Veritas™ Cluster Server Implementation Guide for Microsoft SQL Server 2008

Windows Server 2003

5.1 Application Pack 1

Veritas Cluster Server Implementation Guide for Microsoft SQL Server 2008

Veritas Cluster Server for Windows

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

Third-party legal notices

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

Windows is a registered trademark of Microsoft Corporation.

Licensing and registration

VCS for Windows is a licensed product. See the *Veritas Cluster Server Install and Upgrade Guide* for license installation instructions.

Technical support

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

Contents

Chapter 1 Introducing the VCS agents for SQL Server 2008 and NetApp

About the VCS agents for SQL and NetApp 10
About VCS hardware replication agent for NetApp 11
About the NetApp Filer agent 11
Resource type definition 11
Attribute definitions 12
About the NetApp SnapDrive agent 12
Resource type definition 13
Attribute definitions 13
About the NetApp SnapMirror agent 14
Resource type definition 15
Attribute definitions 15
About the VCS database agent for Microsoft SQL Server 2008 16
About the agent for SQL Server 2008 Database Engine 17
Resource type definition 18
Attribute definitions 18
About the agent for SQL Server 2008 FILESTREAM 21
Resource type definition 21
Attribute definitions 22
About the GenericService agent for SQL Server 2008
Agent and Analysis services 22
About the agent for SQL Server 2008 MSDTC service 23
Resource type definition 23
Attribute definitions 23
SQL Server 2008 sample dependency graph 24
MSDTC sample dependency graph 26
Database monitoring options 27
How the agents make SQL Server highly available 28
Local cluster configuration 28
Disaster recovery configuration 28
Running SQL Server in an Active-Active clustered environment 29
Typical SQL Server 2008 configuration in a VCS cluster 29
Typical disaster recovery configuration 31

Chapter 2 Installing the VCS database agent for SQL Server 2008

About installing the VCS agent for SQL 34
Before installing the agent 34
Installing the agent 35
Configuring the cluster 36
Configuring Web console 48
Configuring notification 49
Configuring Wide-Area Connector process for global clusters 52

Chapter 3 Installing and configuring SQL Server 2008

About installing and configuring SQL 54
Before installing SQL Server 2008 54
Privileges requirements 55
Configuring Microsoft iSCSI Initiator 55
Managing storage using Network Appliance Filer 56
Connecting virtual disks to the cluster node 57
Disconnecting virtual disks from the cluster nodes 58
Managing storage using Windows Logical Disk Manager 59
About LDM support 60
Reserving disks (if you use Windows LDM) 60
Creating volumes (if you use Windows LDM) 60
Mounting volumes (if you use Windows LDM) 61
Unassigning a drive letter 61
Releasing disks (if you use Windows LDM) 61
Installing and configuring SQL Server 2008 on the first cluster node 62
Installing and configuring SQL Server 2008 on the second cluster node ... 63
Assigning ports for multiple SQL Server instances 64

Chapter 4 Configuring the SQL service group

About configuring the SQL service group 66
Before configuring the SQL service group 66
Configuring a SQL Server 2008 service group 68
Running SnapManager for SQL 74
Creating a SQL Server user-defined database 74
Adding storage agent resources to the SQL service group 75
Verifying the service group configuration 76
Bringing the service group online 76
Taking the service group offline 77
Switching the service group 77
Administering a SQL Server service group 78
Modifying a SQL service group configuration 78
Deleting a SQL service group 79

Chapter 5 Configuring an MSDTC service group for high availability

About configuring the MSDTC service group 82
Before configuring the MSDTC service group 82
Reviewing the configuration 83
Creating an MSDTC Server service group 86
About configuring the MSDTC client 88
    Configuring MSDTC client on Windows 2003 89
About using the virtual MMC viewer 90
Viewing DTC transaction information 90
Verifying the installation 91

Chapter 6 Making a standalone SQL Server highly available

About making a standalone SQL Server 2008 highly available 94
Reviewing the configuration 94
    Sample configuration 95
Install and configure VCS on the standalone SQL Server 2008 96
Verify that SQL Server 2008 databases and logs
    are moved to shared storage 97
Install and configure SQL Server on additional nodes 97
Assign ports for multiple SQL Server instances 98
Configure the VCS SQL server service group 98
Create a SQL Server 2008 user-defined database 99
Verify the configuration 99

Chapter 7 Active-Active configuration

About active-active configuration 102
Reviewing the configuration 102
    Sample configuration 103
Installing VCS and configuring the cluster 105
Configure volumes or virtual disks for SQL server 105
Installing and configuring the first instance of SQL Server 106
Configuring the VCS service group for the first SQL Server instance 106
Creating a SQL Server user-defined database 107
Repeating SQL Server installation for additional instances 107
Verify the Active-Active configuration 108

Chapter 8 Disaster recovery configuration

About disaster recovery configuration 110
What is a disaster recovery solution? 110
Why implement a disaster recovery solution? 110
Understanding replication 111
What needs to be protected in a SQL Server environment? 111
Typical disaster recovery configuration 112
Disaster recovery: New SQL Server 2008 installation 113
Reviewing the configuration 113
Installing VCS and configuring the cluster 113
Configuring volumes or LUNs on the shared storage 114
Installing and configuring SQL Server 2008 at the primary site 114
Configuring the VCS SQL service group 114
Create a parallel environment on the secondary site 115
Configuring DR components 116
Configuring replication using NetApp SnapMirror 116
Configure SnapMirror resource in the SQL service group 117
Configure Global Cluster Option for wide-area failover 118
Linking clusters: Adding a remote cluster to a local cluster 119
Converting a local service group to a global service group 120
Bringing a global service group online 122
Administering global service groups 123
Deleting a remote cluster 124

Chapter 9 Removing VCS agents for SQL Server 2008

Prerequisites for removing VCS agents 129
Uninstalling the agents 129

Chapter 10 Troubleshooting the agents

About troubleshooting VCS agents 132
VCS logging 132
VCS Cluster Configuration Wizard logs 133
VCWsilent logs 134
Network Appliance agents error messages 134
SQL agent error messages and descriptions 136
Agent for MSDTC 136
Agent for SQL Server 2008 137
Agent for SQL Server 2008 FILESTREAM 141

Index

143

Chapter 1

Introducing the VCS agents for SQL Server 2008 and NetApp

This chapter contains the following topics:

- "About the VCS agents for SQL and NetApp" on page 10
- "About VCS hardware replication agent for NetApp" on page 11
- "About the NetApp Filer agent" on page 11
- "About the NetApp SnapDrive agent" on page 12
- "About the NetApp SnapMirror agent" on page 14
- "About the VCS database agent for Microsoft SQL Server 2008" on page 16
- "About the agent for SQL Server 2008 Database Engine" on page 17
- "About the agent for SQL Server 2008 FILESTREAM" on page 21
- "About the GenericService agent for SQL Server 2008 Agent and Analysis services" on page 22
- "About the agent for SQL Server 2008 MSDTC service" on page 23
- "SQL Server 2008 sample dependency graph" on page 24
- "MSDTC sample dependency graph" on page 26
- "Database monitoring options" on page 27
- "How the agents make SQL Server highly available" on page 28
- "Running SQL Server in an Active-Active clustered environment" on page 29
- "Typical SQL Server 2008 configuration in a VCS cluster" on page 29

■ "Typical disaster recovery configuration" on page 31

About the VCS agents for SQL and NetApp

The VCS database agent for Microsoft SQL Server 2008 provides high availability to SQL Server 2008. The VCS hardware replication agent for Network Appliance SnapMirror enables configuring Network Appliance 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 Network Appliance filers for shared storage. The agents also support disaster recovery configurations set up using the VCS Global Cluster Option (GCO) and Network Appliance SnapMirror for data replication.

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

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

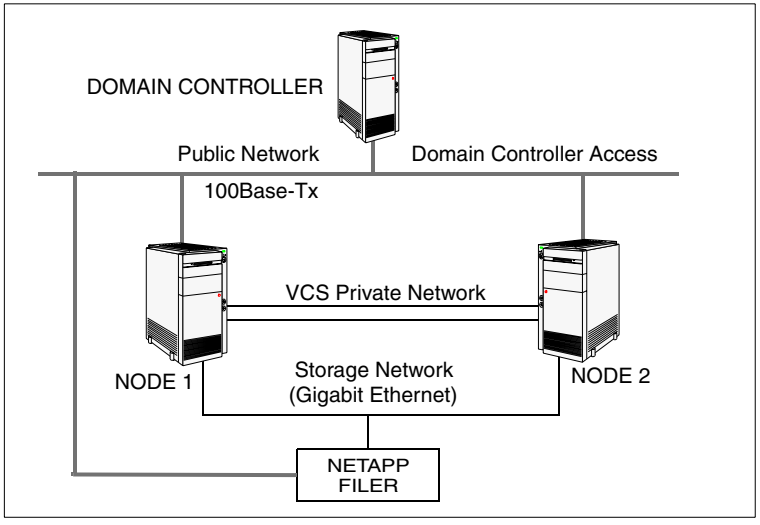

This chapter provides an overview of the agents. You may also want to review the other information about the agents including their resource type definitions and attribute definitions.

See "Resource type definitions" on page 197.

About VCS hardware replication agent for NetApp

The VCS hardware replication agent for Network Appliance provides failover support and recovery in environments employing Network Appliance filers for storage and Network Appliance 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 includes the NetAppFiler agent, the NetAppSnapDrive agent, and the 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.

The agent function includes the following:

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

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
    )
```

Attribute definitions

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

Table 1-1 describes the required attributes associated with the VCS agent for NetApp filer.

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.

Specific agent functions include the following:

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 connectivity 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).

Resource type definition

The 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
)
```

Attribute definitions

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

Table 1-2 describes the required attributes associated with the VCS agent for NetApp SnapDrive.

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

Table 1-2 NetApp SnapDrive agent attributes

Attribute	Description
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.

Specific agent functions include the following:

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 re-synchronize the data manually. The Online entry point touches a lock file if Read Write access is enabled successfully.

Offline — Removes the lock file. The agent does not perform any filer operations because an offline entry point does not necessarily indicate an intention to give up the devices.

Monitor — Verifies the lock file exists. If the lock file exists, the monitor entry point reports the status of the resource as online. If the lock file does not exist, the monitor entry point reports the status of the resource as offline.

Open — Removes the lock file, thereby preventing potential concurrency violation if the group fails over to another node.

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

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

Resource type definition

The 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 }
        str FilerResName
        str VolumeName
        str SnapMirrorArguments
        str SnapMirrorSchedule
        str AppResName
)
```

Attribute definitions

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

Table 1-3 describes the required attributes associated with the VCS agent for NetApp SnapMirror.

Table 1-3 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 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. By default, this attribute is not assigned any value. Type and dimension: string-scalar

Table 1-3 NetApp SnapMirror agent attributes (continued)

Attribute	Description
AppResName	Name of the resource configured to monitor the application being made highly available. Type and dimension: string-scalar

About the VCS database agent for Microsoft SQL Server 2008

The VCS database agent for Microsoft SQL Server 2008 provides high availability for Microsoft SQL Server 2008 in a VCS cluster. The agent monitors Microsoft SQL Server RDBMS and its services on a VCS cluster to ensure high availability. The agent detects an application failure if a configured virtual server becomes unavailable. When this occurs, 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 agents for SQL Server 2008 monitor specific resources within an enterprise application, determines the status of these resources, brings them online, and takes them offline. The database agent also provides "Active-Active" support for SQL Server. In an Active-Active configuration, several SQL server instances are intended to run on a single node when necessary.

The VCS database agent package for SQL Server 2008 includes the following:

- Agent for SQL Server 2008 Database Engine
  The agent provides high availability for SQL Server 2008 Database Engine. This agent also monitors full-text search service, an optional component that is integrated with the Database Engine.
  The SQL Server 2008 agent monitors the SQL Server 2008 service and the integrated Full-Text Search service. If the SQL Server 2008 Database Engine service is not running, the agent returns a failure status and declares the state as OFFLINE. Depending on the detail monitoring configuration, the agent checks the health of critical SQL databases or executes a monitoring script. If the SQL detail monitoring is successful, the agent declares the service group as online.
- Agent for SQL Server 2008 FILESTREAM
  The agent provides high availability for SQL Server 2008 FILESTREAM feature. The agent monitors the Windows FILESTREAM configuration settings for the SQL Server instance.

- GenericService agent for SQL Server 2008 Agent service and Analysis service
  VCS employs the GenericService agent to provide high availability for the SQL Server 2008 Agent service and the Analysis service. The GenericService agent monitors the SQL Server 2008 Agent and Analysis services. If the services are not running, the agent declares the services as OFFLINE.
- Agent for 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 2008 Database Engine

This SQL Server 2008 agent monitors the SQL Server Database Engine service. As Full-text search is an integrated optional component for SQL Server Database Engine, when installed and configured, the agent also monitors the full-text search service. The agent brings the SQL Server 2008 service online, monitors the status, and takes it offline.

Specific agent functions include the following:

Online	Brings the SQL Server 2008 service online.
Offline	Takes the SQL Server 2008 service offline.
Monitor	Queries the Service Control Manager (SCM) for the status of SQL Server 2008 services. Also, if detail monitoring is configured, the agent performs a database health check depending on the configuration. See "Database monitoring options" on page 27.
Clean	Forcibly stops the SQL Server service.

Resource type definition

The agent for SQL Server 2008 is configured as a resource of type SQLServer2008.

```
type SQLServer2008 (
    static i18nstr ArgList[] = { Instance,
    "LanmanResName:VirtualName", SQLOnlineTimeout,
    SQLOfflineTimeout, DetailMonitorInterval,
    SQLDetailMonitorTimeout, Username, Domain, Password, DBList,
    SQLFile, FaultOnDMFailure, "LanmanResName:IPResName" }
    str Instance
    str LanmanResName
    int SQLOnlineTimeout = 90
    int SQLOfflineTimeout = 90
    int DetailMonitorInterval = 0
    int SQLDetailMonitorTimeout = 30
    i18nstr Username
    i18nstr Domain
    str Password
    i18nstr DBList[]
    i18nstr SQLFile
    boolean FaultOnDMFailure = 1
)
```

Attribute definitions

Review the following information to familiarize yourself with the agent attributes for a SQLServer2008 resource type.

Table 1-4 describes the required attributes associated with the VCS agent for SQL Server 2008 Database Engine.

Table 1-4 SQL Server 2008 agent required attributes

Required Attributes	Definition
Instance	Name of SQL Server instance to monitor. If the attribute is blank, the agent monitors the default instance. Type and dimension: string-scalar
LanmanResName	The Lanman resource name on which the SQL Server 2008 resource depends. Type and dimension: string-scalar
SQLOnlineTimeout	Number of seconds that can elapse before online entry point aborts. Default is 90. Type and dimension: integer-scalar

Table 1-4 SQL Server 2008 agent required attributes

Required Attributes	Definition
SQLOfflineTimeout	Number of seconds that can elapse before offline entry point aborts. Default is 90. Type and dimension: integer-scalar

Table 1-5 describes the optional attributes associated with the VCS agent for SQL Server 2008 Database Engine.

Table 1-5 SQL Server 2008 agent optional attributes

Required Attributes	Definition
DetailMonitorInterval	Defines whether the agent performs detail monitoring of SQL Server 2008 database. If set to 0, the agent will not monitor the database in detail. A non-zero value indicates the number of online monitor cycles that the agent must wait before performing detail monitoring. Default = 5 **Note:** If the attribute is set to a non-zero value, and script-based detail monitoring is configured, then the attributes Username, Password, Domain, SQLDetailMonitorTimeOut, and SQLFile must be assigned appropriate values. Type and dimension: integer-scalar
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 not. Default = 1 Type and dimension: boolean
SQLDetailMonitor Timeout	Number of seconds that can elapse before the detail monitor routine aborts. Default is 30. Type and dimension: integer-scalar

Table 1-5 SQL Server 2008 agent optional attributes (continued)

Required Attributes	Definition
Username	The Microsoft Windows authentication name when logging in to a database for detail monitoring. This attribute must not be null if DetailMonitorInterval 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
Domain	Domain for the user account. This attribute is used to create a trusted connection to the SQL Server 2008 instance if DetailMonitorInterval 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
Password	Password for logging in to a database for in-depth monitoring. This attribute must not be null if DetailMonitorInterval attribute is set to a non-zero value and script-based detail monitoring is configured. Type and dimension: string-scalar
SQLFile	The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the DetailMonitorInterval 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

About the agent for SQL Server 2008 FILESTREAM

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

The agent for SQL Server 2008 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 agent for SQL Server 2008 FILESTREAM is configured as a resource of type SQLFilestream.

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

Attribute definitions

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

Table 1-6 describes the required attributes associated with the VCS agent for SQL Server 2008 FILESTREAM.

Table 1-6 SQL Server 2008 Filestream agent required attributes

Required Attributes	Definition
InstanceName	The name of the SQLServer2008 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

About the GenericService agent for SQL Server 2008 Agent and Analysis services

VCS uses the GenericService agent to make the SQL Server 2008 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 2008 services online.
Offline	Takes the configured SQL Server 2008 services offline.
Monitor	Queries the Service Control Manager (SCM) for the status of configured SQL Server 2008 services. See "Database monitoring options" on page 27.
Clean	Forcibly stops the configured SQL Server 2008 services.

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

About the agent for SQL Server 2008 MSDTC service

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

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

Resource type definition

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

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

Table 1-7 describes the required attributes associated with the VCS agent for MSDTC.

Table 1-7 MSDTC agent required attributes

Required Attributes	Definition
LanmanResName	Name of the Lanman resource on which the MSDTC resource depends. Type and dimension: string-scalar

Table 1-7 MSDTC agent required attributes

Required Attributes	Definition
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 2008 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 2008 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 2008 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 2008 resource comes online after each of these resources are brought online. The SQL Server 2008 Analysis service (MSOlap) and SQL Server 2008 Agent service (SQLServerAgent) are configured as GenericService resources.

Figure 1-2 SQL Server 2008 service group dependency graph

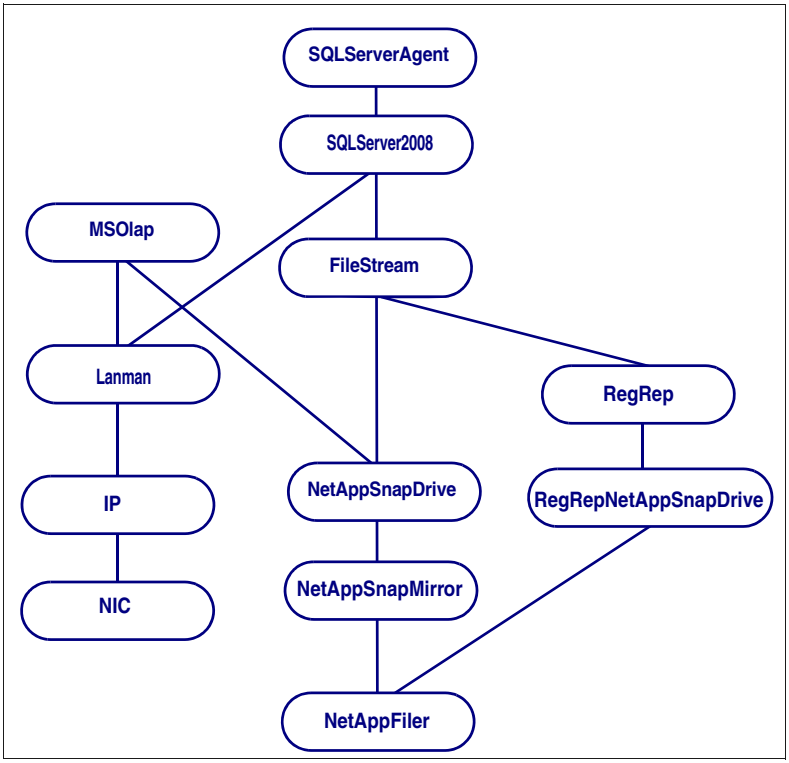

MSDTC sample dependency graph

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 MSDTC service group dependency graph

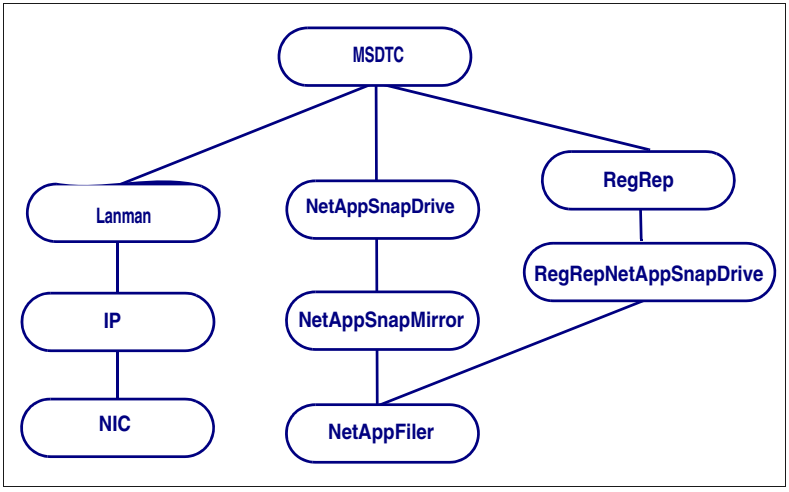

Database monitoring options

Use the detail monitoring capability of the VCS database agent for SQL Server 2008 to monitor the status of a database. The VCS database agent for SQL Server 2008 provides two levels of application monitoring: basic and detail. Basic monitoring queries the Windows Service Control Manager (SCM) to verify whether the configured SQL Server services are continuously active. Detail monitoring queries the database to verify the availability of the database instance.

You can configure detail monitoring in the following ways:

- DBList detail monitoring

  The SQL Server agent monitors only the list of databases specified in the SQL Server 2008 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\SQLServer2005\sample_script.sql`

  Here, the variable `%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.

You can enable and configure detail monitoring by running the SQL Server 2008 Configuration Wizard for VCS. Refer to the instructions for configuring a SQL Server service group for more information.

See "Configuring the VCS SQL Server 2008 service group" on page 146.

Note: If you start the SQL server services from outside VCS, then the SQL resource will go in an UNKNOWN state, because the VCS database agent for Microsoft SQL Server 2008 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.

How the agents make SQL Server highly available

The VCS database agent for Microsoft SQL Server 2008 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

When the VCS database agent for Microsoft SQL Server 2008 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

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.

The fail over involves the following steps:

- Connecting the virtual disks (LUNs) to the target hosts (using the NetAppSnapDrive agent)
- Performing a mirror break, which enables write access to the target (using the NetAppSnapMirror agent)
- Reversing the direction of replication by demoting the original source to a target, and begin replicating from the new source (using the NetAppSnapMirror agent)

- Starting the SQL services on the remote node (using the VCS database agent for SQL Server)

Running SQL Server in an Active-Active clustered environment

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

SQL Server 2008 allows multiple independent instances of SQL Server to run on a single machine. Using this feature, the VCS database agent for Microsoft SQL Server 2008 supports SQL Server in an Active-Active environment by allowing a node to run up as many instances as supported by SQL. A SQL Server instance can fail over to any of the other cluster nodes that you specify when you configure the SQL Server service group.

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

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

Typical SQL Server 2008 configuration in a VCS cluster

A typical SQL Server 2008 configuration in a VCS cluster involves two cluster nodes accessing a shared storage. The SQL Server binaries are installed on the cluster nodes. The shared storage is used to store SQL Server data files and the MSDTC log files. The cluster nodes access the shared storage. The shared storage can be managed using Network Appliance suite of products.

The cluster nodes are configured to host the SQL Server 2008 resource, the SQL Server 2008 FILESTREAM resource, the SQL Server 2008 Analysis and Agent service resources. The MSDTC resource can be configured on the same cluster nodes. You need not configure an MSDTC client if the MSDTC resource is configured on the same nodes that have SQL Server 2008 resource configured. However, if the MSDTC resource is configured on other nodes, you must configure an MSDTC client to point to the virtual server name of the MSDTC resource.

Figure 1-4 on page 30 depicts a two node cluster hosting a SQL Server 2008 service group with the different services configured. MSDTC resource is also configured on the same nodes.

Figure 1-4 Typical SQL Server 2008 configuration in a VCS cluster

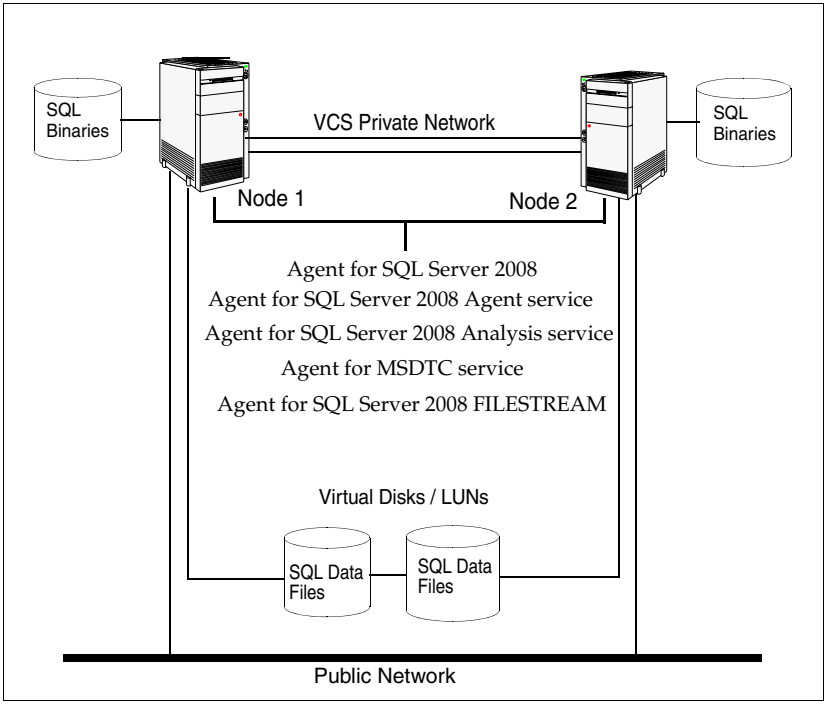

Typical disaster recovery configuration

A Disaster Recovery (DR) configuration enables you to restore application data and services in the event of a catastrophic failure. A typical DR solution requires primary and secondary sites, and clusters within those sites. The cluster at the primary site provides data and services during normal operation, and the cluster at the secondary site provides data and services if the primary site fails.

Figure 1-5 on page 31 displays an environment with a DR solution that is prepared for a disaster. In this case, the primary site consists of two nodes, System1 and System2. Similarly the secondary setup consists of two nodes, System3 and System4. Each site has a clustered setup with the nodes set up appropriately for failover within the site.

Figure 1-5 Typical DR configuration in a VCS cluster

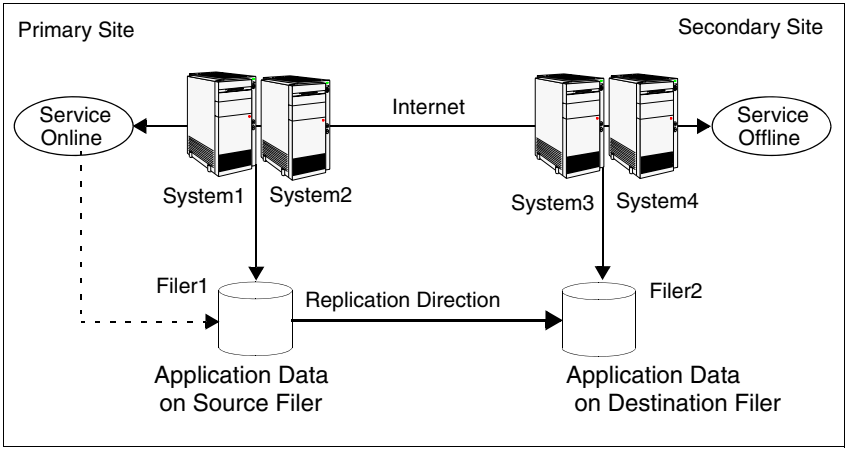

Filer1 in the cluster on the primary site replicates to Filer2 in the cluster on the secondary site. Replication between the filers is set up using NetApp SnapMirror for SQL. Refer to NetApp documentation for more information on replication using NetApp filers.

If the Microsoft SQL Server server on System1 fails, SQL Server comes online on node System2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal.

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

Chapter 2

Installing the VCS database agent for SQL Server 2008

This chapter contains the following topics:

- "About installing the VCS agent for SQL" on page 34
- "Before installing the agent" on page 34
- "Installing the agent" on page 35
- "Configuring the cluster" on page 36

About installing the VCS agent for SQL

This chapter describes how to install the VCS database agent for SQL Server 2008 and then configure the cluster. The agent installation is required only if you have not already installed the agent while installing the Application Pack 1.

Note: This guide contains instructions on setting up a SQL Server 2008 configuration in a VCS for Windows cluster environment. If you already have SQL Server 2000 or SQL Server 2005 configured with VCS 5.1 and wish to upgrade the configuration to Application Pack 1 and SQL 2008, refer to the *Veritas Cluster Server 5.1 for Windows Application Pack 1 Release Notes* for more information.

Before installing the agent

Ensure that the following prerequisites are met before installing the agent:

- Verify that you have installed Veritas Cluster Server 5.1 for Windows on all cluster nodes. Refer to the *Veritas Cluster Server 5.1 Installation and Upgrade Guide* for instructions.
- Ensure that you have installed Veritas Cluster Server 5.1 Application Pack 1 for Windows on all the cluster nodes. Refer to the *Veritas Cluster Server 5.1 Application Pack 1 for Windows Release Notes* for instructions.
- Ensure that you have local administrator privileges on the node where you are installing the agent.

Installing the agent

Complete these steps to install the VCS database agent for SQL Server 2008 on a cluster node. Repeat these steps on all nodes on which you wish to configure SQL Server 2008 with VCS.

Perform these steps only if you did not install the agent while installing VCS 5.1 Application Pack 1 for Windows.

To install the agent

1. Launch the installer for VCS 5.1 for Windows Application Pack 1. In the Windows Add/Remove Programs window, click **Veritas Cluster Server 5.1 for Windows Application Pack 1** and click **Change**.
2. In the VCS 5.1 for Windows dialog box, choose the **Add or Remove** option and click **Next**.
3. In the VCS product options dialog box, check the **Microsoft SQL Server 2008 support** check box and click **Next**.
   The disk space required for the installation is displayed toward the bottom of the screen. When you add or remove an option, the total space changes.
4. The installer validates the system for prerequisites. After the system is accepted, click **Next**.
   If a system is rejected, the Comments column displays the cause for rejecting the system. Highlight the system to view a detailed information about the failure in the Details box. Resolve the error, highlight the system from the list, and click **Validate Again**.
5. Review the summary of your selections and click **Update** to start the installation. The installer displays the status of installation.
6. After the installation is complete, review the installation report and click **Next**.
7. Click **Finish**.
   After the agent is successfully installed, proceed to configuring the cluster.

Configuring the cluster

After installing the VCS database agent for SQL Server 2008 using the product installer, set up the components required to run a cluster. The VCS Cluster Configuration Wizard (VCW) sets up the cluster infrastructure, including LLT and GAB, and configures Symantec Product Authentication Service in the cluster. The wizard also configures the ClusterService group, which contains resources for the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console, notification, and global clusters.

- If you plan to set up a disaster recovery environment, configure the wide-area connector process for global clusters.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Domain Account Operators group.
- When configuring a user account for the VCS Helper service, make sure that the user account is a domain user. The VCS HAD, which runs in the context of the local system built-in account, uses the VCS Helper service user context to access the network. This account does not require domain admin privileges.
- Make sure the VCS Helper service domain user account has “Add workstations to domain” privilege enabled in the Active Directory.
- In case of NetApp, the user account for the VCS Helper service must have administrative privileges on the NetApp filer.

To configure a VCS cluster

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

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

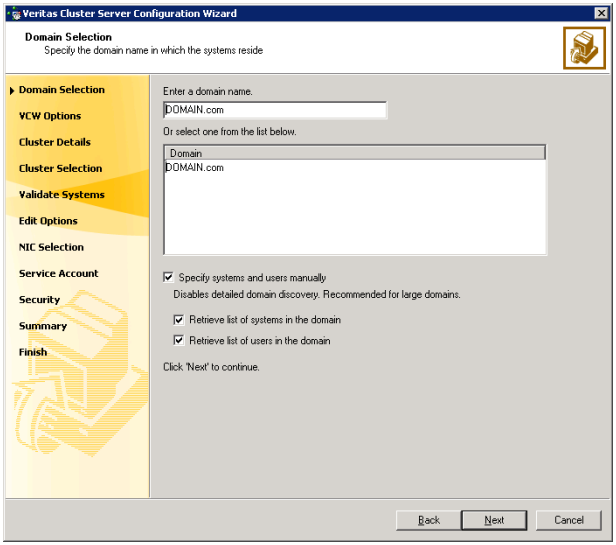

To discover information about all systems and users in the domain:

- Clear the **Specify systems and users manually** check box.
- Click **Next**.
  Proceed to step 8 on page 39.

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

- Check the **Specify systems and users manually** check box.
  Additionally, you may instruct the wizard to retrieve a list of systems and users in the domain by selecting appropriate check boxes.
- Click **Next**.
  If you chose to retrieve the list of systems, proceed to step 6 on page 38. Otherwise, proceed to the next step.

5 On the System Selection panel, type the name of each system to be added, click **Add**, and then click **Next**. Do not specify systems that are part of another cluster.

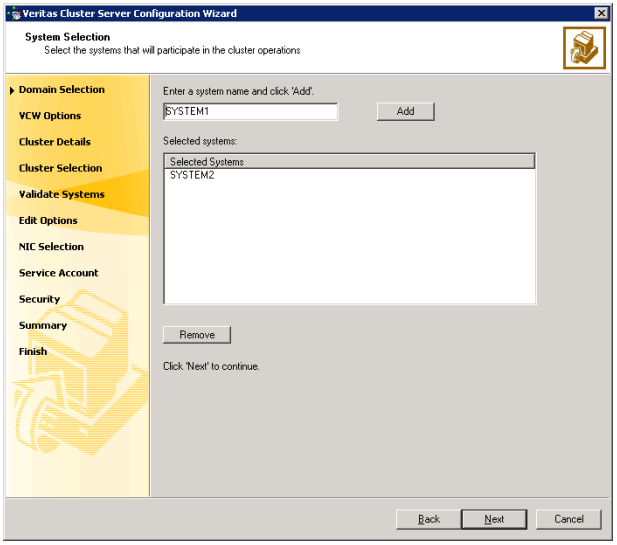

Proceed to step 8 on page 39.

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

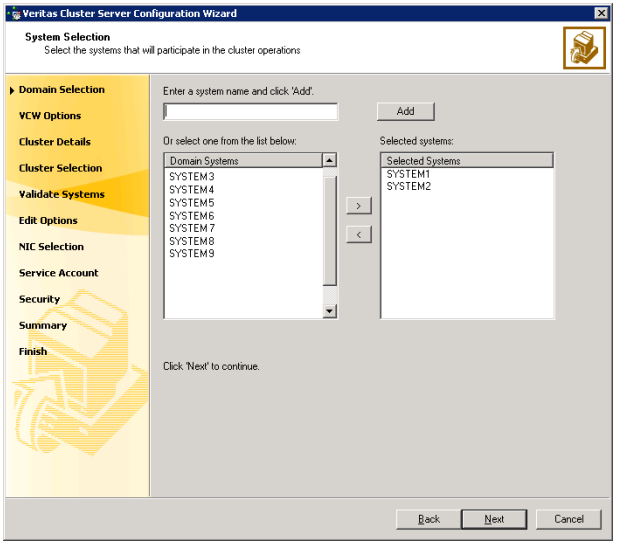

Enter the name of the system and click **Add** to add the system to the **Selected Systems** list, or click to select the system in the Domain Systems list and then click the **>** (right-arrow) button.

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

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

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

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

Click **Next** to proceed.

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

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

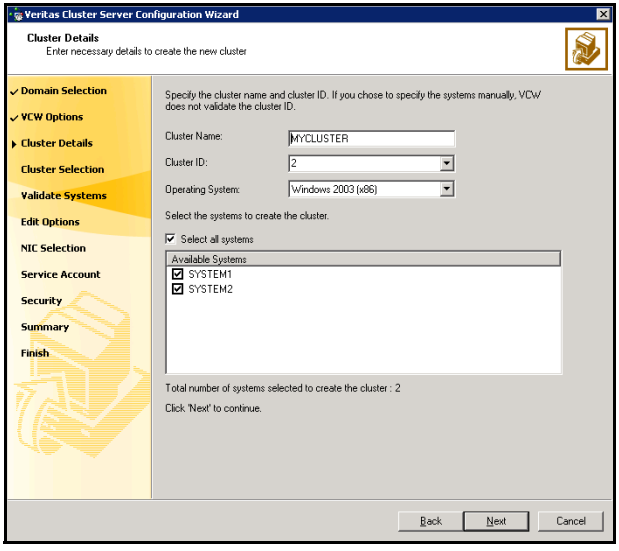

Cluster Name	Type a name for the new cluster. Symantec recommends a maximum length of 32 characters for the cluster name.
Cluster ID	Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 255. **Caution:** If you chose to specify systems and users manually in step 4 or if you share a private network between more than one domain, make sure that the cluster ID is unique.
Operating System	From the drop-down list, select the operating system that the systems are running.
Available Systems	Select the systems that will be part of the cluster. The wizard discovers the NICs on the selected systems. For single-node clusters with the required number of NICs, the wizard prompts you to configure a private link heartbeat. In the dialog box, click **Yes** to configure a private link heartbeat. Check the **Select all systems** check box to select all the systems simultaneously.

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 on page 39, proceed to the next step. Otherwise, proceed to step 12 on page 43.

11 On the Private Network Configuration panel, configure the VCS private network and click **Next**.
Do one of the following:

- To configure the VCS private network over Ethernet

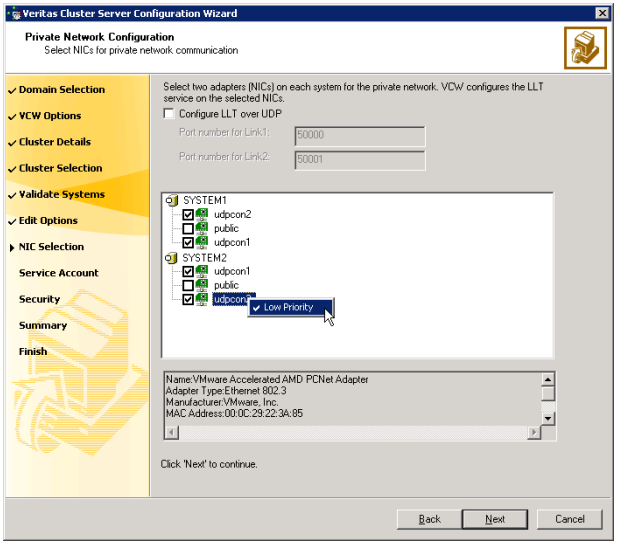

  - Select the check boxes next to the two NICs to be assigned to the private network.

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

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

The wizard will configure the LLT service (over Ethernet) on the selected network adapters.

- To configure the VCS private network over the User Datagram Protocol (UDP) layer

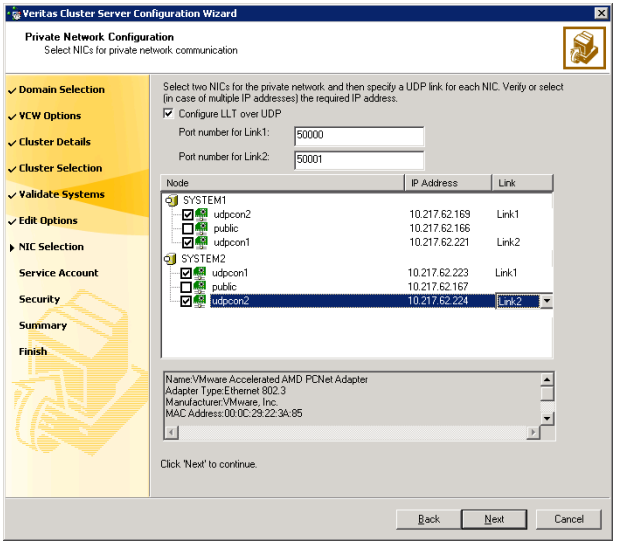

  - Check the **Configure LLT over UDP** check box.
  - Specify a unique UDP port in the **Port number for Link1** and **Port number for Link2** fields. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively.
  - Select the check boxes next to the two NICs to be assigned to the private network. 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. Each IP address can be in a different subnet.

    The IP address is used for the VCS private communication over the specified UDP port.
  - For each selected NIC, double-click the respective field in the Link column and choose a link from the drop-down list. Specify a different link (Link1 or Link2) for each NIC. Each link is associated with a UDP port that you specified earlier.

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

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

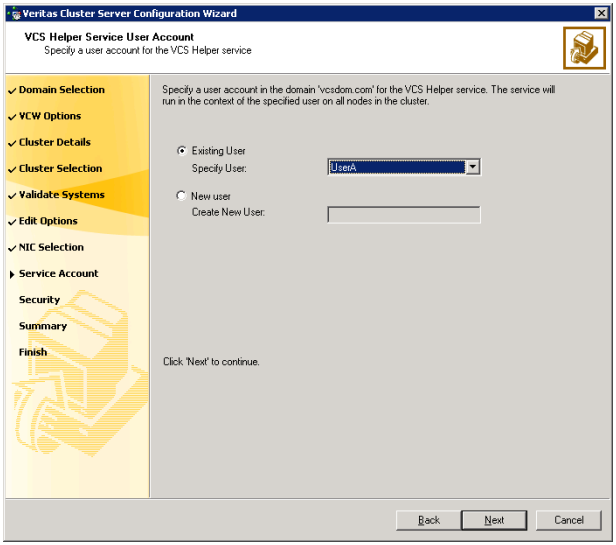

- To specify an existing user, do one of the following:
  - Click **Existing user** and select a user name from the drop-down list,
  - If you chose not to retrieve the list of users in step 4 on page 37, type the user name in the **Specify User** field, and then click **Next**.
- To specify a new user, 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 and then click **Next**.

Do one of the following:

- To use the single sign-on feature

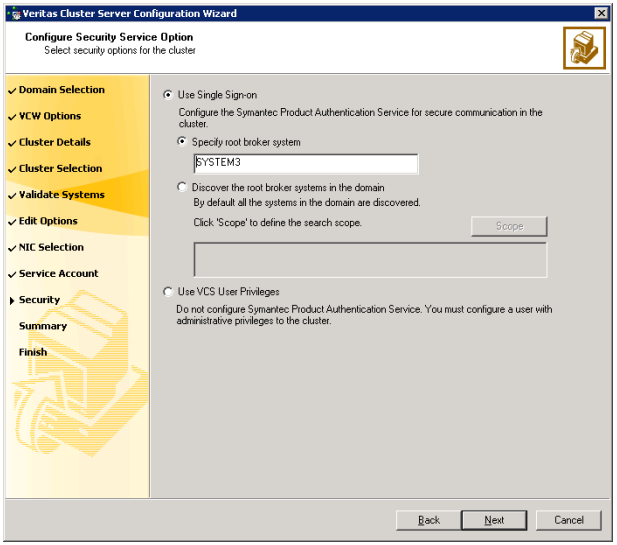

  - Click **Use Single Sign-on**. In this mode, VCS uses SSL encryption and platform-based authentication. The VCS engine (HAD) and Veritas Command Server run in secure mode.

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

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

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

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

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

- If you want to define a search criteria, click **Scope**. In the Scope of Discovery dialog box, click **Entire Domain** to search across the domain, or click **Specify Scope** and select the Organization Unit from the Available Organizational Units list, to limit the search to the specified organization unit. Use the Filter Criteria options to search systems matching a certain condition.

  For example, to search for systems managed by a user *Administrator*, select **Managed by** from the first drop-down list, **is (exactly)** from the second drop-down list, type the user name **Administrator** in the adjacent field, click **Add**, and then click **OK**. Table 2-1 contains some more examples of search criteria.

Table 2-1 Search criteria examples

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

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

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

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

- To use VCS user privilege:

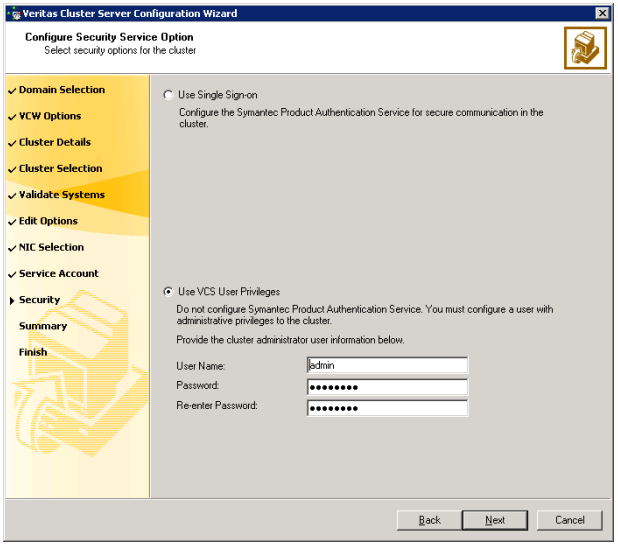

  - Click **Use VCS User Privileges**.
    The default user name for the VCS administrator is *admin* and the default password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password. Symantec recommends that you specify a new user name and password.
    Use this account to log on to VCS using Cluster Management Console (Single Cluster Mode) or Web Console, when VCS is not running in secure mode.
  - Click **Next**.

14 Review the summary information on the Summary panel, and click **Configure**. The wizard configures the VCS private network.
If the selected systems have LLT or GAB configuration files, the wizard displays an informational dialog box before overwriting the files. In the dialog box, click **OK** to overwrite the files. Otherwise, click **Cancel**, exit the wizard, move the existing files to a different location, and rerun the wizard.
The wizard starts running commands to configure VCS services. If an operation fails, click **View configuration log file** to see the log.

15 On the Completing Cluster Configuration panel, click **Next** to configure the ClusterService service group; this group is required to set up components for the Cluster Management Console (Single Cluster Mode) or Web Console, notification, and for global clusters.

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

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

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 the system to the cluster.

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

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

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

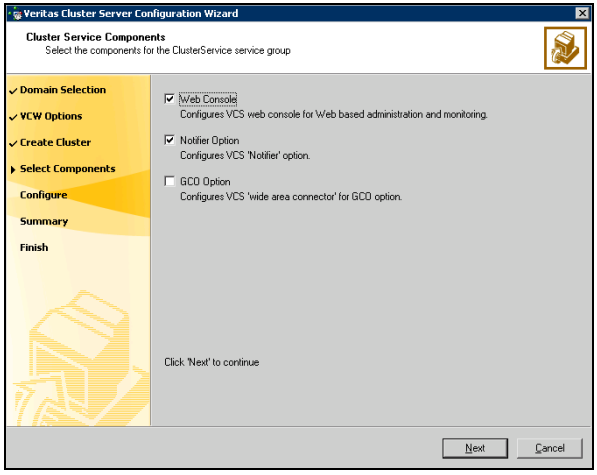

- Check the **Web Console** checkbox to configure the Cluster Management Console (Single Cluster Mode), also referred to as the Web Console.
  See "Configuring Web console" on page 48.

- Check the **Notifier Option** checkbox to configure notification of important events to designated recipients.
  See "Configuring notification" on page 49.

Configuring Web console

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

To configure the Web console

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

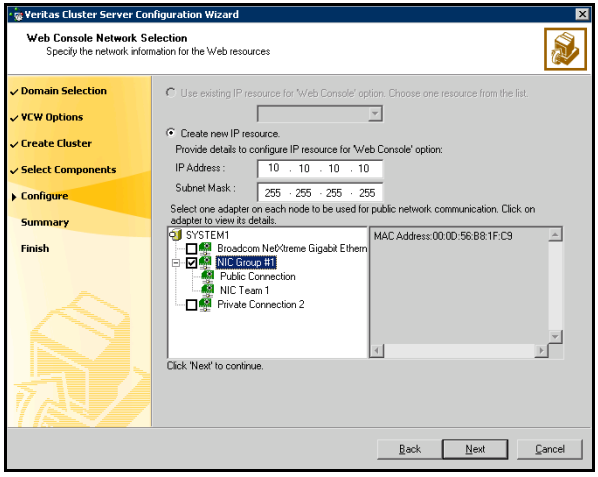

- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address for the Web console.
- If you choose to configure a new IP address, type the IP address and associated subnet mask.
- Select a network adapter for each node in the cluster. Note that the wizard lists the public network adapters along with the adapters that were assigned a low priority.

2 Review the summary information and choose whether you want to bring the Web Console resources online when VCS is started, and click **Configure**.

3 If you chose to configure a Notifier resource, proceed to:
"Configuring notification" on page 49.
Otherwise, click **Finish** to exit the wizard.

Configuring notification

This section describes steps to configure notification.

To configure notification

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

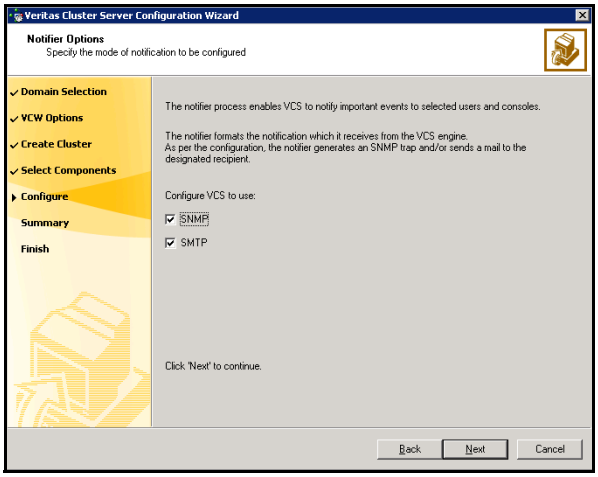

You can configure VCS to generate SNMP (V2) traps on a designated server and send emails to designated recipients in response to certain events.

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

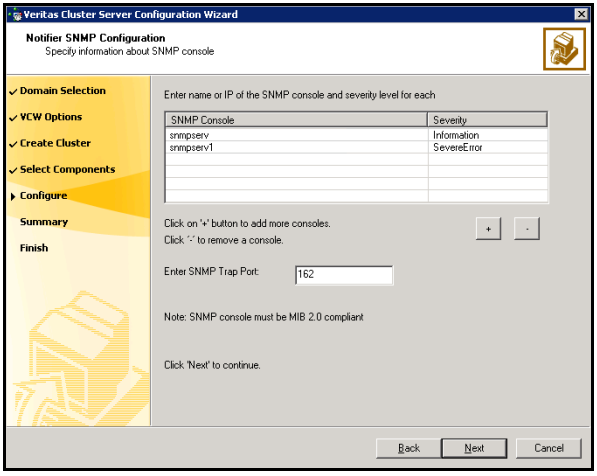

- Click a field in the SNMP Console column and type the name or IP address of the console. The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the Severity column and select a severity level for the console.
- Click '**+**' to add a field; click '**-**' to remove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

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

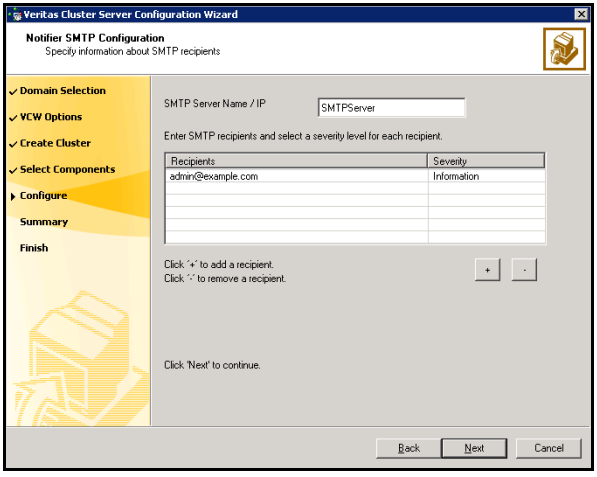

- Type the name of the SMTP server.
- Click a field in the Recipients column and enter a recipient for notification. Enter recipients as admin@example.com.
- Click the corresponding field in the Severity column and select a severity level for the recipient. VCS sends messages of an equal or higher severity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

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

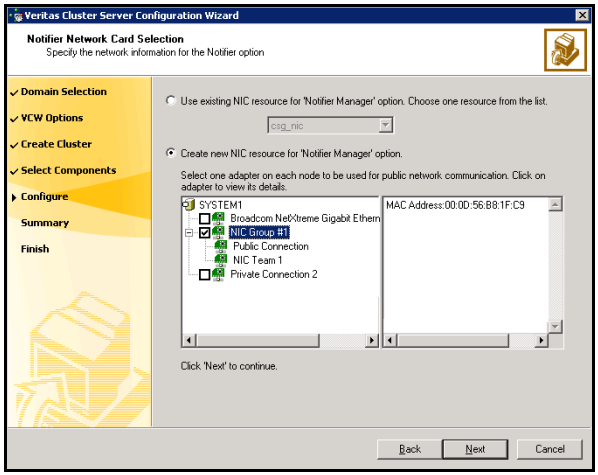

- If the cluster has a ClusterService service group configured, you can use the NIC resource configured in the service group or configure a new NIC resource for notification.
- If you choose to configure a new NIC resource, select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.

5 Review the summary information and choose whether you want to bring the notification resources online when VCS is started.

6 Click **Configure**.

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

Configuring Wide-Area Connector process for global clusters

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

To configure the wide-area connector process for global clusters

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

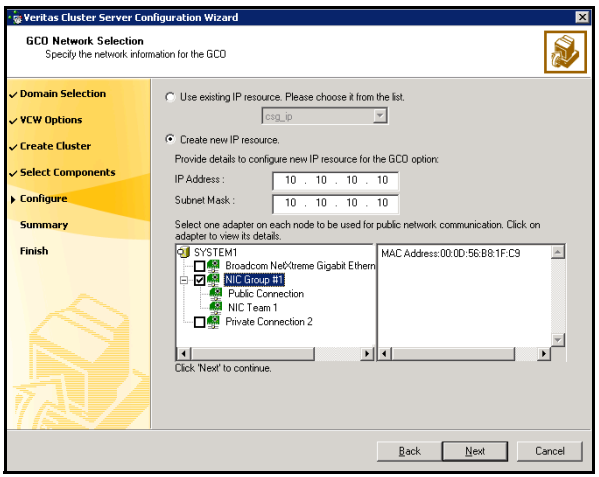

- If the cluster has a ClusterService service group configured, you can use the IP address configured in the service group or configure a new IP address.
- If you choose to configure a new IP address, enter the IP address and associated subnet mask. Make sure that the specified IP address has a DNS entry.
- Select a network adapter for each node in the cluster. The wizard lists the public network adapters along with the adapters that were assigned a low priority.

2 Review the summary information and choose whether you want to bring the resources online when VCS starts and click **Configure**.

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

Chapter 3

Installing and configuring SQL Server 2008

This chapter contains the following topics:

- "About installing and configuring SQL" on page 54
- "Before installing SQL Server 2008" on page 54
- "Managing storage using Network Appliance Filer" on page 56
- "Managing storage using Windows Logical Disk Manager" on page 59
- "Installing and configuring SQL Server 2008 on the first cluster node" on page 62
- "Installing and configuring SQL Server 2008 on the second cluster node" on page 63
- "Assigning ports for multiple SQL Server instances" on page 64

About installing and configuring SQL

This chapter provides information for installing and configuring SQL Server 2008 in a VCS for Windows environment. This environment uses an active-passive configuration with one to one failover capabilities.

To configure MSDTC service groups, see "Configuring an MSDTC service group for high availability" on page 81.

If you already have an existing SQL Server deployment, and you wish to make it highly available using VCS, see "Making a standalone SQL Server highly available" on page 93.

If you are planning to deploy an active-active configuration with multiple SQL instances, see "Active-Active configuration" on page 101.

Before installing SQL Server 2008

Ensure that the following prerequisites are met before you proceed:

- Verify that VCS 5.1 for Windows Application Pack 1 along with the VCS database agent for SQL 2008 are installed on all the nodes in the cluster.
- Verify that you have configured a VCS cluster. You can use the VCS Cluster Configuration Wizard (VCW) for cluster configuration.
  See "Configuring the cluster" on page 36.
- If using iSCSI, verify that the Microsoft iSCSI Initiator is configured to establish a persistent connection between the Network Appliance filer and the cluster nodes. See the Microsoft documentation for instructions.
- Make sure that you have an external, basic disk to create volumes or LUNs (virtual disk).
  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 Network Appliance Filer" on page 56.
  See "Managing storage using Windows Logical Disk Manager" on page 59.
- If your cluster has an Exchange service group configured, make sure to install SQL Server 2008 on a node that is not in the SystemList attribute for the Exchange service group.

Privileges requirements

The following privileges are required:

- You must be a domain user. The domain user must have administrator privileges.
- You must be a member of the Local Administrators group on all nodes on which Microsoft SQL Server 2008 is installed.
- You must have write permissions for the Active Directory objects corresponding to these nodes.
- You must have write permissions on the DNS server to perform DNS updates.

Configuring Microsoft iSCSI Initiator

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

To configure Microsoft iSCSI initiator

1. Make sure the Microsoft iSCSI Initiator software is installed on all cluster nodes. Refer to Microsoft documentation for further information.
2. Double-click the Microsoft iSCSI Initiator icon from the desktop to launch the Microsoft iSCSI configuration manager.
3. On the Target Portals tab, click **Add**.
4. In the Add Target Portals dialog box, specify the IP address or the DNS name for the Network Appliance Filer and then click **OK**.
5. On the Persistent Target tab, verify that the newly added target portal is listed under the Select a target box, and then click **OK**.

Managing storage using Network Appliance Filer

Network Appliance manages data by creating volumes on physical disks. These volumes can further be divided into Logical Unit Numbers (LUNs). The LUNs are accessible from the cluster nodes, provided the nodes have Microsoft iSCSI Initiator and Network Appliance SnapDrive installed. If you plan to use Fibre Channel (FC) for connecting the LUNs, ensure that you install the FCP Attach Kit or Windows Host Utilities on all the cluster nodes. Refer to the NetApp documentation for more information.

Note: Symantec does not support volumes created using qtree.

Figure 3-1 VCS cluster in a NetApp storage environment

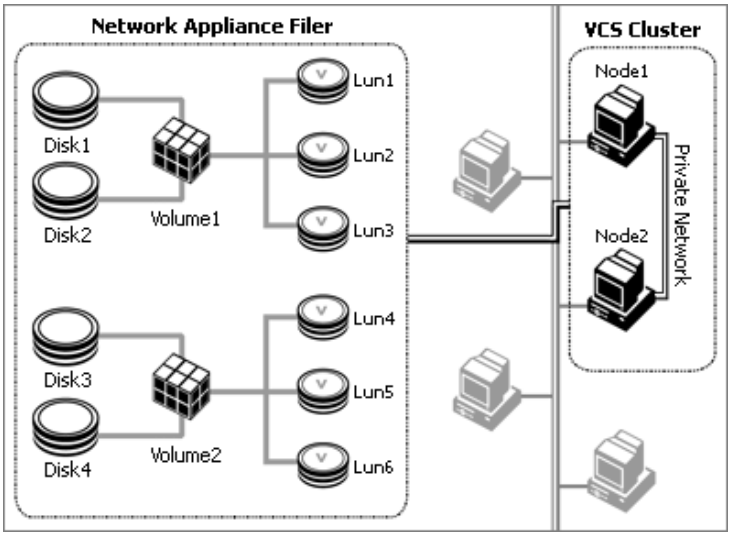

The VCS database agent for Microsoft SQL Server 2008 requires two LUNs to be created on the Network Appliance filer, one for SQL Server data and the other for the registry replication information. If you are using SQL Server 2008 FILESTREAM, create additional LUNs for FILESTREAM enabled database objects. If you are configuring an MSDTC Server service group, create additional LUNs for MSDTC log and MSDTC registry replication. These LUNs must be accessible from all cluster nodes.

For example, you can create the following LUNs to manage your SQL Server 2008 database and logs:

- INST1_DATA_FILES: contains the SQL Server 2008 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

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 Network Appliance documentation for instructions.

Connecting virtual disks to the cluster node

After creating the virtual disks on the NetApp filer, they must be connected (if not connected already) to the cluster nodes using NetApp SnapDrive from the Windows Computer Management Console.

The following steps provide a high level description. Instructions may vary depending on the SnapDrive version. Refer to the NetApp documentation for detailed instructions.

To connect virtual disks to the cluster node

1. Click **Start > All Programs > Administrative Tools > Computer Management** to start the Computer Management MMC on the cluster node where you want to connect the LUN.
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. On the Welcome page, click **Next**.

5 Specify the Storage System Name, and then 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

You must disconnect the virtual disks from the node where you have finished installing SQL and connect them to the additional failover node before installing SQL on that node.

While installing an application on multiple nodes, you must first disconnect the virtual disks from one node, and then connect the volumes using the same drive letters and then install the application on the failover node.

Complete these steps to disconnect the virtual disks from a node. The following steps provide a high level description. Instructions may vary depending on the SnapDrive version. Refer to the NetApp documentation for detailed instructions.

To disconnect virtual disks

1 Click **Start > All Programs > Administrative Tools > Computer Management** to start the Computer Management MMC on the cluster node where you want to disconnect the LUN.

2 From the left pane, expand **Storage > SnapDrive > Disks**.

3 Right-click the LUN you want to disconnect and then click **Disconnect Disk**.

4 In the Disconnect Disk message prompt, click **Yes**.

Managing storage using Windows Logical Disk Manager

If your configuration uses shared disks and volumes managed using Windows Logical Disk Manager (LDM), use the VCS DiskReservation (DiskRes) and 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*.

Symantec recommends 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 2008 FILESTREAM, create additional volumes for FILESTREAM enabled database objects. If you are configuring an MSDTC Server service group, create additional volumes for MSDTC log and MSDTC registry replication. These volumes must be accessible from all cluster nodes.

For example, you could create the following volumes to manage your SQL Server 2008 database and logs:

- INST1_DATA_FILES: contains the SQL Server system data files (including the master, model, msdb, and tempdb databases)
- INST1_REGREP_VOL: contains the list of registry keys that must be replicated among cluster systems for the SQL Service. Create a 100 MB (minimum recommended size) volume for this purpose.
- INST1_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

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

- Reserve disks.
  See "Reserving disks (if you use Windows LDM)" on page 60.
- Create volumes.
  See "Creating volumes (if you use Windows LDM)" on page 60.
- Mount volumes.
  See "Mounting volumes (if you use Windows LDM)" on page 61.

About LDM support

The following restrictions apply in this release:

- Disk Reservation and Mount agents are supported on VCS for Windows only. These agents are not supported in an SFW storage environment.
- LDM support is available on Windows Server 2003 only.
- 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.

Reserving disks (if you use Windows LDM)

Complete the following steps to reserve the disks on the node on which you are going to install SQL Server 2008.

To reserve the disks

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

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

   You will see a table that lists all the disks that are visible from the current system. Make a note of the disk numbers (Disk# column in the table). You will need it in the next step.

2. To reserve a disk, type the following on the command line:

   ```
   C:\>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 while installing SQL Server 2008 and configuring the service group, on additional failover nodes in the cluster.

Creating volumes (if you use Windows LDM)

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. After creating the required volumes on a node, 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)

Use the Windows Disk Management tool to mount the volumes that you created earlier. After mounting the volumes on a cluster node, run the CHKDSK command and verify that there are no errors on the mounted volumes.

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. To open Disk Management, click **Start > Run**, type **diskmgmt.msc** and then click **OK**.
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)

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

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

For example, to release disk #4, type:

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

Make a note of the disk number and the corresponding signature. You will require these details to identify and reserve the disks later.

Installing and configuring SQL Server 2008 on the first cluster node

Run the Microsoft SQL Server 2008 installer to install SQL Server 2008 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 SFW HA 5.1 Application Pack 1 for Windows along with the VCS database agent for SQL 2008, on all the nodes on which you wish to configure SQL Server 2008.
  Refer to the Application Pack 1 release notes for more information.
- Ensure that the volumes or LUNs (virtual disks) on the filer are mounted or connected to the first cluster node for this SQL instance.
  See "Managing storage using Network Appliance Filer" on page 56.
  See "Managing storage using Windows Logical Disk Manager" on page 59.
- Install SQL Server 2008 in the stand-alone installation mode in a non-clustered environment. Also, while installing SQL, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure with SFW HA.
- Install the SQL program files to a local disk and the SQL database files to the shared storage.
- Use the same instance name and instance ID when you install this instance of SQL Server 2008 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name and instance ID.
- While specifying a user name for the SQL Server services account, specify a domain user account.

  Note: If SQL Server services are not installed with a domain user account, the SQL service group may fail to come online on the cluster nodes. It may come online only on the node on which SQL Server was installed last. In such a case, you must perform additional steps after configuring the SQL service group.
  See Technote http://support.veritas.com/docs/281828.

- 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 2008 services to manual. However, set the startup type of the SQL Server 2008 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 and configuring SQL Server 2008 on the second cluster node

Run the Microsoft SQL Server 2008 installer to install SQL Server 2008 on the second or any additional cluster node. Refer to the Microsoft documentation for instructions.

Note the following prerequisites before installing SQL Server 2008 on the second or any additional failover nodes for the SQL instance:

- Ensure that you have installed SFW HA 5.1 Application Pack 1 for Windows along with the VCS database agent for SQL 2008, on all the nodes on which you wish to configure SQL Server 2008.
  Refer to the Application Pack 1 release notes for more information.
- Ensure that the SQL Server services for the SQL instance are stopped on the first node where you installed SQL Server 2008. This allows the installation on the second or additional nodes to manipulate the database files on the shared disks.
- Ensure that the volumes or LUNs (virtual disks) for this SQL instance are deported from the first node and mounted or connected on the second or additional node.
  See "Managing storage using Network Appliance Filer" on page 56.
  See "Managing storage using Windows Logical Disk Manager" on page 59.
- Install SQL Server 2008 in the stand-alone installation mode in a non-clustered environment. Also, while installing SQL, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure with SFW HA.
- Before installing SQL on the second or additional node, open the SQL Server 2008 system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the SQL Server installation on the second or additional nodes.
  If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second or additional node SQL Server installation.

Note: You can delete the renamed folder and its contents after the installation on additional nodes completes successfully.

- Use the same instance name and instance ID when you install this instance of SQL Server 2008 on additional nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name and instance ID.

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 SQL Server documentation for instructions on assigning ports.

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 (up to 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 Microsoft SQL Server documentation for instructions.
- 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.

Chapter 4

Configuring the SQL service group

This chapter contains the following topics:

- "About configuring the SQL service group" on page 66
- "Before configuring the SQL service group" on page 66
- "Configuring a SQL Server 2008 service group" on page 68
- "Running SnapManager for SQL" on page 74
- "Creating a SQL Server user-defined database" on page 74
- "Verifying the service group configuration" on page 76
- "Administering a SQL Server service group" on page 78

About configuring the SQL service group

A VCS SQL Server 2008 service group is used to bring a SQL Server 2008 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. You use the VCS SQL Server 2008 Configuration Wizard to configure the 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 2008 Configuration Wizard.

The SQL Server 2008 Configuration Wizard enables you to create a SQL Server 2008 service group and define the attributes for its resources on all the nodes within the cluster simultaneously.

Before configuring the SQL service group

Prerequisites for configuring the SQL Server 2008 service group are as follows:

- Verify that VCS 5.1 for Windows Application Pack 1, along with the VCS database agent for SQL Server 2008, is installed on all the cluster nodes.
- Verify that you have configured a VCS cluster. You can use the VCS Cluster Configuration Wizard (VCW) for cluster configuration.
- Verify that SQL Server 2008 is installed and configured identically on all the cluster nodes.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must be logged on as a Domain Administrator.
- You must be an Administrator for the NetApp filer containing the LUNs created to store SQL Server components.
- Verify that the Veritas High Availability Daemon (HAD) is running on the system from where you run the wizard.
- Verify the 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

- If you wish to configure high availability for FILESTREAM, ensure that FILESTREAM is configured for the SQL instance, and is enabled for the SQL instance on the node on which you run the wizard and disabled on all the remaining nodes.
  Refer to the Microsoft SQL Server 2008 documentation for instructions.
- Assign a unique virtual IP address for the SQL Server 2008 instance. You specify this IP address when configuring the service group.
- If you wish to use a script for detail monitoring, for example, to create a table and write data to it, note the location(s) of the script to use. Either locate 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\SQLServer2005\sample_script.sql`
  Detailed monitoring is often not necessary.
- Stop the SQL 2008 Server service for the SQL instance that you wish to configure the service group.
- If you have configured Windows Firewall, add the following to the Firewall Exceptions list:
  - Port 14150 or the VCS Command Server service, `%vcs_home%\bin\CmdServer.exe`.
    Here, %vcs_home% is the installation directory for VCS, typically `C:\Program Files\Veritas\Cluster Server`.
  - Port 14141

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

Configuring a SQL Server 2008 service group

This section describes how to create a new SQL service group. To modify an existing service group, see “Modifying a SQL service group configuration” on page 78.

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

See “About VCS hardware replication agent for NetApp” on page 11.

See “About the VCS database agent for Microsoft SQL Server 2008” on page 16.

To create a SQL Server service group on the cluster

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

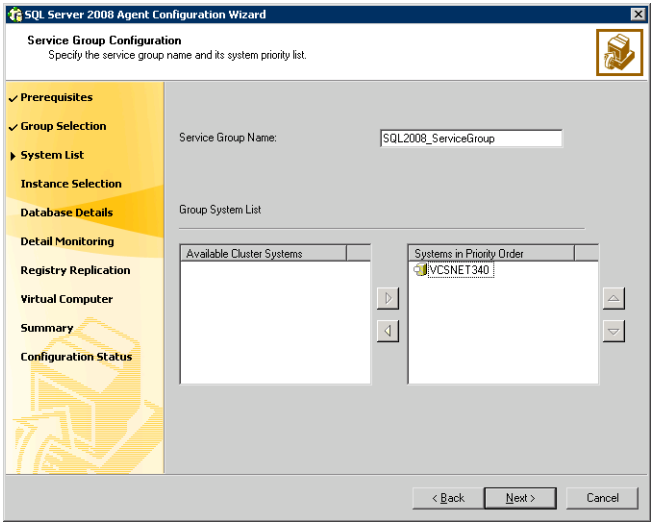

- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
- In the Available Cluster Systems box, select the systems on which to configure the service group and click the 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, complete the following steps and then click **Next**:

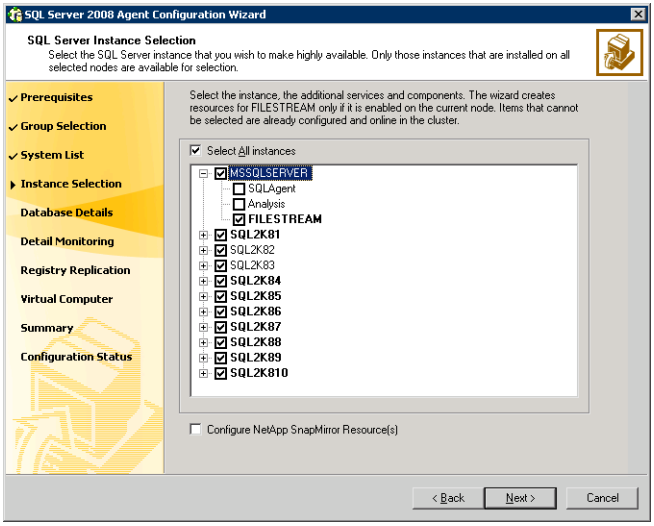

- Select the SQL Server instance(s) that you wish to configure in the service group. The wizard displays only those instances that are installed on all the cluster nodes.
- If required, select the other services that you wish to make highly available. These options are available for selection only if 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 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.

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

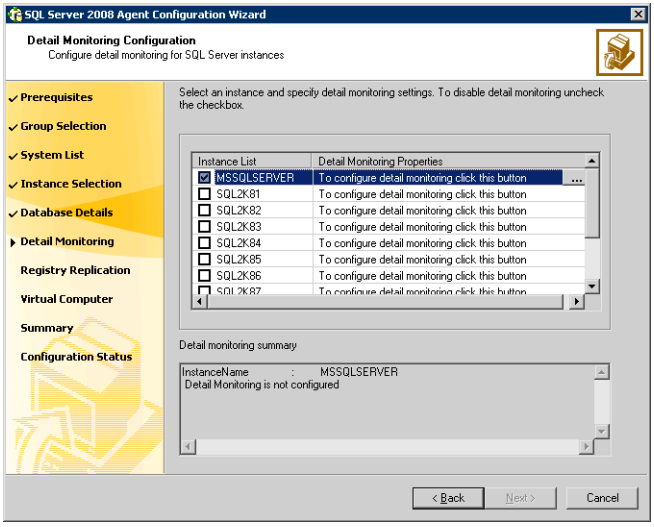

Perform the following steps only if you wish to configure detail monitoring for an instance:

- 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.
  Clear the check box to disable detail monitoring for the instance.
- On the Detail Monitor configuration dialog box, specify the monitoring interval in the **Detail monitoring interval** field.
  This sets the value for the DetailMonitoringInterval attribute of the SQL agent. It indicates the number of online monitor cycles that the agent must wait before performing detail monitoring. The default value is 5. Symantec recommends that you set the monitoring interval between 1 and 12.
- Select **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 **SQLFile Detail Monitoring** if you wish to use a script to monitor SQL databases. In this mode of detail monitoring, the agent executes the script that you specify for detail monitoring.
- 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.
- 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.
- 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.
- Click **Apply**.
- Repeat these steps for each SQL instance that you wish to configure detail monitoring for, and then click **Next**.

9 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.

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

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

This sets the Lanman resource attributes ADUpdateRequired and ADCriticalForOnline to true. This allows the Lanman agent to update Active Directory with the SQL virtual server name.

You can type the OU details in the format "CN=Computers,DC=domainname,DC=com". If you wish to search for the OU, click the ellipsis button and specify the search criteria in the Windows "Find Organizational Units" dialog box.

- Click **Next**.

11 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 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.

12 In 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 as follows:

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

13 Click **Yes** when prompted that the wizard will modify the configuration. The wizard begins to create the service group. Various messages indicate the status of the commands.

14 Check **Bring the service group online** checkbox if you want to bring the service group online and then click **Finish**.

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.

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

To configure an MSDTC service group, see "Configuring an MSDTC service group for high availability" on page 81.

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 2008 database.

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

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

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

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

Creating a SQL Server user-defined database

You can use VCS for Windows to manage user-defined SQL Server 2008 databases. Create the required SQL databases using the SQL Server Management Studio and then make them highly available with VCS.

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

- Create volumes or LUNs for a user-defined SQL Server 2008 database and its transaction log, if you have not already created them.
  In the sample deployment the following volumes are named:
  - INST1_DB1_VOL: contains a user-defined database file
  - INST1_DB1_LOG: contains a user-defined database log file
  - INST1_DB1_FS_VOL: contains FILESTREAM enabled database objects for the user database
- Create a SQL Server user-defined database for the required SQL instance using the SQL Server Management Studio. While creating the database, ensure that you point the database files and transaction log to the paths of the new volumes or LUNs created earlier.
  Refer to the SQL Server documentation for instructions on how to create databases.
- After creating the database, run the SQL Server 2008 Configuration Wizard and modify the SQL Server 2008 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.
See "Adding storage agent resources to the SQL service group" on page 75.

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

Adding storage agent resources to the SQL service group

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

Before you proceed, make sure the volumes for the user-defined 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. Start the SQL Server 2008 Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Symantec > Veritas Cluster Server > Configuration Tools > SQL Server 2008 Configuration Wizard**.
2. Review the prerequisites on the Welcome panel 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.

Verifying the service group configuration

Failover simulation is an important part of configuration testing. This section provides steps to verify the SQL Server 2008 service group configuration by bringing the service group online, taking the service group offline, and switching the service group to another cluster node.

Bringing the service group online

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

To bring a service group online from the Java Console

1 In the Cluster Explorer configuration tree, select the SQL service group to be taken online.

2 Right-click the service group name, and select **Enable Resources**. This enables all resources in the service group.

3 Right-click the service group name, and select the system on which to enable the service group. (Right-click > Enable > *system_name* or Right-click > Enable > All)

4 Save your configuration (**File > Close Configuration**).

5 Right-click the service group and select to online the service group on the system. (Right-click > Online > *system_name*)

Taking the service group offline

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

To take a service group offline from the Java Console

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

Switching the service group

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

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

To switch service groups

1. In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, click the Service Groups tab, and right-click the service group icon in the view panel.
   - Click **Switch To**, and click the appropriate node from the menu.
   - In the dialog box, click **Yes.** The service group you selected is taken offline on the original node and brought online on the node you selected.
     If there is more than one service group, you must repeat this step until all the service groups are switched.
2. Verify that the service group is online on the node you selected to switch to in step 1.
3. To move all the resources back to the original node, repeat step 1 for each of the service groups.

To shut down an active cluster node

1. Gracefully shut down or restart the cluster node where the service group is online.

2 In the Veritas Cluster Manager (Java Console) on another node, connect to the cluster.

3 Verify that the service group has failed over successfully, and is online on the next node in the system list.

4 If you need to move all the service groups back to the original node:
   - Restart the node you shut down in step 1.
   - Click **Switch To**, and click the appropriate node from the menu.
   - In the dialog box, click **Yes.**
     The service group you selected is taken offline and brought online on the node that you selected.

Administering a SQL Server service group

You can dynamically modify the SQL service group configuration in several ways, including the SQL Server 2008 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 2008 Configuration Wizard.

Modifying a SQL service group configuration

Prerequisites for modifying a service group are as follows:

- 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 Configuration Wizard**.

2 Review the prerequisites and click **Next**.

3 In the Wizard Options panel, click **Modify service group**, select the service group and then 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 and click **Next**. The wizard will create NetAppSnapDrive resource for each database.

7 On the Detail Monitoring Configuration panel, specify or modify detail monitoring options for an instance and then click **Next**.
Follow the wizard instructions and make desired modifications to the service group configuration. For more details refer to the service group creation procedure.
See "Configuring a SQL Server 2008 service group" on page 68.

Deleting a SQL service group

The following steps describe how to delete a SQL Server service group using the SQL Server 2008 Configuration Wizard.

To delete 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 Configuration Wizard**.

2 Review the prerequisites and click **Next**.

3 In the Wizard Options panel click **Delete service group**, select the service group, and then click **Next**.

4 In the Service Group Summary dialog box, click **Next**.

5 Click **Yes** on the message that prompts you that the wizard will modify the configuration. The wizard then deletes the service group.

6 Click **Finish**.

Chapter 5

Configuring an MSDTC service group for high availability

This chapter contains the following topics:

- "About configuring the MSDTC service group" on page 82
- "Before configuring the MSDTC service group" on page 82
- "Reviewing the configuration" on page 83
- "Creating an MSDTC Server service group" on page 86
- "About configuring the MSDTC client" on page 88
- "About using the virtual MMC viewer" on page 90
- "Viewing DTC transaction information" on page 90
- "Verifying the installation" on page 91

About configuring the MSDTC service group

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

To configure high availability for MSDTC, you first configure the MSDTC Server service group using the SQL Server Configuration Wizard (not the SQL Server 2008 Configuration Wizard), and then configure the MSDTC client manually. This chapter describes the required steps.

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

Before configuring the MSDTC service group

Review the following prerequisites before configuring the MSDTC service group:

- You must be a cluster administrator. This user classification is required to create and configure a service group.
- You must be a Local Administrator on the node where you run the wizard.
- Verify that the VCS database agent for SQL Server 2008 is installed on all cluster nodes.
- Verify that the VCS cluster is configured. You can use the VCS Cluster Configuration Wizard (VCW) for cluster configuration.
- Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC Server service group.
- Assign a unique virtual server name and virtual IP address to the MSDTC Server.
- Verify that the Distributed Transaction Coordinator service is stopped.
- Verify that you have created the volumes or LUNs (virtual disks) for MSDTC log, and MSDTC registry replication, on the shared disk.

- Verify that the volumes or LUNs containing the MSDTC logs and registry replication directory are mounted or connected to one node (the node on which you are configuring the service group) and dismounted and disconnected from all other nodes.
  See "Managing storage using Network Appliance Filer" on page 56.
  See "Managing storage using Windows Logical Disk Manager" on page 59.
- 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 Installation and Upgrade Guide*.

Reviewing the configuration

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

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

The following configurations are possible:

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

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

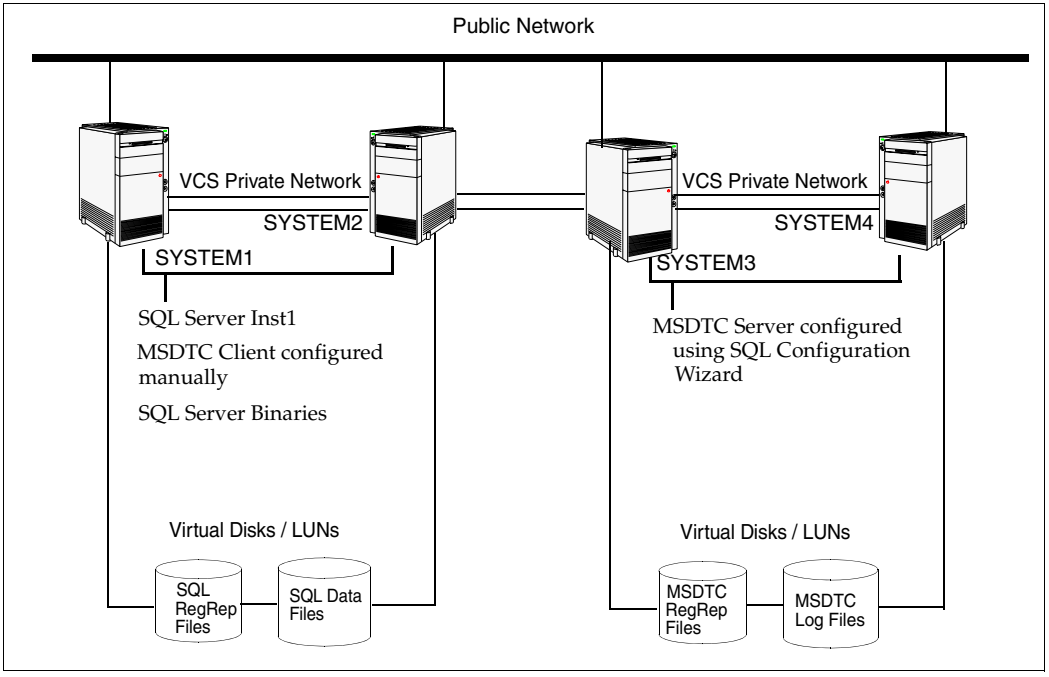

Figure 5-2 MSDTC Server configured on the same node as SQL Server

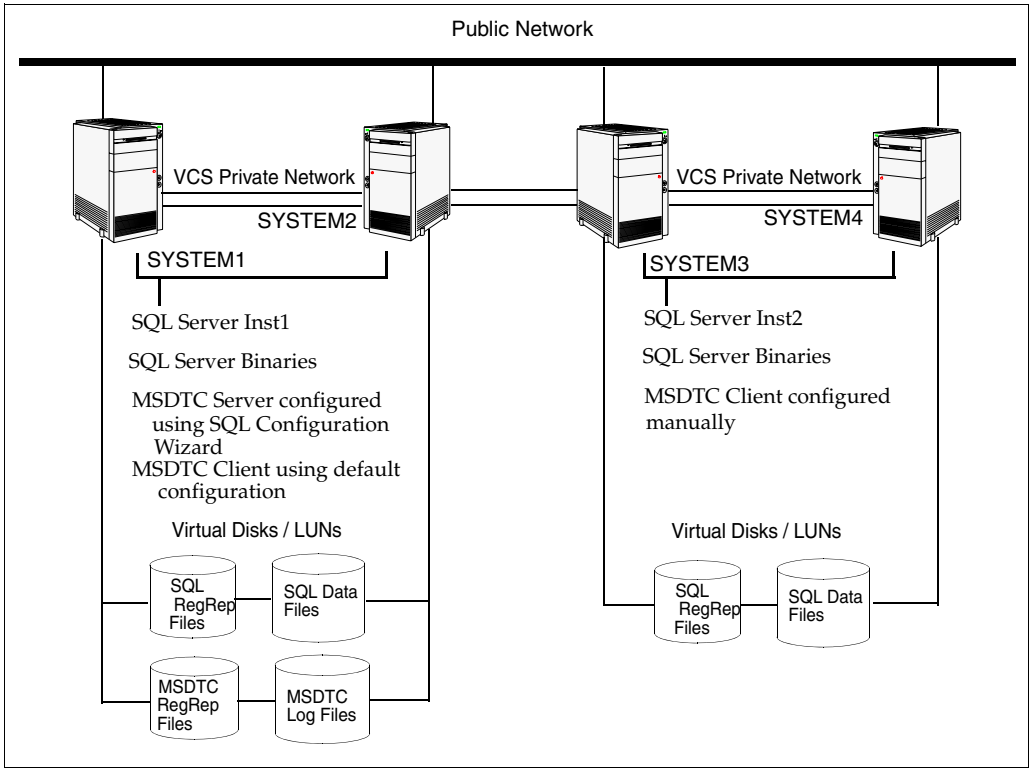

Creating an MSDTC Server service group

MSDTC is a global resource and is accessed by more than one SQL Server service group. Symantec recommends you to configure one MSDTC service group in a VCS cluster. VCS provides 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. After configuring the service group proceed to configuring the MSDTC client.

This section describes the steps required to create a new MSDTC service group using the wizard.

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

To create an MSDTC Server 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 and then click **Next**.
4. On the Service Group Configuration panel, specify the service group name and system list as follows:
   - Enter a name for MSDTC service group.
   - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the service group's system list. Make sure to select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.

     To remove a system from the service group's system list, select the Systems in Priority Order list and click the left arrow.

     To change a system's priority in the service group's system list, select the system from the Systems in Priority Order and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
   - Click **Next**. If the configuration is in read-only mode, the wizards 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:

- Enter the virtual server name for the node on which DTC service is running. Ensure that the virtual server name you enter is unique in the cluster.
- Enter a unique virtual IP address for the MSDTC Server.
- Enter the subnet mask to which the virtual IP address belongs.
- For each system in the cluster, select the public network adapter name. Click the **Adapter Display Name** field to view the adapters associated with a system.
  The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory, click **Advanced Settings**, check the **Active Directory Update Required** check box, and select the Organizational Unit from the drop down list.
  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.
  By default, the Lanman resource adds the virtual server to the default container Computers. The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
- Click **OK** and then click **Next**.

6 On the Specify Data Path panel, specify the drive letter for the MSDTC log and replication directory and then click **Next**.
Symantec recommends using different paths for these directories.
Clear the **Configure NetApp SnapMirror Resource(s)** check box. This option is applicable only in case of a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. If you are setting up a disaster recovery environment, check this check box to configure the SnapMirror resource at the primary site. The SnapMirror resource must be configured only after you have configured the cluster at the secondary site.

7 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 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 name 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 in the Attributes box.
- The wizard assigns unique names to resources. Change names of the resources, if desired. To edit a resource name, select the resource name and either click it or press the **F2** key. Press enter after editing each resource name. To cancel editing a resource name, press 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 to bring the configured service group online and then click **Finish**.

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

About configuring the MSDTC client

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

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

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

Configuring MSDTC client on Windows 2003

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

To configure an MSDTC client on Windows 2003

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

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 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 the virtual MMC viewer will be run.

To view transactions from a node where MSDTC resource is online

1. Run the virtual DTC viewer. Type the following on the command prompt:
   ```
   C:\>VCSVMMCView.exe -target MSDTC -u <username> - -p <password>
   -d <domain> -t <domain type>
   ```
   Running the Virtual DTC viewer will restart the COM+ services.
   The Component Services MMC snap-in is launched.
2. In the console tree of the Component Services administrative tool, expand **Component Services > Computers > My Computer > Distributed Transaction Coordinator > Local DTC**.
3. 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.
4. 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**
   or
   Click **Start > Run**, type **dcomcnfg** and click **OK**.
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.

Verifying the installation

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

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

See "Verifying the service group configuration" on page 76.

Chapter 6

Making a standalone SQL Server highly available

This chapter contains the following topics:

- "About making a standalone SQL Server 2008 highly available" on page 94
- "Reviewing the configuration" on page 94
- "Install and configure VCS on the standalone SQL Server 2008" on page 96
- "Verify that SQL Server 2008 databases and logs are moved to shared storage" on page 97
- "Install and configure SQL Server on additional nodes" on page 97
- "Assign ports for multiple SQL Server instances" on page 98
- "Configure the VCS SQL server service group" on page 98
- "Create a SQL Server 2008 user-defined database" on page 99
- "Verify the configuration" on page 99

About making a standalone SQL Server 2008 highly available

This chapter describes the procedure to convert an existing standalone SQL Server 2008 into a “clustered” SQL Server in a new Veritas Cluster Server environment. This environment involves an active-passive configuration with one to one failover capabilities.

Note: In addition to the information contained in this chapter, refer to the Microsoft SQL Server documentation for instructions on how to move SQL Server databases.

Reviewing the configuration

This section 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

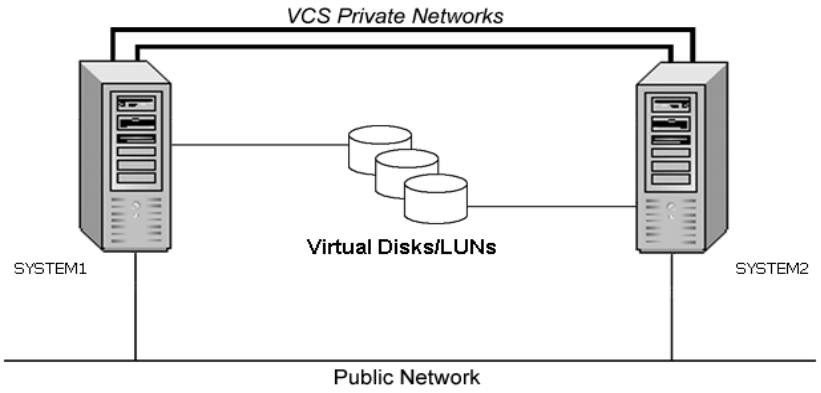

The virtual SQL Server is online on SYSTEM1, serving client requests. The shared LUNs (virtual disks) provide storage for the SQL Server databases. SYSTEM2 waits in a warm standby state as a backup node, prepared to begin handling client requests if SYSTEM1 becomes unavailable. From the user’s perspective 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: used by Veritas Cluster Manager (Web Console)

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

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

Name	Object
SYSTEM1 & SYSTEM2	server names; SYSTEM1 is the existing standalone SQL Server 2008
INST1_SG	Microsoft SQL Server 2008 service group
SQL_CLUS1	virtual SQL server cluster
INST1_DG	Disk group for the volumes for the SQL instance
INST1_DATA_FILES	volume/LUNs for Microsoft SQL Server system data files
INST1_DB1_VOL	volume/LUNs for storing a Microsoft SQL Server user-defined database
INST1_DB1_LOG	volume/LUNs for storing a Microsoft SQL Server user-defined database log file
INST1_REGREP_VOL	volume/LUNs that contains the list of registry keys that must be replicated among cluster systems for the SQL server
INST1_FS_VOL	volume/LUNs for storing FILESTREAM enabled database objects
INST1	SQL Instance Name
INST1-VS	name of the SQL Virtual Server

Install and configure VCS on the standalone SQL Server 2008

Before converting the existing standalone SQL into a clustered server, complete the following procedures on the system on which SQL server is installed:

- Create a backup of the data on the existing standalone SQL server.
- Use the SQL Configuration Manager and set the startup type of all the SQL services for the SQL instance to manual.
- Install Veritas Cluster Server 5.1 for Windows on the standalone SQL server. Refer to the *Veritas Cluster Server 5.1 Installation and Upgrade Guide* for more information.
- Install Veritas Cluster Server 5.1 for Windows Application Pack 1 for Windows along with the VCS agents for SQL Server 2008, on the standalone SQL Server. Refer to the *Veritas Cluster Server 5.1 for Windows Application Pack 1 Release Notes* for more information.
- Configure the cluster using the VCS Cluster Configuration Wizard (VCW).
  See "Configuring the cluster" on page 36.
- Create volumes or LUNs (virtual disks) necessary to manage the SQL Server database.
  See "Managing storage using Network Appliance Filer" on page 56.
  See "Managing storage using Windows Logical Disk Manager" on page 59.

Verify that SQL Server 2008 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.

Complete the following tasks to move the databases.

To move the database and logs to shared storage

1. Stop the SQL Server service.
2. Verify that you have backed up your existing data.
3. Ensure that the volumes or LUNs are imported or connected to the node where the original database files are located on the local drives.
   See "Managing storage using Network Appliance Filer" on page 56.
   See "Managing storage using Windows Logical Disk Manager" on page 59.
4. Modify the SQL Server 2008 data file and user database locations.
   Refer to the Microsoft SQL Server documentation for instructions.
5. Restart SQL Server 2008.

Install and configure SQL Server on additional nodes

Run the Microsoft SQL Server 2008 installer to install SQL Server 2008 on the second or any additional cluster node. Refer to the Microsoft documentation for instructions.

You will need the following information:

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

Note the following prerequisites before installing SQL Server 2008 on the second or any additional failover nodes for the SQL instance:

- Stop the SQL Server services for the SQL instance, on the first node where you installed SQL Server 2008. This allows the installation on the second or additional nodes to manipulate the database files on the shared disks.

- Ensure that the volumes or LUNs (virtual disks) for this SQL instance are deported from the first node and mounted or connected on the second or additional node.
  See "Managing storage using Network Appliance Filer" on page 56.
  See "Managing storage using Windows Logical Disk Manager" on page 59.
- Before installing SQL on the second or additional node, open the SQL Server 2008 system data files volume (INST1_DATA_FILES) and rename or remove the first node SQL Server system data files. The files will be replaced during the SQL Server installation on the second or additional nodes.
  If you rename the folder that contains the system data files, the files are available as backup files in case problems occur during the second or additional node SQL Server installation.

  Note: You can delete the renamed folder and its contents after the installation on additional nodes completes successfully.

- Use the same instance name and instance ID when you install this instance of SQL Server 2008 on failover nodes. If you are installing multiple instances of SQL in the cluster, each instance must have a unique instance name and instance ID.
- Install SQL Server 2008 in the stand-alone installation mode in a non-clustered environment. Also, while installing SQL, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure with SFW HA.

Assign 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.

See "Assigning ports for multiple SQL Server instances" on page 64.

Configure the VCS SQL server service group

The SQL Server 2008 Configuration Wizard enables you to create a SQL Server 2008 service group and define the attributes for its resources on all the nodes within the cluster simultaneously. Create the service group for SQL Server 2008.

See "Configuring a SQL Server 2008 service group" on page 68.

Create a SQL Server 2008 user-defined database

Create databases and configure resources for those databases in the SQL service group.

See "Creating a SQL Server user-defined database" on page 74.

The procedure includes the following tasks:

- Create volumes for a user-defined SQL Server database and its transaction log.
- Create a new SQL Server user-defined database and point the database files and transaction log to the paths of the new volumes.
- Use the SQL Server 2008 Configuration Wizard to add the storage resources (DiskRes and Mount for Windows LDM, or NetAppFiler and NetAppSnapDrive) for the user databases, to the existing SQL Server service group.

Verify the configuration

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

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

See "Verifying the service group configuration" on page 76.

Chapter 7

Active-Active configuration

This chapter contains the following topics:

- "About active-active configuration" on page 102
- "Reviewing the configuration" on page 102
- "Installing VCS and configuring the cluster" on page 105
- "Configure volumes or virtual disks for SQL server" on page 105
- "Installing and configuring the first instance of SQL Server" on page 106
- "Configuring the VCS service group for the first SQL Server instance" on page 106
- "Creating a SQL Server user-defined database" on page 107
- "Repeating SQL Server installation for additional instances" on page 107
- "Verify the Active-Active configuration" on page 108

About active-active configuration

This chapter describes how to install and configure VCS for SQL Server that includes active-active clustering.

Reviewing the configuration

A SQL Server instance is a completely independent SQL Server installation, with its own services, master database, storage, and memory resources. Each instance is defined uniquely by a separate SQL Server virtual server and service group. SQL Server supports multiple independent instances of SQL Server to run on a single machine. A SQL Server instance can fail over to any of the other nodes configured in the SQL Server service group's system list.

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

Figure 7-1 Active-active configuration

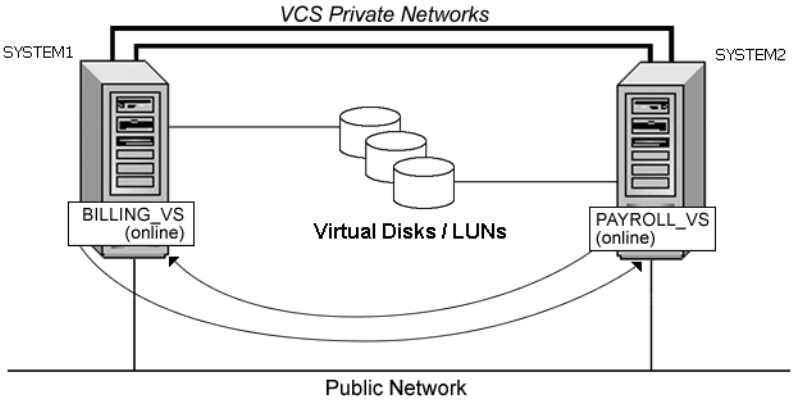

For example, consider a two-node cluster hosting two SQL Server virtual servers, BILLING_VS and PAYROLL_VS. The table below and the sample configuration illustrate that the virtual servers are configured in two separate service groups with BILLING_VS online on SYSTEM1 but able to fail over to

SYSTEM2, and PAYROLL_VS online on SYSTEM2 but able to fail over to SYSTEM1.

Table 7-1 Active-active configuration

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

Sample configuration

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

During the configuration process, create virtual IP addresses for the following:

- Billing virtual server: virtual IP address is the same on all nodes
- Payroll virtual server: virtual IP address is the same on all nodes
- Cluster IP address: used by Web Console

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

Table 7-2 describes the objects created and used during the installation and configuration:

Table 7-2 Active-active configuration objects

Name	Description
SYSTEM1 & SYSTEM2	server names
SQL_CLUS1	virtual SQL server cluster
BILLING_DG	cluster disk group for volumes for the billing instance
PAYROLL_DG	cluster disk group for volumes for the payroll instance
BILLING_VS_SYS_FILES	volume or LUNs for the SQL Server system data files for the billing instance
PAYROLL_VS_SYS_FILES	volume or LUNs for the SQL Server system data files for the payroll instance

Table 7-2 Active-active configuration objects

Name	Description
BILLING_VS_FS_VOL	volume or LUNs for the FILESTREAM enabled database objects for the billing instance
PAYROLL_VS_FS_VOL	volume or LUNs for the FILESTREAM enabled database objects for the payroll instance
BILLING_DATA	volume or LUNs for a SQL Server user-defined database for the billing instance
PAYROLL_DATA	volume or LUNs for a SQL Server user-defined database for the payroll instance
BILLING_LOG	volume or LUNs for a SQL Server user-defined database log file fort he billing instance
PAYROLL_LOG	volume or LUNs for a SQL Server user-defined database log file for the payroll instance
BILLING_REGREP	volume or LUNs for the list of registry keys replicated among the nodes for the billing instance
PAYROLL_REGREP	volume or LUNs for the list of registry keys replicated among the nodes for the payroll instance
BILLING_INST	instance name for the billing instance
PAYROLL_INST	instance name for the payroll instance
BILLING_VS	virtual SQL server name for the billing instance
PAYROLL_VS	virtual SQL server name for the payroll instance
BILLING_SG	SQL Server service group for the billing instance
PAYROLL_SG	SQL Server service group for the payroll instance

Installing VCS and configuring the cluster

Complete the following procedures on the systems on which you wish to configure SQL Server 2008:

- Install Veritas Cluster Server 5.1 for Windows. Refer to the *Veritas Cluster Server Installation and Upgrade Guide* for more information.
- Install Veritas Cluster Server 5.1 for Windows Application Pack 1 for Windows. Refer to the *Veritas Cluster Server 5.1 Application Pack 1 Release Notes* for more information.
- Configure the cluster using the VCS Cluster Configuration Wizard (VCW). See "Configuring the cluster" on page 36.

Configure volumes or virtual disks for SQL server

For each instance of SQL Server 2008, create volumes or LUNs (virtual disks) on the shared storage, as follows:

Referring to the sample configuration described earlier, for the Billing instance, create the following:

- BILLING_DG: a cluster disk group for the volumes related to the Billing instance
- BILLING_VS_SYS_FILES: volume or LUNs for the SQL Server system data files
- BILLING-REGREP: volume or LUNs for the list of registry keys replicated among cluster nodes for the Billing instance
- BILLING_VS_FS_VOL: volume or LUNs for FILESTREAM enabled objects

For the Payroll Instance, create the following:

- PAYROLL_DG: a cluster disk group for the volumes related to the Payroll instance
- PAYROLL_VS_SYS_FILES: volume or LUNs for the SQL Server system data files
- PAYROLL_REGREP: volume or LUNs for the list of registry keys replicated among cluster nodes for the Payroll instance
- PAYROLL_VS_FS_VOL: volume or LUNs for FILESTREAM enabled objects

See "Managing storage using Network Appliance Filer" on page 56.

See "Managing storage using Windows Logical Disk Manager" on page 59.

Installing and configuring the first instance of SQL Server

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

- Do not accept the default instance name. Specify an instance name for each SQL Server installation.
- Install SQL Server 2008 in the stand-alone installation mode in a non-clustered environment. Also, while installing SQL, ensure that you select all the desired features (for example, Full-Text Search, Analysis Services) that you wish to configure with SFW HA.
- Each SQL Server instance must be assigned a different port. The default port is 1433; ports for subsequent instances are generally assigned in descending order (1432, 1431, 1430, etc.).

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

Refer to the following topics:

- See "Installing and configuring SQL Server 2008 on the first cluster node" on page 62.
- See "Installing and configuring SQL Server 2008 on the second cluster node" on page 63.
- See "Assigning ports for multiple SQL Server instances" on page 64.

Configuring the VCS service group for the first SQL Server instance

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

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

See "Configuring a SQL Server 2008 service group" on page 68.

Creating a SQL Server user-defined database

Create a user-defined database and then configure it with VCS.

The procedure includes the following tasks:

- Create volumes or LUNs for a user-defined SQL Server database and its transaction log and for FILESTREAM enabled database objects if configured.
- Create a SQL Server 2008 user-defined database and point the database files and transaction log to the paths of the new volumes or LUNs.
- Add storage agent resources (Mount and DiskRes in case of Windows LDM, or NetAppSnapDrive and NetAppFiler) to the existing SQL Server service group.

See "Creating a SQL Server user-defined database" on page 74.

Repeating SQL Server installation for additional instances

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

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

Complete the tasks to install and configure additional SQL Server instances on all nodes of the cluster. Refer to the topics mentioned below.

See "Configure volumes or virtual disks for SQL server" on page 105.

See "Installing and configuring SQL Server 2008 on the first cluster node" on page 62.

See "Installing and configuring SQL Server 2008 on the second cluster node" on page 63.

See "Assigning ports for multiple SQL Server instances" on page 64.

See "Configuring a SQL Server 2008 service group" on page 68.

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

- Assign a unique name to the SQL Server service group, for example PAYROLL_SG.

- Pay close attention to the priority order of the systems. For example, if the system priority for the first instance is SYSTEM1 then SYSTEM2; reverse the priority order for the second instance, so that SYSTEM2 has a higher priority.
- Create any SQL Server user-defined databases required for your environment.
  See "Creating a SQL Server user-defined database" on page 107.

Verify the Active-Active configuration

See "Verifying the service group configuration" on page 76.

Chapter 8

Disaster recovery configuration

This chapter contains the following topics:

- "About disaster recovery configuration" on page 110
- "What is a disaster recovery solution?" on page 110
- "What needs to be protected in a SQL Server environment?" on page 111
- "Typical disaster recovery configuration" on page 112
- "Disaster recovery: New SQL Server 2008 installation" on page 113
- "Create a parallel environment on the secondary site" on page 115
- "Configuring DR components" on page 116
- "Configuring replication using NetApp SnapMirror" on page 116
- "Configure SnapMirror resource in the SQL service group" on page 117
- "Configure Global Cluster Option for wide-area failover" on page 118

About disaster recovery configuration

This chapter describes how to set up a disaster recovery configuration for SQL Server 2008 using the VCS agents for Network Appliance SnapMirror and Microsoft SQL Server 2008.

What is a disaster recovery solution?

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

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

Why implement a disaster recovery solution?

A DR solution is vital for businesses that rely on the availability of data. A well-designed DR solution prepares a business for unexpected disasters and provides the following benefits in a DR situation:

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

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

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

Understanding replication

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

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

What needs to be protected in a SQL Server environment?

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

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

Typical disaster recovery configuration

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

Figure 8-1 Typical Disaster Recovery configuration

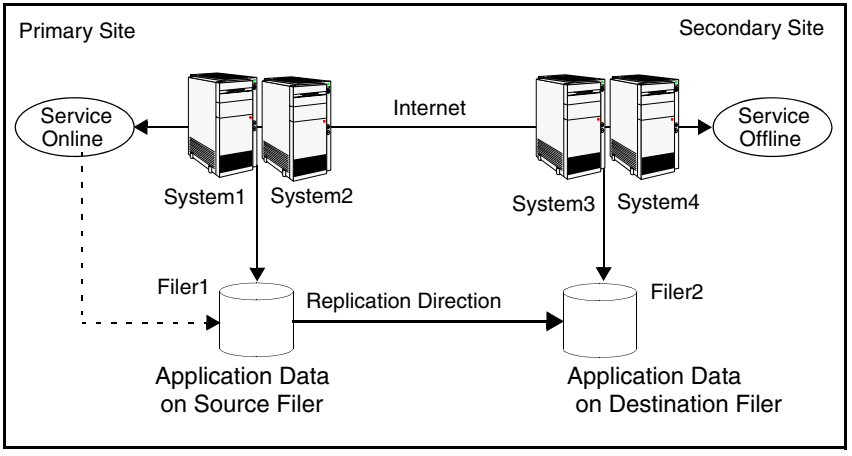

The illustration displays an environment with a DR solution that is prepared for a disaster. The primary site consists of two nodes, System1 and System2. The secondary site consists of two nodes, System3 and System4. Each site has a clustered setup with the nodes set up appropriately for failover within 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 Microsoft SQL Server 2008 on System1 fails, SQL Server comes online on node System2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal.

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

Disaster recovery: New SQL Server 2008 installation

This section provides information on how to install and configure the high availability and SQL Server 2008 components on the primary and secondary sites, with the intent of creating a parallel setup for the SQL service group on both sites. The configuration process is the same for both sites.

Perform the following tasks to set up a disaster recovery environment for SQL Server 2008.

Reviewing the configuration

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

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

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

Installing VCS and configuring the cluster

Perform the following tasks:

- Install Veritas Cluster Server 5.1 for Windows. While installing, ensure that you select the Global Cluster Option to enable wide-area failover.
  Refer to the *Veritas Cluster Server 5.1 Installation and Upgrade Guide* for instructions.
- Install Veritas Cluster Server 5.1 for Windows Application Pack 1 for Windows. Ensure that you select the option to install Veritas Cluster Server Database Agent for Microsoft SQL Server 2008.
  Refer to the *Veritas Cluster Server 5.1 for Windows Application Pack 1 Release Notes* for instructions.
- Configure cluster components including the Wide-Area Connector (WAC) resource for global clusters, using the VCS Cluster Configuration Wizard (VCW).
  See "Configuring the cluster" on page 36.

Configuring volumes or LUNs on the shared storage

Perform the following tasks:

- Create volumes or LUNs (virtual disks) required for SQL Server 2008.
- Ensure that the volumes or LUNs are connected to the first cluster node.

See "Managing storage using Network Appliance Filer" on page 56.

See "Managing storage using Windows Logical Disk Manager" on page 59.

Installing and configuring SQL Server 2008 at the primary site

Perform the following tasks:

- Create volumes or LUNs (virtual disks) and ensure that they are mounted or connected to the first cluster node
- Install and configure SQL Server 2008
- Configure SQL services
- Stop the SQL services
- Install SQL Server 2008 on the additional node

See Chapter 3, "Installing and configuring SQL Server 2008" on page 53.

Configuring the VCS SQL service group

Perform the following tasks:

- Create a SQL Server 2008 service group using the VCS SQL Server 2008 Configuration Wizard
- If required, create volumes or LUNs for a user-defined database and transaction log
- If required, create a user-defined database in SQL Server
- If required, add VCS storage agent resources for a user-defined database in VCS
- Verify the service group configuration

See Chapter 4, "Configuring the SQL service group" on page 65 for instructions.

Create a parallel environment on the secondary site

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

Before you begin creating the SQL Server 2008 service group for the cluster at the secondary site, make sure that the SQL Server service group at the primary site is offline.

Perform the following tasks at the secondary site:

- See "Reviewing the configuration" on page 113.
- See "Configuring the cluster" on page 36.
- See "Managing storage using Network Appliance Filer" on page 56.
  See "Managing storage using Windows Logical Disk Manager" on page 59.
  During the creation of volumes or LUNs (virtual disks) 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 "Installing and configuring SQL Server 2008 on the first cluster node" on page 62.
  Select the same options at the secondary site as you did at the primary site.
- See "Installing and configuring SQL Server 2008 on the second cluster node" on page 63.
  The instance name must be the same on the primary site and secondary site.
- See "Configuring the SQL service group" on page 65.
  The service group name and virtual computer name must be the same on both the primary site and secondary site.

Configuring DR components

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

Perform the following tasks:

- Set up volume replication using Network Appliance SnapMirror.
  See "Configuring replication using NetApp SnapMirror" on page 116.
- Configure SnapMirror resources in the SQL service group at the primary site.
  See "Configure SnapMirror resource in the SQL service group" on page 117.
- Configure GCO option for wide-area failover.
  See "Configure Global Cluster Option for wide-area failover" on page 118.

Configuring replication using NetApp SnapMirror

You can replicate SQL Server 2008 data by establishing a SnapMirror relationship between the Filers at the primary and secondary sites. Before configuring replication, make sure the SQL service group is offline at the secondary site.

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

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

Refer to Network Appliance documentation for more information.

Configure SnapMirror resource in the SQL service group

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

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

This is required in cases where:

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

Use the SQL Server 2008 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.

See "Configuring a SQL Server 2008 service group" on page 68.

Configure Global Cluster Option for wide-area failover

The Global Cluster Option (GCO) is required to manage global clustering for wide-area disaster recovery.

Creating a global cluster environment involves the following tasks:

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

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

Creating a global cluster environment requires the following:

- Wide-area Connector process is configured and the ClusterService group is online at both sites.
- 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.

See "Configuring Wide-Area Connector process for global clusters" on page 52.

Linking clusters: Adding a remote cluster to a local cluster

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

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

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

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

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

To add a remote cluster in Cluster Explorer

1. From Cluster Explorer, click **Add/Delete Remote Cluster** on the **Edit** menu.
   or
   From the Cluster Explorer configuration tree, right-click the cluster name, and click **Add/Delete Remote Cluster**.
2. Review the required information for the Remote Cluster Configuration Wizard and then 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 as follows:
   If the cluster is not running in secure mode:

- 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 and the password.
- Click **Next**.

If the cluster is running in secure mode:

- 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:

- 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:
- 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:
   - 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:

- 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. This operation involves the following tasks:

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

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

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

To take the wide area cluster (wac) resource offline

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

2 In the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the **wac** resource under the **Application** type in the **ClusterService** group.

or

Click a service group in the configuration tree, click the **Resources** tab, and right-click the **wac** resource in the view panel.

3 Click **Offline**, and click the appropriate system from the menu.

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

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

2 Click **Next**.

3 Enter the details of the service group to modify:

- 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:

- 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:

- 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**.

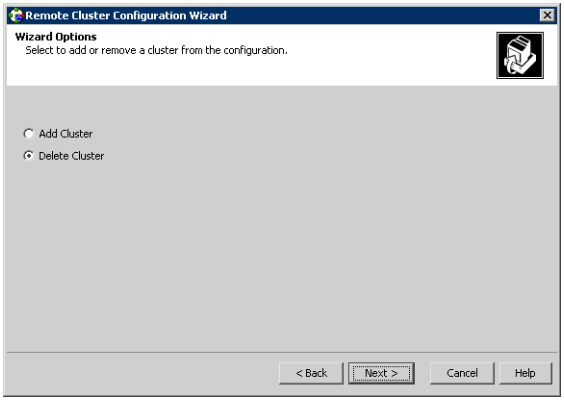

4 In the Delete Cluster panel, click the name of the remote cluster to delete, then click **Next**.

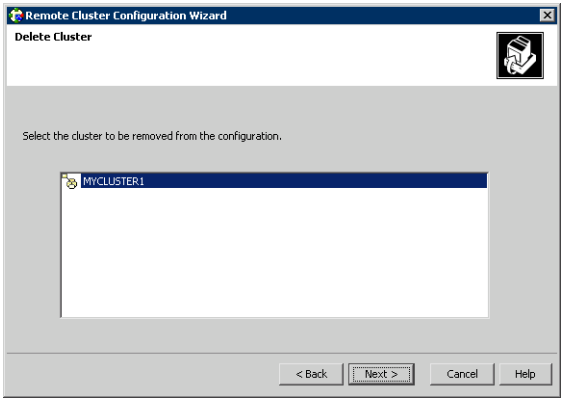

5 Review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:

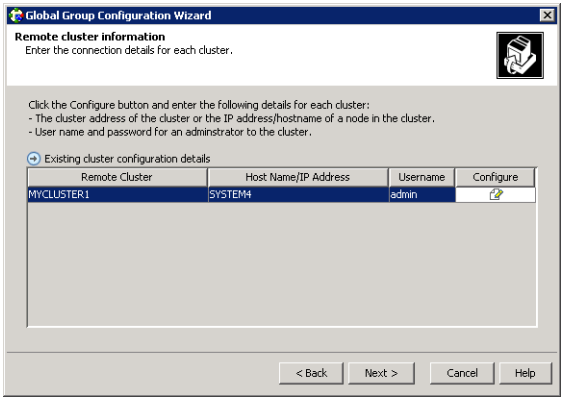

If the cluster is not running in secure mode:

- 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:

- 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 9

Removing VCS agents for SQL Server 2008

This chapter contains the following topics:

- "Prerequisites for removing VCS agents" on page 129
- "Uninstalling the agents" on page 129

Prerequisites for removing VCS agents

Prerequisites for removing VCS are as follows:

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

Uninstalling the agents

This section describes steps for uninstalling VCS. Verify that all SQL service groups are offline on all nodes in the cluster.

To remove the VCS agents

1. In the Windows Add/Remove Programs applet, click Veritas Cluster Server 5.1 for Windows Application Pack 1 and then click **Remove**.
2. Review the Welcome page and click **Next**.
3. Select the check box if you want to remove the client components and click **Next**.

4 Specify the nodes from which you want to remove the agents, as follows:
   - Select the node. This may take some time depending on the size of the domain and network conditions.
   - Click **Add**. To remove a node, highlight the node in the computer list and click **Remove**.
   - Click **Next**.

5 The installer validates the system for uninstallation. After the node is accepted, click **Next**.
   If a node is rejected, the Comments column displays the cause for rejecting the node. Highlight the node to view a detailed information about the failure in the Details box. Resolve the error, highlight the node in the selected systems list, and click **Validate Again**. Once all the nodes are accepted, click **Next**.

6 Review the summary of your selections and click **Uninstall**. The installer displays the status of uninstallation.

7 After the uninstallation is complete, review the report and click **Next**.

8 Click **Finish**.
   After uninstallation you may be required to reboot the system. After rebooting the system, the SnapDrive service may fail to start with a logon failure. In such cases, reset the password for the SnapDrive service account and then start the service.

Chapter 10

Troubleshooting the agents

This chapter contains the following topics:

- "About troubleshooting VCS agents" on page 132
- "VCS logging" on page 132
- "Network Appliance agents error messages" on page 134
- "SQL agent error messages and descriptions" on page 136

About troubleshooting VCS agents

This chapter describes how to troubleshoot common problems in the VCS agents for Network Appliance and Microsoft SQL Server 2008. 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 as follows:

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

The agent log message components are defined as follows:

- Timestamp denotes the date and time when the message was logged.
- Mnemonic denotes which Veritas product logs the message. For VCS enterprise agent for Microsoft SQL, mnemonic is 'VCS'.
- Severity denotes the seriousness of the message. Severity of the VCS error messages is classified into the following types:
  - CRITICAL indicates a critical error within a VCS process. Contact Technical Support immediately.
  - ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action.
  - WARNING indicates a warning or error, but not an actual fault.
  - NOTE informs that VCS has initiated an action.
  - INFO informs about various state messages or comments.

  Of these, CIRITCAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.
- UMI or Unique Message ID is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the SQLServer agent would resemble: `V-16-20020-13`

  Originator ID for all VCS products is 'V-16.' Category ID for SQLServer agent is 20020, for MSDTC is 20021, and that for MSSearch agent is 20022. Message ID is a unique number assigned to the message text.
- Message text denotes the actual message string.

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

A typical agent log resembles the following:

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

VCS Cluster Configuration Wizard logs

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

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

The format of the wizard log is

ThreadID | Message Text

- *ThreadID*: the ID of the thread initiated by the wizard.
- *Message Text*: 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.
```

Network Appliance agents error messages

This section describes the error messages for the NetApp agents.

Table 10-1 NetApp agents error messages

Message	Description
Failed to open connection to filer *%s*.	Make sure that the VCS Helper service account is part of the administrator's group on the local host and the filer. Make sure the private network is functioning properly. Verify you can ping the IP used for the private storage network. This is the IP defined the StorageIP attribute of the NetAppFiler resource.
Failed to initialize ONTAPI on system	The agent could not find the file NTAPADMIN.DLL on the system. Verify the file exists in the *%VCS_HOME%*`\bin` directory
Invalid attributes exist in the configuration	Some agent attributes have not been defined or have been defined incorrectly. Verify the configuration definition for the agent.

Table 10-1 NetApp agents error messages (continued)

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

SQL agent error messages and descriptions

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

Agent for MSDTC

This section describes the error messages for the MSDTC agent.

Table 10-2 MSDTC agent error messages

Message	Description
Lanman attribute has not been configured.	No value specified for the LanmanResName attribute. Solution: Specify a valid value for the LanmanResName attribute.
MountResName attribute has not been configured.	No value specified for MountResName attribute. Solution: Specify a valid value for the MountResName attribute.
LogPath attribute has not been configured.	No value specified for LogPath attribute. Solution: Specify a valid value for the MountResName attribute.
Failed to open the SCM handle. Error = *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.

Table 10-2 MSDTC agent error messages (continued)

Message	Description
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 2008

This section describes the error messages for the SQL Server 2008 agent.

Table 10-3 SQL Server 2008 agents error messages

Message	Description
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.
Failed to convert the argument list. Error = *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.
Invalid value specified for attribute *attribute name*.	No value provided for the specified attribute. Solution: Provide a value for the attribute.
Failed to initialize the SQLServer2008 agent.	The agent failed to initialize the SQLServer2008 agent for SQL Server 2008.

Table 10-3 SQL Server 2008 agents 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.
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 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.
SQL Server Instance name is not specified. Agent will operate on the default instance.	No value spaced-out for SQL Server instance name. Agent would operate on the default SQL Server instance.
Failed to set the virtual computer name in the environment of the service *service name*. Error = *Error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.
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.

Table 10-3 SQL Server 2008 agents 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.
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 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.

Table 10-3 SQL Server 2008 agents error messages (continued)

Message	Description
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.
Sql script has failed with 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 with error *error code*.	This is a VCS internal error. Solution: Contact Symantec Technical Support.
The password attribute has not been configured.	The password attribute used for detail monitoring is not configured.
Failed to start the Sql script. (User = *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.
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.
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.

Table 10-3 SQL Server 2008 agents error messages (continued)

Message	Description
Invalid arglist. The ArgList should contain LanmanResName:IPResName	The Lanman resource name is incorrect. Solution: Verify that the Lanman resource name is valid.

Agent for SQL Server 2008 FILESTREAM

This section describes the error messages for the SQL Server 2008 FILESTREAM agent.

Table 10-4 SQL Server 2008 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.
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.

Table 10-4 SQL Server 2008 FILESTREAM agent error messages (continued)

Message	Description
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.

Index

A

Active/Active
- configuration diagram 102
- configuration overview 102
- creating a user-defined database 107

Active-Active
- configuring VCS SQL Server service group 106

agent functions
- MSDTC agent 23
- NetApp Filer agent 11
- NetApp SnapDrive agent 12
- NetApp SnapMirror agent 14
- SQL Server 2008 agent 17
- SQL Server 2008 Agent service agent 22
- SQL Server 2008 Analysis service agent 22
- SQL Server 2008 FILESTREAM agent 21

application failure, detecting 28

attribute definitions
- MSDTC agent 23
- SQL Server 2008 Filestream agent 22

attributes
- NetApp Filer agent 12
- NetApp SnapDrive agent 13
- NetApp SnapMirror agent 15

C

clusters
- configuring (DR) 113
- configuring (HA) 36

configurations
- Disaster Recovery 31, 112

configure
- LLT over UDP using VCW 42

configuring
- cluster (DR) 113
- cluster (HA) 36
- SQL Server 2008 DR overview 113
- SQL Server service group 2008 114

converting
- standalone SQL conversion to HA 94

converting standalone SQL Server to HA
- moving data files and databases 97

D

database
- new database for SQL Server 2008 74
- user-defined for SQL Server 2008 74

database agent
- installing 35

dependency graph
- MSDTC agent 26
- SQL Server 2008 agent 24
- SQL Server Agent Service agent 24
- SQL Server Analysis Service agent 24
- SQL Server Filestream agent 24

disaster recovery (DR)
- deploying for SQL Server 2008 113
- illustrated 31, 112
- protection in a SQL Server environment 111
- rationale 110
- setting up a secondary SQL Server 2008 site 115
- typical configuration 112

Disaster Recovery configuration 31, 112

G

Global Cluster Option (GCO)
- prerequisites 118

I

installing
- SQL Server 2008 on a second node (VCS) 63
- SQL Server 2008 on the first node 62, 114
- SQL Server 2008 tasks (DR) 113
- SQL Server 2008, secondary site tasks 115
- VCS for Windows (DR) 113
- verifying the installation 91

L

LLT over UDP

configuring using VCW 42
logging
VCW logs 133
VCWsilent logs 134
Logs
VCS 132
VCS Cluster Configuration Wizard 133
VCWsilent 134

M

MSDTC
configuring the client 88
creating an MSDTC service group 86
HA configuration overview 83
service group configuration (VCS) 83
MSDTC agent
attributes 23
dependency graph 26
functions 23
resource type definition 23
multiple SQL Server instances
assigning the port 64, 98

N

NetApp Filer agent
about 11
agent functions 11
attributes 12
resource type definition 11
NetApp SnapDrive agent
about 12
agent functions 12
attributes 13
resource type definition 13
NetApp SnapMirror agent
about 14
agent functions 14
attributes 15
resource type definition 15

P

port assignment for multiple instances 64, 98

R

replication
defined 111
resource type definition
MSDTC agent 23
NetApp Filer agent 11
NetApp SnapDrive agent 13
NetApp SnapMirror agent 15
SQL Server 2008 agent 18
SQL Server 2008 Filestream agent 21

S

sample configurations
SQL Server (HA) 103
standalone SQL Server conversion (HA) 95
Security Services
configuring 43
service groups
administering global groups 123
SQL agent
removing 129
uninstalling 129
SQL cluster
Disaster Recovery setup 31, 112
SQL cluster configuration
Disaster Recovery 31, 112
SQL Server 2008
Active-Active sample configuration 103
configuring the SQL Server service group 114
creating a new database 74
DR configuration overview 113
installing on a second node 63
installing on the first node 62, 114
new installation tasks (DR) 113
overview of user-defined database 74
tasks for setting up a secondary site 115
SQL Server 2008 agent
agent functions 17
attributes 18
dependency graph 24
resource type definition 18
SQL Server 2008 Agent service agent
functions 22
SQL Server 2008 Analysis service agent
functions 22
SQL Server 2008 FILESTREAM agent
agent functions 21
SQL Server 2008 Filestream agent
attributes 22
resource type definition 21
SQL Server Agent Service agent
dependency graph 24
SQL Server Analysis Service agent

dependency graph 24
SQL Server database agent
monitoring options 27
MSDTC agent 23
SQL Server Filestream agent
dependency graph 24
standalone SQL Server
installing and configuring SFW (HA) 96
overview of HA configuration 94
sample HA configuration 95

U

uninstalling
SQL agent 129

V

VCS
configuring the cluster for DR 113
configuring the cluster for HA 36
configuring the SQL Server service group 114
configuring the SQL Server service group (Active-Active) 106
VCS for Windows
uninstalling 129
viewing distributed transactions 90
virtual DTC viewer 90