run the wizard again.

- 12 On the Detail Monitoring Configuration panel, configure detail monitoring for the SQL server instances. This step is optional. If you do not want to configure detail monitoring, click **Next** and proceed to the next step.

Perform the following steps only if you wish to configure detail monitoring for an instance:

- Select the SQL instance, and then select the **Configure Detail Monitoring** check box.
- In the Detail Monitor Cycle box, specify the desired value for the interval in which you want the detail monitoring to happen.

This sets the value for the `DetailMonitoringInterval` attribute of the SQL agent. It indicates the number of online monitor cycles that the agent must wait before performing detail monitoring. The default value is 5. Symantec recommends that you set the monitoring interval between 1 and 12.

- Select **Database monitoring** and then choose the databases from the list of databases available for the instance. The selected databases populate the `DBList` attribute of the SQL agent. In this mode of detail monitoring the agent monitors the health of the databases by connecting to those databases. The agent monitors only the databases specified in the `DBList` attribute.
 - Select **SQL script based monitoring** if you wish to use a script to monitor SQL databases. In this mode of detail monitoring, the agent executes the script that you specify for detail monitoring.
 To select the SQL script,
 - Specify the fully qualified user name and the password for connecting to the SQL Server database. Make sure that the user has SQL Server logon permissions, on all systems in the cluster.
 - Browse to select the SQL script file path. Note that this file should exist on the same path on all the systems.
 - Select **Fault service group if detail monitoring fails** check box, if not already selected. This allows the SQL agent to fail over the service group to another system if the detail monitoring fails.
 - Repeat these steps for each SQL instance that you wish to configure detail monitoring for, and then click **Next**.
- 13 On the Registry Replication Path panel, select the shared drive on which you wish to store the registry replication information and click **Next**.
 Symantec recommends to have the RegRep data and SQL data in separate volumes.
- 14 On the Service Group Summary panel, review the service group configuration and then click **Next**.
 The Service Group Details box lists the configured resources. The wizard assigns unique names to resources based on their respective name rules. Click a resource to view its attributes and their configured values in the Attributes box. Optionally, if desired, change the names of the resources as follows:
- To edit a resource name, select the resource and then either click the resource name or press the **F2** key. Press the **Enter** key after editing each resource name.

Note: You cannot edit the resource name or the SG name, if it already exists in the cluster.

- To cancel editing a resource name, press the **Esc** key.
- 15 The Commit Configuration panel displays the service group configuration progress. The wizard starts running the commands to create the service group. Various messages indicate the status of these commands. After all the tasks are complete, click **Next**.
- 16 On the **Finish** panel, check the **Bring the service group online** check box if you want the wizard to bring the service group online after the service group is configured, and then click **Finish**.
 The wizard issues a command to bring the service group online on the local system.
 This completes the SQL Server service group configuration.
 A service group may fail to come online due to various reasons. This does not mean that the service group was not configured. In such cases, you must probe the resources and then bring the service group online manually from the VCS One console.

Making an existing SQL Server highly available

This chapter contains the following topics:

- [About making existing SQL server highly available](#)
- [Reviewing the configuration](#)
- [Before you begin to install VCS One](#)
- [Converting the existing SQL server into a clustered server](#)

About making existing SQL server highly available

In order to ensure that the SQL resource is always available, you can bring your existing SQL Server installation under the VCS One control. This environment involves an active-passive configuration with one to one failover capabilities.


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

Reviewing the configuration

Bringing an existing SQL Server under VCS One control involves active-passive configuration. The active system of the cluster hosts the SQL Server under virtual server context and the second system is a dedicated redundant server that takes over and hosts the SQL Server running under virtual server context, in the event of a failure on the active system.

Figure 6-1 represents the active-passive configuration

Figure 5-1 Active-Passive configuration



The SQL Server running in virtual server context is online on SYSTEM1, serving client requests. The shared LUNs (virtual disks) provide storage for the SQL Server databases. SYSTEM2 waits in a warm standby state as a backup system. It gets operative when SYSTEM1 fails, and the service group failovers to SYSTEM2.

Before you begin to install VCS One

Complete the following tasks before you begin to install VCS One and create a clustered environment:

- Back up the existing SQL data.
- Set SQL Server services to manual start, using SCM.
- Configure the storage, hardware and network
For details refer to, “[Configuring the storage, hardware and network](#)” on page 82

Configuring the storage, hardware and network

Perform the following to configure the hardware and verify DNS settings.

To configure the hardware, perform the following tasks on all the systems.

- 1 Install the required network adapters, and SCSI controllers.
- 2 Connect the network adapters on each system.
- 3 Using Windows Disk Management for LDM and VEA for SFW, ensure that the attached shared LUNs (virtual disks) are visible on each system.
- 4 Verify that the DNS Server is able to resolve the System Name.

Converting the existing SQL server into a clustered server

Perform the following tasks to convert the existing SQL Server into a clustered server.

- Create a backup of the data from the existing standalone SQL server.
- Use the SQL Configuration Manager and set the startup type of all the SQL services for the SQL instance to manual start.
- Install and configure VCS One client on all the systems on which you wish to bring the SQL server under VCS One control. Refer to the *Veritas Cluster Server One Installation and Upgrade Guide* for more information.
- Install the VCS One agent for SQL Server
 See [“Before installing the agent”](#) on page 35.
- Create volumes or LUNs (virtual disks) necessary to manage the SQL Server storage.
 See [“About managing storage using Windows Logical Disk Management \(LDM\)”](#) on page 49.
 See [“About managing storage using SFW”](#) on page 52.
- Install and configure SQL Server on additional systems, if required, using the SQL Server installer.
 See [“Installing SQL Server on the failover cluster systems”](#) on page 58.
- Move the existing SQL Server data files and user databases from the local drive to the shared drives
 Verify the location of all SQL Server data files for the existing standalone server. If they are located on the local drive, move them to the appropriate volumes or LUNs on shared drive, to ensure proper failover operations in the cluster.
 See [“Moving the existing SQL Server data files and user databases”](#) on page 84.
- Configure the service group using the SQL Server Service Group Configuration Wizard.
 The SQL Server 2008 Configuration Wizard enables you to create a SQL Server service group and simultaneously define the attributes for its resources on all the systems within the cluster. Run the configuration wizard from the system where the shared drives containing the SQL database files are mounted.
 See [“About configuring the failover SQL Server service group”](#) on page 64..

See [“About configuring the start/stop SQL Server service group”](#) on page 64.

- Create and configure SQL Server user-defined databases.
Create databases and configure resources for those databases in the SQL service group.
See [“Creating a SQL Server user-defined database”](#) on page 89.
- Verify the configuration
Verify your configuration by switching the service group or by shutting down the system that is currently online. Either process tests the service group can be smoothly transferred between systems.
If you have opted to shutdown the online system in your cluster, remember to bring it back up after you have confirmed that the service group successfully failed over to another system.
Refer to *Veritas Cluster One User's Guide* for more details.

Moving the existing SQL Server data files and user databases

Before you begin to configure the SQL Server service group, move the existing SQL Server data files and user databases from the local drive to the shared drives to ensure proper failover operations in the cluster.

Perform the following steps, to move the databases from the local drive to the shared drive:

- 1 Stop the SQL Server service, using SCM.
- 2 Verify that you have backed up your existing data.
- 3 Import the volumes or LUNs to the system where the original files are located on the local drives and mount the volumes (add drive letters).
See [“About managing storage”](#) on page 48.
- 4 Move the SQL Server data files and user database locations.
Follow the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach.
<http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>
- 5 Restart SQL Server service, using SCM.

Administering the service group

This chapter contains the following topics:

- [Modifying the SQL Server service group](#)
- [Creating a SQL Server user-defined database](#)
- [Adding storage agent resources to the SQL service group](#)

Modifying the SQL Server service group

You can modify the SQL Server service group configuration using the SQL Server Service Group Configuration Wizard, VCS One Web Console, and the command line.

While modifying the service group, you can perform the following administrative tasks.

- Add or remove systems or change their priority order
- Modify network settings of the service group. For example, the IP address, LANMAN or Subnet mask
- Add or remove the instances
- Add or remove storage for the instances after adding or deleting databases
- Add, remove or modify the detail monitoring settings
- Add, or modify the Regrep resource
- Modify the OU to which the service group is attached.

Additionally, using the VCS One console you can also delete a service group.

For more details refer to, *VCS One User's Guide*.

Prerequisites

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

- Run the wizard from the system on which the service group is active (online or partially online). You can use the wizard to add resources to and remove them from the configuration. However, you cannot change the resource attributes.
- To change the resource attributes, you must take the service group offline. However, the MountV (Mount in case of Windows LDM) and VMDg (DiskRes in case of Windows LDM) resources for the service group should be online on the system where you run the wizard and offline on all other systems.
- If you are running the wizard to remove a system from the service group's system list, do not run the wizard on the system being removed.
- If you are running the wizard to add or remove MountV (Mount in case of Windows LDM) resources for user defined databases, make sure the SQL Server service is online in the selected service group.

Additionally, review the prerequisites required to create a SQL Server service group.

See [“Before you configure a SQL Server service group”](#) on page 66.

Modifying a SQL Server service group

The following steps describe how to modify the service group using the SQL Server Configuration Wizard.

To modify a SQL Server service group

- 1 Click **Start > All Programs > Symantec > VCS One > SQL Server Configuration Wizard**.
- 2 Review the prerequisites on the Welcome panel and then click **Next**.
- 3 On the Policy Master Details panel, specify the login credentials, as follows:

Policy Master: IP Address: Port:	The wizard populates the virtual IP address and port number of the Policy Master (PM) server.
Use Logged-on user credentials	<p>This check box is selected by default, if the Policy master server you are connecting to, runs on the Windows operating system and belongs to the same domain as that of your client system.</p> <p>Continue with the selection if you want the wizard to use the logged-on user account context to perform the SQL service group configuration tasks.</p> <p>The credentials values are populated in the respective fields below and the fields are disabled by default.</p> <p>If you do not want to use the logged-on user account context, clear this check box and enter the relevant credentials in the fields below.</p>
User Name	Type a user name.
Password	Type the password for the specified user.
Domain Type	Select the appropriate domain type from the drop-down list. For a Windows PM, select nt. For a UNIX PM select unixpwd
Domain Name	<p>Type the domain name of the specified user.</p> <p>For example, if Tom is part of the domain <i>mydomain.com</i>, type mydomain or mydomain.com.</p> <p>Note that <i>Tom@mydomain.com</i> must already be configured in the VCS One cluster.</p>

The wizard performs the verification checks to authenticate the user with the Policy Master server. Click **Next**.

- 4 On the Service Group Options panel, click **Modify service group**.
Note that only the service groups that were created using the wizard are displayed in the grid below. Also, if the service group contains any mixed

versions of SQL instances, then the service group validation fails and the wizard hangs.

In case the service group you wish to modify does not appear in the list of service groups displayed in the grid, type the service group name in the **Service group name** text box. When you begin to type the service group name, a list of all the service groups configured with the Policy Master gets sorted and populated.

Select the desired service group and click **Next**.

- 5 On the OUValue Details panel, note that the OUValue under which the selected service group is associated is selected by default. Click **Next** to proceed without changing the OUValue.
However, in case you wish to change the OUValue, select the new OUValue and then click **Next**.
- 6 On the System Selection panel, add or remove the systems or modify their failover priority order, in case you wish to do the same. Click **Next**.
For more details to add a system refer to Step7 from [Creating a SQL Server user-defined database](#)
Note that you cannot remove the system from which you are running the wizard. Also, ensure that you have the administrative privileges and the VCS One Remoting and Client Daemon services are running on the systems that you wish to add.
- 7 On the Network Configuration panel, if desired, modify the network settings of the selected SQL Server service group and then click **Next**.
Note that you can modify the network settings only if the Lanman and IP resource are offline.
For more information on assigning the network settings refer to, step 9 from [Creating a SQL Server user-defined database](#)
- 8 On the SQL Server Instance Selection panel, select the instances and the corresponding services that you wish to configure or remove.
Note that,
 - Only those instances that are installed on the selected systems and are not a part of another service group are displayed.
 - All the instances that are configured are selected by default.
 - You cannot remove the instances that are online. In case you want to remove any of these resources, access the Web Console and bring the resources offline and then relaunch the wizard.
- 9 On the Instance Database List panel, review the database details for the configured SQL instances and click **Next**.

The panel displays the databases and the respective files for which the wizard configures storage resources. Click a database name to view its database files.

Databases that appear with a red cross indicate that the wizard does not configure the storage agent resources for those items. These databases either do not reside on shared storage or the wizard is unable to locate them. If you wish to configure resources for these databases, ensure that the database are located on shared storage and then run the wizard again.

- 10 On the Detail Monitoring Configuration panel, if desired, modify the detail monitoring settings and click **Next**.
For more details on setting the detail monitoring options refer to, step 12 from “[Creating a SQL Server user-defined database](#)” on page 89
- 11 On the Registry Replication Path panel, if desired, modify the location for storing the RegRep information and then click **Next**.
Note that, you can modify the location only if all the RegRep resources in the selected service group are offline.
- 12 On the Service Group Summary panel, review the service group configuration details. If you want to rename any of the resources, you can do so. However, note that only those resources that are newly added can be renamed. The resources that are already a part of the service group can not be edited.
Click **Next** to commit the changes.
- 13 On the Commit Configuration panel, review the configuration progress and click **Next** after the configuration is complete.
- 14 On the Finish panel, select the **Bring the service group online** checkbox, if you want to bring the service group online. Click **Finish**.
You may want to review the service group configuration in the VCS One Web console before bringing the service group online. You can use the VCS One Web Console to bring the service group online later.

Creating a SQL Server user-defined database

Create the required SQL databases using the SQL Server Management Studio and then make them highly available with Veritas Cluster Server One.

Refer to the following guidelines before creating and configuring the user-defined databases:

- Create volumes or LUNs for a user-defined SQL Server databases, if you have not already created them.
In the sample deployment the following volumes are named:
 - INST1_DB1_VOL: contains a user-defined database file

- INST1_DB1_LOG: contains a user-defined database log file
- INST1_DB1_FS_VOL: contains FILESTREAM enabled databases.
- Create a SQL Server user-defined database for the required SQL instance using the SQL Server Management Studio. While creating the database, ensure that you point the database files and log files to the paths of the new volumes or LUNs created earlier.
Refer to the SQL Server documentation for instructions on how to create databases.
- After creating the databases, run the SQL Server Service Group Configuration Wizard and modify the SQL Server service group. This allows the wizard to add the corresponding storage resources for the user databases, to the SQL Server service group.
See [“Adding storage agent resources to the SQL service group”](#) on page 90.

Adding storage agent resources to the SQL service group

Perform the following steps to add the storage resources for the user-defined databases, to the SQL Server service group. You must run the wizard in the modify mode even if you have added or changed volumes in your existing configuration. This allows the wizard to make the necessary changes to the SQL Server service group.

Before you proceed, make sure the volumes for the user-defined databases are mounted on the system.

Note: You are required to add storage resources only if the database is created on a new volume.

To add storage agent resources to the SQL service group

- 1 Start the SQL Server Service Group Configuration Wizard. Click **Start > All Programs > Symantec > VCS One > SQL Server Configuration Wizard**
- 2 Review the prerequisites on the Welcome panel and click **Next**.
- 3 On the Service Group Options panel, click **Modify service group**, select the service group and then click **Next**.
- 4 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
- 5 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**.

- 6 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 7 Click **Yes** to continue when a Veritas Cluster Server Notice indicates the configuration will be modified.
- 8 Click **Finish** to exit the wizard.
The wizard marks all the resources in the service group as CRITICAL. If desired, use VCS One Web Console or the command line to change the state

Upgrading SQL Server

This chapter contains the following topics:

- [About upgrading SQL Server](#)
- [Upgrading from Microsoft SQL Server 2005 to SQL Server 2008](#)
- [Upgrading the SQL Server service packs](#)

About upgrading SQL Server

This chapter provides information for upgrading SQL Server 2005 to SQL Server 2008, in a VCS One cluster.

Upgrading from Microsoft SQL Server 2005 to SQL Server 2008

The following steps describe how to upgrade your existing clustered SQL Server 2005 setup to SQL Server 2008. Complete these steps on all the cluster systems that are part of the SQL service group, one system at a time.

At a high level, upgrading to SQL Server 2008 involves the following tasks:

- Upgrade SQL Server on the first cluster system.
- Upgrade SQL Server on each additional failover system.
- In case of a Disaster Recovery configuration, repeat the SQL upgrade procedures on the systems at the secondary site. First upgrade the first cluster system at the DR site, and then the additional failover systems.
- If the existing SQLServer2005 service group contains the Full text search service resource then delete the resources from the service group.
- Run the SQL Server Configuration wizard in modify mode and select the upgraded service groups to upgrade the cluster configuration. In case of a DR setup, perform this at the secondary site also.

Note: In case of a Disaster Recovery setup, you must first upgrade SQL on the cluster systems at the primary site and then proceed with the systems at the secondary site. You must follow the same upgrade sequence at both sites, upgrade first system and then the additional systems, as described in the procedures in this section.

Ensure that you perform the following before the upgrade:

- Take a backup of the SQL databases.
- In case of a Disaster Recovery environment, ensure that the databases on the primary and secondary sites are synchronized and then stop the replication between the sites.
- Make a note of the SQL virtual server name and all the IP addresses configured at both the primary and the secondary site, for the SQL setup in the DR environment. You will need these details later.

Upgrading SQL on the first cluster system

These steps assume a single SQL Server instance configured in a two cluster system configuration.

To upgrade SQL Server on the first cluster system

- 1 On the system on which the SQL service group is online, take all the resources excluding the Mount resource offline.
To bring the resources offline, access the VCS One Web Console. Right-click the resource and select **Offline**.
- 2 Bring the registry replication (RegRep) resource online.
- 3 Make a backup copy of the SQL Server 2005 and the Registry (RegRep) directories from the shared disk and store them in a temporary location.
You will need the backed-up directories while upgrading SQL Server 2005 on the additional failover systems, later.
- 4 Freeze the SQL Server service group.
From the VCS One Web Console, right-click the SQL Server service group in tree view on the left pane, and click **Freeze > Persistent**.
- 5 Launch the Microsoft SQL Server 2008 installer and install SQL Server 2008 on the system. Make sure that you select the option to upgrade the existing SQL Server instance(s), when prompted to do so. The SQL Server 2008 installer then automatically places the SQL data files in the appropriate location.
Refer to the Microsoft SQL Server 2008 documentation for instructions.
- 6 Unfreeze and then take the SQL Server service group offline. From the VCS One Web Console, right-click the SQL Server service group in tree view on the left pane and click Unfreeze, and then take the entire service group offline on the system.

This completes the upgrade steps on the first cluster system. Proceed to upgrading SQL on the additional failover systems.

Upgrading SQL Server on the failover systems

Perform the following steps on each additional failover system that is part of the SQL service group.

To upgrade SQL Server on each additional systems

- 1 Bring all the SQL service group resources including the storage resources, online. From the VCS One Web Console, right-click the resource and select **Online**. Click **Yes** in the confirmation pop-up box to bring the resource online.

- 2 Rename the SQL Server and the Registry (RegRep) directories on the shared disks. These directories are updated when SQL Server 2008 is installed on the first system. You can also delete these directories, if desired.
- 3 Copy the backed-up Registry (RegRep) directories from the temporary location to the shared disks. The backup directories are the same that you had backed up earlier while upgrading SQL on the first cluster system.
- 4 Bring the RegRep resource online.
From the VCS One Web Console, right-click the resource and select **Online**. Click **Yes** in the confirmation pop-up box to bring the resource online.
- 5 Freeze the SQL service group.
From the VCS One Web Console, right-click the SQL Server service group in tree view on the left pane and click **Freeze > Persistent**.
- 6 Launch the Microsoft SQL Server 2008 installer and install SQL Server 2008 on the system. Make sure that you select the option to upgrade the existing SQL Server instance(s), when prompted to do so. The SQL Server 2008 installer then automatically places the SQL data files in the appropriate location.
Refer to the Microsoft SQL Server 2008 documentation for instructions.
- 7 Unfreeze and then take the SQL Server service group offline. From the VCS One Web Console, right-click the SQL Server service group in tree view on the left pane and click **Unfreeze**, and then take the entire service group offline on the system.
This completes the upgrade steps on the additional failover system.
- 8 Run the SQL Sever Service Group Configuration wizard to configure the upgraded SQL service group in the cluster.

Upgrading the SQL Server service packs

This section describes how to upgrade SQL Server 2005 or SQL Server 2008 to their corresponding service packs.

See [“Upgrading SQL Server 2005 to SQL Server 2005 SP1”](#) on page 97.

See [“Upgrading SQL Server 2005 to 2005 SP2 or later”](#) on page 98.

See [“Upgrading Microsoft SQL Server 2008 to 2008 SP1”](#) on page 103.

The outlined procedures are applicable only if you already have your SQL server setup in a VCS One cluster.

Upgrading SQL Server 2005 to SQL Server 2005 SP1

This section describes how to upgrade SQL 2005 to SQL 2005 SP1.

Ensure that you meet the following pre-requisites before begin to upgrade SQL Server.

- Review your Microsoft documentation for the requirements for a Microsoft SQL 2005 Server SP1 installation.
Make sure your system meets these requirements.
- Make sure that you have a recent backup of your system and user databases.
- Ensure that you are prepared for a server down time during this upgrade.

To install SQL Server 2005 SP1

- 1 In VCS One Web Console, right-click the SQL Server service group and select **Offline on all systems**.
- 2 On the system where the SQL Server service group was taken offline, online the SQL Server 2005 resource for the shared drive containing the SQL Server databases.
- 3 In VCS One Web Console, right-click the SQL Server service group which is now partially online, and select **Freeze > Persistent**.
- 4 Install SQL Server 2005 Service Pack 1 on the active system (where the SQL Server service group is online), using the instructions provided by Microsoft.
- 5 Repeat step 5 for each additional SQL instance in this service group, if you have more than one instance in this service group.
- 6 In VCS One Web Console, right-click the SQL Server service group which is still online and select **Unfreeze**.
- 7 VCS One Web Console, right-click the SQL Server service group and select **Offline** on the system where it was online.

- 8 On the failover system, online the SQL Server 2005 resource for the shared drive containing the SQL databases.
- 9 In VCS One Web Console, right-click the SQL Server service group which is now partially online and select **Freeze >Persistent**
- 10 Install SQL Server 2005 Service Pack 1 on the active system (where the SQL Server service group is online), using the instructions provided by Microsoft SQL Server 2005 Service Pack 1 Setup.
- 11 Repeat step 10 for each additional SQL Server instance in this service group, if you have more than one instance in this service group.
- 12 In VCS One Web Console, right-click the SQL Server service group which is still online and select **Unfreeze**.
- 13 In VCS One Web Console, right-click the SQL Server service group and select **Offline** on the system where it was online.
- 14 Optionally reboot and online each service group to verify the database connect for each system.
- 15 Repeat step 8 through 14 on each additional system is more that two SQL Server 2005 systems are in use.
- 16 For a Disaster Recovery environment, repeat this procedure at the secondary site.
- 17 When SQL Server 2005 Service Pack 1 has been completely installed on all systems, test user connectivity to the instances by bringing it online and failing it over from system to system.
- 18 If more than one SQL Server service group is present, repeat this entire procedure for each SQL Server service group.

Upgrading SQL Server 2005 to 2005 SP2 or later

This section describes how to upgrade SQL 2005 to SQL 2005 SP2 or later.

Note: Do not follow the installation steps provided in this section to install SQL Server 2005 Service Pack 1 and all other hotfixes released before Service Pack 2.

Prerequisites

Consider the following points before applying SQL Server 2005 Server SP2 or later service pack version to a production server:

- You must be a domain user having administrative privileges to the cluster systems.

- You must have administrative privileges to the SQL instance that you wish to upgrade.
- You must back up the SQL Server 2005 databases.
- Refer to the Microsoft documentation for prerequisites related to SQL Server 2005 Service Pack installation.

Preliminary installation information

Typically, multiple SQL instances are configured in a VCS One cluster. Each SQL service group is configured to fail over on one or more systems in the cluster. The system on which the SQL service group is online is called as the Active system for that SQL instance. The system on which the SQL service group is offline is called as the Passive system for that SQL instance. The procedure for applying service packs, patches, or hotfixes for SQL instances varies depending on whether it is an active or a passive system. This document describes procedures for both the cases in detail.

Use the procedure that applies to the type of setup you have.

To provide context, the installation procedures described in this document assume two SQL Server 2005 instances configured in a VCS one cluster with three failover systems.

Table 7-1 SQL Server 2005 SP upgrade configuration objects

Object	Description
System1, System2, System3	Cluster system names
SQLinst1, SQLinst2	SQL Server 2005 instance names

The configuration is as follows:

- SQLinst1 can fail over on System1 and System2, and SQLinst2 can fail over on System3 and System2.
So, System2 is the common failover system for SQLinst1 and SQLinst2.
- The SQL service group for SQLinst1 is online on System1, and the SQL service group for SQLinst2 is online on System3.
So, System1 and System3 are the “active” systems for SQLinst1 and SQLinst2 respectively. System2 is the “passive” system for both SQL instances. We will first install the service pack on System2 (passive system) and then proceed to install on System1 and System3.

Installing the Service Pack on “passive” cluster systems

Perform these steps on all the systems where the SQL service group is configured to fail over but is not online. You can either perform the installation at one time for all the SQL instances that are configured to fail over on the system, or repeat the steps for each SQL instance separately.

Do not run these steps for SQL instances whose corresponding service groups are online on the active systems. For installation on active systems,

See [“Installing the Service Pack on “active” cluster systems”](#) on page 101.

Note: You can install SQL Server 2005 Service Pack in an unattended mode from the command prompt using the /quiet switch to suppress the setup dialog boxes. Refer to the Microsoft documentation for more information.

To install the service pack on passive cluster systems

- 1 Ensure that service groups for SQL instances SQLInst1 and SQLInst2 are offline on System2.

Note: This upgrade procedure will not upgrade the SQL instance whose corresponding service group is online on the system

- 2 On System2, copy the SQL Server 2005 Service Pack installer or map a drive to the directory where the installer is located.
- 3 From the command prompt on System2, navigate to the directory where the installer is located.
- 4 From the command prompt, run the service pack installer command with the appropriate options.

For example,

```
SQLServer2005SP2-KB921896-x64-ENU.exe [options]  
/passive=1
```

You can use the following options for the command:

- /allinstances
This option upgrades all SQL Server 2005 instances and shared components to the desired SQL Server 2005 SP.
- /instancename = “<instance1>, <instance2>, ...”
This option upgrades only the specified SQL Server 2005 instances and shared components to the desired SQL Server 2005 SP.

You can run any of the following commands on System2:

```
SQLServer2005SP2-KB921896-x64-ENU.exe  
/allinstances /passive=1
```

or

```
SQLServer2005SP2-KB921896-x86-ENU.exe
/instancename = SQLinst1, SQLinst2 /passive=1
```

Note that in case of multiple SQL instances, there should be no spaces between instance names in the command.

- 5 Follow the upgrade wizard to complete the installation.
Once the installation is complete on the passive systems, proceed to install on the active systems.

Installing the Service Pack on “active” cluster systems

Perform these steps on all the systems on which the SQL service group is online. You can either perform the installation at one time for all the SQL instances that are configured to fail over and are online on the system, or repeat the steps for each SQL instance separately.

Do not run these steps for SQL instances whose corresponding service groups are offline on the systems (passive systems). For installation on passive systems,

See [“Installing the Service Pack on “passive” cluster systems”](#) on page 100.

Referring to the configuration example described earlier, run these steps on System1 and System3 where the SQL service groups for SQLinst1 and SQLinst2 are online.

To install the service pack on active cluster systems

- 1 Ensure that the SQL service group for SQLinst1 is online on System1.
- 2 In the SQL service group for SQLinst1, take all resources of type SQLServer2005 offline on System1.
If there are other SQL Server 2005 instances configured on the system that you wish to upgrade, take SQLServer2005 resources of the respective service groups offline as well.
- 3 From the Services snap-in, stop the SQL server Full Text Search service and the Analysis service, if they are not configured as part of the SQL service groups.
- 4 Freeze the SQL service group for SQLinst1 on System1.
From the VCS One Web Console, right-click the SQL service group, select **Freeze** and click **Temporary**.

or

Type the following on the command prompt:

```
hagrp -freeze service_group
```

- 5 If the SQL Server Reporting Services is installed for a particular instance, start the SQL Server Database Service of the respective instance using the Services snap-in.
- 6 Run the SQL Server 2005 Service Pack installer.
Double-click SQLServer2005SP2-KB921896-x64-ENU.exe to launch the SP installation wizard.
- 7 Follow the upgrade wizard to complete the installation.
- 8 After the installation is complete, stop the SQL Server services, if started before applying the patch.

Note: SQLServer2005 resources may go in UNKNOWN state if we start the services outside the VCS cluster. Ignore this and probe the resources after installation is completed and all the services are stopped.

- 9 Unfreeze the SQL service group and probe the resources for SQLinst1 on System1.
From the VCS One Web Console, right-click the SQL service group, and select **Unfreeze**.
or
Type the following on the command prompt:
`hagrp -unfreeze service_group`
- 10 From the Services snap-in, start the SQL server Full Text Search service and the Analysis service, if they are not configured as part of the SQL service groups.
- 11 Ensure that all the services related to the SQL Server 2005 instance are in stopped state.
- 12 Apart from the SQL Browser service, set the startup type of all the SQL services to manual.
- 13 Bring the SQLServer2005 resources in the SQL service group for SQLinst1 online on System1.
- 14 Repeat step 1 to step 13 for SQLinst2 on System3.

Upgrading Microsoft SQL Server 2008 to 2008 SP1

This section discusses the procedure to upgrade Microsoft SQL Server 2008 to Microsoft SQL Server 2008 SP1 in a VCS One cluster.

The upgrade procedure considers a two systems in a cluster, SystemA and SystemB. The SQL service group is ONLINE on System A, and System B is the passive System.

Prerequisites

Consider the following points before proceeding with the upgrade:

- You must have administrative privileges to the SQL instance that you wish to upgrade.
- Make sure that you have a recent backup of your system and user databases.
- Ensure that you have referred to the Microsoft documentation for prerequisites related to SQL Server 2008 Service Pack 1 installation.

Upgrading the cluster

- 1 Change the RegRep resource Exclude Keys Attribute on any one of the cluster systems, using HA commands or Web Console.
Add the following key in Exclude Keys:
HKLM\SOFTWARE\Microsoft\Microsoft SQL Server\InstanceName\MSSQLServer\Filestream
- 2 Upgrade the SQL 2008 instance on SystemB (passive system) and reboot the system, if the system prompts for.
- 3 Freeze the service group on SystemA (active node).
- 4 Upgrade the SQL 2008 instance on SystemA.
- 5 Unfreeze the service group.
- 6 If the system asks for a reboot, failover the service group to the passive node and reboot the system.
- 7 Remove the Exclude key from RegRep resource that was added in step 1.

Resource type definitions

This chapter contains the following topics:

- [About resource type definitions](#)
- [Agent for SQL Server](#)
- [Agent for SQL Server FileStream](#)

About resource type definitions

This appendix lists resource type definitions and attribute definitions of the SQL Server agents .

The resource type represents VCS One specifies how the agent is defined in the configuration file. The Attribute Definition table explains the attributes associated with the agent. The sample dependency graph depicts the resource types, resources, and resource dependencies required to create the service groups.

Agent for SQL Server

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

Resource type definition

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

```
<agent name="SQLServer" platform="windows/x86">
  <attributes class="static">
    <attribute name="ArgList" type="str" dimension="vector">
      <default>
        <val>"Instance"</val>
        <val>"LanmanResName:VirtualName"</val>
        <val>"SQLOnlineTimeout"</val>
        <val>"SQLOfflineTimeout"</val>
        <val>"DetailMonitorInterval"</val>
        <val>"SQLDetailMonitorTimeout"</val>
        <val>"Username"</val>
        <val>"Domain"</val>
        <val>"Password"</val>
        <val>"DBList"</val>
        <val>"SQLFile"</val>
        <val>"FaultOnDMFailure"</val>
        <val>"LanmanResName:IPResName"</val>
      </default>
      <no_override>1</no_override>
    </attribute>
    <attribute name="MinCompatAgentVersion" type="str"
dimension="scalar">
      <default><scalar>"5.1.0.0"</scalar></default>
      <editable>0</editable>
      <no_override>1</no_override>
    </attribute>
    <attribute name="MsgCatID" type="int" dimension="scalar">
      <default><scalar>20069</scalar></default>
    </attribute>
```

```

        <attribute name="TypeDefinitionVersion" type="str"
dimension="scalar">
        <default><scalar>"5.1.0.0"</scalar></default>
        <editable>0</editable>
        <no_override>1</no_override>
    </attribute>
</attributes>
<attributes>
    <attribute name="DBList" type="str" dimension="vector">
        <important>1</important>
        <attr_description>34</attr_description>
    </attribute>
    <attribute name="DetailMonitorInterval" type="int"
dimension="scalar">
        <important>1</important>
        <must_configure>1</must_configure>
        <validation>"R:0,65535"</validation>
        <attr_description>33</attr_description>
    </attribute>
    <attribute name="Domain" type="str" dimension="scalar">
        <important>1</important>
        <attr_description>35</attr_description>
    </attribute>
    <attribute name="FaultOnDMFailure" type="boolean"
dimension="scalar">
        <default><scalar>1</scalar></default>
        <important>1</important>
        <attr_description>37</attr_description>
    </attribute>
    <attribute name="Instance" type="str" dimension="scalar">
        <important>1</important>
        <unique>1</unique>
        <attr_description>38</attr_description>
    </attribute>
    <attribute name="LanmanResName" type="str"
dimension="scalar">
        <important>1</important>
        <attr_description>39</attr_description>
    </attribute>
    <attribute name="Password" type="str" dimension="scalar">
        <important>1</important>
        <agentencrypt>1</agentencrypt>
        <attr_description>42</attr_description>
    </attribute>
    <attribute name="SQLDetailMonitorTimeout" type="int"
dimension="scalar">
        <default><scalar>30</scalar></default>
        <important>1</important>
        <must_configure>1</must_configure>
        <validation>"R:0,65535"</validation>
        <attr_description>43</attr_description>
    </attribute>

```

```

        <attribute name="SQLFile" type="str" dimension="scalar">
            <important>1</important>
            <attr_description>44</attr_description>
        </attribute>
        <attribute name="SQLOfflineTimeout" type="int"
dimension="scalar">
            <default><scalar>90</scalar></default>
            <important>1</important>
            <must_configure>1</must_configure>
            <validation>"R:0,65535"</validation>
            <attr_description>45</attr_description>
        </attribute>
        <attribute name="SQLOnlineTimeout" type="int"
dimension="scalar">
            <default><scalar>90</scalar></default>
            <important>1</important>
            <must_configure>1</must_configure>
            <validation>"R:0,65535"</validation>
            <attr_description>46</attr_description>
        </attribute>
        <attribute name="Username" type="str" dimension="scalar">
            <important>1</important>
            <attr_description>47</attr_description>
        </attribute>
    </attributes>
</agent>

```

Agent for SQL Server FileStream

The resource type definitions for the VCS One agent for SQL Server 2008 FileStream are as follows. This information will assist you during the agent configuration.

Resource type definition

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

```

<agent name="SQLFilestream" platform="windows/x86">
    <attributes class="static">
        <attribute name="ArgList" type="str" dimension="vector">
            <default>
                <val>"InstanceName"</val>
            </default>
            <no_override>1</no_override>
        </attribute>
        <attribute name="MinCompatAgentVersion" type="str"
dimension="scalar">
            <default>
                <scalar>"5.1.0.0"</scalar>
            </default>

```

```
<editable>0</editable>
<no_override>1</no_override>
</attribute>
<attribute name="MsgCatID" type="int" dimension="scalar">
  <default>
    <scalar>20070</scalar>
  </default>
</attribute>
<attribute name="TypeDefinitionVersion" type="str"
dimension="scalar">
  <default>
    <scalar>"5.1.0.0"</scalar>
  </default>
  <editable>0</editable>
  <no_override>1</no_override>
</attribute>
</attributes>
<attributes>
  <attribute name="InstanceName" type="str" dimension="scalar">
    <important>1</important>
    <attr_description>7</attr_description>
  </attribute>
</attributes>
</agent>
```


Troubleshooting VCS One agents

This chapter contains the following topics:

- [About troubleshooting the VCS One database agent for SQL](#)
- [Newly added SQL Server 2008 Databases are not accessible on failover](#)
- [SQL Server service group faults upon a failover with an error - Failed to find Disk Group](#)
- [Error messages and descriptions](#)
- [VCS One logging](#)

About troubleshooting the VCS One database agent for SQL

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

Newly added SQL Server 2008 Databases are not accessible on failover

When a new database is created in SQL Server 2008, the NTFS permissions assigned to the database files (*.mdf and *.ldf) are applicable only for the local groups (Local Administrators group and a local SQL Users group that SQL 2008 creates during installation)

This results in database access issues when the database fails over or is moved to another system in the cluster and you attempt to access it from the SQL Management Studio.

This situation may occur in the following cases:

- Domain administrator group is not a part of the local administrator's group.
- The SQL Server user account added is a local user that exists only on the system where the database was created, and not on the other systems in the cluster.

Workaround:

Set the SQLClusterAccount attribute of the SQL Server agent for Database Engine.

Specify a domain group or the SQL Server service name. The agent assigns the account with Full Control privileges to the SQL Server databases and log files. If you specify a domain group, then the SQL service account must be part of this domain group.

For a domain group specify in the format Domain.com\DomainGroup.

For SQL Server service name specify in the format MSSQL\$InstanceName. For the default instance, the service name is MSSQLServer.

SQL Server service group faults upon a failover with an error - Failed to find Disk Group

Consider you have installed SQL Server with its data path on a shared disk for one of the cluster system (system A) and on the system drive for another cluster system (system B). Also, while creating the service groups, you have included both these systems in the failover system lists and have specified the data path of the shared drive.

Now, during the failover, from system A to system B the service group faults with the “Failed to find Disk Group” error.

This issue occurs because system B is unable to locate the shared disks.

Workaround:

Rescan the shared disks from system B.

To rescan the disk, perform the following steps:

- 1 Click **Start > Program > Symantec > Veritas Storage Foundation > Veritas Enterprise Administrator**.
- 2 Login to the system where you want to rescan the disks.
- 3 From the Actions menu, click **Rescan**.

Error messages and descriptions

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

Table B-1 SQL Server 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</i> name.	No value provided for the specified attribute. Solution: Provide a value for the attribute.

Table B-1 SQL Server agent error messages

Message	Description
Lanman resource is not configured	<p>This is the debug message with logdbg attribute value equal to 1.</p> <p>The error occurs because the LanmanResName attribute is not configured. The agent gets the SQL instance online in the physical system name context.</p>
The username attribute has not been configured	<p>The username attribute of the SQLServer resource for the instance is not configured.</p> <p>This error occurs when the detail monitor interval attribute is configured with a non zero value, but no value is assigned to the dblist attribute. Under this situation, SQL Server tries to monitor the SQL File with the specified user context.</p> <p>Solution: Provide the SQL admin user name who has the privilege to connect to the sql instance to perform the SQL File detail monitoring.</p>
Invalid arglist. The ArgList should contain LanmanResName:IPResName	<p>The SQLServer type does not contain the LanmanResName:IPResName in the ArgList attribute.</p> <p>Solution:</p> <ul style="list-style-type: none">■ Verify if the Lanman resource name is valid.■ Contact Symantec helpdesk
Invalid value specified for SQL Domain group account	<p>The configured value for the DomainGroupAccount attribute is invalid.</p> <p>Solution:</p> <p>Specify the DomainGroupAccount attribute value in the format fully_qualified_domain_name\Domain_Group_Name (For example Domain.com\SQL Admins)</p>

Table B-1 SQL Server agent error messages

Message	Description
The password attribute has not been configured	<p>The password attribute of the SQLServer resource for the instance is empty.</p> <p>This error occurs when the detail monitor interval attribute is non zero and the dblist is empty.</p> <p>The SQL Server tries to perform the detail monitoring in the specified user context.</p> <p>Solution: Specify the valid password</p>
Failed to initialize the SQLServer2008 agent	<p>This is a VCS internal error.</p> <p>Contact Symantec Technical Support.</p>
Failed to open SCM handle. Error = Error code	<p>The agent fails to get a handle to the Service Control Manager (SCM). This occurs if the specified SCM database does not exists or the requested access is denied.</p> <p>Solution: Verify if SCM can run on the host system and refer to the associated Windows error code for more information.</p>
Failed to start the service <i>service name</i> . Error = Error code	<p>The agent failed to start the service.</p> <p>Solution: Verify if you can start the service from the Windows Services console. If it fails, to start look at the event log for more information on the reasons for the failure.</p>
Failed to set the virtual computer name in the environment of the service <i>Service name</i> . Error = Error code	<p>This is a VCS One internal error.</p> <p>Solution: Contact Symantec Technical Support.</p>
The service <i>service name</i> did not start within the specified time limit.	<p>The agent failed to start the service within the time limit specified in the SQLOnlineTimeOut attribute.</p> <p>Solution: Monitor the time required to start the service and modify the SQLOnlineTimeout attribute value.</p>

Table B-1 SQL Server agent error messages

Message	Description
SQL Server Instance name is not specified. Agent will operate on the default instance.	No value is specified for the SQL Server instance name. In this situation the agent operates on the default SQL Server instance (MSSQLSERVER).
Failed to wait for the service <i>service name</i> to start. Error = Error code	This is a VCS One internal error. Solution: Contact Symantec Technical Support
The service <i>service name</i> service is not in stopped or running state. State=state name	During the agent monitor function, the agent expects the service to be either in the stopped or running state. The agent accordingly declares that the resource is offline or in the running state. If the service is not in the stopped or started state, this error is logged and the resource goes into an Unknown state. Solution: Probe the resource or wait for the next agent monitor cycle.
Failed to get the password attribute. Error = %d	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 = Error code	This is a VCS One internal error. Solution: Contact Symantec Technical Support.
Failed to open the service <i>service name</i> . Error = Error code.	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 = Error code	The agent failed to query the state of the service. Solution: Check whether the service is present in the Service Control Manager.

Table B-1 SQL Server agent error messages

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 = Error code.	The agent failed to stop the service. Solution: See the associated Windows error code for more information.
The SQL script has failed with error code.	The SQL script for detail monitoring failed. Solution: See the associated Windows error code for more information.
Error occurred while getting the process exit code. Error = Error code	This is a VCS One internal error. Solution: Contact Symantec Technical Support.
WaitForSingleObject failed with error code	This is a VCS One internal error. Solution: Contact Symantec Technical Support.
Failed to start the SQL script (User = user name, Domain = domain name) Error = Error code	The agent failed to execute the script for detail monitoring. Solution: See the associated Windows error code for detail monitoring.
Unable to convert the buffer to UNICODE Error = Error code	This is a VCS One internal error. Solution: Contact Symantec Technical Support.
Sql script failed. Script output : output	The SQL script failed to monitor the SQL Server instance. Solution: See the script output for more information. (The script output is located in %vcsone_home%\log\SQL<InstanceName>-TMP-OUT.tml")

Table B-1 SQL Server agent error messages

Message	Description
Failed to get the temporary file path. Error : Error code	This is a VCS One internal error. Solution: Contact Symantec Technical Support.
Failed to create the temporary file. Error = Error code.	This is a VCS One internal error. Check whether the log folder (%vcsone_home%\log) has Read-Write permission for the user name configured in the user name attribute. If the error occurs inspite of the Read-Write permission for the user name configured in the user name attribute, contact Symantec Technical Support.
Failed to read the temporary file. Error = Error code	This is a VCS One internal error. Check whether the log folder (%vcsone_home%\log) has Read-Write permission for the user name configured in the user name attribute. If the error occurs inspite of the Read-Write permission for the user name configured in the user name attribute, contact Symantec Technical Support.
Failed to remove the virtual name environment for the service <i>service name</i>	This is a VCS One internal error. Solution: Contact Symantec Technical Support.
Failed to add SQL domain group account to domain group account. Error = error code	The agent failed to add the domain group account to the SQL group. Verify if the domain group configured in the DomainGroupAttribute is valid and refer to the associated Windows error code for more information.
Failed to start the 'Veritas Cluster Server One Client Helper Service'	The agent to start the VCSOne remoting service Solution: Attempt to start the VCSOne Remoting service.

Table B-1 SQL Server agent error messages

Message	Description
Failed to look up Service Account 'Service Account' in domain group 'domain group'. Error: Error code	The agent failed to look for the SQL Service account in the configured domain group account attribute value. Solution: Verify the SQL Server service user account for the configured instance is a member of the Group Value configured in the DomainGroupAccount Attribute.

VCS One logging

VCS One 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 %VCSONE_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 One database agent for Microsoft SQL, mnemonic is VCSOne.
- Severity denotes the seriousness of the message. Severity of the VCS One error messages is classified into the following types:
 - CRITICAL indicates a critical error within a VCS One process. Contact Technical Support immediately.
 - ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS One action.
 - WARNING indicates a warning or error, but not an actual fault.
 - NOTE informs the user that VCS One has initiated an action.
 - INFO informs the user of various state messages or comments.
 - Of these, CIRITCAL, 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 One products is 'V-16.<number>' Category ID for SQLServer agent is 20020. Message ID is a unique number assigned to the message text.
- Message text denotes the actual message string.
You can view the 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 ONE NOTICE V-16-20020-10 SQLServer2000 :SQLGrp-SQLServer2000:monitor:SQL Server Instance name is not specified. Agent will operate on the default instance.

Index

A

- agent functions
 - SQL Server 2005 agent 13
 - SQL Server 2008 Agent Service agent 18
 - SQL Server 2008 Analysis Service agent 18
- attributes
 - GenericService agent 18
 - SQL Server 2005 agent 14

C

- converting
 - moving data files and databases (HA) 84
 - standalone SQL conversion to HA 81

D

- database
 - new database for SQL Server 2008 89
 - user-defined for SQL Server 2008 89
- database agent
 - installing 36
- disk groups, importing 53

G

- GenericService agent
 - attributes 18

M

- multiple SQL Server instances
 - assigning the port 61

P

- port assignment for multiple instances 61

R

- resource type definition
 - SQL Server FILESTREAM agent 108

S

- sample configurations
 - standalone SQL Server conversion (HA) 82
- SFW
 - creating volumes 53
 - importing disk groups 53
- SQL Server 2005 agent
 - attributes 14
 - functions 13
- SQL Server 2008
 - creating a new database 89
 - overview of user-defined database 89
- SQL Server 2008 Agent Service agent
 - functions 18
- SQL Server 2008 Analysis Service agent
 - functions 18
- SQL Server database agent
 - monitoring options 23
 - supported operating systems 47
 - supported versions 47
- SQL Server service group
 - modifying configuration 87
- standalone SQL Server
 - configuring storage hardware and network 82
 - installing and configuring SFW (HA) 83
 - overview of HA configuration 81
 - sample HA configuration 82
 - supported software 47

