Veritas™ Cluster Server Implementation Guide for Microsoft SQL Server 2012

Windows Server 2008 (x64), Windows Server 2008 R2 (x64)

6.0

January 2013

Veritas Cluster Server Agent for SQL Server

6.0

6.0.0

Legal Notice

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

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

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

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

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

- Product release level

- Hardware information
- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
    - Error messages and log files
    - Troubleshooting that was performed before contacting Symantec
    - Recent software configuration changes and network changes

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Contents

Chapter 3 Installing SQL Server in a VCS environment 51

Chapter 4 Configuring the SQL Server service group 67

Chapter 5 Configuring an MSDTC service group 85

Chapter 1

Introducing the VCS agents for SQL Server and NetApp

This chapter includes the following topics:

About the VCS agents for SQL and NetApp

The VCS agent for Microsoft SQL Server provides high availability to SQL Server. The VCS hardware replication agent for NetApp SnapMirror enables configuring NetApp filers over an iSCSI or Fibre Channel (FC) connection in a VCS cluster environment. Both the agents work together to provide high availability and disaster recovery to SQL Server in environments that use NetApp filers for shared storage. The agents also support disaster recovery configurations set up using the VCS Global Cluster Option and NetApp SnapMirror for data replication.

In a typical configuration, the agents are installed on each node in the cluster. The nodes are connected to the NetApp filers through a dedicated (private) storage network. VCS nodes are physically attached to the NetApp filer via an ethernet cable supporting iSCSI or FC as the transport protocol.

Figure 1-1 shows a typical VCS configuration in a NetApp storage environment.

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

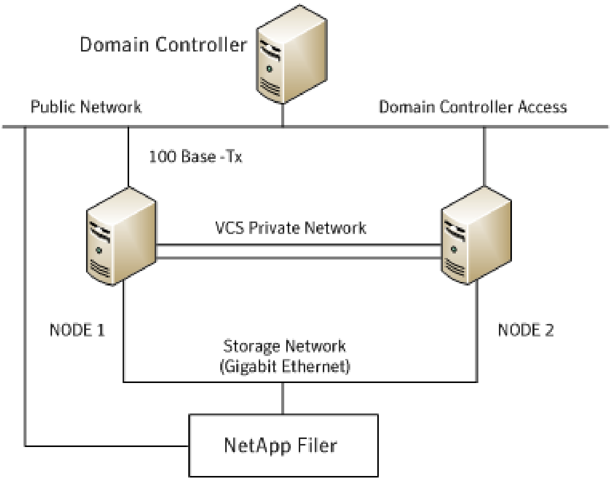

This chapter provides an overview of the agents. For more information about the agents, refer to their resource type definitions and attribute definitions.

About the VCS hardware replication agent for NetApp

The VCS hardware replication agent for NetApp provides failover support and recovery in environments employing NetApp filers for storage and NetApp SnapMirror for replication.

The agent monitors and manages the state of replicated filer devices and ensures that at a time only one system has safe and exclusive access to the configured devices.

The agent can be used in local clusters, single VCS replicated data clusters, and multi-cluster environments set up using the VCS Global Cluster Option (GCO).

The VCS agents for NetApp are as follows:

- NetAppFiler agent
- NetAppSnapDrive agent
- NetAppSnapMirror agent

About the NetApp Filer agent

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

NetApp Filer agent function

The NetApp Filer agent function is as follows:

Monitor	Performs the following tasks: ■ Verifies the state of the filer attached to the host by sending an ICMP ping command to the filer. If the filer does not respond, the agent reports the state of the filer as faulted. ■ Opens a filer connection and checks if ONTAPI version is supported by the filer. If the connection fails or the ONTAPI version is not supported, the agent reports the state as offline.

NetAppFiler agent resource type definition

The NetApp Filer agent is configured as a resource of type NetAppFiler.

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

NetAppFiler agent attribute definitions

Table 1-1 describes the NetApp Filer agent attributes.

Table 1-1 NetApp Filer agent attributes

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

About the NetApp SnapDrive agent

The NetApp SnapDrive agent monitors, connects, and disconnects filer volumes. You can configure the agent to use the iSCSI or the FC protocol.

NetApp SnapDrive agent functions

The NetApp SnapDrive agent functions are as follows:

Online	Connects a virtual disk (LUN) using an iSCSI or an FC initiator. The agent presents the LUN as a locally-attached drive to the host. The agent also removes LUN-host mappings made before the online operation.
Offline	Disconnects the virtual disk (LUN) from the host.
Monitor	Verifies that the specified virtual disk (LUN) is connected to the host.
Open	Verifies that there is connectivitiy to the filer. It also checks that the VCS Helper service is running with the same privileges as the SnapDrive service.
Clean	Attempts to forcibly disconnect a virtual disk (LUN).

NetAppSnapDrive agent resource type definition

NetApp SnapDrive agent is configured as a resource of type NetAppSnapDrive.

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

NetAppSnapDrive agent attribute definitions

Table 1-2 describes the NetApp SnapDrive agent attributes.

Table 1-2 NetApp SnapDrive agent attributes

Attribute	Description
FilerResName	Name of the VCS NetAppFiler-type resource in the service group. Type and dimension: string-scalar
VolumeName	Name of the volume containing the virtual disk. Define the volume name in the same case as on the filer. Type and dimension: string-scalar
ShareName	Name of the CIFS share containing the virtual disk. This attribute is ignored if NetApp SnapDrive version 6.0 is used. Type and dimension: string-scalar
LUN	Name of the LUN (virtual disk) on the filer that is presented to the host for mounting. Define the LUN name in the same case as on the filer. Type and dimension: string-scalar
MountPath	Drive letter to be assigned to the virtual disk. Type and dimension: string-scalar
Initiator	Name of iSCSI or FC initiator the host uses to connect virtual disks. You can retrieve this value from the Disk Management console. Type and dimension: string-vector

About the NetApp SnapMirror agent

The NetApp SnapMirror agent monitors the replication state of filer devices. When a failover occurs, the agent reverses the direction of replication. The agent supports the replication modes supported by NetApp. The agent supports asynchronous, semi-synchronous, and synchronous modes of replication. You can set the mode of replication using the SyncMode agent attribute.

NetApp SnapMirror agent functions

The NetApp SnapMirror agent functions are as follows:

Online	If the state of the local filer device is SOURCE, the agent creates a lock file to indicate that the resource can come online. This effectively makes the devices writable for the application. If the state of the local filer is SNAPMIRRORED, the agent attempts to reverse the direction of replication by changing the state of the local filer to SOURCE and that of the original source to SNAPMIRRORED. If the original source filer is down, the agent performs a mirror breakoff to enable local write access, if the filer is not already broken off. If the original source returns to life, you must resynchronize the data manually. The online function touches a lock file if read-write access is enabled successfully.
Offline	Removes the lock file. The agent does not perform any filer operations because an offline entry point does not necessarily indicate an intention to give up the devices.
Monitor	Verifies that the lock file exists. If the lock file exists, the monitor function reports the status of the resource as online. If the lock file does not exist, the monitor function reports the status of the resource as offline.
Open	Removes the lock file thereby preventing potential concurrency violation if the group fails over to another node. **Note:** The agent does not remove the lock file if the agent is started after an `hastop -force` command.
Clean	Removes the lock file. No filer operations are performed as taking the resource offline does not indicate a pending role swap.

Action function

Use the Action function to perform predefined actions on a resource. To perform an action on a resource, type the following command:

```
hares -action <SnapMirror_resname> <token> [-actionargs <arg1> ...]
[-sys <system>] [-clus <cluster> ]
```

Table 1-3 lists the action supported by the NetAppSnapMirror agent.

Table 1-3 Actions supported by NetAppSnapMirror agent

Token for Action	Description
fbsync	Resynchronises an original source volume with a broken-off volume. After synchronization, the original source volume becomes the target volume. The broken-off volume was initially the target volume, but was broken off as a result of a take over.

To synchronize volumes, type the following at the command prompt:

```
hares -action SnapMirror_resname fbsync -sys node_name
```

Where, *SnapMirror_resname* represents the name of the SnapMirror resource and *node_name* represents the node on which the service group is online.

Run the action for each SnapMirror resource.

You can also add custom actions for the agents. Refer to the *Veritas Cluster Server Agent Developer's Guide* for more information.

NetAppSnapMirror agent resource type definition

NetApp SnapMirror agent is configured as a resource of type NetAppSnapMirror.

```
type NetAppSnapMirror (
     static keylist SupportedActions = { fbsync }
     static int MonitorInterval = 300
     static int NumThreads = 1
     static i18nstr ArgList[] = { FilerResName,
     "FilerResName:FilerName",
     "FilerResName:StorageIP",VolumeName, SnapMirrorArguments,
     SnapMirrorSchedule, AppResName, VisibilityFrequency, SyncMode }
     str FilerResName
     str VolumeName
     str SnapMirrorArguments
     str SnapMirrorSchedule
     str AppResName
     int VisibilityFrequency = 180
     str SyncMode = async
 )
```

NetAppSnapMirror agent attribute definitions

Table 1-4 describes the NetApp SnapMirror agent attributes.

Table 1-4 NetApp SnapMirror agent attributes

Attribute	Description
FilerResName	Name of the VCS NetAppFiler-type resource in the group. Type and dimension: string-scalar
VolumeName	Name of the filer volume containing the virtual disk. This is the volume that is to be mounted. Define the volume name in the same case as on the filer. Type and dimension: string-scalar
SnapMirrorArguments	Specifies the SnapMirror arguments such as maximum transfer speed and restart mode. Type and dimension: string-scalar
SnapMirrorSchedule	Specifies the schedule the destination uses for updating data. Do not assign a value for this attribute if you use SnapManager. The schedule is in the following format: **minute hour dayofmonth dayofweek** Each field is separated by a space. Refer to the NetApp documentation for more details on the rules for each of these schedule fields. By default, this attribute does not have any value. Type and dimension: string-scalar
AppResName	Name of the resource configured to monitor the application being made highly available. Type and dimension: string-scalar
SyncMode	Specifies the mode of replication for the mirror. This attribute can have the following values: ■ async: Indicates that the mirror should be configured in the asynchronous mode. ■ semi-sync: Indicates that the mirror should be configured in the semi-synchronous mode. ■ sync: Indicates that the mirror should be configured in the synchronous mode. The default is async (asynchronous) mode. Type and dimension: string-scalar

Table 1-4 NetApp SnapMirror agent attributes *(continued)*

Attribute	Description
VisibilityFrequency	Specifies how often the source snapshot will be visible on the destination mirror. It controls the value of visibility_interval in the snapmirror.conf file. The default value is 180 seconds. This attribute is applicable only if the mirror is configured in synchronous or semi-synchronous mode. Type and dimension: string-scalar

About the VCS agent for Microsoft SQL Server 2012

The VCS agent for Microsoft SQL Server provides high availability for SQL Server in a VCS cluster. The agent monitors the SQL Server instance and its services on a VCS cluster to ensure high availability.

The VCS agent package for SQL Server 2012 includes the following:

Agent for SQL Server 2012 Database Engine	The agent provides high availability for SQL Server Database Engine. If the SQL Server Database Engine service is not running, the agent returns a failure status and declares the state as offline.
Agent for SQL Server FILESTREAM	The agent provides high availability for the SQL Server FILESTREAM feature. The agent monitors the Windows FILESTREAM configuration settings for the SQL Server instance.
GenericService agent for SQL Server Agent service and Analysis service	VCS employs the GenericService agent to provide high availability for the SQL Server Agent service and the Analysis service. The VCS GenericService agent monitors the SQL Server Agent and Analysis service. If the services are not running, the agent declares the services as offline.
Agent for SQL Server MSDTC	The VCS database agent for MSDTC provides high availability for the Microsoft Distributed Transaction Coordinator (MSDTC) service used in distributed transactions. The MSDTC agent monitors the MSDTC service to detect failure. The agent detects an MSDTC failure if the MSDTC service is not running.

About the agent for SQL Server 2012 Database Engine

This VCS agent for SQL Server 2012 monitors the SQL Server Database Engine service. The agent brings the SQL Server 2012 service online, monitors the status, and takes it offline.

Specific agent functions include the following:

Online	Brings the SQL Server service online.
Offline	Takes the SQL Server service offline.
Monitor	Queries the Service Control Manager (SCM) for the status of SQL Server services. Also, if detail monitoring is configured, the agent performs a database health check depending on the configuration.
Clean	Forcibly stops the SQL Server service.

Resource type definition for SQL Server 2012 Database Engine agent

The agent for SQL Server 2012 is configured as a resource of type SQLServer.

```
type SQLServer (
static int IMF{} = { Mode=3, MonitorFreq=5,
RegisterRetryLimit=3 }
static i18nstr IMFRegList[] = { Instance }
static keylist LogDbg = { DBG_1 }
static int AgentReplyTimeout = 999999999
static i18nstr ArgList[] = { Instance,
"LanmanResName:VirtualName", SQLOnlineTimeout,
SQLOfflineTimeout, Username, Domain, Password,
"IPResName:Address", DBist, SQLFile, FaultOnDMFailure,
SQLDetailMonitorTimeout, SQLClusterAccount }
str Instance
str LanmanResName
str IPResName
int SQLOnlineTimeout = 90
int SQLOfflineTimeout = 90
i18nstr Username
i18nstr Domain
str Password
i18nstr DBList[]
i18nstr SQLFile
boolean FaultOnDMFailure = 1
str SQLClusterAccount
int SQLDetailMonitorTimeout = 30
)
```

Attribute definitions for VCS agent for SQL Server 2012

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

Table 1-5 describes the required attributes associated with the VCS agent for SQL Server 2012 Database Engine.

Table 1-5 SQL Server 2012 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. Type and dimension: string-scalar
LanmanResName	Lanman resource name on which the SQLServer resource depends Type and dimension: string-scalar
SQLOnlineTimeout	Number of seconds that can elapse before the online entry point aborts Default = 90 Type and dimension: integer-scalar
SQLOfflineTimeout	Number of seconds that can elapse before the offline entry point aborts Default = 90 Type and dimension: integer-scalar

Table 1-6 describes the optional attributes associated with the VCS agent for SQL Server 2012 Database Engine.

Table 1-6 SQL Server 2012 agent optional attributes

Optional attributes	Definition
LevelTwoMonitorFreq	Defines whether the agent performs detail monitoring of SQL Server database. If set to 0, the agent only performs the basic monitoring of the instance service. A non-zero value indicates the number of online monitor cycles that the agent must wait before performing detail monitoring. **Note:** You can either configure script-based detail monitoring or DBList-based detail monitoring. In either case, the attributes Username, Password, and Domain must be assigned appropriate values. Default = 5 Type and dimension: integer-scalar **Note:** This is not a SQL Server agent-specific attribute, but a common type-level attribute. The value of this attribute can only be set through the wizard. If you configure the service group manually, you need to remember to specify its value manually. See "Configuring the detail monitoring frequency using the CLI" on page 31.
FaultOnDMFailure	Defines whether the agent fails over the service group if the detail monitoring script execution fails. The value 1 indicates that the agent fails over the service group if detail monitoring script fails to execute. The value 0 indicates that it does notfail over, but goes into the UNKNOWN state. Default = 1 Type and dimension: boolean
SQLDetailMonitor Timeout	Number of seconds that can elapse before the detail monitor routine aborts. Default = 30 Type and dimension: integer-scalar

Table 1-6 SQL Server 2012 agent optional attributes *(continued)*

Optional attributes	Definition
Username	The Microsoft Windows authentication name when logging in to a database for detail monitoring. This attribute must not be null if the LevelTwoMonitorFreq attribute is set to a non-zero value. The user must have the necessary privileges to connect to the database and execute the appropriate query. **Note:** This attribute can take localized values. Type and dimension: string-scalar
Domain	Domain for the user account. This attribute is used to create a trusted connection to the SQL Server instance if the LevelTwoMonitorFreq attribute is set to a non-zero value. **Note:** This attribute can take localized values. Type and dimension: string-scalar
Password	Password for logging in to a database for in-depth monitoring. This attribute must not be null if the LevelTwoMonitorFreq attribute is set to a non-zero value. Type and dimension: string-scalar
SQLFile	The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the LevelTwoMonitorFreq attribute is set to a non-zero value and script-based detail monitoring is configured. **Note:** This attribute can take localized values. Type and dimension: string-scalar
DBList	List of databases for which the agent will perform detail monitoring. **Note:** This attribute can take localized values. Type and dimension: string-vector

Table 1-6 SQL Server 2012 agent optional attributes *(continued)*

Optional attributes	Definition
SQLClusterAccount	Use this attribute if the user account that you specify for the SQL Server service and the SQL Server Agent service is not a member of the local Administrators group on all the cluster nodes that are part of the service group. Specify a domain group or the SQL Server service name. If you specify a domain group, then the SQL service account must be part of this domain group. The agent assigns the account with Full Control privileges to the SQL Server databases and log files. 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.
IPResName	The IP resource on which the Lanman resource for the SQLServer resource depends. Type and dimension: string-scalar

About the agent for SQL Server FILESTREAM

FILESTREAM in SQL Server enables SQL-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. The agent for SQL Server FILESTREAM enables FILESTREAM, monitors the status, and disables it.

The agent makes FILESTREAM highly available in a clustered environment.

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 or if FILESTREAM is disabled on the node, the FILESTREAM resource in the service group faults.

Resource type definition the SQL Server FILESTREAM agent

The agent for SQL Server FILESTREAM is configured as a resource of type SQLFilestream.

```
type SQLFilestream (
static i18nstr ArgList[] = { InstanceName }
str InstanceName
)
```

Attribute definitions the SQL Server FILESTREAM agent

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

Table 1-7 SQL Server Filestream agent required attributes

Required attributes	Definition
InstanceName	The name of the SQLServer resource to which the FILESTREAM is bound. If this attribute is blank, the agent monitors the default SQL server instance (MSSQLSERVER). Type and dimension: string-scalar **Note:** This attribute can take localized values.

About the agent for SQL Server Agent and Analysis services

VCS uses the GenericService agent to make the SQL Server Agent service and Analysis service highly available. The GenericService agent brings these services online, monitors their status, and takes them offline.

Specific agent functions include the following:

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

Clean	Forcibly stops the configured SQL Server services.

Refer to *Veritas Cluster Server Bundled Agents Reference Guide* for more information about the GenericService agent.

About the agent for MSDTC service

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

Specific agent functions include the following:

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

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

Resource type definition for MSDTC agent

The MSDTC agent is configured as a resource of type MSDTC.

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

Attribute definitions for MSDTC agent

Review the following information to familiarize yourself with the agent attributes for an MSDTC resource type.

Table 1-8 MSDTC agent required attributes

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

SQL Server sample dependency graph

A sample configuration graphically depicts the resources and their dependencies within the service group. The following example illustrates a typical service group configured to make SQL Server highly available in a VCS cluster.

The shared disk group is configured using the NetApp Filer (NetAppFiler) resource. The virtual name for the SQL Server is created using the Lanman resource. The service group IP address for the SQL Server is configured using the IP and NIC resources. The NetApp SnapDrive mount point is created using the NetAppSnapDrive resource. SQL Server registry is replicated using the RegRep and RegRepNetAppSnapDrive resources. The FileStream resource monitors the Windows FILESTREAM configuration settings for the SQL Server instance. The SQL Server resource comes online after each of these resources are brought online. The SQL Server Analysis service (MSOlap) and SQL Server Agent service (SQLServerAgent) are configured as GenericService resources.

Figure 1-2 shows the dependencies in the SQL Server service group.

Figure 1-2 SQL Server service group dependency graph

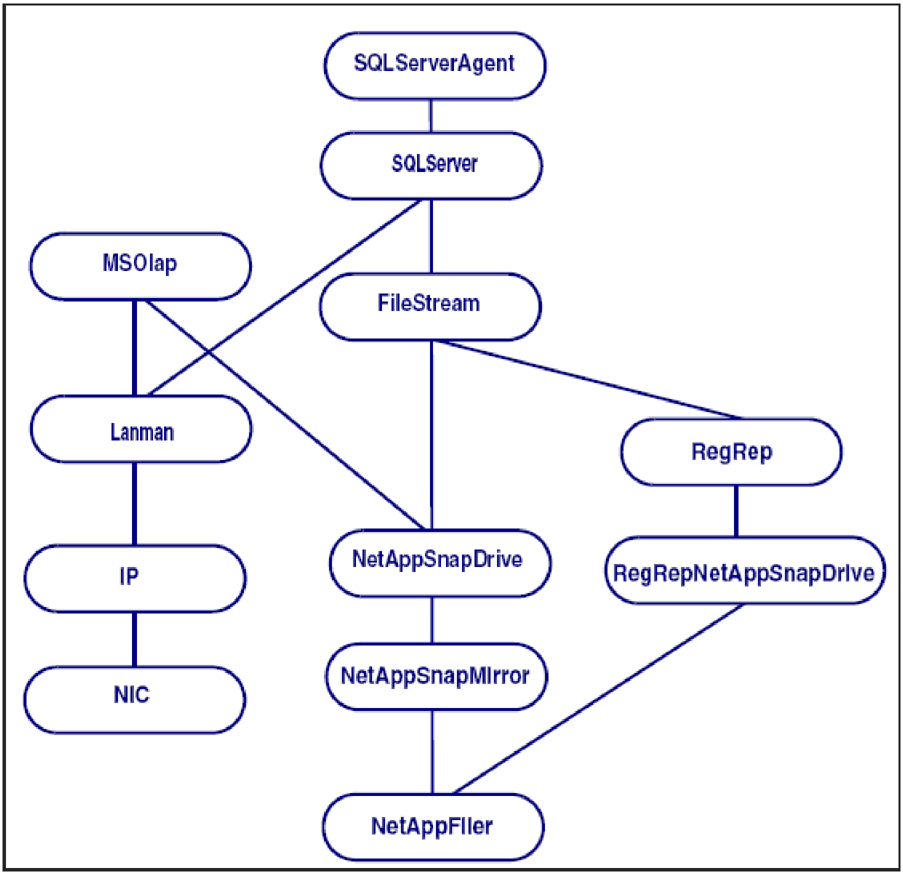

MSDTC sample dependency graph

A sample configuration graphically depicts the resources and their dependencies within the service group. The following example describes a typical MSDTC service group configured to monitor the state of the MSDTC services in a VCS cluster.

In the sample configuration shown in the dependency graph below, the shared disk group is configured using the Volume Manager Diskgroup (VMDg) resource. The virtual name for the MSDTC Server is created using the Lanman resource. The service group IP address for the MSDTC Server is configured using the IP and NIC resources. The MountV mount point is created using the MountV resource. MSDTC registry is replicated using the RegRep and RegRepMountV resources.

The MSDTC resource comes online after each of these resources are brought online.

Figure 1-3 shows the dependencies in the MSDTC service group.

Figure 1-3 MSDTC service group dependency graph

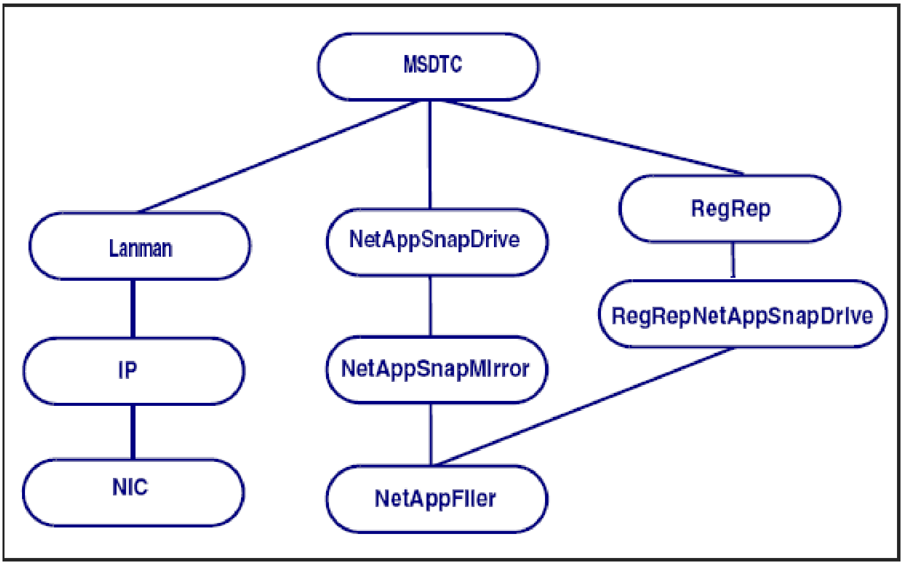

Detail monitoring options

The VCS agent for Microsoft SQL Server provides two levels of application monitoring: basic and detail. Basic monitoring queries the SCM to verify whether the SQL Server services are continuously active. Detail monitoring updates a temporary table in the SQL Server database to verify the availability of the database instance.

Table 1-9 Methods of configuring detail monitoring for SQL databases

Method	Description
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 is configured, the service group fails over to the failover nodes.
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 is configured, the service group fails over to the failover nodes. A sample script is provided with the agent for the purpose. You can customize the script to meet your configuration requirements. The script is located at: %*VCS_HOME*%\bin\SQLServer\sample_script.sql Here, %*VCS_HOME*% is the default installation directory for VCS, typically it is C:\Program Files\Veritas\Cluster Server. You should use a separate script for each SQL Server service group that exists in the cluster. The script should exist on all the nodes in the service group.

Note: If you provide input for both types of detail monitoring, DBList monitoring takes precedence, and SQL script-based monitoring is not performed.

You can enable and configure detail monitoring by running the SQL Server 2008 Agent Configuration Wizard for VCS. 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, then the SQL resource will go in an UNKNOWN state, because the VCS agent monitors the computer context of the services. If the SQL service is not started in the virtual server context the resource goes in an UNKNOWN state. You must ensure that you start all the SQL related services from within VCS.

Configuring the detail monitoring frequency using the CLI

The LevelTwoMonitorFreq attribute defines whether the agent performs detail monitoring of the SQL Server database. If set to 0, the agent only performs the basic monitoring of the instance service. A non-zero value indicates the number of online monitor cycles that the agent must wait before performing detail monitoring. If you configure a service group manually, you must also specify the value of this attribute manually. The following procedures describe how to manually configure the detail monitoring frequency using this attribute.

Note: To configure detail monitoring for SQL Server 2012, you must apply the VCS 6.0 Agent Pack Q2 2012 or later.

To enable the LevelTwoMonitorFreq option

◆ From any of the cluster nodes that have the agent pack installed, run the following commands sequentially:

```
haconf -makerw
hares -override SQLServer-MSSQLSERVER LevelTwoMonitorFreq
hares -modify SQLServer-MSSQLSERVER LevelTwoMonitorFreq 1
haconf -dump -makero
```

To remove the LevelTwoMonitorFreq option from the configuration

◆ From any of the cluster nodes that have the agent pack installed, run the following commands sequentially:

```
haconf -makerw
hares -undo_override SQLServer-MSSQLSERVER LevelTwoMonitorFreq
haconf -dump -makero
```

How the agents make SQL Server highly available

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

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

Local cluster configuration after a failover

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

Disaster recovery configuration after a failover

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

This involves the following steps:

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

Typical SQL Server configuration in a VCS cluster

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

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

Figure 1-4 Typical SQL Server configuration in a VCS cluster

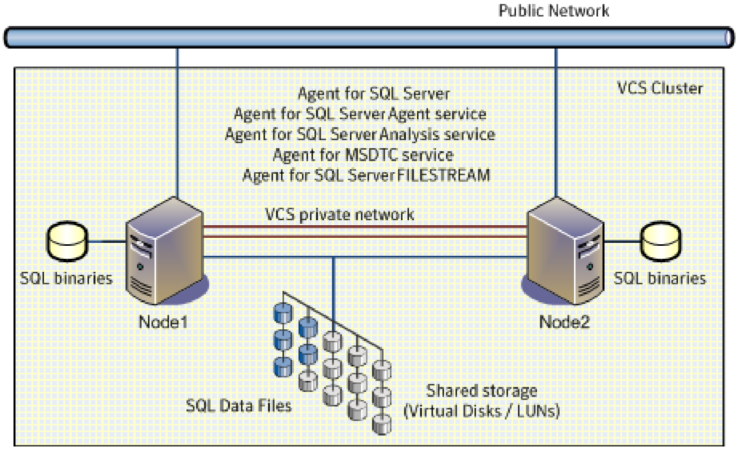

Chapter 2

Installing the VCS agents for SQL Server and configuring the cluster

This chapter includes the following topics:

- About installing the VCS agents for SQL Server
- Configuring the cluster using the Cluster Configuration Wizard

About installing the VCS agents for SQL Server

Install Veritas Cluster Server (VCS) on all the systems where you want to configure the application. During installation, the product installer installs the VCS agents required for making the applications highly available.

You must install the VCS agents before configuring the application with VCS.

Refer to the *Veritas Cluster Server for Windows Installation and Upgrade Guide* for instructions.

Configuring the cluster using the Cluster Configuration Wizard

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

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

Note the following prerequisites before you proceed:

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

  To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Symantec recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Symantec recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.
- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Symantec recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Symantec recommends that you disable TCP/IP from private NICs to lower system overhead.
- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which SQL will be installed and ensure that the public adapter is the first adapter in the Connections list.

  When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.
- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.

- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the VCS Helper Service user context to access the network. This account does not require Domain Administrator privileges.
- Make sure the VCS Helper Service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.

  Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.
- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

1. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** to start the VCS Cluster Configuration Wizard.
2. Read the information on the Welcome panel and click **Next**.
3. On the Configuration Options panel, click **Cluster Operations** and click **Next**.
4. On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.

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

   - Clear **Specify systems and users manually**.
   - Click **Next**.

     Proceed to step 8.

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

   - Select **Specify systems and users manually**.

     Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.

- Click **Next**.

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

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

   Do not specify systems that are part of another cluster.

   Proceed to step 8.

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

   Do not select systems that are part of another cluster.

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

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

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

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

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

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

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

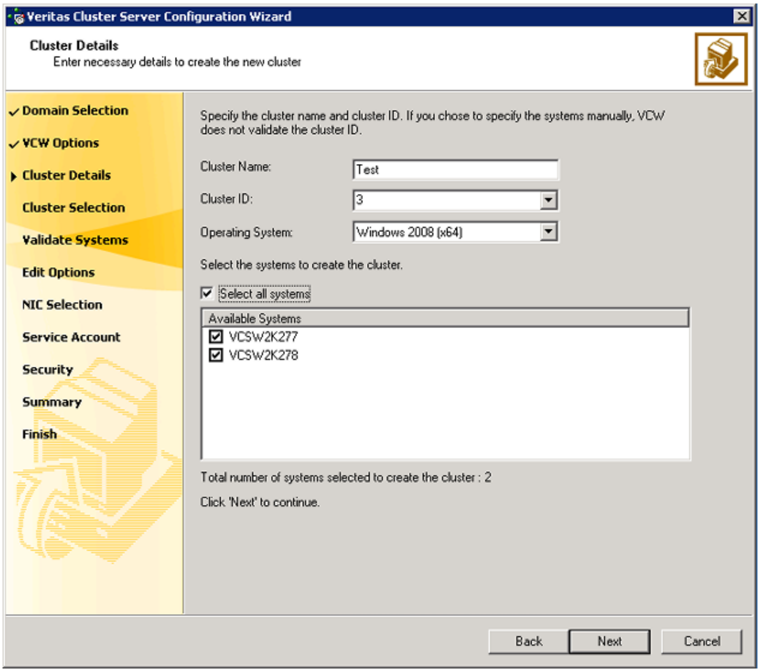

Specify the cluster name and cluster ID. If you chose to specify the systems manually, VCW does not validate the cluster ID.

Cluster Name: Test

Cluster ID: 3

Operating System: Windows 2008 (x64)

Select the systems to create the cluster.

☑ Select all systems

Available Systems
☑ VCSW2K277
☑ VCSW2K278

Total number of systems selected to create the cluster : 2

Click 'Next' to continue.

Specify the cluster details as follows:

Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 65535. **Caution:** If you chose to specify systems and users manually in step 4 or if you share a private network between more than one domain, make sure that the cluster ID is unique.
Operating System	From the drop-down list, select the operating system. The Available Systems box then displays all the systems that are running the specified operating system. All the systems in the cluster must have the same operating system and architecture. You cannot configure a Windows Server 2008 and a Windows Server 2008 R2 system in the same cluster.

Available Systems	Select the systems that you wish to configure in the cluster. Check the **Select all systems** check box to select all the systems simultaneously. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat.

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

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

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

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

Do one of the following:

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

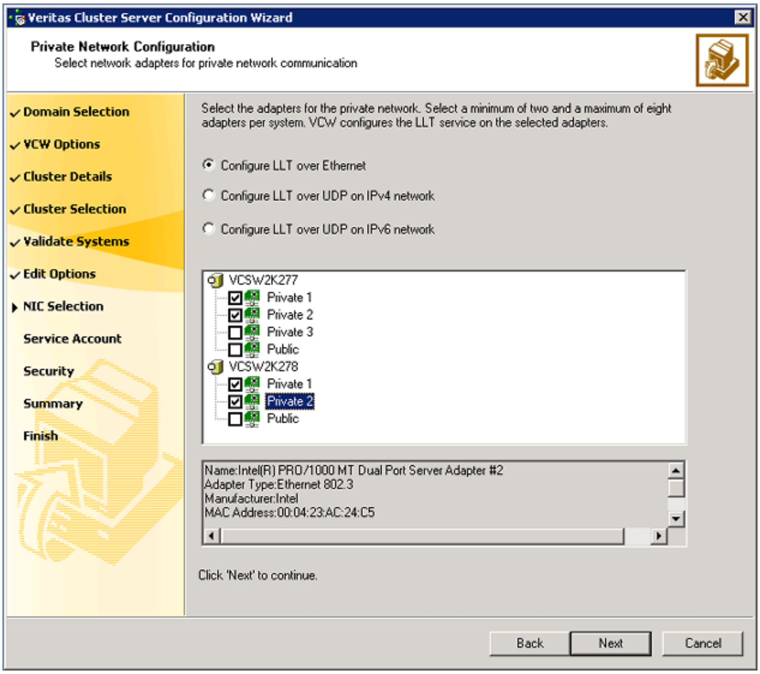

Private Network Configuration
Select network adapters for private network communication

Select the adapters for the private network. Select a minimum of two and a maximum of eight adapters per system. VCW configures the LLT service on the selected adapters.

● Configure LLT over Ethernet
○ Configure LLT over UDP on IPv4 network
○ Configure LLT over UDP on IPv6 network

VCSW2K277
- ☑ Private 1
- ☑ Private 2
- ☐ Private 3
- ☐ Public

VCSW2K278
- ☑ Private 1
- ☑ Private 2
- ☐ Public

Name:Intel(R) PRO/1000 MT Dual Port Server Adapter #2
Adapter Type:Ethernet 802.3
Manufacturer:Intel
MAC Address:00:04:23:AC:24:C5

Click 'Next' to continue.

- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links.

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

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

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

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

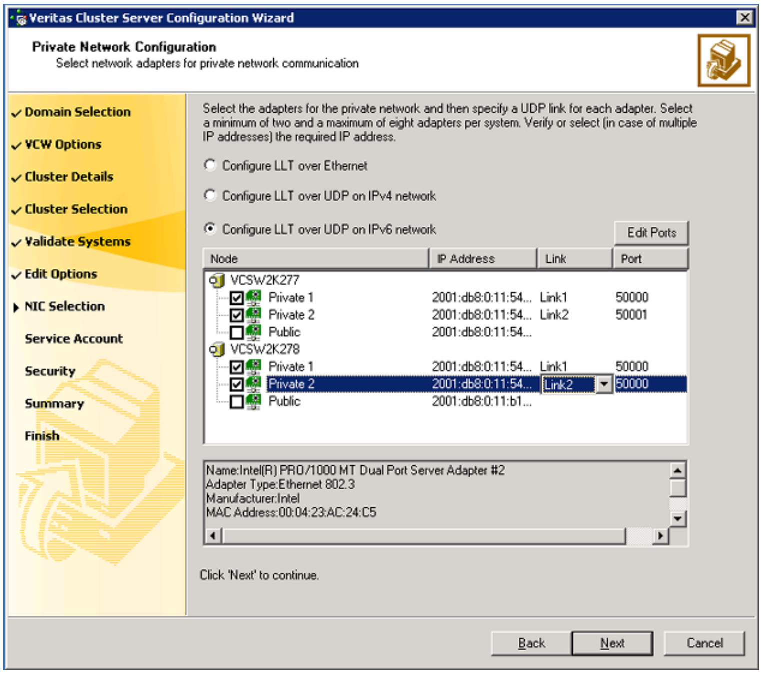

Private Network Configuration

Select network adapters for private network communication

Select the adapters for the private network and then specify a UDP link for each adapter. Select a minimum of two and a maximum of eight adapters per system. Verify or select (in case of multiple IP addresses) the required IP address.

○ Configure LLT over Ethernet

○ Configure LLT over UDP on IPv4 network

● Configure LLT over UDP on IPv6 network

Node	IP Address	Link	Port
VCSW2K277			
☑ Private 1	2001:db8:0:11:54...	Link1	50000
☑ Private 2	2001:db8:0:11:54...	Link2	50001
☐ Public	2001:db8:0:11:54...		
VCSW2K278			
☑ Private 1	2001:db8:0:11:54...	Link1	50000
☑ Private 2	2001:db8:0:11:54...	Link2	50000
☐ Public	2001:db8:0:11:b1...		

Name:Intel(R) PRO/1000 MT Dual Port Server Adapter #2
Adapter Type:Ethernet 802.3
Manufacturer:Intel
MAC Address:00:04:23:AC:24:C5

Click 'Next' to continue.

- Select **Configure LLT over UDP on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use.

  The IPv6 option is disabled if the network does not support IPv6.

- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Symantec recommends reserving two NICs exclusively for the VCS private network.

- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. In case of IPv4, each IP address can be in a different subnet.

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

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

For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.
The wizard configures the LLT service (over UDP) on the selected network adapters. The specified UDP ports are used for the private network communication.

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

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

Specify the domain user details as follows:

- To specify an existing user, do one of the following:
    - Click **Existing user** and select a user name from the drop-down list.
    - If you chose not to retrieve the list of users in step 4, type the user name in the Specify User field and then click **Next**.
- To specify a new user, click **New user** and type a valid user name in the Create New User field and then click **Next**.
  Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

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

Do one of the following:

- To use VCS cluster user privileges, click **Use VCS User Privileges** and then type a user name and password.
  The wizard configures this user as a VCS Cluster Administrator. In this mode, communication between cluster nodes and clients, including Cluster Manager (Java Console), occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password.
  The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.

Symantec recommends that you specify a new user name and password.

- To use the single sign-on feature, click **Use Single Sign-on**.
  In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The VCS high availability engine (HAD) and Veritas Command Server run in secure mode.
  The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

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

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

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

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

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

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

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

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
  See "Configuring notification" on page 45.

- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters.The WAC process is required for inter-cluster communication.

  Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.

  You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.

  See "Configuring Wide-Area Connector process for global clusters" on page 47.

Configuring notification

This section describes steps to configure notification.

To configure notification

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

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

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

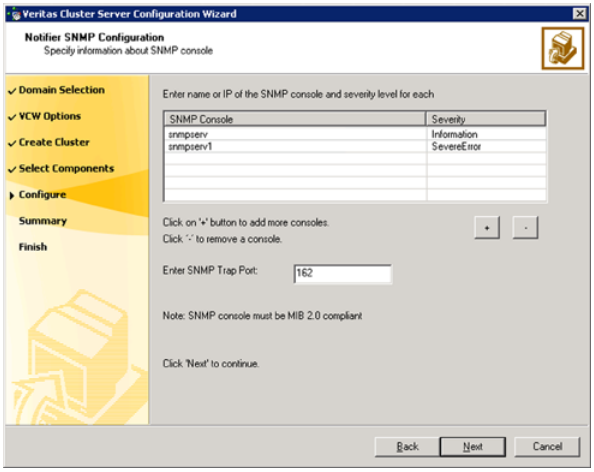

Notifier SNMP Configuration
Specify information about SNMP console

Enter name or IP of the SNMP console and severity level for each

SNMP Console	Severity
snmpserv	Information
snmpserv1	SevereError

Click on '+' button to add more consoles.
Click '-' to remove a console.

Enter SNMP Trap Port: 162

Note: SNMP console must be MIB 2.0 compliant

Click 'Next' to continue.

Do the following:

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

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

Do the following:

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

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

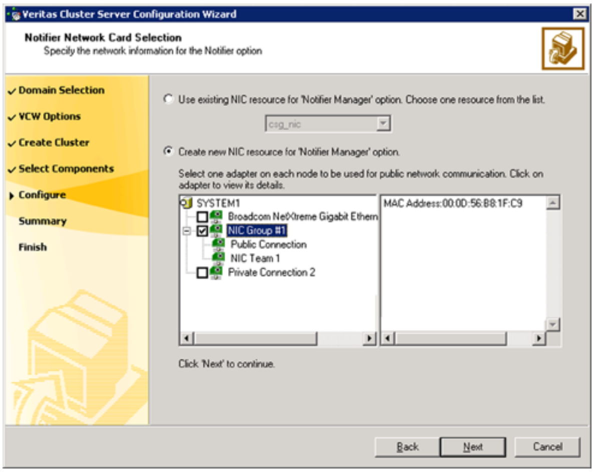

Veritas Cluster Server Configuration Wizard

Notifier Network Card Selection

Specify the network information for the Notifier option

○ Use existing NIC resource for 'Notifier Manager' option. Choose one resource from the list.

csg_nic

● Create new NIC resource for 'Notifier Manager' option.

Select one adapter on each node to be used for public network communication. Click on adapter to view its details.

SYSTEM1
- ☐ Broadcom NetXtreme Gigabit Ethern
- ☑ NIC Group #1
  - Public Connection
  - NIC Team 1
- ☐ Private Connection 2

MAC Address:00:0D:56:B8:1F:C9

Click 'Next' to continue.

Do the following:

- If the cluster has a ClusterService group configured, you can use the NIC resource configured in that service group or configure a new NIC resource for notification.
- If you choose to configure a new NIC resource, select a network adapter for each node in the cluster.

  The wizard lists the public network adapters along with the adapters that were assigned a low priority.

5. Review the summary information and choose whether you want to bring the notification resources online when VCS starts and click **Configure**.
6. Click **Finish** to exit the wizard.

Configuring Wide-Area Connector process for global clusters

Configure the Wide-Area Connector process only if you are configuring a disaster recovery environment. The GCO option configures the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication. Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.

You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.

To configure the wide-area connector process for global clusters

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

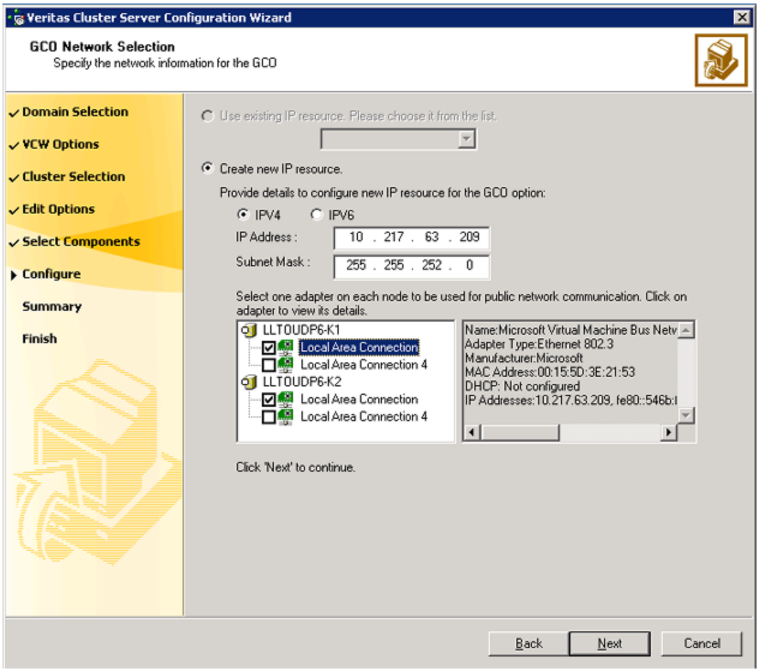

GCO Network Selection

Specify the network information for the GCO

○ Use existing IP resource. Please choose it from the list.

● Create new IP resource.

Provide details to configure new IP resource for the GCO option:

● IPV4 ○ IPV6

IP Address : 10 . 217 . 63 . 209

Subnet Mask : 255 . 255 . 252 . 0

Select one adapter on each node to be used for public network communication. Click on adapter to view its details.

- LLTOUDP6-K1
  - ☑ Local Area Connection
  - ☐ Local Area Connection 4
- LLTOUDP6-K2
  - ☑ Local Area Connection
  - ☐ Local Area Connection 4

Name:Microsoft Virtual Machine Bus Netv
Adapter Type:Ethernet 802.3
Manufacturer:Microsoft
MAC Address:00:15:5D:3E:21:53
DHCP: Not configured
IP Addresses:10.217.63.209, fe80::546b:I

Click 'Next' to continue.

If the cluster has a ClusterService group configured, you can use the IP address configured in the service group or configure a new IP address.

Do the following:

- To specify an existing IP address, select **Use existing IP resource** and then select the IP address from the drop-down list.
- To use a new IP address, do the following:
  - In case of IPv4, select **IPV4** and then enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
  - In case of IPv6, select **IPV6** and select the IPv6 network from the drop-down list.
    The wizard uses the network prefix and automatically generates a unique IPv6 address that is valid on the network.
    The IPv6 option is disabled if the network does not support IPv6.
- Select a network adapter for each node in the cluster.

The wizard lists the public network adapters along with the adapters that were assigned a low priority.

2. Review the summary information and choose whether you want to bring the WAC resources online when VCS starts and then click **Configure**.
3. Click **Finish** to exit the wizard.

Chapter 3

Installing SQL Server in a VCS environment

This chapter includes the following topics:

About installing SQL Server in a VCS environment

Installing and configuring SQL Server in a VCS environment involves the following tasks. This environment uses an active-passive configuration with one to one failover capabilities.

- Installing SQL Server on the first cluster node.
  You must install the SQL instance on the local disk and install the SQL database files and analysis service files on the shared storage.The shared storage must be accessible from all the nodes where you wish to install and configure SQL Server.
- Installing SQL Server on additional cluster nodes.

You must install the SQL instances, database files, and analysis service files on the local disk. You need not install the database files and the analysis service files on shared storage.
The advantage of this method is that while you are installing SQL Server on the first cluster node, you can run parallel installations on the remaining cluster nodes.

- Configuring the SQL Server service group using the SQL Server Configuration Wizard and bringing it online on the first node.
  You must run the wizard from the first cluster node where you installed the SQL Server database and analysis service files on shared storage. This is required as the wizard configures the resources for the SQL Server database and registry information installed on the shared storage and propagates this information to the remaining nodes that are part of the SQL service group. When the service group fails over to any additional cluster node, this information moves to that node and the SQL instance is brought online on that node.

For setting up an active-active SQL Server clustered environment refer to,

For setting up a disaster recovery clustered environment refer to,

See "Configuring a disaster recovery set up for SQL Server" on page 104.

Making a standalone SQL Server highly available refer to,

Prerequisites for installing SQL Server

Ensure the following before you install SQL Server:

- Verify that VCS is installed on all the systems where you want to install and configure SQL Server.
  Refer to the *VCS for Windows Install and Upgrade Guide* for more information.
- Verify that you have configured a VCS cluster using VCS Cluster Configuration Wizard (VCW).
  See "Configuring the cluster using the Cluster Configuration Wizard " on page 35.
- If using iSCSI, verify that the Microsoft iSCSI Initiator is configured to establish a persistent connection between the NetApp filer and the cluster nodes. See the Microsoft documentation for instructions.
- Verify that Microsoft Internet Information Services (IIS) is installed and enabled on the system. Otherwise, you will be prompted to install or enable IIS and then restart the installation.

- Ensure that you have created volumes on an external, basic disk, or LUNs (virtual disks) on a NetApp filer. Mount or connect them to the first cluster node where you will install SQL.
  Ensure that the shared disks or filer LUNs are accessible from all the cluster nodes where you will install SQL Server.

  Symantec recommends that you create volumes for the following:
  - SQL Server data
  - Registry replication
  - User defined database
  - User defined database logs
  - FILESTREAM enabled database objects

  See "Managing storage using NetApp filer " on page 54.
  See "Managing storage using Windows Logical Disk Manager " on page 57.
- If your cluster has an Exchange service group configured, make sure to install SQL Server on a node that is not in the SystemList attribute for the Exchange service group.

Privileges required for installing SQL Server

The following privileges are required for installing SQL Server:

- Ensure that the [NT AUTHORITY\SYSTEM] account is granted the sysadmin server role (from SQL Management Studio Console) on each node.
- The logged-on user must be a domain user with local Administrator privileges.
- The logged-on user must be a member of the local Administrators group on all nodes where you will install Microsoft SQL Server.
- The logged-on user or the VCS Helper Service user account must have write permissions for the Active Directory objects corresponding to these nodes.
- The logged-on user or the VCS Helper Service user account must have write permissions on the DNS server to perform DNS updates.

Configuring Microsoft iSCSI initiator

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

To configure Microsoft iSCSI initiator on a Windows Server 2008 system

1. Start the Microsoft iSCSI initiator.
2. On the Discovery tab, click **Add Portal**.
3. On the Add Target Portal dialog box, specify the DNS name for the NetApp filer and then click **OK**.
4. On the Targets tab, click **Log On**.
5. On the Log On to Target dialog box, clear the **Automatically restore this connection when the system reboots** check box and then click **OK**.
6. On the Targets tab, verify that the newly added portal is listed under the Select a target box and the status shows "connected". Click **OK**.

To configure Microsoft iSCSI initiator on a Windows Server 2008 R2 system

1. Start the Microsoft iSCSI initiator.
2. On the Discovery tab, click **Discover Portal**.
3. On the Discover Target Portal dialog box, specify the DNS name for the NetApp filer and then click **OK**.
4. On the Target tab, click **Connect**.
5. On the Connect to Target dialog box, clear the **Add this connection to list of Favorite Targets** check box and then click **Ok**.
6. On the Targets tab, verify that the newly added portal is listed under the Select a target box and the status shows "connected". Click **OK**.

Managing storage using NetApp filer

NetApp manages data by creating volumes on physical disks. These volumes can further be divided into LUNs (Logical Unit Numbers). The LUNs are accessible from the cluster nodes, provided the nodes have Microsoft iSCSI Initiator and NetApp SnapDrive installed. However, if you plan to use Fibre Channel (FC) for connecting the LUNs, ensure that filer is connected to the nodes and the LUNs are shared between all the cluster nodes.

Refer to the NetApp documentation for more information.

Figure 3-1 illustrates a typical VCS cluster in a NetApp storage environment.

Figure 3-1 VCS cluster in a NetApp storage environment

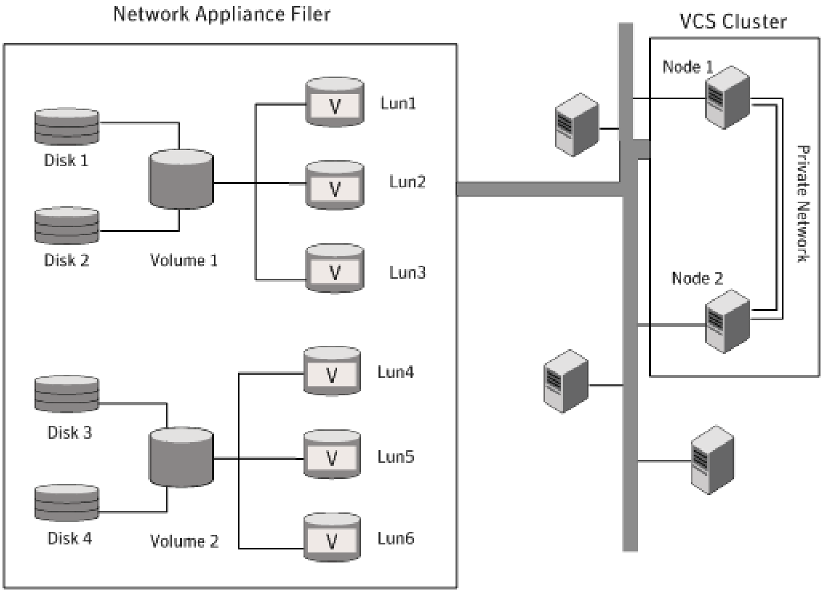

The VCS agent for Microsoft SQL requires two LUNs to be created on the NetApp filer, one for SQL Server data and the other for the registry replication information.

If you are using SQL Server FILESTREAM, create additional LUNs for FILESTREAM enabled database objects.

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

Symantec recommends that you create separate LUNs (virtual disks) for the following:

- INST1_DATA_FILES
  Contains the SQL Server system data files (including the master, model, msdb, and tempdb databases).
- INST1_REGREP_VOL
  Contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB (minimum recommended size) volume for this purpose.
- INST1_FS_VOL
  Contains FILESTREAM enabled database objects for the SQL database.
- INST1_DB1_VOL

Contains the user database files.

- INST1_DB1_LOG
  Contains the user database log files.
- INST1_DB1_FS_VOL
  Contains FILESTREAM enabled database objects for the user database

These LUNs must be accessible from all cluster nodes.

Perform the following tasks to create LUNs on the NetApp filer and to make them accessible from cluster nodes:

- Add the filer storage system to the SnapDrive Storage System Management snap-in on the cluster nodes.
- Create volumes on the NetApp filer.
- Share the volumes.
- Create LUNs or virtual disks on the shared volumes.
  Refer to NetApp documentation for instructions on performing these tasks.

Connecting virtual disks to the cluster node

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

To connect virtual disks to the cluster node

1. On the cluster node where you want to connect the LUN, click **Start > All Programs > Administrative Tools > Computer Management** to start the Computer Management MMC.
2. From the left pane, expand **Storage** and double-click **SnapDrive**.
3. Right-click **Disks** and then click **Connect Disk** to launch the Connect Disk wizard.
4. Click **Next** on the Welcome page.
5. Specify the path of the virtual disk that you wish to connect to the cluster node and then click **Next**.
6. Select **Dedicated** as the Virtual Disk Type and then click **Next**.
7. Click **Assign a Drive Letter** and then choose a drive letter from the drop-down list.
8. On the Select Initiator panel, specify the initiator(s) for the virtual disk and then click **Next**.

9. On the igroup Management Type panel, choose the option that allows SnapDrive to perform igroup management automatically and then click **Next**.
10. Click **Finish** to begin connecting the specified virtual disk to the cluster node.

Disconnecting virtual disks from the cluster nodes

Perform the following steps to disconnect the virtual disks from a cluster node.

To disconnect virtual disks

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

Managing storage using Windows Logical Disk Manager

If your configuration uses shared disks and volumes that are managed using Windows Logical Disk Manager (LDM), use the VCS DiskReservation (DiskRes) and Mount (Mount) agents. Before configuring shared storage, review the resource types and attribute definitions of the Disk Reservation and Mount agents described in the *Veritas Cluster Server Bundled Agents Reference Guide*.

Note the following restrictions before you proceed to manage the storage using LDM:

- Disk Reservation and Mount agents are supported on VCS for Windows only. These agents are not supported in an SFW storage environment.
- For using LDM, your storage devices must be configured to use SCSI-2 disk reservations. SCSI-3 is not supported.
- LDM support is not applicable for Disaster Recovery configurations. Currently only HA configurations are supported.

The VCS SQL Server agent requires that you create two volumes on the shared disk, one for SQL Server data and the other for the registry replication information.

If you are using SQL Server FILESTREAM, create additional volumes for FILESTREAM enabled database objects.

If you will plan to configure an MSDTC service group, create additional volumes for MSDTC log and MSDTC registry replication.

Symantec recommends that you create separate volumes for the following:

- INST1_DATA_FILES
  Contains the SQL Server system data files (including the master, model, msdb, and tempdb databases).
- INST1_REGREP_VOL
  Contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB (minimum recommended size) volume for this purpose.
- INST1_FS_VOL
  Contains FILESTREAM enabled database objects for the SQL database.
- INST1_DB1_VOL
  Contains the user database files.
- INST1_DB1_LOG
  Contains the user database log files.
- INST1_DB1_FS_VOL
  Contains FILESTREAM enabled database objects for the user database.

These volumes must be on the shared storage and should be accessible from all cluster nodes.

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

- Reserve disks.
  See "Reserving disks (if you use Windows LDM)" on page 58.
- Create volumes.
  See "Creating volumes (if you use Windows LDM)" on page 59.
- Mount volumes.
  See "Mounting volumes (if you use Windows LDM)" on page 59.

Reserving disks (if you use Windows LDM)

Complete the following steps to reserve the disks on the node on which you are going to perform the application installation.

To reserve the disks

1 To display all the disks, type the following on the command line:

```
C:\>havol -scsitest /l
```

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:\>havol -scsitest /RES:<disk #>
```

For example, to reserve disk #4, type:

```
C:\>havol -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 during installation and while configuring the service group, on additional nodes in the cluster.

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 cluster 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 (if you use Windows LDM)” on page 60.

3 Rescan the disks on all the remaining nodes in the cluster.

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 cluster node.

To mount a volume

1. Use the Windows Disk Management tool to mount the volumes that you created earlier.
2. After mounting the volumes on a cluster node, 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 cluster nodes.

   Refer to Microsoft Windows documentation for more information about the CHKDSK command and the Disk Management tool.

Unassigning a drive letter

While installing an application on multiple nodes, you must first unassign drive letters and release the disks from one node, and then reserve the disks, 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 at 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**.

Releasing disks (if you use Windows LDM)

Perform the following steps to release reserved disks from a cluster node.

To release disks

1. To display all the disks, type the following on the command line:

   ```
   C:\>havol -scsitest /l
   ```

   Make a note of the disk numbers (Disk# column in the table) of the disk that you wish to release. You will need it in the next step.

2. 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 may require these details to identify and reserve the disks later.

Installing SQL Server on the first cluster node

Run the Microsoft SQL Server installer to install SQL Server on the first cluster node. Refer to the Microsoft documentation for instructions.

Note the following requirements while installing and configuring SQL Server:

- Ensure that you have installed and configured VCS, on all the nodes on which you wish to install and configure SQL Server.
  Refer to *Veritas Cluster Server Installation and Upgrade Guide* for instructions.
  Ensure that you have installed VCS 6.0 Agent Pack Q2 2012 or later.
  For more information, see the *Veritas Cluster Server 6.0 Agent Pack Readme*.
- Make sure that the volumes or LUNs (virtual disks) required for SQL Server are mounted or connected to the first cluster node where you install SQL.
- Install SQL Server in the standalone installation mode in a non-clustered environment.
  From the SQL Server Installation Center, on the Installation panel, choose the **New SQL Server stand-alone installation or add features to an existing installation** option.
- While installing SQL Server, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure for high availability.
- Install the SQL Server instance on the local disk.
  On the SQL Server Installer's Instance Configuration panel, ensure that the Instance root directory resides on the local disk.
- Install the SQL instance data directories on the shared storage.

On the SQL Server Installer’s Database Engine Configuration panel, ensure that all the components displayed on the **Database Directories** tab reside on shared storage.

The components include the following:

- Data root directory
- User database directory
- User database log directory
- Temp DB directory
- Temp DB log directory
- Backup directory

- Install the SQL Server Analysis Services data directories on shared storage. On the SQL Server Installer’s Analysis Services Configuration panel, ensure that all the components displayed on the Data Directories tab reside on the shared storage.

  The components include the following:

  - Data directory
  - Log file directory
  - Temp directory
  - Backup directory

- 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, each instance must have a unique instance name and instance ID. Use the same instance name and instance ID when you install the SQL instances on the additional nodes.
- While specifying a user name for the SQL Server services account, specify a domain user account.
- If the domain user account specified for the SQL Server services is not part of the local Administrators group on all the SQL Server nodes, then you must configure the VCS SQL agent's SQLClusterAccount attribute while configuring the SQL Server service group later.
- Apart from the SQL Browser service, make sure that the other SQL Server services are not set to start at the end of the SQL installation. While installing SQL Server on the first node, set the startup type of all the SQL Server services

to manual. However, set the startup type of the SQL Server Browser service to automatic. You must do this only for the instance which you have installed.

You can change the services startup type either during the installation or using the SQL Server Configuration Manager after the installation. Refer to the Microsoft documentation for instructions.

Installing SQL Server on the additional cluster nodes

Run the Microsoft SQL Server installer to install SQL Server on the second or any additional cluster node. Refer to the Microsoft documentation for instructions.

Note the following prerequisites before installing SQL Server on the second or any additional failover nodes:

- Ensure that you have installed and configured VCS, on all the nodes on which you wish to install and configure SQL Server.
  Refer to the *Veritas Cluster Server Installation and Upgrade Guide* for instructions.
  Ensure that you have installed VCS 6.0 Agent Pack Q2 2012 or later.
  For more information, see the *Veritas Cluster Server 6.0 Agent Pack Readme*.
- Install SQL Server in the standalone installation mode in a non-clustered environment.

  From the SQL Server Installation Center, on the Installation panel, choose the **New SQL Server stand-alone installation or add features to an existing installation** option.
- While installing SQL Server on additional cluster nodes, install the SQL instance, the data directories, and Analysis Services data directories on a local disk. You do not have to install these files on the shared storage.

  If you choose to install the SQL database files to a shared storage, ensure that the shared storage location are not the same as that used while installing SQL on the first cluster node.

  Ensure that you do not overwrite the database directories created by the SQL installation on the first cluster node.
- While installing SQL Server, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure for high availability.
- While installing the SQL instances on additional nodes, ensure that you specify the same instance names and instance IDs that you used while installing the instances on the first cluster node.
- If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name and instance ID. Use the same instance name and instance ID when you install the SQL instances on the additional nodes.

- While specifying a user name for the SQL Server services account, specify a domain user account.
- If the domain user account specified for the SQL Server services is not a part of the local Administrators group on all the SQL Server nodes, then you must configure the SQLClusterAccount attribute while configuring the SQL Server service group later.

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

See Technote: http://support.microsoft.com/kb/823938/en-us.

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

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

Enabling IPv6 support for the SQL Server Analysis Service

This is applicable only if SQL Server is configured in an IPv6 network environment.

The SQL Analysis Services server properties, IPv4 Support and IPv6 Support, determine which protocol is used by the Analysis Server. You must manually modify these properties to enable IPv6 support for Analysis Service.

These steps are required only if you have configured named SQL Server instances. Perform the following steps for each named SQL Server instance. Repeat these steps on all the cluster nodes that will host the SQL service group.

To enable IPv6 support for SQL Server Analysis Service

1. Start the Analysis Service.
2. Open SQL Server Management Studio and connect to the Analysis Server.
3. In the Object Explorer pane, right-click the server to which you have connected and click **Properties**.
4. On the General page, check the **Show Advanced (All) Properties** check box.
5. Locate Network \ Listener \ IPV4Support property and in the Value field type **0**.

   This means that IPv4 is disabled. Analysis Server does not listen on the IPv4 port, and clients will not be able to connect using IPv4.
6. Locate Network \ Listener \ IPV6Support property and in the Value field type **2**.

   This means that IPv6 is optional. The Analysis Server tries to listen on the IPv6 port, but will silently ignore errors and continue to start if IPv6 is not available.
7. Click **OK** to save the changes.
8. Stop the Analysis Service.
9. Perform these steps for each named instance and on all the cluster nodes where SQL Server is installed.

Chapter 4

Configuring the SQL Server service group

This chapter includes the following topics:

About configuring the SQL service group

Configuring the SQL Server service group involves creating resources for the NetApp and SQL agents. VCS provides several ways of configuring a service group, which include the service group configuration wizard, Cluster Manager (Java Console), and the command line. This chapter provides instructions on configuring a SQL service group using the SQL Server Configuration Wizard.

A SQL service group is used to bring a SQL Server instance online on another node if the active node fails. If you have set up multiple cluster nodes, you specify the priority of the failover node while configuring the service group. The SQL Server Configuration Wizard enables you to create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

You use the VCS SQL Server 2008 Agent Configuration Wizard to configure a service group for SQL Server 2008, 2008 R2, or 2012. You can configure a service group for only one SQL Server version in a single wizard workflow. To configure another SQL Server version, you must run the wizard again.

Before configuring the SQL service group

Ensure the following before configuring the SQL service group:

- Verify that VCS, along with the VCS database agent for SQL Server, is installed on all the cluster nodes.
- Verify that you have configured a VCS cluster using VCS Cluster Configuration Wizard (VCW).
- Verify that SQL Server is identically installed on all the cluster nodes that will participate in the service group.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- The logged-on user account must be a local Administrator on the node where you run the wizard. If you wish to configure detail monitoring for a SQL instance, the logged-on user must have the permission to log on to that SQL instance.
- You must be an Administrator for the NetApp filer containing the LUNs created to store SQL Server components.
- Run the SQL Server 2008 Agent Configuration Wizard from the first cluster node where you installed SQL Server. Do not run the wizard from the additional nodes.

  This is required as the wizard configures the resources for the SQL Server database and registry information installed on the shared storage and propagates this information to the remaining nodes that are part of the SQL service group.
- Verify that the Veritas High Availability Engine (HAD) is running on the system from where you run the wizard.
- Verify that the volumes or LUNs (virtual disks) created to store the following data components are mounted or connected to the node where you run the wizard and dismounted or disconnected from other nodes in the cluster:
  - SQL Server system data files
  - Registry replication information.
  - User database files
  - User database log files

  - FILESTREAM database objects

  For creating a service group, this must be the first cluster node where you installed SQL Server.

- If you wish to configure high availability for FILESTREAM, ensure that FILESTREAM is configured and enabled for the SQL instance on the first cluster node where you installed SQL, and disabled on all the remaining nodes.

  Refer to the Microsoft SQL Server documentation for more information.

- In case of IPv4, assign a unique virtual IPv4 address to the SQL Server instance. You specify this IP address when configuring the service group.

  In case of IPv6, the configuration wizard automatically generates an IPv6 address based on the network selected. The IPv6 address is valid and unique on the network.

- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.

- If you wish to use a script for detail monitoring, either save the script file in shared storage or ensure that the same file exists in the same location on all the cluster nodes.

  A sample script is supplied in C:\Program Files\Veritas\cluster server\bin\SQLServer\sample_script.sql.

  A sample script is supplied in C:\Program Files\Veritas\cluster server\bin\SQLServer2008\sample_script.sql. The same script can be used to monitor SQL Server 2008 and SQL Server 2008 R2.

  If the script is successfully executed during monitoring, the agent considers the database instance available. If the execution fails, the database instance is considered not available and the service group faults and fails over to the failover nodes. You can customize the script to meet your configuration requirements.

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

- Make sure that the following services are stopped on the first cluster node where you are running the wizard:

  - SQL Server
  - SQL Server Agent

  - SQL Server Analysis Services

  Stop these services for the SQL instances that you wish to configure in the service group.

- Review the resource types and the attribute definitions of the agents.
- If you have configured Windows Firewall, add the following to the Firewall Exceptions list:
  - Port 14150 or the VCS Command Server service, `%vcs_home%\bin\CmdServer.exe`.

    Here, %vcs_home% is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
  - Port 14141

  For a detailed list of services and ports used by VCS, refer to the *Veritas Cluster Server for Windows Installation and Upgrade Guide*.

Configuring a SQL Server service group

This section describes how to configure a SQL service group.

To modify an existing service group.

The VCS SQL Server 2008 Agent Configuration Wizard is used to configure a service group for SQL Server 2008, 2008 R2, and 2012. You can configure a service group for only one SQL Server version at a time. To configure a service group for another SQL Server version, you must run the wizard again.

To create a SQL Server service group on the cluster

1. Ensure that you have stopped the SQL Server service for the instance and are running the wizard from the first cluster node.
2. Start the SQL Server 2008 Agent Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server 2008 Agent Configuration Wizard**.
3. Review the prerequisites on the Welcome panel and then click **Next**.
4. On the Options panel, select **Create service group** and then click **Next**.
5. On the Service Group Configuration panel, specify the service group name and system list.

   Complete the following:

- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
- To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.

  For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
- Click **Next**.

6 On the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that needs to be configured for high availability in your environment. Complete the following steps and then click **Next**.

- From the **SQL Server version** drop-down list, select the SQL Server version for which you wish to configure the service group.

  You can configure a service group for only one SQL Server version in a single wizard workflow. To configure another SQL Server version, you must run the wizard again.

  The wizard displays instances of the selected SQL Server version that satisfy the following criteria:
  - Instances installed identically on all the systems
  - Instances not configured in other SQL service groups
- Select the SQL Server instance(s) that you wish to configure in the service group.
- If required, select the other services that you wish to make highly available. These options are available for selection only if the corresponding services are installed.

  Note that you can choose only one instance of the Analysis service per service group. If you have selected an instance of Analysis service, you must uncheck it before you can select another instance of the Analysis service.

  Note that services that are already configured and online in the cluster appear in bold and are not available for selection. You have to offline the service group and run the wizard in the modify mode to edit the service resources.

- Select SQLFILESTREAM if you wish to configure high availability for FILESTREAM enabled database objects. The wizard configures a resource only if FILESTREAM is enabled for the instance on the current node. Note that FILESTREAM option will not appear for selection if it is not enabled on the node.
- Clear the **Configure NetApp SnapMirror Resource(s)** check box. This option is applicable only in case of a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. If you are setting up a disaster recovery environment, check this check box to configure the SnapMirror resource at the primary site. Note that you must configure the SnapMirror resource only after you have configured replication between the NetApp filers.

7. Click **Yes** on the dialog box that prompts you whether you wish to allow the wizard to reconfigure the database paths for the selected instances using the current cluster node as a reference.
8. On the User Databases List panel, view the summary of the databases for the selected instance and then click **Next**.

   In case of multiple instances, select the required instance from the SQL Instance dropdown list. The panel displays the databases and the respective files for which the wizard configures 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.
9. On the SQL Server Cluster Account Configuration panel, specify the SQL cluster account details and then click **Next**.

   The SQL Cluster account must be configured if the SQL Server service and the SQL Server Agent service accounts do not have local administrator privileges on all the SQL Server nodes in the service group.

   Complete the following steps for each SQL Server instance that you wish to configure in the service group:

   - Select a SQL instance from the **Instance Name box**.
   - Check the **Configure SQL Server Cluster Account** check box.

- Click **Use service SIDs to set the SQL Server service name as the SQL cluster account**.

  This option is not applicable on Windows Server 2003 systems.
- Click **Use Domain Group Account**and then click the adjacent ellipsis button to launch the Windows Select Users, Computers, or Groups dialog box.

  Then specify a domain group and click **OK** to set the domain group as the SQL cluster account.

  If you specify a domain group as the SQL cluster account, ensure that the SQL Server service and SQL Server Agent service accounts are part of the specified domain group.

The SQL agent assigns the specified account with Full Control privileges to the SQL Server databases and log files. This ensures that they are accessible upon failover.

10 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.

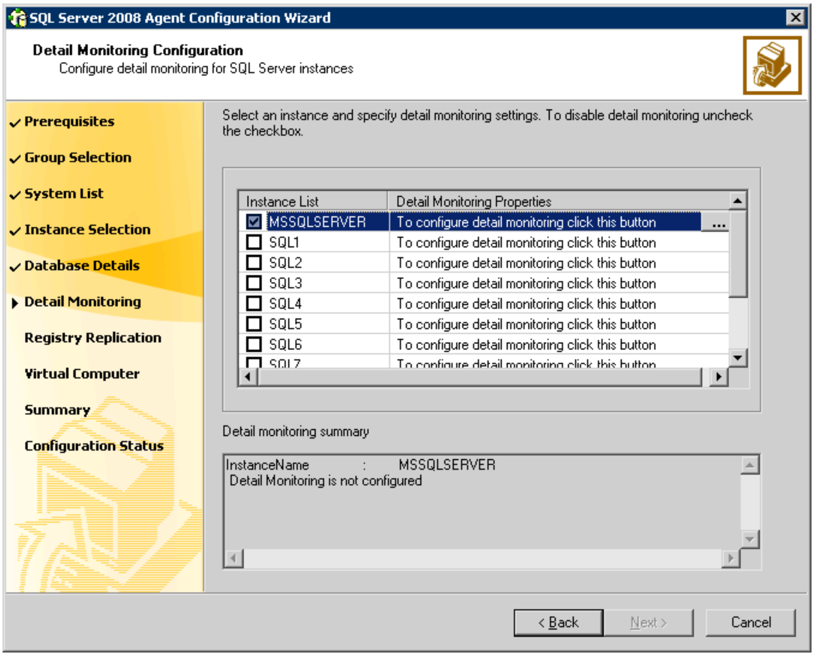

Detail Monitoring Configuration

Configure detail monitoring for SQL Server instances

Select an instance and specify detail monitoring settings. To disable detail monitoring uncheck the checkbox.

Instance List	Detail Monitoring Properties
☑ MSSQLSERVER	To configure detail monitoring click this button
☐ SQL1	To configure detail monitoring click this button
☐ SQL2	To configure detail monitoring click this button
☐ SQL3	To configure detail monitoring click this button
☐ SQL4	To configure detail monitoring click this button
☐ SQL5	To configure detail monitoring click this button
☐ SQL6	To configure detail monitoring click this button
☐ SQL7	To configure detail monitoring click this button

Detail monitoring summary

InstanceName : MSSQLSERVER
Detail Monitoring is not configured

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

- Check the check box for a SQL instance, and then click the button from the Detail Monitoring Properties column to specify the detail monitoring settings.
  See "To configure detail monitoring for a SQL Server instance" on page 76.
- Repeat these steps for each SQL Server instance that you wish to configure detail monitoring for.

Clear the check box to disable detail monitoring for the instance.

Click **Next**.

11 On the Registry Replication Path panel, specify the mount path to the registry replication volume (INST1_REGREP_VOL) and click **Next**.

Symantec recommends that RegRep resources and SQL data be in separate volumes.

12 On the Virtual Server Configuration panel, configure the virtual server as follows:

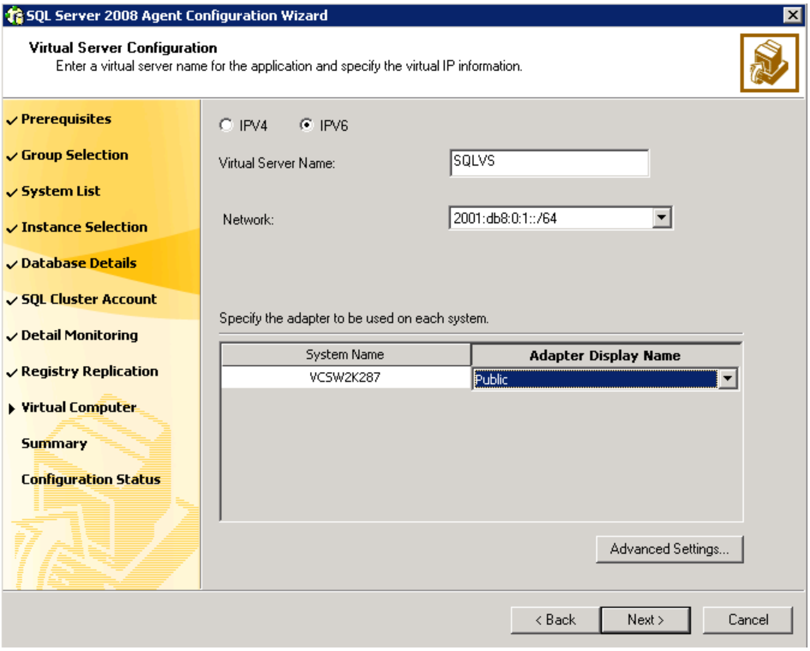

SQL Server 2008 Agent Configuration Wizard

Virtual Server Configuration

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

○ IPV4 ● IPV6

Virtual Server Name: SQLVS

Network: 2001:db8:0:1::/64

Specify the adapter to be used on each system.

System Name	Adapter Display Name
VCSW2K287	Public

- Select **IPv4** to configure an IPv4 address for the virtual server.
  - In the Virtual IP Address field, type a unique virtual IPv4 address that is currently not being used on your network, but is in the same subnet as the current node.

  - In the Subnet Mask field, type the subnet to which the virtual IPv4 address belongs.
- Select **IPv6** to configure an IPv6 address for the virtual server. The IPv6 option is disabled if the network does not support IPv6.
  - Select the network from the drop-down list. The wizard uses the network prefix and automatically generates an IPv6 address that is valid and unique on the network.
- Enter the virtual name for the server, for example `INST1-VS`. Ensure that the virtual server name you enter is unique in the cluster.
- For each system in the cluster, select the public network adapter name.
  The Adapter Display Name field displays the TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory, click **Advanced Settings**, check the **Active Directory Update required** check box, specify the desired Organizational Unit (OU) in the domain and then click **OK**.
  This sets the Lanman resource attributes ADUpdateRequired and ADCriticalForOnline to true. It allows the Lanman agent to update the Active Directory with the virtual server name.
  You can type the OU details in the format **CN=Computers,DC=domainname,DC=com**. To search for an OU, click on the ellipsis button and specify the search criteria in the Windows Find Organization Unit dialog box.
  By default, the Lanman resource adds the virtual server to the default container "Computers." The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- Click **Next**.

13 On the Initiator Selection panel, select the initiator for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.

If you are configuring MPIO over FC, you must select at least 2 FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

14 On the Service Group Summary panel, review the service group configuration and then click **Next**. The Resources 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.

To edit a resource name, click the resource name or press the **F2** key. Press **Enter** after editing each resource name.

To cancel editing a resource name, press **Esc**.

15 Click **Yes** when prompted to confirm creating the service group. Messages indicate the status of the commands.

16 Select the **Bring the service group online** check box, if you want to bring the service group online.

You may want to review the service group configuration in the Cluster Manager (Java Console) before bringing the service group online. You can use the Cluster Manager to bring the service group online later.

You must bring the SQL service group online on the node from where you ran the configuration wizard. This is the first cluster node where you installed SQL Server. This allows the wizard to configure the resources required for SQL Server services.

The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.

If you have created a new SQL Server database, you must modify the SQL Server service group to add the required storage agent resources to the service group. Run the service group configuration wizard to modify the service group.

Click **Finish** to exit the wizard or click **Next** to configure another SQL service group or an MSDTC service group.

The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.

You can also configure an MSDTC service group.

To configure detail monitoring for a SQL Server instance

1 On the Detail Monitor configuration dialog box, specify the monitoring frequency in the **Detail monitoring interval** field.

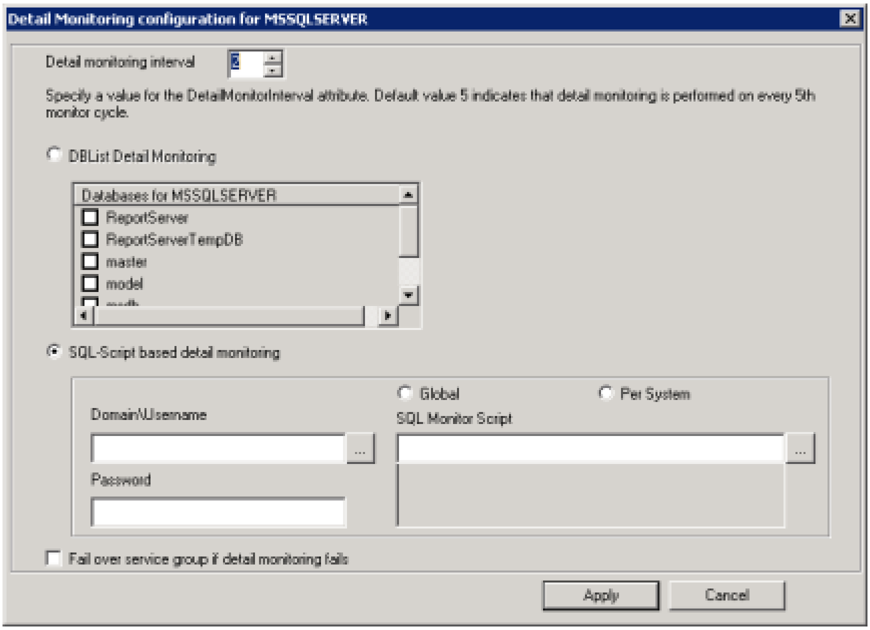

Detail Monitoring configuration for MSSQLSERVER

Detail monitoring interval: 5

Specify a value for the DetailMonitorInterval attribute. Default value 5 indicates that detail monitoring is performed on every 5th monitor cycle.

○ DBList Detail Monitoring

Databases for MSSQLSERVER
- ☐ ReportServer
- ☐ ReportServerTempDB
- ☐ master
- ☐ model

● SQL-Script based detail monitoring

○ Global ○ Per System

☐ Fail over service group if detail monitoring fails

This sets the value for the LevelTwoMonitorFreq 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 **DBList Detail 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 detail monitoring** if you wish to use a script to monitor SQL server databases. In this mode of detail monitoring, the agent executes the script that you specify for detail monitoring.

2 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.

Note: These credentials are required for both, DBList as well as SQLFile detail monitoring.

---

3. Select **Global** or **Per System** depending on whether the monitoring script location is the same for all the nodes or is unique for each cluster node, and then specify the path of the script appropriately.
4. Check **Fail over service group if detail monitoring fails** check box, if not already checked. This allows the SQL agent to fail over the service group to another node if the detail monitoring fails.
5. Click **Apply**.

Assigning privileges to the existing SQL Server databases and logs

Note: The following steps are required only if you have configured the SQL cluster account while creating the SQL Server service group earlier.

While installing SQL Server, if the user account specified for the SQL Server services is not a member of the local administrators group, then the SQL services and databases may not be accessible after a service group failover. For such a case, you configure the SQL cluster account while creating the SQL Server service group.

The SQL cluster account gets full control privileges to all the new databases and log files that are created after the service group is configured.

However, if databases were created before the service group is configured, you have to manually assign the SQL cluster account with full control privileges to the existing databases and log files associated with the instances in the service group.

To assign privileges to the existing SQL databases and logs

1. On the node where the SQL Server service group is online, navigate to the following directory from Windows explorer:

   <*Datarootdirectory*>\<*SQLinstancename*>\MSSQL\

   The directory contains various directories including DATA, FTData, JOBS, Log, repldata.

   Here, <*Datarootdirectory*> is the path that you specified while installing SQL Server.
2. Assign the SQL Cluster account with full control privileges to the following directories:
   - DATA
   - Log

3. Navigate inside the DATA folder and then assign the SQL cluster account with full control privileges to the following files in that directory:
    - tempdb.mdf
    - templog.ldf
4. Repeat these steps for all the instances that are configured in the SQL Server service group.

    This ensures the existing SQL databases are accessible after a service group failover.

Running SnapManager for SQL

After configuring the service group, you may want to run the SnapManager Configuration Wizard on the node on which the service group is online, to schedule backups of SQL Server database.

You must adhere to the following requirements while running SnapManager for SQL:

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

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

See the NetApp documentation for more information about running SnapManager for SQL.

Note: To use SnapManager with SQL Server 2012, make sure that you apply the latest cumulative patch. For more information, see the *Veritas Cluster Server for Windows Agent Pack Readme*.

Making SQL Server user-defined databases highly available with VCS

You can use VCS to manage user-defined SQL Server databases. Create the required SQL databases using the SQL Server Management Studio and then make them highly available with VCS.

Perform the following tasks to configure user-defined databases with VCS:

- Create volumes or LUNs for a user-defined SQL Server database and its transaction log.
- Create a SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes or LUNs.
- Modify the SQL service group using the SQL Server 2008 Agent Configuration Wizard to add the NetAppFiler and NetAppSnapDrive resources for the user databases.

Create volumes or LUNs for SQL user-defined databases

You must create volumes or LUNs for a user-defined SQL Server database and its transaction log.

In the sample deployment these volumes are named as follows:

- 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 database objects for the user database

Creating SQL Server databases

Use the SQL Server Management Studio to create a SQL Server user-defined database for the required SQL instance. While creating the database, ensure that you point the database files and transaction log to the paths of the new volumes or LUNs created earlier.

Refer to the Microsoft SQL Server documentation for instructions on how to create databases.

Adding storage agent resources to the SQL service group

After creating the database, run the SQL Server 2008 Agent Configuration Wizard and modify the SQL Server service group. This allows the wizard to add the NetAppFiler and NetAppSnapDrive (Mount and DiskRes in case of Windows LDM) storage resources for the user databases, to the SQL Server service group.

You must run the SQL Server 2008 Agent Configuration Wizard in the modify mode only if you create user-defined databases after creating the SQL Server service group.

Note: You must run the wizard in the modify mode even if you have added or changed volumes in your existing configuration.

Before running the configuration wizard to add the storage agent resources, do the following:

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

Note: Mount or NetAppSnapDrive resources are required only if the database is created on a new volume.

To add storage agent resources to the SQL service group

1. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server 2008 Agent Configuration Wizard** to start the configuration wizard.
2. Review the Prerequisites page and click **Next**.
3. On the Wizard Options panel, click **Edit service group**, select the service group and then click **Next**.
4. Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
5. In the Service Group Configuration page, click **Next**.
6. In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.
7. In the User Databases List page, make sure the databases are shown with correct paths and click **Next**.
8. If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
9. In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
10. Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
11. Click **Finish** to exit the wizard.

    The wizard marks all the resources in the service group as CRITICAL. If desired, use Cluster Manager (Java Console) or the command line to change the state.

Administering a SQL Server service group

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

Modifying a SQL service group configuration

Note the following prerequisites before modifying the SQL service group:

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

To modify a SQL Server service group

1. Start the SQL Server 2008 Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server 2008 Agent Configuration Wizard**.
2. Review the prerequisites and click **Next**.
3. In the Service Group Selection panel, select the service group to modify and click **Next**.
4. In the Service Group Configuration panel, add or remove systems from the service group's SystemList and click **Next**.
5. In the SQL Server Instance Selection panel, select the SQL Server instance to be made highly available and click **Next**.

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

Deleting a SQL service group

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

To delete a SQL Server service group

1. Start the SQL Server 2008 Agent Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server 2008 Agent Configuration Wizard**.
2. Review the prerequisites and click **Next**.
3. In the Service Group Selection dialog box, select the service group to delete and click **Next**.
4. In the Service Group Summary dialog box, click **Next**.
5. A message appears informing you that the wizard will run commands to delete the service group. Click **Yes** to delete the service group.
6. Click **Finish**.

Chapter 5

Configuring an MSDTC service group

This chapter includes the following topics:

About configuring the MSDTC service group

Microsoft Distributed Transaction Coordinator (MSDTC) service enables you to perform distributed transactions. A distributed transaction updates data on more than one computer in a network. The MSDTC service ensures that a transaction is successfully committed on each computer. A failure to commit on a single system aborts the transaction on all systems in the network. If a transaction spans across more than one computer in the network, you must ensure that the MSDTC service is running on all the computers. Also, all the computers must be able to communicate with each other.

Configuring the MSDTC service group involves the following tasks:

- Creating an MSDTC Server service group using the SQL Configuration Wizard
- Configuring the MSDTC client manually

Note: You have to use the SQL Server Agent Configuration Wizard to configure the MSDTC Server service group. You cannot use the SQL Server 2008 Agent Configuration Wizard to perform this task.

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

Typical MSDTC service group configuration

MSDTC servers can co-exist with SQL servers on the same cluster nodes. If the MSDTC Server and the SQL Server are running on the same node, the MSDTC client is configured in the default configuration. If the MSDTC Server is not configured on the same node as the SQL Server, then the MSDTC client must be configured on that node. In general, you must configure the MSDTC client on all nodes except the node on which the MSDTC Server is configured to fail over. The MSDTC client and the MSDTC Server must not run on the same cluster node.

For example, a SQL Server configuration in a VCS cluster might span four nodes and two sets of shared storage.

The following configurations are possible:

- MSDTC Server and SQL Server are configured on different nodes in the same cluster
- MSDTC Server and SQL Server are configured on the same nodes in a cluster
- MSDTC Server and SQL Server are configured on nodes in different clusters

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

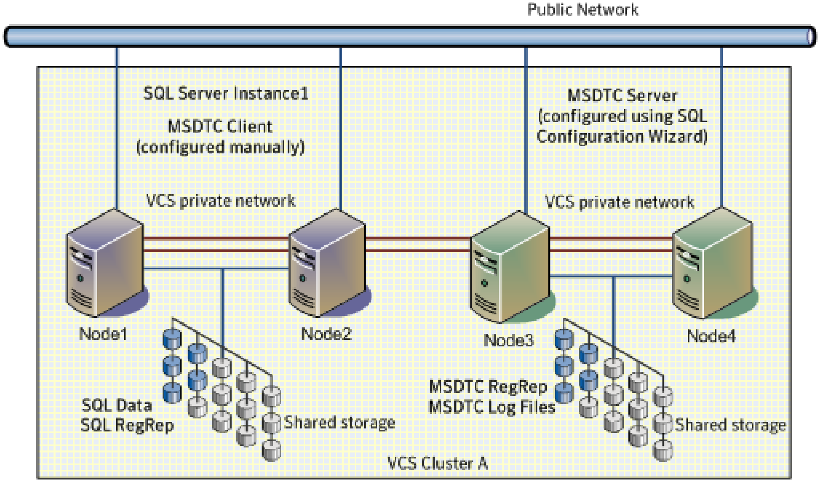

Figure 5-2 MSDTC Server configured on the same node as SQL Server

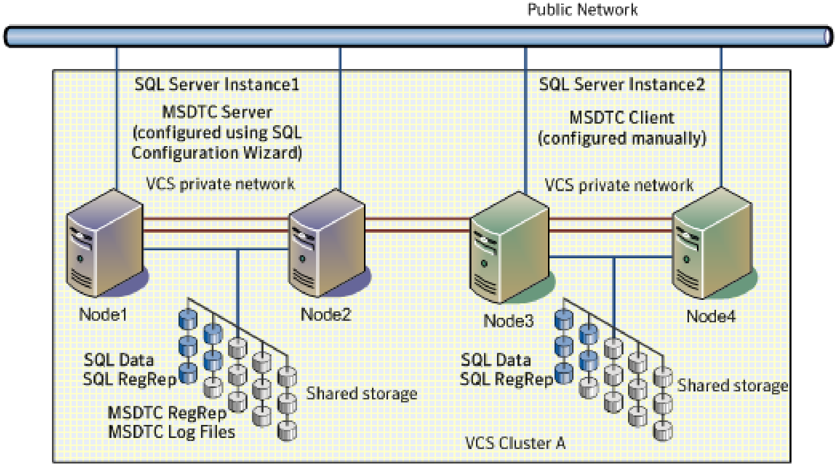

Figure 5-3 MSDTC Server and SQL Server configured on nodes in different clusters

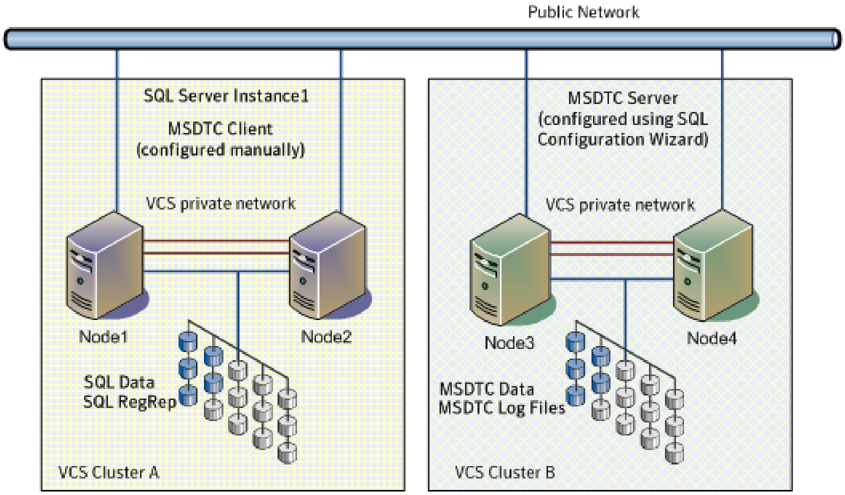

Before configuring the MSDTC service group

Note the following prerequisites before you configure the MSDTC service group:

- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- You must be a local Administrator on the node where you run the wizard.
- Verify that the VCS agent for SQL Server is installed on all cluster nodes.
- Verify that the VCS cluster is configured using the VCS Cluster Configuration Wizard (VCW).
- Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC Server service group.
- Verify that the Distributed Transaction Coordinator (MSDTC) service is stopped.
- Verify that you have created the volumes or LUNs (virtual disks) for storing MSDTC log and MSDTC registry replication information, on a shared disk.
  See “Managing storage using NetApp filer ” on page 54.
  See “Managing storage using Windows Logical Disk Manager ” on page 57.
- Verify that the volumes or LUNs created for the MSDTC logs and registry replication information are mounted or connected to the node where you run the wizard and dismounted or disconnected from all other nodes.

- If you have configured a firewall, add the following to the firewall exceptions list:
    - Port 14150 or the VCS Command Server service, `%vcs_home%\bin\CmdServer.exe`.
      Here, %vcs_home% is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
    - Port 14141

  For a detailed list of services and ports used by VCS, refer to the *Veritas Cluster Server Installation and Upgrade Guide*.
- Keep the following information ready with you; the wizard prompts you for these details:
    - A unique virtual name for the MSDTC Server. This is the name that is used by MSDTC clients to connect to the MSDTC Server. The DTC service runs under this virtual name.
    - A unique virtual IP address for the for the MSDTC Server.
      The virtual IP address is required only if you wish to configure an IPv4 address. In case of IPv6, the wizard prompts you to select the IPv6 network and automatically generates an IPv6 address that is valid and unique on the network. The wizard uses the prefix that is advertised by the router on the IPv6 network.

Creating an MSDTC service group

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

Note: Symantec recommends that you create only one MSDTC Server service group in a cluster.

This section describes the steps required to create an MSDTC Server service group using the SQL Configuration Wizard.

You have to use the SQL Server Agent Configuration Wizard to configure the MSDTC Server service group. You cannot use the SQL Server 2008 Agent Configuration Wizard to perform this task.

To create an MSDTC service group

1. Start the SQL Server Configuration Wizard. Click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
2. In the Select Configuration Option panel, click **MSDTC Server - Service Group Configuration**, click **Create**, and then click **Next**.
3. Review and verify that you have met the prerequisites for configuring an MSDTC Server service group and then click **Next**.
4. On the Service Group Configuration panel, specify the service group name and select the systems for the service group as follows:
   - Type a name for MSDTC service group.
   - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order box. The systems listed in the Systems in Priority Order box represent the SystemList attribute of the service group. While selecting systems, make sure to select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.
     To remove a system from the service group's system list, select the Systems in Priority Order list and click the left arrow.
     To change a system's priority in the service group's system list, select the system from the Systems in Priority Order and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
   - Click **Next**. If the cluster configuration is in read-only mode, the wizard prompts you before changing it to read-write mode. The wizard starts validating your configuration. Various messages indicate the validation status.
5. On the Virtual Server Configuration panel, specify the information related to the virtual server as follows:
   - Type a virtual name for the MSDTC Server. This is the name that is used by MSDTC clients to connect to the MSDTC Server. The DTC service runs under this virtual name. Ensure that the virtual server name is unique in the cluster.
   - Select **IPv4** to configure an IPv4 address for the virtual server.
     - In the Virtual IP Address field, type a unique virtual IPv4 address for the MSDTC server.

  - In the Subnet Mask field, type the subnet to which the virtual IPv4 address belongs.
- Select **IPv6** to configure an IPv6 address for the virtual server. The IPv6 option is disabled if the network does not support IPv6.
  - Select the network from the drop-down list. The wizard uses the network prefix and automatically generates an IPv6 address that is valid and unique on the network.
- For each system, select the public network adapter name. The Adapter Display Name field displays the TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. To view the adapters associated with a system, click the Adapter Display Name field and then click the arrow. Make sure that you select the adapters assigned to the public network, not the private.
- Click **Advanced Settings** to configure the Lanman agent to perform Windows Active Directory (AD) update. These settings are applicable to the Lanman resource in the service group. On the Lanman Advanced Configuration dialog box, complete the following:
  - Check the **Active Directory Update required** check box to enable the Lanman agent to update the Active Directory with the virtual name. This sets the Lanman agent attributes ADUpdateRequired and ADCriticalForOnline to true.
  - In the Organizational Unit field, type the distinguished name of the Organizational Unit for the virtual server in the format **CN=containername,DC=domainname,DC=com**. To browse for an OU, click the ellipsis (...) button and search for the OU using the Windows Find Organization Units dialog box. By default, the Lanman resource adds the virtual server to the default container "Computers."
  - Click **OK**.
    The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- Click **Next**.

6 On the Specify Data Path panel, specify the drive letter for the MSDTC log and registry replication directory and click **Next**. If the directory does not exist, the wizard creates it. Symantec recommends using different paths for these directories.

Clear the **Configure NetApp SnapMirror Resource(s)** check box. This option is applicable only in case of a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. If you are setting up a disaster recovery environment, check this check box to configure the SnapMirror resource at the primary site. The SnapMirror resource must be configured only after you have configured the cluster at the secondary site.

7 On the Initiator Selection panel, select the initiator for the virtual disk from the list of available initiators displayed for each cluster node, and then click **Next**.

If you are configuring Multipath I/O (MPIO) over Fibre Channel (FC), you must select at least two FC initiators for each cluster node. Note that the node from which you run this wizard already has an initiator selected by default. This is the initiator that was specified when you connected the LUNs to this cluster node.

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

- The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values.
- The wizard assigns unique names to resources. Change names of the resources, if desired.
  To edit a resource name, select the resource name and either click it or press the F2 key. Press the Enter key after editing each resource name. To cancel editing a resource name, press the Esc key.

9 Click **Yes** on the message that informs you that the wizard will run commands to create the service group.

Various messages indicate the status of these commands.

10 In the Completing the MSDTC Configuration Wizard panel, check **Bring the service group online** check box if you wish to bring the configured service group online on the local system. To bring the service group online later, clear this check box.

11 Click **Finish** to exit the wizard.

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

About configuring an MSDTC client

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

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

Ensure the following before you configure the MSDTC client:

- Verify that the MSDTC Server service group is online in the cluster.
- Configure the MSDTC client on systems where a SQL instance is configured to run.

Note: You have to configure the MSDTC client manually. You cannot use the service group configuration wizard to configure the MSDTC client.

Configuring an MSDTC client

Complete the following steps to configure the MSDTC client.

To configure an MSDTC client

1. Ensure that the MSDTC Server service group is online in the cluster.
2. Launch the Windows Component Services Administrative tool.

   Click **Start > Programs > Administrative Tools > Component Services**

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

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

Chapter 6

Configuring the standalone SQL Server

This chapter includes the following topics:

- Typical high availability configuration for a standalone SQL Server setup
- Configuring a standalone SQL Server for high availablility

Typical high availability configuration for a standalone SQL Server setup

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

It also describes the tasks necessary to create a virtual server in an active-passive SQL configuration. The active node of the cluster hosts the virtual server. The second node is a dedicated redundant server able to take over and host the virtual server in the event of a failure on the active node.

Figure 6-1 Active-Passive configuration

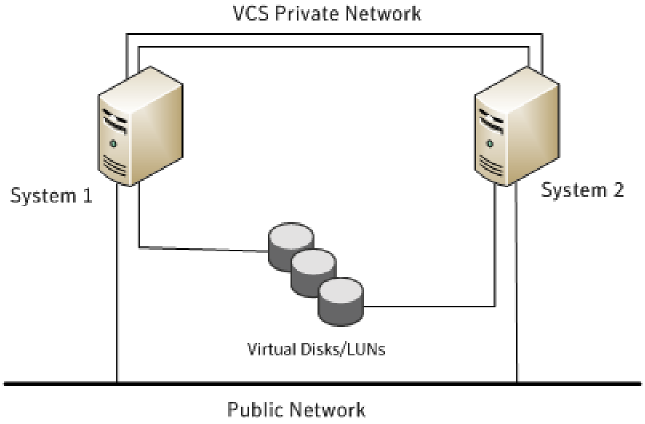

The virtual SQL Server is online on SYSTEM1, serving client requests. The shared LUNs (virtual disks) provide storage for the SQL Server databases. SYSTEM2 waits in a warm standby state as a backup node, prepared to begin handling client requests if SYSTEM1 becomes unavailable. From the user's perspective there will be a small delay as the backup node comes online, but the interruption in effective service is minimized.

Sample configuration

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

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

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

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

Table 6-1 Objects used for clustering standalone server

Name	Object
`SYSTEM1 & SYSTEM2`	server names; SYSTEM1 is the existing standalone SQL server

Table 6-1 Objects used for clustering standalone server *(continued)*

Name	Object
`INST1_SG`	Microsoft SQL Server service group
`SQL_CLUS1`	virtual SQL server cluster
`INST1_DG`	Disk group for the volumes for the SQL instance
`INST1_DATA_FILES`	volume for Microsoft SQL Server system data files
`INST1_DB1_VOL`	volume for storing a Microsoft SQL Server user-defined database
`INST1_DB1_LOG`	volume for storing a Microsoft SQL Server user-defined database log file
`INST1_REGREP_VOL`	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL server
`INST1`	SQL Instance Name
`INST1-VS`	name of the SQL Virtual Server

Configuring a standalone SQL Server for high availablility

Perform the following tasks to configure the standalone SQL Server for high availability, in an active-passive configuration with one to one failover capabilities:

Note: In addition to the tasks mentioned below, 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.

Table 6-2 lists the tasks to configure the standalone SQL Server for high availability.

Table 6-2 Tasks to configure the standalone SQL Server for high availability

Task	Description
Prepare the standalone SQL Server	Complete the following tasks before you begin the process of installing VCS and creating a clustered environment: - Back up the existing SQL data - From the SQL Server Service Manager, set the SQL Server services to manual start. While you set the SQL Server services to manual start, you must select the standalone server that you plan to incorporate, select the service from the Services list, and then clear the **Auto-start service when OS starts** check box. **Note:** Repeat these steps for all other SQL Server services that are running on the server.
Install and configure VCS on the standalone SQL server	Install Veritas Cluster Server (VCS) on all the systems where you want to configure the application. During installation, the product installer installs the VCS agents required for making the applications highly available. Refer to the Veritas Cluster Server Installation and Upgrade Guide for more information on installing VCS. For details on configuring the VCS cluster, refer to, See "Configuring the cluster using the Cluster Configuration Wizard " on page 35.
Create volumes or LUNs (virtual disks) necessary to manage the SQL Server storage.	See "Managing storage using NetApp filer " on page 54. See "Managing storage using Windows Logical Disk Manager " on page 57.
Install and configure SQL Server on additional nodes, if required	Perform the following tasks to install Microsoft SQL Server on additional nodes. - Ensure that the shared volumes or LUNs are imported on the node. - Ensure that the SQL Server configuration is identical on all nodes in the cluster. To have identical configuration, ensure that the instance name (if applicable), destination folder for Program Files and Data Files and the Authentication Mode are same on all the nodes. - Perform the SQL Server installation See "Installing SQL Server on the additional cluster nodes " on page 63.

Table 6-2 Tasks to configure the standalone SQL Server for high availability *(continued)*

Task	Description
Verify that the existing SQL Server databases and logs are moved to shared storage	Verify the location of all SQL Server databases and logs for the existing standalone server. If they are located on local storage, move them from the local drive to the appropriate volumes or LUNs on shared storage to ensure proper failover operations in the cluster.
Configure the SQL Server service group	See "Configuring a SQL Server service group " on page 70.
Create and manage SQL Server user-defined database	See "Making SQL Server user-defined databases highly available with VCS " on page 79.
Verify the configuration	

Chapter 7

Configuring a disaster recovery setup

This chapter includes the following topics:

Setting up the Active/Active cluster

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

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

Why implement a disaster recovery solution?

A DR solution is vital for businesses that rely on the availability of data.

A well-designed DR solution prepares a business for unexpected disasters and provides the following benefits in a DR situation:

- Minimizes economic loss due to the unavailability or loss of data.
- Provides a plan for the safe and orderly recovery of data in the event of a disaster.
- Ensures safe and efficient recovery of data and services.
- Minimizes any decision making during DR.
- Reduces the reliance on key individuals.

Strategically planning a DR solution provides businesses with affordable ways to meet their service level agreements, comply with government regulations, and minimize their business risks.

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

Understanding replication

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

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

What needs to be protected in a SQL Server environment?

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

- User Databases
  The most critical component in any SQL Server implementation is the user data that is stored in user-defined databases.
- Logins
  Logins allow clients to connect to SQL Server and execute queries on user data. Logins are stored in the master database and each of the user-defined databases.

- Jobs

  Jobs are a set of scheduled tasks that maintain SQL Server databases. The job configuration is stored in the msdb system database.
- Alerts

  Alerts are actions that are taken when a specific event occurs. They are used to respond to and correct errors that occur in SQL Server. The alert configuration is stored in the msdb system database.
- Operators

  Operators are contacts that address problems occurring in SQL Server. They are notified in the event of errors. The operator configuration is stored in the msdb system database.
- Extended Stored Procedures

  Extended stored procedures are external routines that are called from within SQL Server. They are typically stored in DLL files on the file system.
- Other Server Extensions

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

Typical disaster recovery configuration

A Disaster Recovery (DR) configuration enables you to restore application data and services in the event of a catastrophic failure. A typical DR solution requires primary and secondary sites, and clusters within those sites. The clusters at the primary and secondary sites are a part of the global cluster. The cluster at the primary site provides data and services during normal operation, and the cluster at the secondary site provides data and services if the primary site fails. VCS continuously monitors and communicates events between clusters. Inter-cluster communication ensures that the global cluster is aware of the state of the global service group at all times.

Figure 7-1 Typical Disaster Recovery configuration

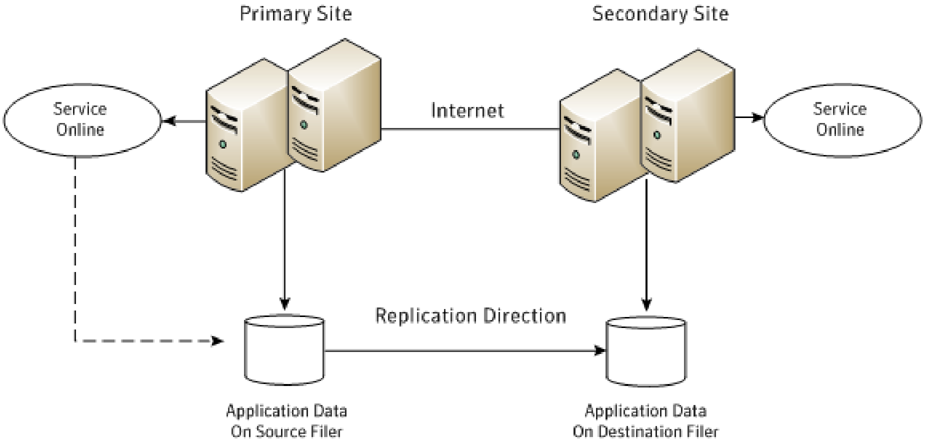

The illustration displays an environment with a DR solution that is prepared for a disaster. The primary site consists of two nodes, System1 and System2. The secondary site consists of two nodes, System3 and System4. Each site has a clustered setup with the nodes set up appropriately for failover within the site.

Filer1 in the cluster on the primary site replicates to Filer2 in the cluster on the secondary site. Replication between the filers is set up using NetApp SnapMirror for SQL. If the Microsoft SQL Server server on System1 fails, SQL Server comes online on node System2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal.

When a failure occurs, such as an earthquake that destroys the data center in which the primary site resides, the DR solution is activated. VCS fails over the entire service group to the cluster at the secondary site. System3 at the secondary site takes over, and the data that was replicated to the secondary site is used to restore the application services to clients.

Configuring a disaster recovery set up for SQL Server

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

Note: You must perform the tasks at the primary site first. After the configuration is complete at the primary site, proceed to perform the tasks at the secondary site.

Before you begin to create the SQL Server service group for the cluster at the secondary site, make sure that the SQL Server service group at the primary site is offline.

Table 7-1 lists the tasks to set up a disaster recovery environment for SQL Server.

Table 7-1 Tasks for SQL Server disaster recovery set up

Task	Description
Review the configuration	Review the system configuration before you start deploying VCS and SQL Server in your environment. See "About installing the VCS agents for SQL Server " on page 35. See "Prerequisites for installing SQL Server " on page 52. See "Privileges required for installing SQL Server " on page 53. See "Before configuring the SQL service group " on page 68.
Install VCS and configure the cluster	Ensure that you perform the following tasks while you install VCS and configure the cluster: ■ Select the Global Cluster Option for VCS to enable wide-area failover. ■ Configure cluster components including the Wide-Area Connector (WAC) resource for global clusters, using the VCS Cluster Configuration Wizard (VCW). See "About installing the VCS agents for SQL Server " on page 35. See "Configuring the cluster using the Cluster Configuration Wizard " on page 35.

Table 7-1 Tasks for SQL Server disaster recovery set up *(continued)*

Task	Description
Configure volumes or LUNs on the shared storage	Create volumes or LUNs required for SQL Server and ensure that the volumes or LUNs (virtual disks) are connected to the first cluster node. During the creation of virtual disks and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site: ■ Volume sizes ■ Volume names ■ Drive letters See "Managing storage using Windows Logical Disk Manager " on page 57.
Install and configure SQL Server on the first node	See "Installing SQL Server on the first cluster node " on page 61.
Install and configure SQL Server on the additional nodes	See "Installing SQL Server on the additional cluster nodes " on page 63. **Note:** The instance name must be the same on the primary site and secondary site.
Configure the SQL Server service group	See "Configuring a SQL Server service group " on page 70. **Note:** The service group name and virtual computer name must be same on both, the primary site and secondary site.
Configure DR components	After configuring the high availability and SQL Server components on the primary and secondary sites, complete the disaster recovery solution by configuring the disaster recovery components for both sites. See "Prerequisites " on page 108. See "Linking clusters: Adding a remote cluster to a local cluster " on page 108. See "Converting a local service group to a global service group " on page 110. See "Bringing a global service group online " on page 112. See "Administering global service groups " on page 112. See "Deleting a remote cluster " on page 114.

Configuring replication using NetApp SnapMirror

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

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

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

Refer to NetApp documentation for more information.

Configuring SnapMirror resources at the primary site

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

You may want to repeat this procedure and create a NetAppSnapMirror resource at the secondary site.

This is required in cases such as the following:

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

Use the SQL Server 2008 Agent Configuration Wizard to add the SnapMirror resource. Verify that the volumes or LUNs created to store the registry replication information and the SQL Server database are connected to the node on which you run the wizard, and disconnected from other nodes in the cluster.

Configuring the Global Cluster Option for wide-area failover

The Global Cluster option is required to manage global clustering for wide-area disaster recovery.

Creating a global cluster environment involves the following:

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

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

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

Prerequisites

Creating a global cluster environment requires the following:

- Wide-area Connector process is configured and the ClusterService group is online at both sites.
  See "Configuring Wide-Area Connector process for global clusters" on page 47.
- All service groups properly configured and able to come online.
- The service group serving as the global group has the same unique name across all applicable clusters.
- The clusters use the same version of VCS.
- The clusters use the same operating system.
- The clusters are standalone and do not already belong to a global cluster environment.
- The names of the clusters at the primary and secondary sites and the virtual IP addresses associated with them are registered in the DNS with reverse lookup.

Linking clusters: Adding a remote cluster to a local cluster

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

You can run the wizard from the following locations:

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

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

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

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

To add a remote cluster in Cluster Explorer

1. From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.

   or

   From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.
2. Review the required information for the Remote Cluster Configuration Wizard and click **Next**.
3. In the Wizard Options panel, click **Add Cluster**, then click **Next**.
4. In the New Cluster Details panel, enter the details of the new cluster.

   If the cluster is not running in secure mode, do the following:

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

- Click **Next**.

If the cluster is running in secure mode, do the following:

- Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
- Verify the port number.
- Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.
- If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
- Click **Next**.

5 Click **Finish**. After running the wizard, the configurations on all the relevant clusters are in read-write mode; the wizard does not close the configurations.

6 Verify that the heartbeat connection between clusters is alive. From the command window enter `hahb -display`. The state attribute in the output should show `alive`.

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

Converting a local service group to a global service group

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

To convert a local service group to a global group

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

or

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

or

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

2 Review the information required for the Global Group Configuration wizard and click **Next**.

3 Enter the details of the service group to modify as follows:

- Click the name of the service group that will be converted from a local group to a global group, or vice versa.
- From the Available Clusters box, click the clusters on which the group can come online. Click the right arrow to move the cluster name to the Clusters for Service Group box; for global to local cluster conversion, click the left arrow to move the cluster name back to the Available Clusters box. A priority number (starting with 0) indicates the cluster on which the group will attempt to come online. If necessary, double-click the entry in the Priority column and enter the new value.
- Select the policy for cluster failover as follows:

Manual	Prevents a group from automatically failing over to another cluster.
Auto	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster fails.
Connected	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster.

- Click **Next**.

4 Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:

Cluster not in secure mode	■ Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system. ■ Verify the port number. ■ Enter the user name. ■ Enter the password. ■ Click **OK**. ■ Repeat these steps for each cluster in the global environment.

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

5. Click **Next**, then click **Finish**.

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

Bringing a global service group online

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

To bring a remote global service group online from Cluster Explorer

1. In the Service Groups tab of the configuration tree, right-click the service group.

   or

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

Administering global service groups

Administering global groups requires the following conditions:

- A group that will serve as the global group must have the same name across all applicable clusters.

- You must know the user name and password for the administrator to each cluster in the configuration.

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

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

Taking a remote global service group offline

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

To take a remote global service group offline from Cluster Explorer

1. In the Service Groups tab of the configuration tree, right-click the service group.

   or

   Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.

2. Click **Offline**, and click **Remote offline**.
3. In the Offline global group dialog box do the following:
   - Click the remote cluster to take the group offline.
   - Click the specific system, or click **All Systems**, to take the group offline.
   - Click **OK**.

Switching a remote service group

Use Cluster Explorer to switch a remote service group.

To switch a remote service group from Cluster Explorer

1. In the Service Groups tab of the configuration tree, right-click the service group.

   or

   Click a cluster in the configuration tree, click the **Service Groups** tab, and right-click the service group icon in the view panel.

2. Click **Switch To**, and click **Remote switch**.
3. In the Switch global group dialog box do the following:
   - Click the cluster to switch the group.
   - Click the specific system, or click **Any System**, to take the group offline.
   - Click **OK**.

Deleting a remote cluster

If necessary, use the Remote Cluster Configuration wizard to delete a remote cluster.

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

Deleting a remote cluster involves the following tasks:

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

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

To take the wide area cluster (wac) resource offline

1. From Cluster Monitor, log on to the cluster that will be deleted from the global cluster environment.
2. In the Service Groups tab of the Cluster Explorer configuration tree, right-click the **wac** resource under the Application type in the ClusterService group.

   or

   Click a service group in the configuration tree, click the **Resources** tab, and right-click the **wac** resource in the view panel.
3. Click **Offline**, and click the appropriate system from the menu.

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

1. From Cluster Explorer, click **Configure Global Groups** on the **Edit** menu.
2. Click **Next**.
3. Enter the details of the service group to modify, as follows:
   - Click the name of the service group.
   - For global to local cluster conversion, click the left arrow to move the cluster name from the cluster list back to the Available Clusters box.
   - Click **Next**.
4. Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:

   If the cluster is not running in secure mode, do the following:
   - Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
   - Verify the port number.
   - Enter the user name.
   - Enter the password.
   - Click **OK**.

   If the cluster is running in secure mode, do the following:
   - Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system.
   - Verify the port number.
   - Choose to connect to the remote cluster using the connected cluster’s credentials, or enter new credentials, including the user name, password, and domain.

- Click **OK**.

5 Click **Next**.

6 Click **Finish**.

To delete a remote cluster from the local cluster

1 From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.

or

From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Clusters**.

2 Review the required information for the Remote Cluster Configuration wizard and click **Next**.

3 On the Wizard Options panel, click **Delete Cluster**, then click **Next**.

4 In the Delete Cluster panel, click the name of the remote cluster to delete, then click **Next**.

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

If the cluster is not running in secure mode do the following:

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

If the cluster is running in secure mode do the following:

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

6 Click **Finish**.

Chapter 8

Troubleshooting VCS agents for NetApp and Microsoft SQL Server

This chapter includes the following topics:

About troubleshooting VCS agents for NetApp and Microsoft SQL Server

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

VCS logging

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

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

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

Table 8-1 describes the agent log message components and their descriptions.

Table 8-1 Log message components and their description

Log message component	Description
Timestamp	Denotes the date and time when the message was logged.
Mnemonic	Denotes which Symantec product logs the message. For Veritas Cluster Server, the mnemonic is ‘VCS’.
Severity	Denotes the severity of the message. Severity is classified into the following types: ■ CRITICAL indicates a critical error within a VCS process. Contact Technical Support immediately. ■ ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action. ■ WARNING indicates a warning or error, but not an actual fault. ■ NOTE informs the user that VCS has initiated an action. ■ INFO informs the user of various state messages or comments. Among these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.

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

Log message component	Description
UMI or Unique Message ID	UMI 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-xxxxx-yy`. Originator ID for all VCS products is 'V-16.' Category ID for the agents is as follows: ■ SQL Server 2000: 20020 ■ SQL Server 2005: 20050 ■ SQL Server 2005 Agent Service: 20051 ■ SQL Server 2005 OLAP: 20052 ■ SQL Server 2008: 20069 ■ SQL Server 2012: 20093 ■ SQL Server Filestream: 20070 ■ MSTDC: 20021 Message ID is a unique number assigned to the message text.
Message Text	Denotes the actual message string.

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

A typical agent log resembles:

```
2010/01/12 11:22:47 VCS NOTICE V-16-20020-10
    SQLServer2005:SQL-Group-SQLServer2005:monitor:SQL Server
    Instance name is not specified. Agent will operate on default
    instance.
```

VCS Cluster Configuration Wizard (VCW) logs

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

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

The format of the wizard log is of the format *ThreadID | Message Text*.

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

A typical wizard log resembles the following:

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

VCWsilent logs

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

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

A typical VCWsilent log resembles the following:

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

NetApp agents error messages

Table 8-2 contains a list of error messages for the VCS agents for NetApp.

Table 8-2 NetApp agents error messages

Message	Description
Failed to open connection to filer %s.	Make sure that the VCS Helper Service account has is a domain user and is part of the administrator's group on the local host and the filer. Make sure the private network is functioning properly. Verify you can ping the IP used for the private storage network. This is the IP defined the StorageIP attribute of the NetAppFiler resource.
Failed to initialize ONTAPI on system	The agent could not find the file NTAPADMIN.DLL on the system. Verify the file exists in the %VCS_HOME%\bin directory
Invalid attributes exist in the configuration	Some agent attributes have not been defined or have been defined incorrectly. Verify the configuration definition for the agent.
ONTAP API called failed for object_name on filer_name.	The specified API failed on the specified object. See the NetApp ONTAP API documentation for information about the associated error message
Volume %s on filer %s is not a SnapMirror replicated volume	Verify replication is set up on the specified volume.
Multiple snapmirror destinations for a volume is not supported by this agent. 'snapmirror status' for volume %s on filer %s returned multiple status entries. Administrative intervention required	There should be only one destination per source volume.
Initialize VLibNetAppHost::Initialize() failed. (error_type: %s, error_code: 0x%s)	The agent could not detect the iSCSI or the FC Initiator on the host. Make sure that you have installed and configured Microsoft iSCSI Initiator or an FC Initiator on each node.

Table 8-2 NetApp agents error messages *(continued)*

Message	Description
Failed to connect/disconnect virtual disk. (error_type: %s, error_code: 0x%s. error_message: %s)	This could occur because one or more of the following parameters are defined incorrectly in the VCS configuration: ■ Filer name ■ Volume name/LUN name ■ Share name ■ Storage IP Verify the configuration definition of the resource. Make sure each attribute is defined correctly.
Unable to create/delete online lock file %s. Error code %s,	Make sure you have write permissions on the specified directory.

SQL Server agent error messages and descriptions

This section lists the messages of type ERROR and WARNING. Each message includes a description and a recommended solution, if applicable.

Agent for MSDTC error messages

Table 8-3 describes the error messages for the MSDTC agent.

Table 8-3 MSDTC agent error messages

Message	Description
Lanman attribute has not been configured.	No value specified for the LanmanResName attribute. Solution: Specify a valid value for the LanmanResName attribute.
MountResName attribute has not been configured.	No value specified for MountResName attribute. Solution: Specify a valid value for the MountResName attribute.
LogPath attribute has not been configured.	No value specified for LogPath attribute. Solution: Specify a valid value for the MountResName attribute.

Table 8-3 MSDTC agent error messages *(continued)*

Message	Description
Failed to open the SCM handle. Error = *Error code*.	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
Failed to open the MSDTC service. Error = *Error code*.	The agent failed to open the MSDTC service from the Service Control Manager (SCM). Solution: Check whether the service is present in the Service Control Manager.
Failed to start the MSDTC service. Error = *Error code*.	The agent failed to start the MSDTC service. See the associated Windows error code for more information.
The MSDTC log path is '*path name*'. Configured one is '*path name*'.	The specified path for the MSDTC logs is different from the actual path. Solution: Specify the correct MSDTC log path.
The MSDTC service is not in running state. Offline might be unsuccessful.	The MSDTC service could be in PAUSE, PAUSE PENDING, or START PENDING state. Solution: Resume the service and then attempt to stop it.
Failed to stop the MSDTC service. Error = *Error code*.	The MSDTC service could not be stopped. See the associated Windows error code for more information.
Failed to wait for the MSDTC service to stop. Error = *Error code*.	The agent could not stop the service within the specified time limit of 20 seconds. See the associated Windows error code for more information.

Agent for SQL Server 2012 error messages

This section describes the error messages for the VCS agent for SQL Server 2012.

Table 8-4 VCS agents for SQL Server 2012 error messages

Message	Description
Failed to initialize the SQLServer agent.	The agent failed to initialize the SQLServer agent for SQL Server 2012.
Failed to open the SCM handle. Error = *Error code*.	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
Failed to open the service *service name*. 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 *service name*. 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.
The service *service name* is not in stopped state.	The agent is trying to start the service. But the service is in an invalid state. Solution: Check the state of the service.
Failed to set the virtual computer name in the environment of the service *service name*. Error = *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to start the service *service name*. Error = *Error code*.	The agent failed to start the service. Solution: Verify if you can start the service from the Windows Services console. If the service starts successfully, stop the service. If the service does not start, see the associated Windows error code for more information.
The service *service name* did not start within the specified time limit.	The agent failed to start the service within the time limit as specified in the SQLOnlineTimeout attribute. Solution: If the system is slow, you can modify the SQLOnlineTimeout attribute value to accomodate the time that the service takes to start.

Table 8-4 VCS agents for SQL Server 2012 error messages *(continued)*

Message	Description
Failed to wait for the service *service name* to start. Error = *Error code.*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Sql script failed. Enable DBG1 Logging for the Script output.	The SQL script failed to monitor the SQL Server instance. See the script output for more information.
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. See the associated Windows error code for more information.
Sql script has failed. Error : *Error code.*	The SQL script for detail monitoring failed. See the associated Windows error code for more information.
Error occurred while getting the process exit code. Error : *Error code.*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
WaitForSingleObject () failed. Error : *Error code*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to convert the argument list. Error = *Error code.*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to convert the password attribute to UNICODE. Error = *Error code.*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Configuration mismatch: Detail monitoring is configured, but user credentials are not provided.	When configuring detail monitoring for the VCS agent for SQL Server 2012, authentication is required for regardless of the monitoring option (database list or SQL script). Solution: Provide the appropriate user credentials to perform detail monitoring.

Table 8-4 VCS agents for SQL Server 2012 error messages *(continued)*

Message	Description
Configuration mismatch: Detail monitoring is configured,but neither a list of databases to be monitored nor a SQL monitor script is provided.	When configuring the VCS agent for SQL Server 2012, if you specify that detail monitoring should be performed, you also need to specify the method to be used. Solution: Select either the database list-based or the SQL script-based monitoring option, and provide the appropriate user credentials.

Table 8-5 VCS agents for SQL Server 2008 error messages

Message	Description
Failed to initialize the SQLServer agent.	The agent failed to initialize the SQLServer agent for SQL Server 2012.
Failed to open the SCM handle. Error = *Error code*.	The agent fails to get a handle to the Service Control Manager (SCM). This could occur if the specified SCM database does not exist or the requested access is denied. Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.
Failed to open the service *service name*. 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 *service name*. 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.
The service *service name* is not in stopped state.	The agent is trying to start the service. But the service is in an invalid state. Solution: Check the state of the service.
Failed to set the virtual computer name in the environment of the service *service name*. Error = *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.

Table 8-5 VCS agents for SQL Server 2008 error messages *(continued)*

Message	Description
Failed to start the service *service name*. Error = *Error code*.	The agent failed to start the service. Solution: Verify if you can start the service from the Windows Services console. If the service starts successfully, stop the service. If the service does not start, see the associated Windows error code for more information.
The service *service name* did not start within the specified time limit.	The agent failed to start the service within the time limit as specified in the SQLOnlineTimeout attribute. Solution: If the system is slow, you can modify the SQLOnlineTimeout attribute value to accomodate the time that the service takes to start.
Failed to wait for the service *service name* to start. Error = *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Sql script failed. Enable DBG1 Logging for the Script output.	The SQL script failed to monitor the SQL Server instance. See the script output for more information.
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. See the associated Windows error code for more information.
Sql script has failed. Error : *Error code*.	The SQL script for detail monitoring failed. See the associated Windows error code for more information.
Error occurred while getting the process exit code. Error : *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.
WaitForSingleObject () failed. Error : *Error code*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to convert the argument list. Error = *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.

Table 8-5 VCS agents for SQL Server 2008 error messages *(continued)*

Message	Description
Failed to convert the password attribute to UNICODE. Error = *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Configuration mismatch: Detail monitoring is configured, but user credentials are not provided.	When configuring detail monitoring for the VCS agent for SQL Server 2012, authentication is required for regardless of the monitoring option (database list or SQL script). Solution: Provide the appropriate user credentials to perform detail monitoring.
Configuration mismatch: Detail monitoring is configured,but neither a list of databases to be monitored nor a SQL monitor script is provided.	When configuring the VCS agent for SQL Server 2012, if you specify that detail monitoring should be performed, you also need to specify the method to be used. Solution: Select either the database list-based or the SQL script-based monitoring option, and provide the appropriate user credentials.
The *service name* service is in STARTED state but is not running under the context of Virtual Server *virtual server name*	The service has started from outside VCS control. Solution: Stop the service and then bring the service group online.
Invalid value specified for attribute *attribute name*.	No value provided for the specified attribute. Solution: Provide a value for the attribute.
SQL Server Instance name is not specified. Agent will operate on the default instance.	No value spaced-out for SQL Server instance name. Agent would operate on the default SQL Server instance.

Table 8-5 VCS agents for SQL Server 2008 error messages *(continued)*

Message	Description
The *service name* service is not in stopped or running state. State=*state name*.	During the agents monitor entry point, the agent expects the service to be either in stopped state [to declare that the resource is offline] or in running/started state [to declare that the resource is online]. If the service is not in a stopped or started state then this error is logged and resource goes into unknown state. Solution: Probe the resource or wait for the next agent monitor cycle.
Failed to get the password attribute. Error = *Error code*.	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 internal error. Solution: Contact Symantec Technical Support.
The service *service name* 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 *service name* did not stop. Error = *Error code*.	The agent failed to stop the service. See the associated Windows error code for more information.
The service *service name* did not stop within the specified timeout. Error = *Error code*.	The agent failed to stop the service within the time limit as specified in the SQLOfflineTimeout attribute. Solution: If the system is slow, you can modify the SQLOfflineTimeout attribute value to accomodate the time that the service takes to stop.
The password attribute has not been configured.	The password attribute used for detail monitoring is not configured.
Unable to convert the buffer to UNICODE. Error = *Error code*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Sql script failed. Script output : *output*	The SQL script failed to monitor the SQL Server instance. See the script output for more information.

Table 8-5 VCS agents for SQL Server 2008 error messages *(continued)*

Message	Description
Failed to get the temporary file path. Error : *Error code*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to create the temporary file. Error = *Error code.*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed read the temporary file. Error = *Error code.*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Failed to remove the virtual name environment for the service *service name.*	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Invalid arglist. The ArgList should contain LanmanResNameIPResName	The Lanman resource name is incorrect. Solution: Verify that the Lanman resource name is valid.

Agent for SQL Server FILESTREAM error messages

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

Table 8-6 SQL Server FILESTREAM agent error messages

Message	Description
Check Filestream is enabled in MSSQL-Configuration Manager if not enable filestream with appropriate enable level [206]	If FILESTREAM is not enabled on the node and the VCS Filestream resource is created manually, the resource fails to discover FILESTREAM settings on the node. Solution: Enable FILESTREAM for that SQL instance and then probe the VCS Filestream resource.

Table 8-6 SQL Server FILESTREAM agent error messages *(continued)*

Message	Description
SQLFilestream Resource will be in UNKNOWN [Actual :Offline] State : Filestream Fileshare exists even filestream is disabled [406]	When the VCS Filestream resource is taken offline, the respective FILESTREAM fileshares on the node are also deleted. If the agent is unable to delete the fileshares the VCS Filestream resource goes in to an unknown state. Solution: Delete the FILESTREAM fileshares from the command line manually, and then probe the resource.
SQLFilestream Resource is in UNKNOWN[Actual:online] : Filestream Fileshare exists even filestream is enabled for local access only [407	The FILESTREAM access level is set to 0 (local access) but the FILESTREAM fileshares exists. This causes the VCS Filestream resource to go in to an unknown state. Solution: Delete the FILESTREAM fileshares from the command line manually, probe the resource, take the resource offline and then bring it online.
Filestream will be in offline [Actual:Online] : Filestream Fileshare doesn't exists even filestream is Enabled [409]	Either the FILESTREAM fileshares do not exist or the agent failed to create them. The VCS Filestream resource goes offline. Solution: From the SQL Configuration Manager, enable FILESTREAM for that instance, and then probe the resource.

Agent for SQL Server Analysis Service error messages

This section describes the error messages for the GenericService agent used to make SQL Server Analysis Service highly available.

Table 8-7 GenericService agent error messages

Message	Description
VCS ERROR V-16-10051-6012 GenericService:MSOlap-*resource name* Online:Failed to wait for the service *service name* to start. Error = 258.	This error may occur if the Analysis Service takes a long time to start. The configured GenericService resource may go into an unknown state. Solution: The GenericService agent attributes DelayAfterOffline and DelayAfterOnline determine the number of seconds the agent waits for the service to start or stop. Modify these attribute values depending on the time the configured service takes to start or stop once the resource is taken online or offline in the environment.

Using the virtual MMC viewer

This appendix includes the following topics:

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

About using the virtual MMC viewer

VCS starts the MSDTC service in the cluster under the context of the virtual server. Because the MMC snap-in is not aware of such a configuration, it is not possible to view the transactions on the DTC virtual server from a node where the MSDTC resource is online. VCS provides a virtual MMC viewer, the VCS Application Manager (VAM) utility, that enables you to view the distributed transaction statistics on the DTC virtual server from a node where the MSDTC resource is online.

Viewing DTC transaction information

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

Prerequisites for viewing DTC transaction information are as follows:

- An MSDTC service group must be configured and online in the cluster.

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

To view transactions from a node where MSDTC resource is online

1. Start the VCS Application Manager utility.

   Click **Start > Programs > Symantec > Veritas Cluster Server > Configuration Tools > Application Manager**.

   The VCS Application Manager displays a list of supported application service groups configured in the cluster. For each service group it also displays the state of the service group, the name of the virtual server resource (Lanman resource) and the corresponding management tools used for that application.
2. Select **MSDTC** from the Select the resource type drop-down list.
3. Select the MSDTC resource that is online and then click **Manage**, or double-click the MSDTC resource name.

   VAM launches the Component Services snap-in in the virtual server context.
4. In the console tree of the Component Services administrative tool, expand **Component Services > Computers > My Computer > Distributed Transaction Coordinator > Local DTC**.
5. Click **Transaction List** to view all transactions, their status, and their identifiers. Right-click a transaction and click **View > Properties** to list the parent transaction and its children.
6. Click **Transaction Statistics** to view statistical information about the transactions in which a server participated.

   You can use transaction statistics to get an overview of DTC performance. Refer to the Microsoft documentation for further information.

The following steps describe how to view DTC transactions from nodes that are not part of the MSDTC Server service group.

To view transactions from any node in the domain

1. Launch the Windows Component Services Administrative tool.

   Click **Start > Programs > Administrative Tools > Component Services**.
2. In the console tree of the Component Services administrative tool, double-click **Component Services**, right-click **Computers**, click **New > Computer**.

3. In the Add Computer dialog box, specify the virtual server name that you specified while creating the MSDTC Server service group. If you are unsure of the exact name, click **Browse** to search from a list of all computers on the network and select the virtual computer name from the list.
4. Click **OK**. The virtual computer entry is added to the Computers container.
5. Expand the newly added virtual computer entry and double-click **Distributed Transaction Coordinator**.
6. Click **Transaction List** to view all transactions, their status, and their identifiers. Right-click a transaction and click **View > Properties** to list the parent transaction and its children.
7. Click **Transaction Statistics** to view statistical information about the transactions in which a server participated.

   You can use transaction statistics to get an overview of DTC performance. Refer to the Microsoft documentation for further information.

Index