

Cluster Server Implementation Guide for Microsoft SQL Server

Windows

7.2

Cluster Server Implementation Guide for Microsoft SQL Server

Document version: 7.2 Rev 0

Last updated: 2016-10-25

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third party software for which Veritas 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. Refer to the third party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

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

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

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

Veritas Technologies LLC

500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

Technical Support maintains support centers globally. All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policies. For information about our support offerings and how to contact Technical Support, visit our website:

<https://www.veritas.com/support>

You can manage your Veritas account information at the following URL:

<https://my.veritas.com>

If you have questions regarding an existing support agreement, please email the support agreement administration team for your region as follows:

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

Make sure that you have the current version of the documentation. Each document displays the date of the last update on page 2. The latest documentation is available on the Veritas website:

<https://sort.veritas.com/documents>

Documentation feedback

Your feedback is important to us. Suggest improvements or report errors or omissions to the documentation. Include the document title, document version, chapter title, and section title of the text on which you are reporting. Send feedback to:

doc.feedback@veritas.com

You can also see documentation information or ask a question on the Veritas community site:

<http://www.veritas.com/community/>

Veritas Services and Operations Readiness Tools (SORT)

Veritas Services and Operations Readiness Tools (SORT) is a website that provides information and tools to automate and simplify certain time-consuming administrative tasks. Depending on the product, SORT helps you prepare for installations and upgrades, identify risks in your datacenters, and improve operational efficiency. To see what services and tools SORT provides for your product, see the data sheet:

https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf

Contents

Section 1	Introducing Veritas InfoScale solutions for application high availability	10
Chapter 1	Understanding the InfoScale solutions for application high availability	11
	About the Veritas InfoScale solutions for monitoring SQL Server	11
	About the VCS agents for SQL Server	12
	About the VCS agent for SQL Server Database Engine	12
	About the VCS agent for SQL Server FILESTREAM	18
	About the VCS GenericService agent for SQL Server Agent service and Analysis service	19
	About the agent for MSDTC service	19
	About the monitoring options	21
	How VCS monitors storage components	22
	Shared storage—if you use NetApp filers	22
	Shared storage—if you use SFW to manage cluster dynamic disk groups	23
	Shared storage—if you use Windows LDM to manage shared disks	23
	Non-shared storage—if you use SFW to manage dynamic disk groups	24
	Non-shared storage—if you use Windows LDM to manage local disks	24
	Non-shared storage—if you use VMware storage	25
	How is application availability achieved in a physical environment	25
	Typical SQL Server cluster configuration using shared storage	26
	Typical SQL Server disaster recovery cluster configuration	27
	SQL Server sample dependency graph	28
	MSDTC sample dependency graph	30
	How is application availability achieved in a VMware virtual environment	31
	How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host	33

	Typical VCS cluster configuration in a virtual environment	34
Chapter 2	Managing storage and installing the VCS agents	36
	Managing storage using NetApp filer	36
	Connecting virtual disks to the cluster node	38
	Disconnecting virtual disks from the cluster nodes	39
	Managing storage using Windows Logical Disk Manager	39
	Reserving disks (if you use Windows LDM)	41
	Creating volumes (if you use Windows LDM)	41
	Mounting volumes (if you use Windows LDM)	42
	Unassigning a drive letter	42
	Releasing disks (if you use Windows LDM)	43
	Managing storage using VMware virtual disks	43
	About installing the VCS agents	45
Chapter 3	Installing SQL Server	46
	About installing SQL Server for a high availability (HA) configuration	46
	Configuring Microsoft iSCSI initiator	48
	About installing SQL Server on the first system	48
	About installing SQL Server on additional systems	50
	Assigning ports for multiple SQL Server instances	51
	Enabling IPv6 support for the SQL Server Analysis Service	52
Section 2	Configuring SQL Server in a physical environment	53
Chapter 4	Overview	54
	About configuring SQL Server in physical environment	54
Chapter 5	Configuring the VCS cluster	57
	Configuring the cluster using the Cluster Configuration Wizard	57
	Configuring notification	66
	Configuring Wide-Area Connector process for global clusters	69
Chapter 6	Configuring the SQL Server service group	72
	About configuring the SQL Server service group	72
	Before configuring the SQL Server service group	73

Configuring a SQL Server service group using the wizard	75
Configuring detail monitoring for a SQL Server instance	81
Assigning privileges to the existing SQL Server databases and logs	83
Configuring the service group in a non-shared storage environment	84
Running SnapManager for SQL Server	87
About the modifications required for tagged VLAN or teamed network	88
Making SQL Server user-defined databases highly available	89
Create volumes or LUNs for SQL Server user-defined databases	89
Creating SQL Server databases	89
Adding storage agent resources to the SQL service group	90
Verifying the service group configuration	91
Bringing the service group online	91
Taking the service group offline	92
Switching the service group	92
Administering a SQL Server service group	93
Modifying a SQL service group configuration	93
Deleting a SQL service group	94
 Chapter 7	
Configuring an MSDTC service group	96
About configuring the MSDTC service group	96
Typical MSDTC service group configuration using shared storage	97
Before configuring the MSDTC service group	99
Creating an MSDTC service group	100
About configuring an MSDTC client	103
Configuring an MSDTC client	104
Verifying the installation	104
 Chapter 8	
Configuring the standalone SQL Server	106
Typical high availability configuration for a standalone SQL Server setup	106
Sample configuration	107
Configuring a standalone SQL Server for high availability	108
Moving the existing SQL Server data files and user databases	110

Chapter 9	Configuring an Active/Active cluster	111
	About running SQL Server in an active-active clustered environment	111
	Sample configuration	112
	Setting up the Active/Active cluster	114
Chapter 10	Configuring a disaster recovery setup	117
	Setting up the disaster recovery cluster	117
	Why implement a disaster recovery solution	117
	Understanding replication	118
	What needs to be protected in a SQL Server environment	118
	Configuring a disaster recovery set up for SQL Server	119
	Configuring replication using NetApp SnapMirror	121
	Configuring SnapMirror resources at the primary site	121
	Configuring the Global Cluster Option for wide-area failover	122
	Prerequisites	122
	Linking clusters: Adding a remote cluster to a local cluster	123
	Converting a local service group to a global service group	124
	Bringing a global service group online	126
	Administering global service groups	127
	Taking a remote global service group offline	127
	Switching a remote service group	127
	Deleting a remote cluster	128
Section 3	Configuring SQL Server in a VMware environment	131
Chapter 11	Configuring application monitoring using the Veritas High Availability solution	132
	Deploying the Veritas High Availability solution for configuring application monitoring	132
	Notes and recommendations	134
	Assigning privileges for non-administrator ESX/ESXi user account	136
	Configuring application monitoring	139
	Configuring the VCS cluster	140
	Configuring the application	143
	Modifying the ESXDetails attribute	148

Chapter 12	Administering application monitoring	151
	About the various interfaces available for performing application monitoring tasks	151
	Administering application monitoring using the Veritas High Availability tab	152
	Understanding the Veritas High Availability tab work area	152
	To view the status of configured applications	155
	To configure or unconfigure application monitoring	155
	To start or stop applications	156
	To suspend or resume application monitoring	158
	To switch an application to another system	158
	To add or remove a failover system	159
	To clear Fault state	163
	To resolve a held-up operation	164
	To determine application state	164
	To remove all monitoring configurations	164
	To remove VCS cluster configurations	165
	Administering application monitoring settings	165
	Administering application availability using Veritas High Availability dashboard	166
	Understanding the dashboard work area	167
	Monitoring applications across a data center	171
	Monitoring applications across an ESX cluster	171
	Searching for application instances by using filters	171
	Selecting multiple applications for batch operations	172
	Starting an application using the dashboard	172
	Stopping an application by using the dashboard	173
	Entering an application into maintenance mode	173
	Bringing an application out of maintenance mode	174
	Switching an application	174
Section 4	Appendixes	176
Appendix A	Troubleshooting	177
	VCS logging	177
	VCS Cluster Configuration Wizard (VCW) logs	179
	VCWsilent logs	179
	NetApp agents error messages	180
	Error and warning messages from VCS agent for SQL Server	181
	Messages from the VCS agent for SQL Server Database Engine	181

Messages from the VCS agent for SQL Server FILESTREAM	184
Messages from the VCS agent for SQL Server Agent service and Analysis service	184
SQL Server Analysis service (MSOLAP) service fails to come online with "invalid context of address" error	185
Messages from the VCS agent for MSDTC	185
Troubleshooting application monitoring configuration issues	187
Running the 'hastop -all' command detaches virtual disks	187
Validation may fail when you add a failover system	187
Adding a failover system may fail if you configure a cluster with communication links over UDP	187
Troubleshooting Veritas High Availability view issues	188
Veritas High Availability tab not visible from a cluster node	188
Veritas High Availability tab does not display the application monitoring status	188
Veritas High Availability tab may freeze due to special characters in application display name	189
Veritas High Availability view may fail to load or refresh	189
Operating system commands to unmount resource may fail	190

Appendix B	Using the virtual MMC viewer	191
	About using the virtual MMC viewer	191
	Viewing DTC transaction information	191

Introducing Veritas InfoScale solutions for application high availability

- [Chapter 1. Understanding the InfoScale solutions for application high availability](#)
- [Chapter 2. Managing storage and installing the VCS agents](#)
- [Chapter 3. Installing SQL Server](#)

Understanding the InfoScale solutions for application high availability

This chapter includes the following topics:

- [About the Veritas InfoScale solutions for monitoring SQL Server](#)
- [About the VCS agents for SQL Server](#)
- [How VCS monitors storage components](#)
- [How is application availability achieved in a physical environment](#)
- [How is application availability achieved in a VMware virtual environment](#)

About the Veritas InfoScale solutions for monitoring SQL Server

Veritas InfoScale provides application monitoring capability for SQL Server in the physical and VMware virtual environment.

The application monitoring capability is based on the VCS application agent and storage agent framework that combine together to monitor the application and storage components respectively.

In a physical environment, the application monitoring configuration employs shared or local storage. The shared storage employs NetApp filers over an iSCSI or Fibre Channel (FC) connection and NetApp SnapMirror for replication. The classic VCS infrastructure is used to configure an application for high availability in physical environment.

In a VMware virtual environment, the storage components employ non-shared virtual disks created on a data store or Raw Device Mappings (RDM)/SAN storage. The Veritas High Availability solution is used to configure an application for high availability in VMware virtual environment.

About the VCS agents for SQL Server

This agent provides high availability for Microsoft SQL Server by monitoring the SQL Server instance and its services in a VCS cluster.

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 Server services and virtual server are started on the new system. The agent monitors specific resources within an enterprise application, determines the status of these resources, brings them online, and takes them offline. The 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 agent package for SQL Server includes the following:

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

About the VCS agent for SQL Server Database Engine

The VCS agent for SQL Server Database Engine monitors the Database Engine service and all the optional components that are integrated with the service. For example, if the Full-Text Search component is configured, the agent also monitors

that component too. The agent brings the service online, monitors its status, and takes it offline. If the service is not running, the agent declares its state as Offline.

If detail monitoring is configured, the agent checks the health of the selected SQL Server databases or executes a monitoring script. If detail monitoring is successful, the agent declares the service group as Online.

Functions

Online	Brings the SQL Server Database Engine service online.
Offline	Takes the Database Engine service offline.
Monitor	Monitors the status of SQL Server services. It queries the Windows Service Control Manager (SCM) for the status of the services. If detail monitoring is configured, then depending on the configuration, the agent performs a database health check or executes a monitoring script. See "About the monitoring options" on page 21.
Clean	Forcibly stops the Database Engine service. Cleans up after a resource fails to come online, fails to go offline, or fails to detect the state as Online even when the resource is online.

Note: If you start the SQL Server services from outside VCS, then the SQL resource will go in an unknown state because the VCS agent monitors the computer context of the services. If the SQL Server 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 Server-related services from within VCS.

Resource type definitions

The resource type represents the VCS configuration definition of the agent, and specifies how the agent is defined in the configuration file, `main.cf`.

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

```
type SQLServer (
    static int IMF{} = { Mode=3, MonitorFreq=5, RegisterRetryLimit=3 }
    static il8nstr IMFRegList[] = { Instance }
    static il8nstr ArgList[] = { Instance, "LanmanResName:VirtualName",
        DetailMonitorInterval, SQLOnlineTimeout, SQLOfflineTimeout,
        SQLDetailMonitorTimeout, Username, Domain, Password, DBList,
        "IPResName:Address", SQLFile, FaultOnDMFailure }
    str Instance
```

```
    str LanmanResName
    str IPResName
    int DetailMonitorInterval = 0
    int SQLOnlineTimeout = 90
    int SQLOfflineTimeout = 90
    int SQLDetailMonitorTimeout = 30
    i18nstr Username
    i18nstr Domain
    str Password
    i18nstr SQLFile
    i18nstr DBList[]
    boolean FaultOnDMFailure = 1
)
```

Attributes

The following table lists the attributes that must be configured for the agent to function properly.

Table 1-1 VCS agent for SQL Server Database Engine—required attributes

Required Attribute	Description
Name: Instance Type: String Dimension: Scalar	Name of the SQL Server instance to monitor. If the value of this attribute is not set, the agent monitors the default SQL Server instance (MSSQLSERVER).
Name: LanmanResName Type: String Dimension: Scalar	Lanman resource name on which the SQLServer resource depends.

Table 1-2 VCS agent for SQL Server Database Engine—optional attributes

Optional attribute	Description
Name: SQLOnlineTimeout Type: Integer Dimension: Scalar	Number of seconds that may elapse before the Online function aborts. Default value: 90

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

Optional attribute	Description
<p>Name: SQLOfflineTimeout</p> <p>Type: Integer</p> <p>Dimension: Scalar</p>	<p>Number of seconds that may elapse before the Offline function aborts.</p> <p>Default value: 90</p>
<p>Name: DetailMonitorInterval</p> <p>Type: Integer</p> <p>Dimension: Scalar</p>	<p>Defines whether the agent performs detail monitoring of the SQL Server instance.</p> <p>The value 0 (zero) indicates that the agent performs only the basic monitoring of the instance service. A non-zero value indicates the number of online monitor cycles that the agent must wait before performing detail monitoring.</p> <p>Veritas recommends that you set this value to an integer between 1 and 12. The value 1 would make the agent perform detail monitoring at each monitor cycle. The value 2 would make the agent perform detail monitoring at every other monitor cycle. This interpretation may be extended to other values.</p> <p>If this attribute is set to a non-zero value, then the following attributes must be assigned appropriate values:</p> <ul style="list-style-type: none"> ■ For script-based monitoring: SQLFile, Username, Password, Domain, and SQLDetailMonitorTimeOut ■ For databases list-based monitoring: Username, Password, Domain, and DBList <p>Default value: 5</p> <p>Note: This is not a SQL Server agent-specific attribute, but a common type-level attribute. The value of this attribute can only be set through the wizard. If you configure the service group manually, you must remember to specify this attribute and its value.</p> <p>See “About the monitoring options” on page 21.</p>

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

Optional attribute	Description
<p>Name: FaultOnDMFailure</p> <p>Type: Boolean</p> <p>Dimension: Scalar</p>	<p>Defines whether the agent fails over the service group if the detail monitoring script execution fails.</p> <p>The value 1 indicates that the agent fails over the service group if detail monitoring script fails to execute. The value 0 indicates that it does not fail over, but goes into the unknown state.</p> <p>Defines whether the agent faults the SQL Server resource when the detail monitoring fails.</p> <ul style="list-style-type: none"> ■ If the value is set to True, the agent faults the configured SQL Server resource when the detail monitoring fails. In this case, the SQL Server resource may go into OFFLINE state. ■ If the value is set to False, the agent does not fault the SQL Server resource when the detail monitoring fails. In this case, the SQL Server resource may go into the UNKOWN state. <p>Default value: True</p>
<p>Name: SQLDetailMonitorTimeout</p> <p>Type: Integer</p> <p>Dimension: Scalar</p>	<p>Number of seconds that may elapse before the detail monitoring routine aborts.</p> <p>Default value: 30</p>
<p>Name: Username</p> <p>Type: String</p> <p>Dimension: Scalar</p>	<p>User account in whose context the detail monitoring is performed.</p> <p>If the DetailMonitorInterval attribute is set to a non-zero value, the value of this attribute must not be empty.</p> <p>The Microsoft Windows authentication name when logging in to a database for detail monitoring. This attribute must not be null if the DetailMonitorInterval attribute is set to a non-zero value. The user must have the necessary privileges to connect to the database and execute the appropriate query.</p> <p>Note: This attribute can take localized values.</p>

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

Optional attribute	Description
Name: Domain Type: String Dimension: Scalar	<p>Domain of the user account specified in the Username attribute.</p> <p>Domain for the user account. This attribute is used to create a trusted connection to the SQL Server instance if the DetailMonitorInterval attribute is set to a non-zero value.</p> <p>Note: This attribute can take localized values.</p>
Name: Password Type: String Dimension: Scalar	<p>Password for the user account specified in the Username attribute.</p> <p>If the DetailMonitorInterval attribute is set to a non-zero value, the value of this attribute must not be empty.</p> <p>Password for logging in to a database for in-depth monitoring. This attribute must not be null if the DetailMonitorInterval attribute is set to a non-zero value.</p>
Name: SQLFile Type: String Dimension: Scalar	<p>Location of the SQL file that is executed during a monitor cycle.</p> <p>If the DetailMonitorInterval attribute is set to a non-zero value, then either the script-based detail monitoring or the database list-based detail monitoring must be configured.</p> <p>This attribute must not be null if DetailMonitorInterval attribute is set to a non-zero value and script-based detail monitoring is configured.</p> <p>Note: This attribute can take localized values.</p>
Name: DBList Type: String Dimension: Vector	<p>List of databases for which the agent must perform detail monitoring.</p> <p>If the DetailMonitorInterval attribute is set to a non-zero value, then either the script-based detail monitoring or the database list-based detail monitoring must be configured.</p> <p>If both the SQLFile attribute and the DBList attribute are configured, then the DBList attribute takes precedence.</p> <p>Note: This attribute can take localized values.</p>
Name: IPResName Type: String Dimension: Scalar	<p>IP resource on which the Lanman resource for the SQLServer resource depends.</p>

About the VCS agent for SQL Server FILESTREAM

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

FILESTREAM 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.

Functions

Online	Enables FILESTREAM on the system.
Offline	Disables FILESTREAM on the system.
Monitor	Monitors FILESTREAM status on the system. If the agent is unable to query the status of FILESTREAM or if FILESTREAM is disabled on the system, the FILESTREAM resource in the application monitoring configuration faults.
Clean	Cleans up the resource state after the resource fails to come online, fails to go offline, or fails to be detected as online even though it is in the ONLINE state.

Resource type definition

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

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

Attributes

Table 1-3 VCS agent for SQL Server FILESTREAM—required attribute

Required attribute	Description
Name: InstanceName	Name of the SQL Server instance for which FILESTREAM is configured.
Type: String	If the value of this attribute is empty, the agent monitors the default SQL Server instance (<code>MSSQLSERVER</code>).
Dimension: Scalar	Note: This attribute can take localized values.

About the VCS GenericService agent for SQL Server Agent service and Analysis service

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

Agent functions

Online	Brings the configured SQL Server services online.
Offline	Takes the configured SQL Server services offline.
Monitor	Monitors the status of the configured SQL Server services. The agent queries the Service Control Manager (SCM) for the status of configured services.
Clean	Forcibly stops the configured SQL Server services. Cleans up after a resource fails to come online, fails to go offline, or fails to detect as ONLINE when resource is in an ONLINE state.

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

About the agent for MSDTC service

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

Agent functions

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. Cleans up after a resource fails to come online, fails to go offline, or fails to detect as ONLINE when resource is in an ONLINE state.

Note: The agent for MSDTC comprises of two component; the MSDTC client and the MSDTC server. These components must not be configured on the same cluster node.

Resource type definition

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

```
type MSDTC (
    static int IMF{} = { Mode=3, MonitorFreq=5, RegisterRetryLimit=3 }
    static il8nstr ArgList[] = { "LanmanResName:VirtualName",
    "MountResName:MountPath", LogPath }
    str LanmanResName
    str MountResName
    il8nstr LogPath
)
```

Attribute definitions

Table 1-4 MSDTC agent required attributes

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

About the monitoring options

Use the detail monitoring capability of the VCS agent for SQL Server to monitor the status of a SQL Server instance.

The agent provides the following levels of application monitoring:

- Basic monitoring queries the Windows Service Control Manager (SCM) to verify whether the configured SQL Server services are continuously active.
- Detail monitoring queries the databases to verify their availability.

Table 1-5 Methods of configuring detail monitoring

Method	Description
Database list-based monitoring	<p>The agent monitors only the list of databases that are specified in the its DBList attribute. It uses Microsoft ActiveX Data Objects (ADO) to establish a connection with the selected databases to verify the health of those databases.</p> <ul style="list-style-type: none">■ If the connection is successful, the agent considers the SQL Server instance as available.■ If the connection fails, the SQL Server instance is considered not available. Then, if the FaultOnDMFailure attribute is configured, the service group fails over to the failover nodes.
Script-based detail monitoring	<p>The agent uses a script to monitor the status of the SQL Server instance.</p> <ul style="list-style-type: none">■ If the script is successfully executed during monitoring, the agent considers the instance available.■ If the execution fails, the instance is considered not available. Then, if the FaultOnDMFailure attribute is configured, the service group fails over to the failover nodes. <p>A sample SQL script is provided with the agent for this purpose. Customize the script to meet your configuration requirements, or use your own script, which can be placed at any other location.</p> <p>The script is located at:</p> <pre>%VCS_HOME%\bin\SQLServer\sample_script.sql</pre> <p>where, %VCS_HOME% points to the default installation directory, typically: C:\Program Files\Veritas\Cluster Server.</p> <p>Use a separate script for each SQL Server service group in the cluster and place it on every node in the service group.</p>

You can enable and configure detail monitoring by running the SQL Server Agent Configuration Wizard.

See [“About configuring the SQL Server service group”](#) on page 72.

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

Note: The VCS agent monitors the computer context of the services that are configured for high availability. If the SQL Server service is not started in the virtual server context—that is, outside VCS—the resource goes into the UNKNOWN state. To avoid this situation, ensure that you start all the SQL Server-related services from within VCS.

How VCS monitors storage components

VCS provides specific agents that monitor storage components and ensure that the shared disks, disk groups, LUNs, volumes, and mounts are accessible on the system where the application is running. Separate agents are available for shared and non-shared storage and for third-party storage arrays such as NetApp filers. Your storage configuration determines which agent should be used in the high availability configuration.

For details on the various VCS storage agents, refer to the *Cluster Server Bundled Agents Reference Guide*.

Shared storage—if you use NetApp filers

The VCS hardware replication agents for NetApp provide failover support and recovery in environments that employ NetApp filers for storage and NetApp SnapMirror for replication. The agents enable configuring NetApp filers over an iSCSI or Fibre Channel (FC) connection in a VCS cluster environment.


The VCS agents for NetApp are as follows:

- NetAppFiler
- NetAppSnapDrive
- NetAppSnapMirror

These agents monitor and manage the state of replicated filer devices and ensure that only one system has safe and exclusive access to the configured devices at a time. The agents can be used in local clusters, single VCS replicated data clusters,

and multi-cluster environments that are set up using the VCS Global Cluster Option (GCO).

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



VCS also provides agents for other third-party hardware arrays. For details on the supported arrays, refer to the product Software Compatibility List (SCL).

Shared storage—if you use SFW to manage cluster dynamic disk groups

The VCS MountV and VMDg agents are used to monitor shared storage that is managed using Storage Foundation (SFW). SFW manages storage by creating disk groups from physical disks. These disk groups are further divided into volumes that are mounted on the cluster systems.

The MountV agent monitors volumes residing on disk groups. The VMDg agent monitors cluster dynamic disk groups and is designed to work using SCSI reservations. Together the MountV and VMDg agents ensure that the shared cluster dynamic disk groups and volumes are available.

Shared storage—if you use Windows LDM to manage shared disks

The VCS Mount and DiskReservation (DiskRes) agents are used to monitor shared disks that are managed using Windows Logical Disk Management (LDM).

The Mount agent monitors basic disks and mount points and ensures that each system is able to access the volume or mount path in the same way. The DiskRes agent monitors shared disks and uses persistent reservation to ensure that only one system has exclusive access to the disks. During failovers, these agents ensure that the disks and volumes are deported and imported on the node where the application is running.

Non-shared storage—if you use SFW to manage dynamic disk groups

VCS introduces the Volume Manager Non-Shared Diskgroup (VMNSDg) agent to support local non-shared storage configurations that are managed using SFW. The VMNSDg agent works without SCSI reservations and is designed for locally attached storage devices that do not support SCSI.

The VMNSDg agent monitors and manages the import and deport of dynamic disk groups created on local storage. The only difference between the VMDg agent and the VMNSDg agent is that the VMDg agent is designed for shared cluster dynamic disk groups and uses SCSI reservations, whereas the VMNSDg agent supports only non-shared local dynamic disk groups and works without SCSI reservations.

The VMNSDg agent can be used to set up single node Replicated Data Clusters (RDC) or Disaster Recovery (DR) configurations with replication set up between the sites.

During a failover, the VCS MountV and VMNSDg agents deport the locally attached storage from the affected node and then import the locally attached storage of the target node. Replication ensures that the data is consistent and the application is up and running successfully.

Note: The VMNSDg agent does not support fast failover and Intelligent Monitoring Framework (IMF).

Non-shared storage—if you use Windows LDM to manage local disks

VCS introduces the NativeDisks agent to support local non-shared storage configurations managed using Windows LDM. The NativeDisks agent works without SCSI reservations and is designed for local storage that does not support SCSI.

Together with the Mount agent, the NativeDisks agent monitors and manages the import and deport of basic local disks on the system. The only difference between the DiskRes agent and the NativeDisks agent is that the DiskRes agent is designed for shared disks and uses SCSI reservations, whereas the NativeDisks agent supports only non-shared local disks and works without SCSI reservations.

Note: The NativeDisks agent does not support fast failover and Intelligent Monitoring Framework (IMF).

Non-shared storage—if you use VMware storage

VCS introduces the VMwareDisks agent to support storage configurations in a VMware virtual environment. The agent is platform independent and supports VMware Virtual Machine Disk (VMDK), Raw Device Mapping (RDM) disk files (virtual), and storage that is configured using Network File System (NFS). The VMwareDisks agent works without SCSI reservations and supports locally attached non-shared storage.

VMware features such as snapshots, vMotion, and DRS do not work when SCSI disks are shared between virtual machines. The VMwareDisks agent is designed to address this limitation. With this agent, the disks can now be attached to a single virtual machine at a time in the VCS cluster. On failover, along with the service group, the VMwareDisks agent moves the disks to the target virtual machine.

The VMwareDisks agent communicates with the host ESXi server to configure storage. This agent manages the disk attach and detach operations on a virtual machine in the VCS cluster. The agent is VMware HA aware. During failovers, the agent detaches the disk from one system and then attaches it to the system where the application is actively running. The VMwareDisks agent presents the virtual disks to the operating system. On Windows, the agent relies on the VMNSDg agent (in case of SFW-managed local storage) and the NativeDisks agent (in case of LDM-managed local storage) for initializing and managing the virtual disks. On Linux, the agent relies on the LVM and VxVM agents.

Note: The VMwareDisks agent does not support fast failover and Intelligent Monitoring Framework (IMF).

How is application availability achieved in a physical environment

The VCS agents continuously monitor the application, storage, and network components that the application uses in the cluster. The agents are able to detect failures in all of these components. For example, an application-level failure such as a configured application virtual server or application service becoming unavailable, a fault in the storage such as a configured disk becoming inaccessible, or a network failure.

When a fault occurs, VCS fails over the application service group to the next available system in the application service group's system list. A service group failover means that the VCS storage agents deport and import the disks or LUNs on the new system. The VCS network agents bring the network components online and the application-specific agents then start the application services on the new system.

In a disaster recovery cluster configuration, VCS first attempts to failover the application service group within the local cluster. If all the systems in the local cluster are unavailable, VCS attempts to failover the service group to a system at the remote site.

In a NetApp environment, the VCS NetApp agents perform the following actions in that order:

- Connect the virtual disks (LUNs) to the target hosts (NetAppSnapDrive agent).
- Perform a mirror break that enables write access to the target (NetAppSnapMirror agent).
- Reverse the direction of replication by demoting the original source to a target, and begin replicating from the new source (NetAppSnapMirror agent).

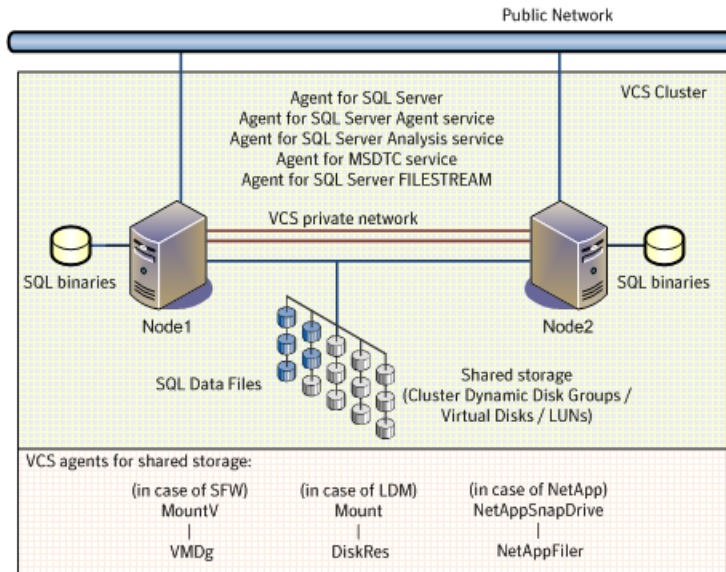
If replication is set up using Volume Replicator (Volume Replicator), the Volume Replicator replication agents make the Secondary RVG at the remote site write-enabled so that it becomes the new Primary. After the storage is connected, VCS starts the application services on the new system at the remote site. The data that is replicated to the remote site is used to restore the application services to the clients.

Typical SQL Server cluster configuration using shared storage

A typical SQL Server cluster configuration using shared storage 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 shared storage can be virtual disks or LUNs managed using NetApp suite of products, or shared cluster dynamic disk groups managed using SFW, or shared disks managed using Windows LDM. Appropriate VCS storage agent resources are configured depending on how the shared storage is managed.

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


Figure 1-1 Typical SQL Server cluster configuration using shared storage



Typical SQL Server disaster recovery cluster 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 1-2 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.

Note: The figure depicts a typical configuration. The number of systems at the primary and secondary site clusters need not be the same.

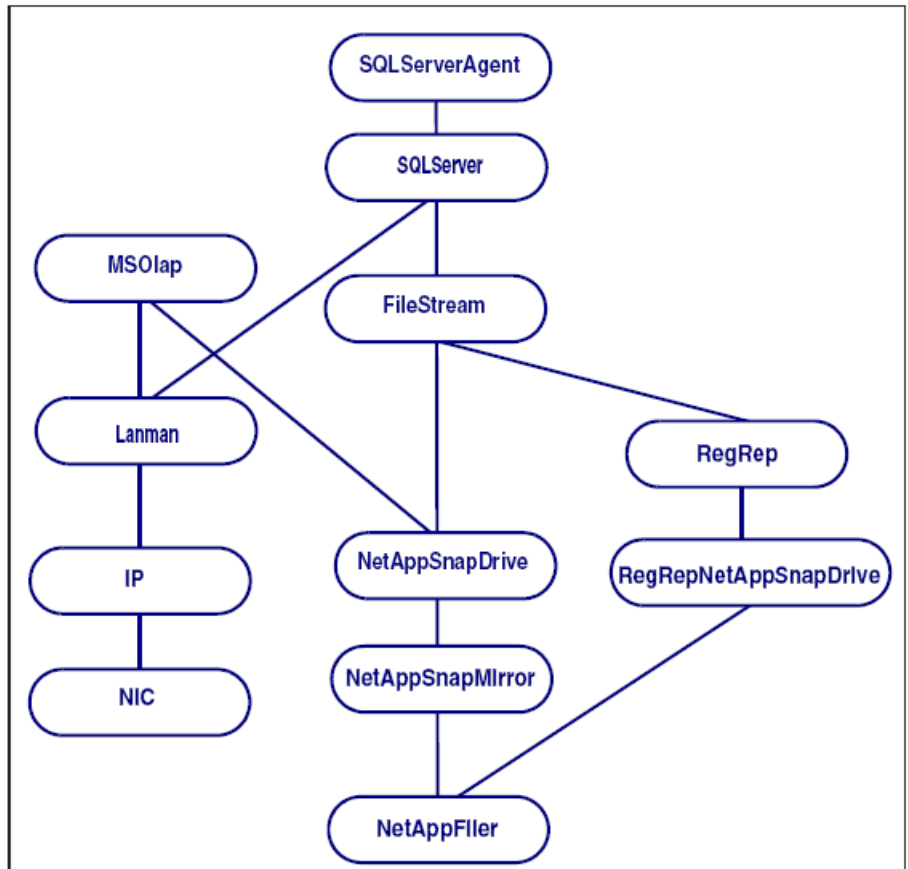
Data is replicated from the primary site to the secondary site. Replication between the storage is set up using a replication software. In case of a NetApp environment, replication between the filers at the primary and secondary sites is set up using NetApp SnapMirror for SQL. If the Microsoft SQL Server server on System1 fails, SQL Server comes online on node System2 and begins servicing requests. From the user's perspective there might be a small delay as the backup node comes online, but the interruption in effective service is minimal.

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

SQL Server sample dependency graph

The following figure is a sample dependency graph, which depicts the resources and their dependencies a typical SQL Server service group in a NetApp environment.

Figure 1-3 Dependency graph for a SQL Server service group



A SQL Server service group is configured to make a SQL Server application highly available in a VCS cluster. The SQL Server Analysis service (MSOLAP) and the SQL Server Agent service (SQLServerAgent) are configured as GenericService resources.

As the dependency graph depicts:

- The virtual name for the SQL Server is configured using the Lanman resource, and the service group IP address is configured using the IP and NIC resources.
- The FileStream resource monitors the Windows FILESTREAM configuration settings for the SQL Server instance.
- This graphic depicts shared storage that is configured in a NetApp environment.

If the shared storage is managed using Windows LDM, the VCS Mount and DiskRes resources replace the NetAppSnapDrive, NetAppSnapMirror, and NetAppFiler resources.

In case of non-shared storage managed using Windows LDM, the VCS Mount and NativeDisks resources should be configured instead.

In case the storage belongs to a VMware virtual environment, the VCS Mount, NativeDisks, and VMwareDisks resources are configured instead of the NetApp resources.

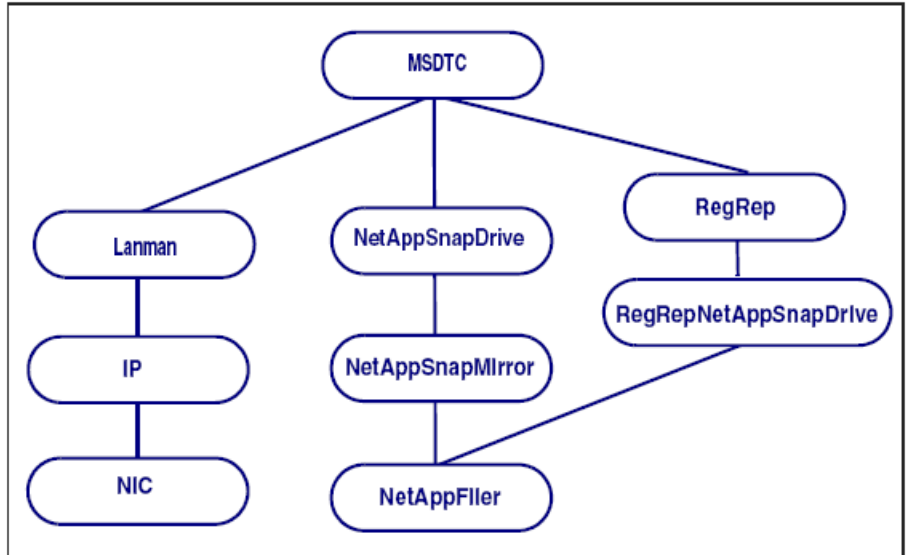
- The SQL Server registry is replicated using the RegRep and RegRepNetAppSnapDrive resources.

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

MSDTC sample dependency graph

The following figure is a sample dependency graph, which depicts the resources and their dependencies a typical MSDTC service group.

Figure 1-4 Dependency graph for an MSDTC service group



An MSDTC service group is configured to monitor the state of the MSDTC services in a VCS cluster.

As the dependency graph depicts:

- The virtual name for the MSDTC server is configured using the Lanman resource, and the service group IP address is configured using the IP and NIC resources.
- This graphic depicts shared storage configured in a NetApp environment. If the shared storage is managed using Windows LDM, the VCS Mount and DiskRes resources replace the NetAppSnapDrive, NetAppSnapMirror, and NetAppFile resources.
In case of non-shared storage managed using Windows LDM, the VCS Mount and NativeDisks resources should be configured instead.
In case the storage belongs to a VMware virtual environment, the VCS Mount, NativeDisks, and VMwareDisks resources are configured instead of the NetApp resources.
- The MSDTC registry is replicated using the RegRep and RegRepNetAppSnapDrive resources.

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

How is application availability achieved in a VMware virtual environment

The Veritas High Availability solution for VMware employs Cluster Server (VCS) and its agent framework to monitor the state of applications and their dependent components running on the virtual machines that use non-shared storage. Specific agents are available to monitor the application, storage, and network components. Together, these agents monitor the overall health of the configured applications by running specific commands, tests, or scripts.

The storage configuration in the VMware virtual environment determines how VCS functions differently in a non-shared virtual environment. The non-shared storage configuration in the VMware virtual environment involves the VMware VMDK and RDM disks that reside on the shared datastore. This datastore is accessible to multiple virtual machines. However, the disks are attached to a single virtual machine at any given point of time. VCS provides a new storage agent “VMwareDisks” that communicates with the VMware ESX/ESXi hosts to perform the disk detach and attach operations to move the storage disk between the virtual machines, in a VCS cluster.

Note: By default the VMwareDisks agent communicates with the ESX/ESXi host to perform the disk detach and attach operations. However, instead of the ESX/ESXi hosts you can choose to communicate with the vCenter Server to perform these operations.

See [“How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host”](#) on page 33.

In event of an application failure, the agents attempt to restart the application services and components for a configurable number of times. If the application fails to start, they initiate an application fail over to the failover target system. During the fail over, the VMwareDisks agent moves the storage disk to the failover target system, the network agents bring the network components online, and the application-specific agents then start the application services on the failover target system.

In case of a virtual machine fault, the VCS agents begin to fail over the application to the failover target system. The VMwareDisks agent sends a disk detach request to the ESX/ESXi host. After the detach operation is successful, the agent proceeds to attach the disks to the new failover target system.

In a scenario where the ESX/ESXi host itself faults, the VCS agents begin to fail over the application to the failover target system that resides on another host. The VMwareDisks agent communicates with the new ESX/ESXi host and initiates a disk detach operation on the faulted virtual machine. The agent then attaches the disk to the new failover target virtual machine.

In event of a failure in a site recovery configuration, the following tasks are performed for application monitoring continuity:

- The virtual machines at the protected site are failed over to the recovery site.
- The pre-online script defined in the form of a command in the SRM recovery plan applies the specified attribute values for the application components.
- The status monitoring script retrieves the application status.
- The network agents bring the network components online and the application-specific agents start the application services on the failover target system.

For details on the VCS configuration concepts and clustering topologies, refer to the *Cluster Server Administrator's Guide*.

For details on the application agents, refer to the application-specific agent guide.

For details on the storage agents, refer to the *Cluster Server Bundled Agents Reference Guide*.

How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host

In addition to the ESX hosts the VMwareDisks agent can also communicate the disk detach and attach operations with the vCenter Server to which the virtual machines belong.

In this scenario, in event of a failure, the VMwareDisks agent sends the disk detach and attach requests to the vCenter Server (instead of the ESX hosts). The vCenter Server then notifies the ESX host for these operations. Since the communication is directed through the vCenter Server, the agent successfully detaches and attaches the disks even if the ESX host and the virtual machines reside in a different network.

In a scenario where the host ESX/ESXi itself faults, the VMareDisks agent from the target virtual machine sends a request to the vCenter Server to detach the disks from the failed virtual machine. However, since the host ESX has faulted, the request to detach the disks fails. The VMwareDisks agent from the target virtual machine now sends the disk attach request. The vCenter Server then processes this request and disks are attached to the target virtual machine. The application availability is thus not affected.

See [“Modifying the ESXDetails attribute”](#) on page 148.

Limitation

The configuration of VMwareDisks agent to communicate with the vCenter Server has the following limitation:

If VMHA is not enabled and the host ESX faults, then even after the disks are attached to the target virtual machine they remain attached to the failed virtual machine. This issue occurs because the request to detach the disks fails since the host ESX itself has faulted. The agent then sends the disk attach request to the vCenter Server and attaches the disks to the target virtual machine.

Even though the application availability is not impacted, the subsequent power ON of the faulted virtual machine fails. This issue occurs because of the stale link between the virtual machine and the disks attached. Even though the disks are now attached to the target virtual machine the stale link with the failed virtual machine still exists.

Workaround

As a workaround, you must manually detach the disks from the failed virtual machine and then power ON the machine.

About the vCenter Server user account privileges

You must have the administrative privileges or must be a root user to communicate the disk detach and attach operations through the vCenter Server. If the vCenter

Server user account fails to have the administrative privileges or is not a root user, then the disk detach and attach operation may fail, in event of a failure.

If you do not want to use the administrator user account or the root user, then you must create a role and add the following privileges to the created role:

- "Low level file operations" on datastore
- "Add existing disk" on virtual machine
- "Change resource" on virtual machine
- "Remove disk" on virtual machine

After you create a role and add the required privileges, you must add a local user to the created role. You can choose to add an existing user or create a new user.

Refer to the VMware product documentation for details on creating a role and adding a user to the created role.

Typical VCS cluster configuration in a virtual environment

A typical VCS cluster configuration in a VMware virtual environment involves two or more virtual machines. The virtual machine on which the application is active, accesses a non-shared VMware VMDK or RDM disk that resides on a VMware datastore.

The virtual machines involved in the VCS cluster configuration may belong to a single ESX host or could reside on separate ESX hosts. If the virtual machines reside on separate ESX hosts, the datastore on which the VMware VMDK or RDM disks (on which the application data is stored) reside must be accessible to each of these ESX hosts.

The application binaries are installed on the virtual machines and the data files are installed on the VMware disk drive. The VCS agents monitor the application components and services, and the storage and network components that the application uses.

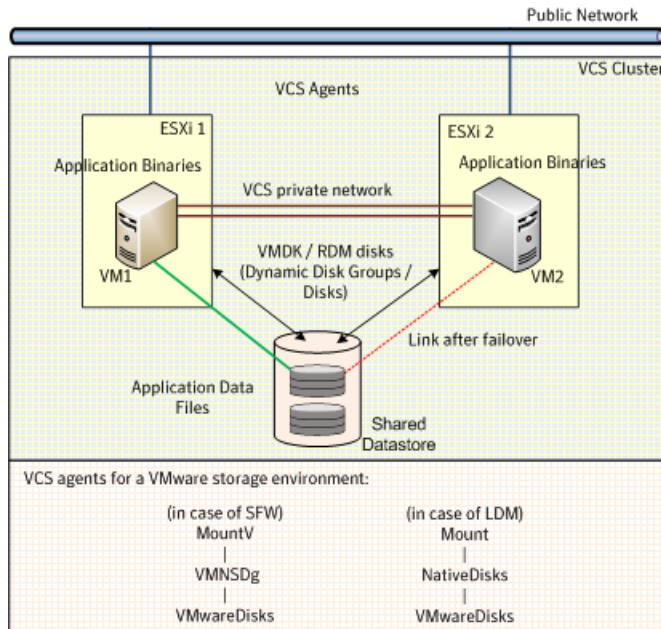
During a failover, the VCS storage agents (MountV-VMNSDg-VMwareDisks in case of SFW storage, Mount-NativeDisks-VMwareDisks in case of LDM storage) move the VMware disks to the new system. The VCS network agents bring the network components online, and the application-specific agents then start the application services on the new system.

In a site recovery environment, Veritas High Availability solution additionally provides script files for the following tasks. These files are invoked when the SRM recovery plan is executed.

- Set up communication between the vCenter Server and the SRM Server at the recovery site and the virtual machines at the protected site.

- Assign a SiteID to both the sites.
- Specify attribute values for the application components at the respective site.
- Retrieve the application status in the SRM recovery report, after the virtual machine is started at the recovery site.

Figure 1-5 Typical cluster configuration in a VMware virtual environment



Managing storage and installing the VCS agents

This chapter includes the following topics:

- [Managing storage using NetApp filer](#)
- [Managing storage using Windows Logical Disk Manager](#)
- [Managing storage using VMware virtual disks](#)
- [About installing the VCS agents](#)

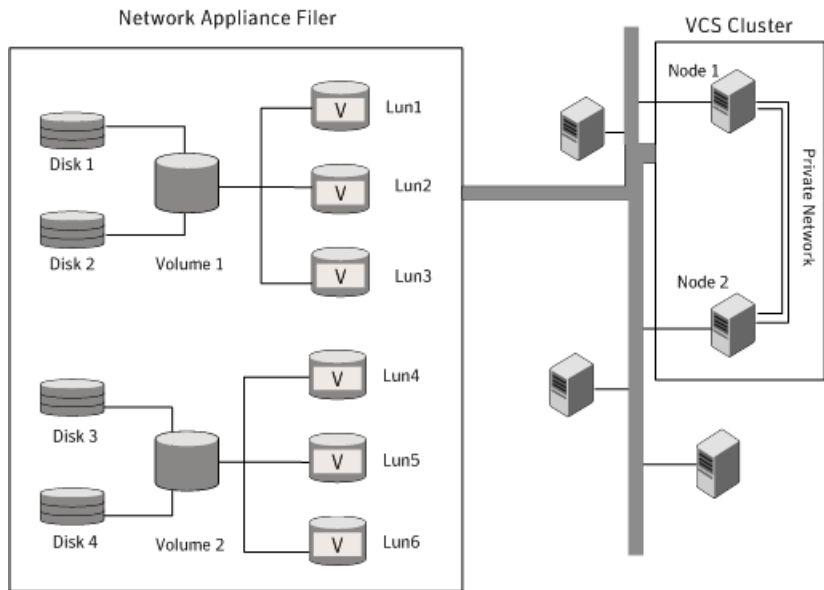
Managing storage using NetApp filer

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

Refer to the NetApp documentation for more information.

The following figure illustrates a typical VCS cluster in a NetApp storage environment.

Figure 2-1 VCS cluster in a NetApp storage environment



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

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

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

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

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

Contains the user database files.

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

These LUNs must be accessible from all cluster nodes.

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

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

Connecting virtual disks to the cluster node

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

To connect virtual disks to the cluster node

- 1** On the cluster node where you want to connect the LUN, launch the Computer Management MMC from **Start > All Programs > Administrative Tools > Computer Management** or, on Windows Server 2012 operating systems, click **Administrative Tools** from the **Start** screen.
- 2** From the left pane, expand **Storage** and double-click **SnapDrive**.
- 3** Right-click **Disks** and then click **Connect Disk** to launch the Connect Disk wizard.
- 4** Click **Next** on the Welcome page.
- 5** Specify the path of the virtual disk that you wish to connect to the cluster node and then click **Next**.
- 6** Select **Dedicated** as the Virtual Disk Type and then click **Next**.
- 7** Click **Assign a Drive Letter** and then choose a drive letter from the drop-down list.
- 8** On the Select Initiator panel, specify the initiator(s) for the virtual disk and then click **Next**.

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

Disconnecting virtual disks from the cluster nodes

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

To disconnect virtual disks

- 1 On the cluster node where you want to disconnect the LUNs, launch the Computer Management MMC from **Start > All Programs > Administrative Tools > Computer Management** or, on Windows Server 2012 operating systems, click **Administrative Tools** from the **Start** screen .
- 2 From the left pane, expand **Storage** and double-click **SnapDrive**.
- 3 Double-click **Disks** to see the LUNs that are connected to the node.
- 4 Right-click the LUN you want to disconnect and then click **Disconnect Disk**.
- 5 In the Disconnect Disk alert box, click **OK**.

Managing storage using Windows Logical Disk Manager

If your configuration uses shared disks and volumes that are managed using Windows Logical Disk Manager (LDM), use the VCS Mount and DiskReservation (DiskRes) agents. If you use LDM to manage non-shared local storage, use the VCS Mount and NativeDisks agents.

Before configuring the storage, review the resource types and attribute definitions of these VCS storage agents (Mount, DiskRes, NativeDisks) described in the *Cluster Server Bundled Agents Reference Guide*.

The following restrictions apply for storage managed using LDM:

- Mount, DiskRes, and NativeDisks agents are supported on VCS for Windows only. These agents are not supported if the storage is managed using Storage Foundation (SFW).
- If you are using shared storage, your storage devices must be configured to use SCSI-2 disk reservations. SCSI-3 is not supported.
SCSI support is not required if you are using non-shared storage.
- LDM support is not applicable for Disaster Recovery configurations. Currently only HA configurations are supported.

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

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

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

Veritas recommends that you create separate volumes for the following:

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

If you are using a shared storage configuration, ensure that these volumes are created on shared storage and are accessible from all cluster nodes.

If you are using a non-shared storage configuration, create these volumes separately on the local storage attached to each cluster node.

Perform the following tasks to configure your storage:

- **Reserve disks**
See [“Reserving disks \(if you use Windows LDM\)”](#) on page 41.
- **Create volumes**
See [“Creating volumes \(if you use Windows LDM\)”](#) on page 41.
- **Mount volumes**
See [“Mounting volumes \(if you use Windows LDM\)”](#) on page 42.
- **Unassign the drive letter**
See [“Unassigning a drive letter”](#) on page 42.

- Release the disks
 See [“Releasing disks \(if you use Windows LDM\)”](#) on page 43.

Reserving disks (if you use Windows LDM)

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

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

To reserve the disks

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

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

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

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

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

For example, to reserve disk #4, type:

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

Make a note of the disk number and the corresponding signature. You will require these details to identify and reserve the disks during installation and while configuring the service group, on additional nodes in the cluster.

Creating volumes (if you use Windows LDM)

Perform the following steps to create volumes.

To create volumes

- 1 Use the Windows Disk Management tool to verify that the disks are visible on the cluster nodes, and then create volumes on the disks.
- 2 In case of shared storage, after creating the required volumes on a node, release the reserved disks from that node.
 See [“Releasing disks \(if you use Windows LDM\)”](#) on page 43.
- 3 In case of shared storage, rescan the disks on all the remaining nodes in the cluster.

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

Mounting volumes (if you use Windows LDM)

Perform the following steps to mount volumes on a cluster node.

To mount a volume

- 1** Use the Windows Disk Management tool to mount the volumes that you created earlier.
- 2** After mounting the volumes on a cluster node, run the CHKDSK command and verify that there are no errors on the mounted volumes.
- 3** Make a note of the drive letters that you assign to the mounted volumes.

Use the same drive letters while mounting these volumes on the remaining cluster nodes.

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

Unassigning a drive letter

In case of a shared storage configuration, 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.

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

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

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

To unassign drive letter

- 1** Log in as Administrator.
- 2** Open Disk Management. Type the following at the command prompt:

```
C:\> diskmgmt.msc
```

- 3** Right-click the partition or logical drive and click **Change Drive Letter and Path**.
- 4** In the **Change Drive Letter and Paths** dialog box, click the drive letter and click **Remove**.

Releasing disks (if you use Windows LDM)

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

These steps are required only if you are configuring shared storage. Skip these steps for a non-shared storage configuration.

To release disks

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

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

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

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

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

For example, to release disk 4, type:

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

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

Managing storage using VMware virtual disks

Configure the storage disks to save the application data.

VMware virtualization manages the application data by storing it on SAN LUNs (RDM file), or creating virtual disks on a local or networked storage attached to the ESX host using iSCSI, network, or Fibre Channel. The virtual disks reside on a datastore or a raw disk that exists on the storage disks used.

For more information, refer to the VMware documentation.

The application monitoring configuration in a VMware environment requires you to use the RDM or VMDK disk formats. During a failover, these disks can be deported from a system and imported to another system.

Consider the following to manage the storage disks:

- Use a networked storage and create virtual disks on the datastores that are accessible to all the ESX servers that hosts the VCS cluster systems.
- In case of virtual disks, create non-shared virtual disks (Thick Provision Lazy Zeroed).

- Add the virtual disks to the virtual machine on which you want to start the configured application.
- Create volumes on the virtual disks.

Note: If your storage configuration involves NetApp filers that are directly connected to the systems using iSCSI initiator, you cannot configure application monitoring in a virtual environment with non-shared disks.

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

If you use SQL Server FILESTREAM, create additional volumes for the FILESTREAM-enabled database objects.

Veritas recommends that you create separate volumes for the following:

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

The following VCS storage agents are used to monitor the storage components involving non-shared storage:

- If the storage is managed using SFW, the MountV, VMNSDg, and VMwareDisks agents are used.
- If the storage is managed using LDM, the Mount, NativeDisks, and VMwareDisks agents are used.

Before configuring the storage, you can review the resource types and attribute definitions of these VCS storage agents. For details refer to the *Cluster Server Bundled Agents Reference Guide*.

About installing the VCS agents

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

Refer to the *Veritas InfoScale Installation and Upgrade Guide* for instructions.

For the latest information about supported operating systems and software, see the Software Compatibility List at:

<https://sort.veritas.com/documents>

Installing SQL Server

This chapter includes the following topics:

- [About installing SQL Server for a high availability \(HA\) configuration](#)
- [About installing SQL Server on the first system](#)
- [About installing SQL Server on additional systems](#)
- [Assigning ports for multiple SQL Server instances](#)
- [Enabling IPv6 support for the SQL Server Analysis Service](#)

About installing SQL Server for a high availability (HA) configuration

Review the following considerations for installing SQL Server:

- If you use shared NetApp storage, ensure that the following configurations are in place:
 - The Microsoft iSCSI Initiator is configured to establish a persistent connection between the NetApp filer and the systems on which you plan to install SQL Server.
 - The volumes are created on an external, basic disk, or LUNs (virtual disks) on a NetApp filer. These disks must be mounted on or connect to the first system on which you plan to install SQL Server.
- If you use Windows LDM storage, ensure that the disks are accessible from all the systems on which you plan to install SQL Server .
Veritas recommends that you create volumes for the following:
 - SQL Server data
 - Registry replication

- User defined database
- User defined database logs
- FILESTREAM enabled database objects
- When installing SQL Server on the first system in a physical environment, install the SQL Server instance on the local disk.
 - In case of shared storage, install the database files and the analysis service files on shared disks. These disks must be accessible from all the systems on which you plan to install SQL Server.
 - In case of non-shared storage, install the database files and analysis service files on the disks that reside on a datastore. This datastore must be accessible from all the systems on which you plan to install SQL Server. These disks are deported and imported during a failover.
- When installing SQL Server on the first system in a VMware virtual environment, install the SQL Server instance at the default path. Ensure that you configure the data folders and log folders on a VMDK.
- When installing SQL Server on additional systems in a physical environment, it does not matter whether the storage is shared or non-shared. Install the SQL Server instances, database files, and analysis service files on the local disk.
- When installing SQL Server on additional systems in a VMware virtual environment, install the SQL Server instance at the default path. The data folders and log folders may be placed anywhere. However, ensure that you use the same instance name that you used when installing SQL Server on the first node.
- If you are installing multiple instances of SQL Server on the same system, ensure the following:
 - Assign a unique name and a unique instance ID to each SQL Server instance. When installing SQL Server on additional systems for the same instance, ensure that you specify the same instance name and ID.
 - The order of the instance installation does not matter. You must ensure that the instances are installed with the same name and ID on all the systems.
 - Assign a unique port number for each instance.
- (SQL Server 2012 or 2014 only) Ensure that the [NT AUTHORITY\SYSTEM] account is granted the sysadmin server role (from SQL Management Studio Console).
- The logged-on user must be a domain user with local Administrator privileges.
- The logged-on user must be a member of the local Administrators group on all there systems where you plan to install SQL Server .

Configuring Microsoft iSCSI initiator

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

To configure Microsoft iSCSI initiator on a Windows Server system

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

About installing SQL Server on the first system

Review the following considerations for installing SQL Server on the first system:

- Install SQL Server in the standalone installation mode in a non-clustered environment.
- Make sure that the volumes or LUNs (virtual disks) required for SQL Server are mounted or connected to the system.
- When installing SQL Server on the first system in a VMware virtual environment, install the SQL Server instance at the default path. Place the data folders and log folders on a VMDK. The VMware disk agent detaches and attaches the disk between nodes during a failover.

Open the SQL Server Installation Center using the appropriate Microsoft SQL Server installer.

To launch the SQL Server Setup wizard, click **New SQL Server stand-alone installation or add features to an existing installation** on the Installation panel.

On the SQL Server Setup wizard, provide the following information:

- On the Feature Selection panel, ensure that you select all the features (for example, Full-Text Search, Analysis Services) that you plan to configure for high availability.
- On the Instance Configuration panel, make the following selections:

- Specify the instance name and instance ID, and make a note of these values. You must use the same name and ID when you install the SQL Server instance on additional systems.
If you plan to install multiple instances of SQL Server in a cluster, each instance must have a unique instance name and instance ID. Later, use these same values when installing the instance on the corresponding additional systems.
- Ensure that the **Instance root directory** resides on the local system disk. Doing so installs the SQL Server instance on the local disk. The location of data directories is to be specified later on a different panel of the installation wizard.
- On the Service Accounts tab of the Server Configuration panel, specify the following values:
 - Specify a domain user name for the SQL Server services account.

Note: This user may not be a part of the local Administrators group on all the systems. In this case, you must configure the **SQLClusterAccount** attribute of the VCS agent while configuring the SQL Server service group later.

- Set the Startup Type of the SQL Server Browser service to Automatic.
- Set the Startup Type of all the other services to Manual.
Alternatively, use the SQL Server Configuration Manager to change the Startup Type of each service it after the installation is complete.
- On the Data Directories tab of the Database Engine Configuration panel, verify the following:
 - In case of shared storage, all the directories must be located on the shared disks.
 - In case of non-shared storage, all the directories must be located on the disks that reside on a shared datastore. These include Data root, User database, User database log, Temp DB, Temp DB log, and Backup.
- On the Data Directories tab of the Analysis Services Configuration panel, verify the following:
 - In case of shared storage, all the directories must be located on the shared disks.

- In case of non-shared storage, all the directories must be located on the disks that reside on a shared datastore. These directories include Data, Log file, Temp, and Backup.

Refer to the Microsoft documentation for the application-specific instructions.

About installing SQL Server on additional systems

Note: This topic is not applicable in case of non-shared storage configurations. Instead, follow the prerequisites and procedure for installing SQL Server on the first system.

See [“About installing SQL Server on the first system”](#) on page 48.

Review the following considerations for installing SQL Server on the additional systems:

- Install SQL Server in the standalone installation mode in a non-clustered environment.
- Ensure that all the SQL Server services other than the SQL Server Browser service are stopped on the first system where this instance is installed.
- When installing SQL Server on additional systems in a VMware virtual environment, install the SQL Server instance at the default path. The data folders and log folders may be placed anywhere. However, ensure that you use the same instance name that you used when installing SQL Server on the first node.

Open the SQL Server Installation Center using the appropriate Microsoft SQL Server installer.

To launch the SQL Server Setup wizard, click **New SQL Server stand-alone installation or add features to an existing installation** on the Installation panel.

On the SQL Server setup wizard, provide the following information:

- On the Feature Selection panel, select the same features that you selected during the installation on the first system.
- On the Instance Configuration panel, make the following selections:
 - Specify the same instance name and instance ID that you used while installing this instance on the first system.
 - Ensure that the **Instance root directory** resides on the local system disk. Doing so installs the SQL Server instance on the local disk. The location of data directories is to be specified later on a different panel of the installation wizard.

- On the Service Accounts tab of the Server Configuration panel, specify a domain user name for the SQL Server services account.

Note: This user may not be a part of the local Administrators group on all the systems. In this case, you must configure the VCS agent's **SQLClusterAccount** attribute while configuring the SQL Server service group later.

- On the Data Directories tabs of the Database Engine Configuration panel and the Analysis Services Configuration panel, do the following:
 - Set the locations of database files and the analysis service files to the local disk.

You do not need to install these files to the shared storage that cluster disk group manages. The SQL Server Agent Configuration Wizard configures the SQL Server instance to use the files from the shared storage.
 - If you choose a shared storage location, ensure that you do not overwrite the data directories that were created when installing SQL Server on the first system.

Refer to the Microsoft documentation for application-specific instructions.

Assigning ports for multiple SQL Server instances

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

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

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

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

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

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

Enabling IPv6 support for the SQL Server Analysis Service

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

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

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

To enable IPv6 support for SQL Server Analysis Service

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

This means that IPv4 is disabled. Analysis Server does not listen on the IPv4 port, and clients will not be able to connect using IPv4.

- 6 Locate Network \ Listener \ IPV6Support property and in the Value field type **2**.

This means that IPv6 is optional. The Analysis Server tries to listen on the IPv6 port, but will silently ignore errors and continue to start if IPv6 is not available.

- 7 Click **OK** to save the changes.
- 8 Stop the Analysis Service.
- 9 Perform these steps for each named instance and on all the cluster nodes where SQL Server is installed.

Configuring SQL Server in a physical environment

- [Chapter 4. Overview](#)
- [Chapter 5. Configuring the VCS cluster](#)
- [Chapter 6. Configuring the SQL Server service group](#)
- [Chapter 7. Configuring an MSDTC service group](#)
- [Chapter 8. Configuring the standalone SQL Server](#)
- [Chapter 9. Configuring an Active/Active cluster](#)
- [Chapter 10. Configuring a disaster recovery setup](#)

Overview

This chapter includes the following topics:

- [About configuring SQL Server in physical environment](#)

About configuring SQL Server in physical environment

You can configure SQL Server in a physical environment following five simple steps:

- Manage storage
- Install Veritas InfoScale Availability
- Install SQL Server
- Configure VCS cluster
- Configure SQL Server service group

The following table provides the workflow for installing and configuring SQL Server in VCS environment, involving physical systems.

Table 4-1 Installing and configuring SQL Server in a VCS environment on physical systems or virtual machines involving shared or non-shared storage

Task	Details
Manage storage	<p>Depending on the type of storage used, perform the steps to manage it for VCS cluster configuration:</p> <ul style="list-style-type: none"> ■ Shared storage employing NetApp filers over an iSCSI or Fibre Channel (FC) connection and NetApp SnapMirror for replication See Table 4-1 on page 55. See “Managing storage using NetApp filer” on page 36. ■ Shared or non-shared storage managed using Windows Logical Disk Manager See “Managing storage using Windows Logical Disk Manager” on page 39.
Install Veritas InfoScale Availability	<p>Install Veritas InfoScale Availability on all the systems where you want to configure the application.</p> <p>See “About installing the VCS agents” on page 45.</p>
Install SQL Server	<p>See “About installing SQL Server for a high availability (HA) configuration” on page 46.</p>
Configure VCS Cluster	<p>Set up the components required to configure the cluster and the cluster service group.</p> <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 57.</p>
Configure SQL Server service group	<p>Configure the SQL Server service group to create resources for the storage and application agents.</p> <p>See “About configuring the SQL Server service group” on page 72.</p> <p>For configuring an MSDTC service group:</p> <p>See “About configuring the MSDTC service group” on page 96.</p> <p>For configuring an active/active cluster:</p> <p>See “About running SQL Server in an active-active clustered environment” on page 111.</p> <p>For configuring a disaster recovery set up:</p> <p>See “Setting up the disaster recovery cluster” on page 117.</p>

Table 4-1 Installing and configuring SQL Server in a VCS environment on physical systems or virtual machines involving shared or non-shared storage *(continued)*

Task	Details
Configure a standalone SQL Server server	<p>You can configure a standalone SQL Server server in the following two cases:</p> <ul style="list-style-type: none">■ You have a standalone SQL Server server and want to add it to the VCS cluster■ You have already installed the application on the systems where you want to configure the application for high availability <p>See “Configuring a standalone SQL Server for high availability” on page 108.</p>

Configuring the VCS cluster

This chapter includes the following topics:

- [Configuring the cluster using the Cluster Configuration Wizard](#)
- [Configuring notification](#)
- [Configuring Wide-Area Connector process for global clusters](#)

Configuring the cluster using the Cluster Configuration Wizard

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

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

Note the following prerequisites before you proceed:

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

To prevent lost heartbeats on the private networks, and to prevent VCS from mistakenly declaring a system down, Veritas recommends disabling the Ethernet auto-negotiation options on the private network adapters. Contact the NIC manufacturer for details on this process. Veritas recommends removing Internet Protocol TCP/IP from private NICs to lower system overhead.

- Verify that the public network adapters on each node use static IP addresses (DHCP is not supported) and name resolution is configured for each node.
- Veritas recommends that you use three network adapters (two NICs exclusively for the VCS private network and one for the public network) per system. You can implement the second private link as a low-priority link over a public interface. Route each private NIC through a separate hub or switch to avoid single points of failure. Veritas recommends that you disable TCP/IP from private NICs to lower system overhead.

Note: If you wish to use Windows NIC teaming, you must select the Static Teaming mode. Only the Static Teaming mode is currently supported.

- Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. GAB supports hub-based or switch network paths, or two-system clusters with direct network links.
- Verify the DNS settings for all systems on which the application is installed and ensure that the public adapter is the first adapter in the Connections list. When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- The logged on user must have local Administrator privileges on the system where you run the wizard. The user account must be a domain user account.
- The logged on user must have administrative access to all systems selected for cluster operations. Add the domain user account to the local Administrator group of each system.
- If you plan to create a new user account for the VCS Helper service, the logged on user must have Domain Administrator privileges or must belong to the Domain Account Operators group.
- When configuring a user account for the Veritas VCS Helper service, make sure that the user account is a domain user. The Veritas High Availability Engine (HAD), which runs in the context of the local system built-in account, uses the Veritas VCS Helper service user context to access the network. This account does not require Domain Administrator privileges.

- Make sure the VCS Helper service domain user account has "Add workstations to domain" privilege enabled in the Active Directory.
- Verify that each system can access the storage devices and each system recognizes the attached shared disk.
Use Windows Disk Management on each system to verify that the attached shared LUNs (virtual disks) are visible.
- If you plan to set up a disaster recovery (DR) environment, you must configure the wide-area connector process for global clusters.
- If you are setting up a Replicated Data Cluster configuration, add only the systems in the primary zone (zone 0) to the cluster, at this time.

To configure a VCS cluster using the wizard

- 1** Start the VCS Cluster Configuration Wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard** or, on Windows Server 2012 operating systems, from the **Apps** menu in the **Start** screen.
- 2** Read the information on the Welcome panel and click **Next**.
- 3** On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4** On the Domain Selection panel, select or type the name of the domain in which the cluster resides and select the discovery options.

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

- Clear **Specify systems and users manually**.
- Click **Next**.
Proceed to step [8](#).

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

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

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

Do not specify systems that are part of another cluster.

Proceed to step [8](#).

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

Do not select systems that are part of another cluster.

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

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

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

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

- System is not pingable.
- WMI access is disabled on the system.
- Wizard is unable to retrieve the system architecture or operating system.
- Product is either not installed or there is a version mismatch.

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

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

Veritas Cluster Server Configuration Wizard

Cluster Details
Enter necessary details to create the new cluster

Domain Selection

Cluster Details

Cluster Selection

Validate Systems

Edit Options

NIC Selection

Service Account

Security

Summary

Finish

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

Cluster Name:

Cluster ID:

Operating System:

Select the systems to create the cluster.

☒ **Select all systems**

Available Systems

- ☒ ROGER
- ☒ SCOOPYDU

Total number of systems selected to create the cluster : 2

Click 'Next' to continue.

VERITAS

Specify the cluster details as follows:

- | | |
|------------------|---|
| Cluster Name | Type a name for the new cluster. Veritas recommends a maximum length of 32 characters for the cluster name. |
| Cluster ID | <p>Select a cluster ID from the suggested cluster IDs in the drop-down list, or type a unique ID for the cluster. The cluster ID can be any number from 0 to 65535.</p> <p>Note: 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.</p> |
| Operating System | <p>From the drop-down list, select the operating system.</p> <p>All the systems in the cluster must have the same operating system and architecture.</p> |

Available Systems Select the systems that you wish to configure in the cluster.

Check the **Select all systems** check box to select all the systems simultaneously.

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

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


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

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

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

Do one of the following:

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



- Select **Configure LLT over Ethernet**.
- Select the check boxes next to the two NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one of the NICs and use the low-priority NIC for both public and as well as private communication.
- If there are only two NICs on a selected system, Veritas recommends that you lower the priority of at least one NIC that will be used for private as well as public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
- If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. Veritas recommends that you do not select teamed NICs for the private network.

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

- To configure the VCS private network over the User Datagram Protocol (UDP) layer, complete the following steps:
- Select **Configure LLT over UDP on IPv4 network** or **Configure LLT over UDP on IPv6 network** depending on the IP protocol that you wish to use. The IPv6 option is disabled if the network does not support IPv6.
- Select the check boxes next to the NICs to be assigned to the private network. You can assign a maximum of eight network links. Veritas recommends reserving two NICs exclusively for the VCS private network.
- For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. In case of IPv4, each IP address can be in a different subnet.

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

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

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

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

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

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

Specify the domain user details as follows:

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

- To specify a new user, click **New user** and type a valid user name in the Create New User field and then click **Next**.
Do not append the domain name to the user name; do not type the user name as Domain\user or user@domain.
- In the Password dialog box, type the password for the specified user and click **OK**, and then click **Next**.

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

Do one of the following:

- To use VCS cluster user privileges, click **Use VCS User Privileges** and then type a user name and password.
The wizard configures this user as a VCS Cluster Administrator. In this mode, communication between cluster nodes and clients, including Cluster Manager (Java Console), occurs using the encrypted VCS cluster administrator credentials. The wizard uses the VCSEncrypt utility to encrypt the user password.
The default user name for the VCS administrator is *admin* and the password is *password*. Both are case-sensitive. You can accept the default user name and password for the VCS administrator account or type a new name and password.
Veritas recommends that you specify a new user name and password.
- To use the single sign-on feature, click **Use Single Sign-on**.
In this mode, the VCS Authentication Service is used to secure communication between cluster nodes and clients by using digital certificates for authentication and SSL to encrypt communication over the public network. VCS uses SSL encryption and platform-based authentication. The Veritas High Availability Engine (HAD) and Veritas Command Server run in secure mode.
The wizard configures all the cluster nodes as root brokers (RB) and authentication brokers (AB). Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates signed by the root. These brokers can authenticate clients such as users and services. The wizard creates a copy of the certificates on all the cluster nodes.

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

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

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

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

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

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

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

Do the following:

- Check the **Notifier Option** check box to configure notification of important events to designated recipients.
See [“Configuring notification”](#) on page 66.
- Check the **GCO Option** check box to configure the wide-area connector (WAC) process for global clusters. The WAC process is required for inter-cluster communication.
Configure the GCO Option using this wizard only if you are configuring a Disaster Recovery (DR) environment and are not using the Disaster Recovery wizard.
You can configure the GCO Option using the DR wizard. The Disaster Recovery chapters in the application solutions guides discuss how to use the Disaster Recovery wizard to configure the GCO option.
See [“Configuring Wide-Area Connector process for global clusters”](#) on page 69.

Configuring notification

This section describes steps to configure notification.

To configure notification

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

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

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

The screenshot shows the 'Notifier SNMP Configuration' panel of the Veritas Cluster Server Configuration Wizard. The panel has a blue header with the title and a close button. Below the header is a sub-header 'Notifier SNMP Configuration' with the instruction 'Specify information about SNMP console.' and a printer icon. On the left is a sidebar with navigation links: 'Domain Selection', 'Create Cluster', 'Select Components', 'Configure', 'Summary', and 'Finish'. The main area contains instructions: 'Enter the name or the IP address of the SNMP console and then select the desired severity level.' Below this is a table with two columns: 'SNMP Console' and 'Severity Information'. The 'SNMP Console' column has a text input field with the placeholder 'Click here to change the text..'. The 'Severity Information' column has a dropdown menu. Below the table are instructions: 'Click on '+' button to add more consoles.' and 'Click '-' to remove a console.' with corresponding '+' and '-' buttons. Below these is a text input field for 'SNMP Trap Port' with the value '162'. A note states: 'Note: SNMP console must be MIB 2.0 compliant.' and instructions: 'Click 'Next' to continue.' At the bottom are 'Back', 'Next', and 'Cancel' buttons. The Veritas logo is in the bottom left corner.

SNMP Console	Severity Information
Click here to change the text..	

Do the following:

- Click a field in the **SNMP Console** column and type the name or IP address of the console.
The specified SNMP console must be MIB 2.0 compliant.
- Click the corresponding field in the **Severity** column and select a severity level for the console.
- Click the + icon to add a field; click the - icon to remove a field.

- Enter an SNMP trap port. The default value is 162.
- 3** If you chose to configure SMTP, specify information about SMTP recipients and then click **Next**.

The screenshot shows the 'Notifier SMTP Configuration' window of the Veritas Cluster Server Configuration Wizard. The window has a blue title bar and a light blue sidebar on the left with the following menu items: Domain Selection, Create Cluster, Select Components, Configure, Summary, and Finish. The main area is white and contains the following elements:

- Notifier SMTP Configuration**
Specify information about SMTP recipients.
- SMTP Server Name / IP**: A text input field.
- Enter SMTP recipients and select a severity level for each recipient.**
- Recipients Table**: A table with two columns: 'Recipients' and 'Severity'. The first row contains the text 'Click here to change the text..' and 'Information' respectively. Below this are three empty rows.
- Buttons**: A '+' button to add a recipient and a '-' button to remove a recipient.
- Instructions**: Text indicating to click '+' to add a recipient, '-' to remove a recipient, and 'Next' to continue.
- VERITAS logo**: Located in the bottom left of the main area.
- Navigation Buttons**: 'Back', 'Next', and 'Cancel' buttons at the bottom right.

Do the following:

- Type the name of the SMTP server.
 - Click a field in the **Recipients** column and enter a recipient for notification. Enter recipients as admin@example.com.
 - Click the corresponding field in the **Severity** column and select a severity level for the recipient.
VCS sends messages of an equal or higher severity to the recipient.
 - Click the + icon to add fields; click the - icon to remove a field.
- 4** On the Notifier Network Card Selection panel, specify the network information and then click **Next**.

Do the following:

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


Configuring Wide-Area Connector process for global clusters

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

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

To configure the wide-area connector process for global clusters

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



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

Do the following:

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

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

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

Configuring the SQL Server service group

This chapter includes the following topics:

- [About configuring the SQL Server service group](#)
- [Before configuring the SQL Server service group](#)
- [Configuring a SQL Server service group using the wizard](#)
- [Configuring the service group in a non-shared storage environment](#)
- [Running SnapManager for SQL Server](#)
- [About the modifications required for tagged VLAN or teamed network](#)
- [Making SQL Server user-defined databases highly available](#)
- [Verifying the service group configuration](#)
- [Administering a SQL Server service group](#)

About configuring the SQL Server service group

A SQL Server service group is used to bring a SQL Server instance online on another node if the active node fails. If you have set up multiple cluster nodes, specify the priority of the failover node when you configure the service group. Configuring the SQL Server service group involves creating VCS resources for the agents SQL Server, storage agents (NetApp, LDM) and network 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. The following topics provide instructions on configuring a SQL Server service group using the wizard.

The SQL Server Configuration Wizard lets you create a SQL Server service group and define the attributes for its resources on all the nodes within the cluster simultaneously. You can configure a service group for only one SQL Server version in a single wizard workflow. To configure another SQL Server version, you must run the wizard again.

For a shared storage configuration, ensure that you use the wizard to configure the service group.

Before configuring the SQL Server service group

Perform the following activities before you configure the SQL Server service group:

- Verify that you have configured a VCS cluster using VCS Cluster Configuration Wizard (VCW).
- Verify that SQL Server is identically installed on all the cluster nodes that will participate in the service group.
- Verify that you have VCS Administrator privileges. This user classification is required to create and configure a service group.
- The logged-on user account must be a local Administrator on the node where you run the wizard. If you wish to configure detail monitoring for a SQL instance, the logged-on user must have the permission to log on to that SQL instance.
- You must be an Administrator for the NetApp filer containing the LUNs created to store SQL Server components.
- In case of a shared storage configuration, run the SQL Server Configuration Wizard from the first cluster node where you installed SQL Server. Do not run the wizard from the additional nodes.
This is required as the wizard configures the resources for the SQL Server database and registry information installed on the shared storage and propagates this information to the remaining nodes that are part of the SQL service group.
- Verify that the Veritas High Availability Engine (HAD) is running on the system from where you run the wizard.
- Verify that the volumes or LUNs (virtual disks) created to store the following data components are mounted or connected to the node where you run the wizard and dismounted or disconnected from other nodes in the cluster:
 - SQL Server system data files
 - Registry replication information

- User database files
- User database log files
- FILESTREAM database objects

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

- If you wish to configure high availability for FILESTREAM, ensure that FILESTREAM is configured and enabled for the SQL instance on the first cluster node where you installed SQL, and disabled on all the remaining nodes. Refer to the Microsoft SQL Server documentation for more information.
- In case of IPv4, assign a unique virtual IPv4 address to the SQL Server instance. You specify this IP address when configuring the service group.
 In case of IPv6, the configuration wizard automatically generates an IPv6 address based on the network selected. The IPv6 address is valid and unique on the network.
- In an IPv6 environment, the Lanman agent relies on the DNS records to validate the virtual server name on the network. If the virtual servers configured in the cluster use IPv6 addresses, you must specify the DNS server IP, either in the network adapter settings or in the Lanman agent's AdditionalDNSServers attribute.
- If you wish to use a script for detail monitoring, either save the script file in shared storage or ensure that the same file exists in the same location on all the cluster nodes.

A sample script is supplied in C:\Program Files\Veritas\cluster server\bin\SQLServer\sample_script.sql. The same script can be used to monitor SQL Server 2012 and 2014.

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

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

- Make sure that the following services are stopped on the first cluster node where you are running the wizard:
 - SQL Server
 - SQL Server Agent
 - SQL Server Analysis Services

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

- Review the resource types and the attribute definitions of the agents.
- If you have configured Windows Firewall, add the required ports and services to the Firewall Exception list.

For a detailed list of services and ports, refer to the *Veritas InfoScale Installation and Upgrade Guide*.

Configuring a SQL Server service group using the wizard

This section describes how to configure a SQL service group using the configuration wizard.

The SQL Server Configuration Wizard is used to configure a service group for SQL Server 2012 and 2014. You can configure a service group for only one SQL Server version at a time. To configure a service group for another SQL Server version, you must run the wizard again.

For a shared storage configuration, use the service group configuration wizard to configure the service group.

To create a SQL Server service group on the cluster

- 1 Ensure that you have stopped the SQL Server service for the instance and are running the wizard from the first cluster node.

Note: If the SQL Server service is running when you launch the wizard in the create mode, the wizard fails to reconfigure the service to start under Lanman context.

- 2 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 3 Review the prerequisites on the Welcome panel and then click **Next**.
- 4 On the Options panel, select **Create service group** and then click **Next**.
- 5 On the Service Group Configuration panel, specify the service group name and system list.

Complete the following:

- In the Service Group Name field, specify a name for the SQL Server service group, for example, INST1_SG. If there are multiple instances, ensure that the name is unique within the cluster.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order list.
 - To change the priority of a system in the Systems in Priority Order list, select the system and click the up and down arrow icons. Arrange the systems in priority order as failover targets for the group. The server that needs to come online first must be at the top of the list.
 For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - To enable the service group to automatically come online on one of the systems, select the **Include selected systems in the service group's AutoStartList attribute** checkbox.
 For information about the AutoStartList attribute, see the *Cluster Server Administrator's Guide*.
 - Click **Next**.
- 6** On the SQL Server Instance Selection panel, select the SQL Server instance and any other services that were installed and that needs to be configured for high availability in your environment. Complete the following steps and then click **Next**.
- From the **SQL Server version** drop-down list, select the SQL Server version for which you wish to configure the service group.
 You can configure a service group for only one SQL Server version in a single wizard workflow. To configure another SQL Server version, you must run the wizard again.
 The wizard displays instances of the selected SQL Server version that satisfy the following criteria:
 - Instances installed identically on all the systems
 - Instances not configured in other SQL service groups
 - Select the SQL Server instances that you wish to configure in the service group.
 - If required, select the other services that you wish to make highly available. These options are available for selection only if the corresponding services are installed.

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

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

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

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

In case of multiple instances, select the required instance from the SQL Instance dropdown list. The panel displays the databases and the respective files for which the wizard configures resources. Click a database name to view its database files.

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

- 9 On the SQL Server Cluster Account Configuration panel, specify the SQL cluster account details and then click **Next**.

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

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

- Select a SQL instance from the **Instance Name box**.
- Check the **Configure SQL Server Cluster Account** check box.
- Click **Use service SIDs to set the SQL Server service name as the SQL cluster account**.
- Click **Use Domain Group Account** and then click the adjacent ellipsis button to launch the Windows Select Users, Computers, or Groups dialog box.
 Then specify a domain group and click **OK** to set the domain group as the SQL cluster account.
 If you specify a domain group as the SQL cluster account, ensure that the SQL Server service and SQL Server Agent service accounts are part of the specified domain group.

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

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

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

- Check the check box for a SQL instance, and then click the button from the Detail Monitoring Properties column to specify the detail monitoring settings. See [“Configuring detail monitoring for a SQL Server instance”](#) on page 81.
- Repeat these steps for each SQL Server instance that you wish to configure detail monitoring for.

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

Click **Next**.

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

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

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

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

Note: If you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that have the same MAC address, then you must edit the “MACAddress” attribute of the NIC agent and the IP agent, after you configure the application service group.

See [“About the modifications required for tagged VLAN or teamed network”](#) on page 88.

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

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

- 14** On the Service Group Summary panel, review the service group configuration and then click **Next**. The Resources box lists the configured resources. The wizard assigns unique names to resources based on their respective name rules. Click a resource to view its attributes and their configured values in the Attributes box. Optionally, if desired, change the names of the resources.

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

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

- 15** Click **Yes** when prompted to confirm creating the service group. Messages indicate the status of the commands.
- 16** Select the **Bring the service group online** check box, if you want to bring the service group online.

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

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

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

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

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

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

You can also configure an MSDTC service group.

Configuring detail monitoring for a SQL Server instance

You can configure detail monitoring for a SQL Server instance while configuring its service group. This dialog box appears when you click a button in the Detail Monitoring Properties column for a SQL Server instance.

See [“Configuring a SQL Server service group using the wizard”](#) on page 75.

To configure detail monitoring for a SQL Server instance

- 1** On the Detail Monitoring configuration dialog box, specify the monitoring frequency in the **Detail monitoring frequency** field.

Detail Monitoring configuration for COYOTESQL_INST1

Detail monitoring frequency:

Default value 5 indicates that detail monitoring is performed on every 5th monitor cycle.

☒ **DBList detail monitoring** ☒ Select all

Databases for COYOTESQL_INST1

- ☐ ReportServer\$COYOTESQL_INST1
- ☐ ReportServer\$COYOTESQL_INST1TempDB
- ☒ master
- ☒ model
- ☐ msdb
- ☐ tempdb

Domain\Username:

Password:

☐ SQL-script based detail monitoring

☒ Global ☐ Per system

SQL monitor script

☒ Fail over service group if detail monitoring fails

Apply Cancel

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

- Select **DBList Detail Monitoring** and then choose the databases from the list of databases available for the instance. The selected databases populate the DBList attribute of the SQL agent. In this mode of detail monitoring the agent monitors the health of the databases by connecting to those databases. The agent monitors only the databases specified in the DBList attribute.

- Select **SQL-Script based detail monitoring** if you wish to use a script to monitor SQL server databases. In this mode of detail monitoring, the agent executes the script that you specify for detail monitoring.
- 2 Specify the fully qualified user name and the password for connecting to the SQL Server database. Make sure that the user has SQL Server logon permissions.

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

- 3 Select **Global** or **Per System** depending on whether the monitoring script location is the same for all the nodes or is unique for each cluster node, and then specify the path of the script appropriately.
- 4 If the detail monitoring fails, the SQL Server agent can fail over the service group to another node. If you want the agent to initiate the failover in such a scenario, make sure that the **Fail over service group if detail monitoring fails** check box is checked.
- 5 Click **Apply**.

Assigning privileges to the existing SQL Server databases and logs

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

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

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

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

To assign privileges to the existing SQL databases and logs

- 1 On the node where the SQL Server service group is online, navigate to the following directory from Windows explorer:
`<Datarootdirectory>\<SQLinstancename>\MSSQL\`
 The directory contains various directories including DATA, FTData, JOBS, Log, repldata.
 Here, `<Datarootdirectory>` is the path that you specified while installing SQL Server.
- 2 Assign the SQL Cluster account with full control privileges to the following directories:
 - DATA
 - Log
- 3 Navigate inside the DATA folder and then assign the SQL cluster account with full control privileges to the following files in that directory:
 - tempdb.mdf
 - templog.ldf
- 4 Repeat these steps for all the instances that are configured in the SQL Server service group.
 This ensures the existing SQL databases are accessible after a service group failover.

Configuring the service group in a non-shared storage environment

If you are using a non-shared storage configuration, you have to use the VCS MountV – VMNSDg agents to monitor your local storage. Currently, the service group configuration wizards do not support configuring these agents in the service group. You have to configure these agents manually by using the Cluster Manager (Java Console) or the VCS commands.

VCS provides templates for configuring service groups that use non-shared storage agent resources.

The Java Console templates are located in the following directory:

`%VCS_HOME%\Templates`

Here, %VCS_HOME% is the default product installation directory, typically,
 C:\Program Files\Veritas\Cluster Server.

For information about adding a service group using templates from the Java Console, refer to the *Cluster Server Administrator's Guide*.

The following steps describe how to create a service group using the Cluster Manager (Java Console).

To configure the service group in a non-shared storage environment

- 1** Open the Cluster Manager (Java Console) from **Start > All Programs > Veritas > Veritas Cluster Server** and then click **Veritas Cluster Manager - Java Console** or, on Windows Server 2012 operating systems, from the **Apps** menu.
- 2** Log on to the cluster. On the Cluster Monitor window click **File > New Cluster**, then on the New Cluster window type **localhost** in the Host name field, and then click **OK**.
- 3** Launch the service group configuration wizard. From the Cluster Explorer window menu, click **Tools > Configuration Wizard**.
- 4** On the Service Group Configuration Wizard Welcome panel, click **Next**.
- 5** Fill in the following information and then click **Next**:
 - Specify a name for the service group.
 - Select the systems for the service group. Click a system in the Available Systems box and then click the right arrow to move the systems to Systems for Service Group.
 - Leave the service group type as the default, Failover.
- 6** Click **Next** again.

- 7
- In the Templates list, select the desired service group template depending on the configuration and then click **Next**.

Template name	Description
SQLServer-VMNSGroup (SQL Server Agent service)	Use these templates to create a single-node high availability service group that uses non-shared storage.
SQLServer-OlapVMNSGroup (SQL Server Agent service, and SQL Server Analysis service)	These templates include resources for configuring MountV and VMNSDg agents.
SQLServer-VMNSFilestreamGroup (SQL Server Agent service, and SQL Server FILESTREAM)	
MSDTCVMNSGroup (SQL Server MSDTC service)	
SQLServer-VirtVMNSGroup (SQL Server Agent service)	Use these templates to create a single-node high availability service group in a VMware virtual environment.
SQLServer-OlapVirtVMNSGroup (SQL Server Agent service, and SQL Server Analysis service)	These templates include resources for configuring MountV, VMwareDisks, and VMNSDg agents.
SQLServer-VirtVMNSFilestreamGroup (SQL Server Agent service, and SQL Server FILESTREAM)	
MSDTCVirtVMNSGroup (SQL Server MSDTC service)	
VvrRvgVMNSRVGGroup	Use this template to create a VVR replication service group on a single node that uses non-shared storage.

The Templates box lists the templates available on the system to which Cluster Manager is connected. The resource dependency graph of the templates, the number of resources, and the resource types are also displayed.

- 8
- Click **Next**. The wizard starts creating the service group.
- 9
- After the service group is successfully created, click **Next** to edit attributes using the wizard.

- 10** The wizard lists the resources and their attributes. You must specify values for the mandatory attributes that appear in bold. The remaining resources listed in the window are preconfigured by the template and do not require editing.

To modify an attribute, do the following:

- Click the resource.
- Click the attribute to be modified.
- Click the **Edit** icon at the end of the table row.
- In the Edit Attribute dialog box, enter the attribute values.
- Click **OK**.

For details on application-specific agent attributes, refer to the application-specific agent or solutions guide.

For details on the storage and network agent attributes, refer to the *Cluster Server Bundled Agents Reference Guide*.

- 11** Click **Finish**.
- 12** Right-click the newly created service group and select **Enable Resources**.
- 13** Right-click the newly created service group, select **Online** from the context menu, and then select a system on which to bring the service group online.

If you are configuring the service group on a node at the secondary site in a DR environment, bring the service group online only after completing all the DR configuration steps.

Running SnapManager for SQL Server

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

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

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

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

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

About the modifications required for tagged VLAN or teamed network

Perform this task only if you have a tagged VLAN network configuration having multiple logical network interfaces or a teamed network interface that share the same MAC address.

After you configure the application service group, you must edit the “MACAddress” attribute of the VCS NIC agent and the IP agent.

During the application service group configuration, you are required to select a network adapter for each cluster system and specify the virtual IP address for the virtual server. The application configuration wizard internally retrieves the MAC address of the specified interface and the MAC address of the interface to which the specified IP address is assigned. It then sets these MAC Addresses as the value of the “MACAddress” attribute of the VCS NIC and IP agent respectively.

If the selected interface or the interface to which the specified IP is assigned shares the MAC address with other logical interfaces, then the following issues may occur:

- NIC agent may begin to monitor an interface other than the one selected.
- The IP agent may assign the specified virtual IP address or the virtual server name to an interface other than the one selected. As a result, the IP agent may monitor an IP address other than the one specified.

As a workaround, use the VCS Java Console to edit the “MACAddress” attribute and specify its value as the interface name instead of the MAC address. You must enter the interface name in double quotes. For example, MACAddress = “InterfaceName”

Notes:

- After you specify the interface name as the “MACAddress” attribute value, if you want to use the VCS wizards to modify any settings, then you must first reset the value of the “MACAddress” attribute to the MAC address of the interface. Failing this, the VCS wizard may fail to identify and populate the selected interface. Use the VCS Java Console to edit the attribute values.
- If you change the interface name, you must update the “MACAddress” attribute value to specify the new name. Failing this, the NIC resource will go in an UNKNOWN state.
- While editing the “MACAddress” attribute to specify the interface name, you must specify the name of only one interface.

Making SQL Server user-defined databases highly available

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

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

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

Create volumes or LUNs for SQL Server user-defined databases

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

In the sample deployment these volumes are named as follows:

- INST1_DB1_VOL
Contains a user-defined database file
- INST1_DB1_LOG
Contains a user-defined database log file
- INST1_DB1_FS_VOL
Contains FILESTREAM-enabled database objects for the user database

Creating SQL Server databases

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

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

Adding storage agent resources to the SQL service group

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

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

For a shared storage configuration, use the service group configuration wizard to modify the service group. For a non-shared storage configuration, use the Cluster Manager (Java Console) to add the required storage resources to the service group manually.

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

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

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

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

To add storage agent resources to the SQL service group

- 1 Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** to start the configuration wizard.
- 2 Review the Prerequisites page and click **Next**.
- 3 On the Wizard Options panel, click **Edit service group**, select the service group and then click **Next**.
- 4 Click **Yes** on the VCS Notice informing you that the service is not completely offline. No adverse consequences are implied.
- 5 In the Service Group Configuration page, click **Next**.
- 6 In the SQL Server Instance Selection page, make sure the correct instance of SQL Server is selected and click **Next**.

- 7 In the User Databases List page, make sure the databases are shown with correct paths and click **Next**.
- 8 If a database is not configured correctly, a VCS warning appears indicating potential problems. Click **OK** to continue.
- 9 In the Detail Monitoring and succeeding pages, review the information and click **Next** to continue.
- 10 Click **Yes** to continue when a VCS Notice indicates the configuration will be modified.
- 11 Click **Finish** to exit the wizard.

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

Verifying the service group configuration

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

Bringing the service group online

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

To bring a service group online from the Java Console

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

Taking the service group offline

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

To take a service group offline from the Java Console

- 1 On the **Service Groups** tab of the Cluster Explorer configuration tree, right-click the service group.

or

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

Switching the service group

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

- 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. Then do the following:
 - 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 the earlier step.
- 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, do the following:
 - Restart the node you shut down in step [1](#).
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**.
The service group you selected is taken offline and brought online on the node that you selected.

Administering a SQL Server service group

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

Modifying a SQL service group configuration

Note the following prerequisites before modifying the SQL service group:

- If the SQL Server service group is online, you must run the wizard from a node on which the service group is online. You can then use the wizard to add resources to and remove them from the configuration. You cannot change resource attributes.
- To change the resource attributes, you must take the service group offline. However, the NetAppFile 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.

- After the application service group configuration, if you have manually edited any of the resource attributes, then you must reset them to their default values. Failing this, the wizard may fail to identify and populate the resources involved in the service group configuration. After you modify the service group configuration you can again edit the resource attributes to set the desired value.

To modify a SQL Server service group

- 1 Start the SQL Server Configuration Wizard. Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 2 Review the prerequisites and click **Next**.
- 3 In the Service Group Selection panel, select the service group to modify and click **Next**.
- 4 In the Service Group Configuration panel, perform the following actions as necessary:
 - Add or remove systems from the service group's system list.
 - Change the priority order of the system list.
 - To add (or remove) the selected systems from the service group's AutoStartList attribute, select (or deselect) the **Include selected systems in the service group's AutoStartList attribute** checkbox.
 For information about the AutoStartList attribute, see the *Cluster Server Administrator's Guide*.

click **Next**.
- 5 In the SQL Server Instance Selection panel, select the SQL Server instance to be made highly available and click **Next**.
- 6 In the User Databases List panel, verify the master and user defined databases configured for the SQL instance. The wizard will create NetAppSnapDrive resource for each database. Click **Next**.
- 7 Follow the wizard instructions and make desired modifications to the service group configuration.

Deleting a SQL service group

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

To delete a SQL Server service group

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

Configuring an MSDTC service group

This chapter includes the following topics:

- [About configuring the MSDTC service group](#)
- [Typical MSDTC service group configuration using shared storage](#)
- [Before configuring the MSDTC service group](#)
- [Creating an MSDTC service group](#)
- [About configuring an MSDTC client](#)
- [Configuring an MSDTC client](#)
- [Verifying the installation](#)

About configuring the MSDTC service group

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

Configuring the MSDTC service group involves the following tasks:

- Creating an MSDTC Server service group using the MSDTC Configuration Wizard

- Configuring the MSDTC client manually

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

Typical MSDTC service group configuration using shared storage

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

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

The following configurations are possible:

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

Figure 7-1 MSDTC Server and SQL Server configured on different nodes




Figure 7-2 MSDTC Server configured on the same node as SQL Server

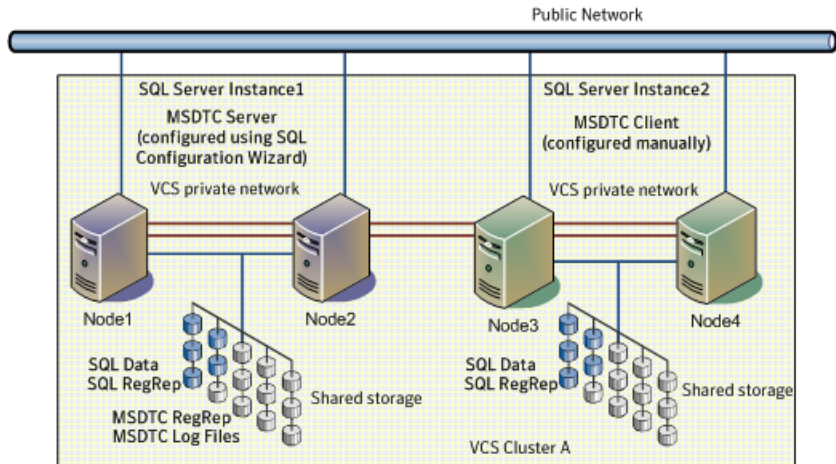
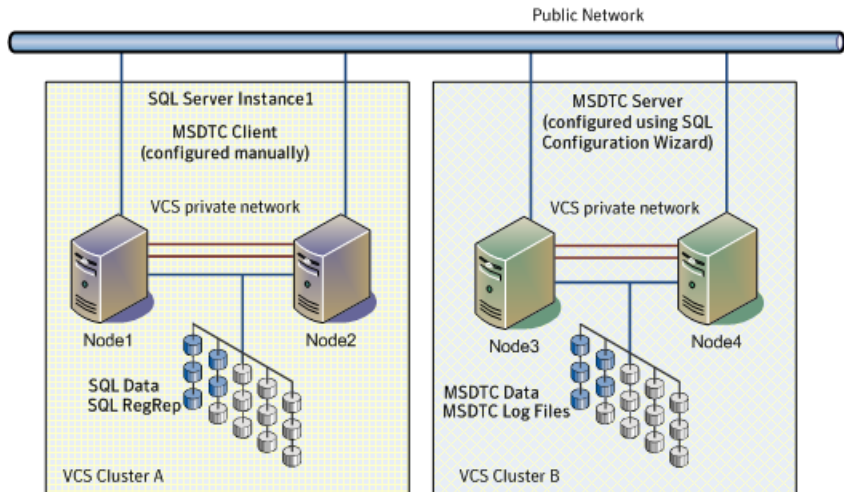


Figure 7-3 MSDTC Server and SQL Server configured on nodes in different clusters



Before configuring the MSDTC service group

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

- You must be a Cluster Administrator. This user classification is required to create and configure a service group.
- You must be a local Administrator on the node where you run the wizard.
- Verify that the VCS agent for SQL Server is installed on all cluster nodes.
- Verify that the VCS cluster is configured using the VCS Cluster Configuration Wizard (VCW).
- Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC Server service group.
- Verify that the Distributed Transaction Coordinator (MSDTC) service is stopped.
- Verify that you have created the volumes or LUNs (virtual disks) for storing MSDTC log and MSDTC registry replication information.
See [“Managing storage using NetApp filer”](#) on page 36.
See [“Managing storage using Windows Logical Disk Manager”](#) on page 39.
- Verify that the volumes or LUNs created for the MSDTC logs and registry replication information are mounted or connected to the node where you run the wizard. In case of a shared storage configuration, ensure that they are dismounted or disconnected from all other nodes.

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

Creating an MSDTC service group

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

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

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

To create an MSDTC service group

- 1 Start the MSDTC Configuration Wizard. Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > MSDTC Configuration Wizard** or on Windows Server 2012 operating systems, from the **Apps** menu in the Start screen.
- 2 In the Select Configuration Option panel, click **MSDTC Server - Service Group Configuration**, click **Create**, and then click **Next**.
- 3 Review and verify that you have met the prerequisites for configuring an MSDTC Server service group and then click **Next**.

- 4 On the Service Group Configuration panel, specify the service group name and select the systems for the service group as follows:
 - Type a name for MSDTC service group.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the Systems in Priority Order box. The systems listed in the Systems in Priority Order box represent the SystemList attribute of the service group. While selecting systems, make sure to select the systems that are not in the SystemList attribute for an Exchange service group configured in the cluster.

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

To change a system's priority in the service group's system list, select the system from the Systems in Priority Order and click the up and down arrows. The system at the top of the list has the highest priority while the system at the bottom of the list has the lowest priority.
 - To enable the service group to automatically come online on one of the systems, select the **Include selected systems in the service group's AutoStartList attribute** checkbox.

For information about the AutoStartList attribute, see the *Cluster Server Administrator's Guide*.
 - Click **Next**. If the cluster configuration is in read-only mode, the wizard prompts you before changing it to read-write mode. The wizard starts validating your configuration. Various messages indicate the validation status.
- 5 On the Virtual Server Configuration panel, specify the information related to the virtual server as follows:
 - Type a virtual name for the MSDTC Server. This is the name that is used by MSDTC clients to connect to the MSDTC Server. The DTC service runs under this virtual name. Ensure that the virtual server name is unique in the cluster.
 - Select **IPv4** to configure an IPv4 address for the virtual server.
 - In the Virtual IP Address field, type a unique virtual IPv4 address for the MSDTC server.
 - In the Subnet Mask field, type the subnet to which the virtual IPv4 address belongs.
 - Select **IPv6** to configure an IPv6 address for the virtual server. The IPv6 option is disabled if the network does not support IPv6.

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

The user account for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.
 - Click **Next**.
- 6** On the Specify Data Path panel, specify the drive letter for the MSDTC log and registry replication directory and click **Next**. If the directory does not exist, the wizard creates it. Veritas recommends using different paths for these directories.
- Clear the **Configure NetApp SnapMirror Resource(s)** check box. This option is applicable only in case of a disaster recovery configuration. The SnapMirror resource is used to monitor replication between filers at the primary and the secondary site, in a disaster recovery configuration. If you are setting up a disaster recovery environment, check this check box to configure the SnapMirror resource at the primary site. The SnapMirror resource must be configured only after you have configured the cluster at the secondary site.

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

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

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

- The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values.

- The wizard assigns unique names to resources. Change names of the resources, if desired.

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

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

Various messages indicate the status of these commands.

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

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

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

About configuring an MSDTC client

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

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

Ensure the following before you configure the MSDTC client:

- Verify that the MSDTC Server service group is online in the cluster.
- Verify that the systems on which you plan to install the MSDTC client, have a SQL instance configured on them.

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

Configuring an MSDTC client

Complete the following steps to configure the MSDTC client.

To configure an MSDTC client

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

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.

You must complete the procedure to verify the service group configuration.

Configuring the standalone SQL Server

This chapter includes the following topics:


- [Typical high availability configuration for a standalone SQL Server setup](#)
- [Configuring a standalone SQL Server for high availability](#)

Typical high availability configuration for a standalone SQL Server setup

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

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

Figure 8-1 Active-Passive configuration



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

Sample configuration

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

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

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

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

Table 8-1 Objects used for clustering standalone server

Name	Object
SYSTEM1 & SYSTEM2	server names; SYSTEM1 is the existing standalone SQL server
INST1_SG	Microsoft SQL Server service group

Table 8-1 Objects used for clustering standalone server (continued)

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

Configuring a standalone SQL Server for high availability

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

Note: In addition to the tasks mentioned below, the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach are required.

Refer to: <http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>.

The following table lists the tasks to configure the standalone SQL Server for high availability.

Table 8-2 Tasks to configure the standalone SQL Server for high availability

Task	Description
Prepare the standalone SQL Server	<p>Complete the following tasks before you begin the InfoScale installation and configuring a VCS cluster:</p> <ul style="list-style-type: none"> ■ Back up the existing SQL data ■ From the SQL Server Service Manager, set the SQL Server services to manual start. <p>While you set the SQL Server services to manual start, you must select the standalone server that you plan to incorporate, select the service from the Services list, and then clear the Auto-start service when OS starts check box.</p> <p>Note: Repeat these steps for all other SQL Server services that are running on the server.</p>
Install the InfoScale product and configure VCS cluster on the standalone SQL server	<p>Install Veritas InfoScale Availability on all the systems where you want to configure the application. During installation, the product installer installs the VCS agents required for making the applications highly available.</p> <p>For details refer to the <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p> <p>For details on configuring the VCS cluster, refer to, See “Configuring the cluster using the Cluster Configuration Wizard” on page 57.</p>
Create volumes or LUNs (virtual disks) necessary to manage the SQL Server storage.	<p>See “Managing storage using NetApp filer” on page 36.</p> <p>See “Managing storage using Windows Logical Disk Manager” on page 39.</p>
Install and configure SQL Server on additional nodes, if required	<p>Perform the following tasks to install Microsoft SQL Server on additional nodes.</p> <ul style="list-style-type: none"> ■ Ensure that the shared volumes or LUNs are imported on the node. ■ Ensure that the SQL Server configuration is identical on all nodes in the cluster. To have identical configuration, ensure that the instance name (if applicable), destination folder for Program Files and Data Files and the Authentication Mode are same on all the nodes. ■ Perform the SQL Server installation <p>See “About installing SQL Server on additional systems” on page 50.</p>

Table 8-2 Tasks to configure the standalone SQL Server for high availability
(continued)

Task	Description
Verify that the existing SQL Server databases and logs are moved to shared storage	Verify the location of all SQL Server databases and logs for the existing standalone server. If they are located on local storage, move them from the local drive to the appropriate volumes or LUNs on shared storage to ensure proper failover operations in the cluster.
Configure the SQL Server service group	See “About configuring the SQL Server service group” on page 72.
Create and manage SQL Server user-defined database	See “Making SQL Server user-defined databases highly available” on page 89.

Moving the existing SQL Server data files and user databases

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

- 1 From the SQL Server Service Manager, stop the SQL Server service.
- 2 Verify that you have backed up your existing data.
- 3 Import the volumes or LUNs to the node where the original files are located on the local drives and mount the volumes (add drive letters).

See [“Connecting virtual disks to the cluster node”](#) on page 38.
- 4 Move the SQL Server data files and user databases to the shared volumes or LUNs. Follow the procedures described in Microsoft Knowledge Base Article - 224071: INF: Moving SQL Server databases to a New Location with Detach/Attach.

Refer to: <http://support.microsoft.com/default.aspx?scid=kb;en-us;224071>.
- 5 Restart SQL Server.

Configuring an Active/Active cluster

This chapter includes the following topics:

- [About running SQL Server in an active-active clustered environment](#)
- [Setting up the Active/Active cluster](#)

About running SQL Server in an active-active clustered environment


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.

A SQL Server instance can fail over to any of the other nodes configured nodes on its system list.

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

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 9-1 Active-active configuration



For example, consider a two-node cluster hosting two SQL Server Virtual Servers, BILLING_VS and PAYROLL_VS.

Table 9-1 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 node)
- Cluster IP address

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

The IP addresses are required only in case of IPv4. In an IPv6 network environment, the configuration wizards automatically generate a unique and valid IPv6 address based on the network prefix advertised by the routers.

The following table describes the objects created and used during the installation and configuration.

Table 9-2 Active-active configuration objects

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

Setting up the Active/Active cluster

Perform the following tasks to configure an active-active SQL Server cluster.

Table 9-3 Tasks to set up an active-active SQL Server cluster

Task	Description
Install Veritas InfoScale Availability and configure the VCS cluster	<p>Install Veritas InfoScale Availability on all the systems where you want to configure an application for availability and configure a VCS cluster.</p> <p>For details refer to the <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p> <p>See “About installing the VCS agents” on page 45.</p> <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 57.</p>
Configure volumes or virtual disks for SQL Server	<p>For each instance of SQL Server (SQL Server system data files and the registry keys replicated among cluster nodes), create volumes or LUNs (virtual disks) on the shared storage.</p> <p>See “Managing storage using NetApp filer” on page 36.</p> <p>See “Managing storage using Windows Logical Disk Manager” on page 39.</p>
Install the first instance of SQL Server	<p>Consider the following points while you install the first instance of SQL Server:</p> <ul style="list-style-type: none"> ■ Do not accept the default instance name. Specify an instance name for each SQL Server installation. ■ Each SQL Server instance must be assigned a different port. The default port is 1433; ports for subsequent instances are generally assigned in descending order (1432, 1431, 1430, etc.). ■ Set a unique internal name for each instance. ■ Install SQL Server in the standalone installation mode in a non-clustered environment. From the SQL Server Installation Center, on the Installation panel, choose the New SQL Server stand-alone installation or add features to an existing installation option. ■ While installing SQL, ensure that you select all the desired features that you wish to configure with VCS. For example, Full-Text Search, Analysis Services, and so on. <p>See “About installing SQL Server on the first system” on page 48.</p>

Table 9-3 Tasks to set up an active-active SQL Server cluster (*continued*)

Task	Description
Configure the service group for the first SQL Server instance	<p>Consider the following points as you configure the SQL Server service group for the first instance:</p> <ul style="list-style-type: none"> ■ Assign a unique name to the SQL Server service group. ■ 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. <p>See “About configuring the SQL Server service group” on page 72.</p>
Configure the service group for user-defined databases, if any	<p>You can use VCS to manage user-defined SQL Server databases. Create the required SQL databases using the SQL Server Management Studio and then make them highly available with VCS.</p> <p>See “Making SQL Server user-defined databases highly available” on page 89.</p>

Table 9-3 Tasks to set up an active-active SQL Server cluster (*continued*)

Task	Description
Repeat the SQL Server installation and service group configuration for additional instances	<p>To install and configure one or more additional SQL Server instances, follow the same procedures as when installing the first instance.</p> <p>Make the same changes to the process, as follows:</p> <ul style="list-style-type: none"> ■ Do not accept the default instance name. Specify a unique instance name for each SQL Server installation. ■ Each SQL Server instance must be assigned a different port. The default port is 1433; ports for subsequent instances are generally assigned in descending order (1432, 1431, 1430, etc.). ■ Set a unique internal name for each instance. <p>See “About installing SQL Server on the first system” on page 48.</p> <p>Consider the following points as you configure the SQL Server service groups for the additional instances:</p> <ul style="list-style-type: none"> ■ Assign a unique name to the SQL Server service group ■ 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. <p>See “About configuring the SQL Server service group” on page 72.</p> <p>Configure the service group for user-defined databases, if any.</p> <p>See “Making SQL Server user-defined databases highly available” on page 89.</p>
Verify the configuration	See “Verifying the service group configuration” on page 91.

Configuring a disaster recovery setup

This chapter includes the following topics:

- [Setting up the disaster recovery cluster](#)
- [Configuring a disaster recovery set up for SQL Server](#)
- [Configuring the Global Cluster Option for wide-area failover](#)
- [Administering global service groups](#)

Setting up the disaster recovery cluster

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

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

Why implement a disaster recovery solution

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

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

- Minimizes economic loss due to the unavailability or loss of data.

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

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

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

Understanding replication

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

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

What needs to be protected in a SQL Server environment

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

- User Databases
The most critical component in any SQL Server implementation is the user data that is stored in user-defined databases.
- Logins
Logins allow clients to connect to SQL Server and execute queries on user data. Logins are stored in the master database and each of the user-defined databases.
- Jobs
Jobs are a set of scheduled tasks that maintain SQL Server databases. The job configuration is stored in the msdb system database.
- Alerts

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

- Operators
Operators are contacts that address problems occurring in SQL Server. They are notified in the event of errors. The operator configuration is stored in the msdb system database.
- Extended Stored Procedures
Extended stored procedures are external routines that are called from within SQL Server. They are typically stored in DLL files on the file system.
- Other Server Extensions
SQL Server is a very flexible database engine and it is possible to extend its functionality in several ways. These extensions are also important to the operation of the SQL Server.

Configuring a disaster recovery set up for SQL Server

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

Note: You must perform the tasks at the primary site first. After the configuration is complete at the primary site, proceed to perform the tasks at the secondary site. Before you begin to create the SQL Server service group for the cluster at the secondary site, make sure that the SQL Server service group at the primary site is offline.

The following table lists the tasks to set up a disaster recovery environment for SQL Server.

Table 10-1 Tasks for SQL Server disaster recovery set up

Task	Description
Review the configuration	Review the system configuration before you start deploying VCS and SQL Server in your environment. See “Before configuring the SQL Server service group ” on page 73.

Table 10-1 Tasks for SQL Server disaster recovery set up (*continued*)

Task	Description
Install the product and configure the VCS cluster	<p>Install Veritas InfoScale Availability on all the systems where you want to configure an application for availability and configure a VCS cluster:</p> <p>For details refer to the <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p> <p>See “About installing the VCS agents” on page 45.</p> <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 57.</p>
Configure volumes or LUNs on the shared storage	<p>Create volumes or LUNs required for SQL Server and ensure that the volumes or LUNs (virtual disks) are connected to the first cluster node.</p> <p>During the creation of virtual disks and volumes for the secondary site, make sure the following is exactly the same as the cluster on the primary site:</p> <ul style="list-style-type: none"> ■ Volume sizes ■ Volume names ■ Drive letters <p>See “Managing storage using NetApp filer” on page 36.</p> <p>See “Managing storage using Windows Logical Disk Manager” on page 39.</p>
Install and configure SQL Server on the first node	See “About installing SQL Server on the first system” on page 48.
Install and configure SQL Server on the additional nodes	<p>See “About installing SQL Server on additional systems” on page 50.</p> <p>Note: The instance name must be the same on the primary site and secondary site.</p>
Configure the SQL Server service group	<p>See “About configuring the SQL Server service group” on page 72.</p> <p>Note: The service group name and virtual computer name must be same on both, the primary site and secondary site.</p>

Table 10-1 Tasks for SQL Server disaster recovery set up (*continued*)

Task	Description
Configure DR components	<p>After configuring the high availability and SQL Server components on the primary and secondary sites, complete the disaster recovery solution by configuring the disaster recovery components for both sites.</p> <p>See “Prerequisites” on page 122.</p> <p>See “Linking clusters: Adding a remote cluster to a local cluster” on page 123.</p> <p>See “Converting a local service group to a global service group” on page 124.</p> <p>See “Bringing a global service group online” on page 126.</p> <p>See “Administering global service groups” on page 127.</p> <p>See “Deleting a remote cluster” on page 128.</p>

Configuring replication using NetApp SnapMirror

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

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

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

Refer to NetApp documentation for more information.

Configuring SnapMirror resources at the primary site

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

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

This is required in cases such as the following:

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

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

Configuring the Global Cluster Option for wide-area failover

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

Creating a global cluster environment involves the following:

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

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

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

Prerequisites

Creating a global cluster environment requires the following:

- Wide-area Connector process is configured and the ClusterService group is online at both sites.
See [“Configuring Wide-Area Connector process for global clusters”](#) on page 69.
- All service groups properly configured and able to come online.
- The service group serving as the global group has the same unique name across all applicable clusters.
- The clusters use the same version of VCS.
- The clusters use the same operating system.

- The clusters are standalone and do not already belong to a global cluster environment.
- The names of the clusters at the primary and secondary sites and the virtual IP addresses associated with them are registered in the DNS with reverse lookup.

Linking clusters: Adding a remote cluster to a local cluster

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

You can run the wizard from the following locations:

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

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

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

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

To add a remote cluster in Cluster Explorer

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

- 3 In the Wizard Options panel, click **Add Cluster**, then click **Next**.
- 4 In the New Cluster Details panel, enter the details of the new cluster.
 If the cluster is not running in secure mode, do the following:
 - Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
 - If necessary, change the default port number.
 - Enter the user name.
 - Enter the password.
 - Click **Next**.
 If the cluster is running in secure mode, do the following:
 - Enter the host name of a cluster system, an IP address of a cluster system, or the IP address of the cluster that will join the global environment.
 - Verify the port number.
 - Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and the domain.
 - If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection.
 - Click **Next**.
- 5 Click **Finish**. After running the wizard, the configurations on all the relevant clusters are in read-write mode; the wizard does not close the configurations.
- 6 Verify that the heartbeat connection between clusters is alive. From the command window enter `hahb -display`. The state attribute in the output should show `alive`.
 If the state is `unknown`, then offline and online the ClusterService group.

Converting a local service group to a global service group

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

To convert a local service group to a global group

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

or

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

or

From the Cluster Explorer configuration tree, right-click the service group, click **Configure As Global**, and proceed to step 3b.
- 2
- Review the information required for the Global Group Configuration wizard and click **Next**.
- 3
- Enter the details of the service group to modify as follows:
 - Click the name of the service group that will be converted from a local group to a global group, or vice versa.
 - From the Available Clusters box, click the clusters on which the group can come online. Click the right arrow to move the cluster name to the Clusters for Service Group box; for global to local cluster conversion, click the left arrow to move the cluster name back to the Available Clusters box. A priority number (starting with 0) indicates the cluster on which the group will attempt to come online. If necessary, double-click the entry in the Priority column and enter the new value.
 - Select the policy for cluster failover as follows:

Manual	Prevents a group from automatically failing over to another cluster.
Auto	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster fails.
Connected	Enables a group to automatically fail over to another cluster if it is unable to fail over within the cluster.
 - Click **Next**.
- 4
- Enter or review the connection details for each cluster. Click the **Configure** icon to review the remote cluster information for each cluster:

- | | |
|----------------------------|--|
| Cluster not in secure mode | <ul style="list-style-type: none"> ■ 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 | <ul style="list-style-type: none"> ■ Enter the IP address of the remote cluster, the IP address of a cluster system, or the host name of a cluster system. ■ Verify the port number. ■ Choose to connect to the remote cluster with the credentials used for the current cluster connection, or enter new credentials, including the user name, password, and domain. ■ If you connected to the remote cluster earlier through the wizard, you can use the credentials from the previous connection. ■ Click OK. ■ Repeat these steps for each cluster in the global environment. |

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

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

Bringing a global service group online

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

To bring a remote global service group online from Cluster Explorer

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

Administering global service groups

Administering global groups requires the following conditions:

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

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

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

Taking a remote global service group offline

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

To take a remote global service group offline from Cluster Explorer

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

Switching a remote service group

Use Cluster Explorer to switch a remote service group.

To switch a remote service group from Cluster Explorer

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

Deleting a remote cluster

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

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

Deleting a remote cluster involves the following tasks:

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

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

To take the wide area cluster (wac) resource offline

- 1 From Cluster Monitor, log on to the cluster that will be deleted from the global cluster environment.
- 2 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the **wac** resource under the Application type in the ClusterService group.
or
Click a service group in the configuration tree, click the **Resources** tab, and right-click the **wac** resource in the view panel.
- 3 Click **Offline**, and click the appropriate system from the menu.

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

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

- Click **OK**.

5 Click **Next**.

6 Click **Finish**.

To delete a remote cluster from the local cluster

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

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

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

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

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

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

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

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

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

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

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

- Click **OK**.

6 Click **Finish**.

Configuring SQL Server in a VMware environment

- [Chapter 11. Configuring application monitoring using the Veritas High Availability solution](#)
- [Chapter 12. Administering application monitoring](#)

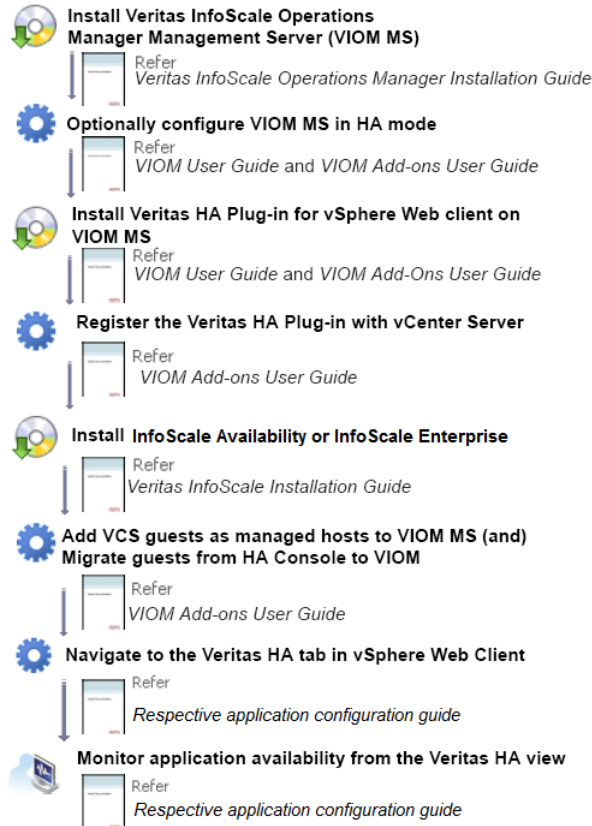
Configuring application monitoring using the Veritas High Availability solution

This chapter includes the following topics:

- [Deploying the Veritas High Availability solution for configuring application monitoring](#)
- [Notes and recommendations](#)
- [Configuring application monitoring](#)
- [Modifying the ESXDetails attribute](#)

Deploying the Veritas High Availability solution for configuring application monitoring

The following figure depicts the workflow for deploying the Veritas High Availability solution for configuring application monitoring.



Notes:

- You must follow this workflow only if you plan to configure application monitoring using the vSphere Web Client infrastructure. Alternatively, you can configure application monitoring and perform the further administrative tasks directly from a browser window using the following URL:
https://<virtualmachineNameorIPaddress>:5634/vcs/admin/application_health.html?priv=ADMIN
- For the configuring application monitoring in a VMware SRM environment, see:
<http://www.symantec.com/docs/DOC7958>
- If you plan to migrate from an existing Veritas High Availability vSphere Desktop Client integrated environment to a vSphere Web Client integrated environment, see:
https://sort.symantec.com/public/vom/Symantec_HA_Plug-in_for_vSphere_Web_Client_README.pdf

Notes and recommendations

Note the following prerequisites before configuring application monitoring:

- Verify that the boot sequence of the virtual machine is such that the boot disk (OS hard disk) is placed before the removable disks.
If the sequence places the removable disks before the boot disk, the virtual machine may not reboot after an application failover. The reboot may halt with an "OS not found" error.
This issue occurs because during the application failover the removable disks are detached from the current virtual machine and are attached on the failover target system.
- Verify that VMware Tools is installed on the virtual machine.
Install the version that is similar to or later than that available with VMware ESX 4.1.
- Verify that all the systems on which you want to configure application monitoring belong to the same domain.
- Verify that the ESX/ESXi host user account has administrative privileges or is a root user.
If the ESX/ESXi user account fails to have the administrative privileges or is not a root user, then in event of a failure the disk deattach and attach operation may fail.
If you do not want to use the administrator user account or the root user, then you must create a role, add the required privileges to the created role and then add the ESX user to that role.
See ["Assigning privileges for non-administrator ESX/ESXi user account"](#) on page 136.
- Verify that the SQL Server instances that you want to monitor are installed on the non-shared local disk that can be deported from the system and imported to another system.
- If you have configured a firewall, ensure that your firewall settings allow access to ports used by Veritas High Availability installer, wizard, and services.
- You must run the Veritas High Availability Configuration wizard from the system to which the disk residing on the shared datastore is attached (first system on which you installed SQL Server).
- After configuring SQL Server databases for monitoring, if you create another database or service, then these new components are not monitored as part of the existing configuration.

In this case, you can either use the VCS commands to add the components to the configuration or unconfigure the existing configuration and then run the wizard again to configure the required components.

- In case the VMwareDisks agent resource is configured manually, care should be taken not to add the operating system disk in the configuration. The VMwareDisks agent does not block this operation. This might lead to a system crash during failover.
- If VMware vMotion is triggered at the same time as an application fails over, the VMwareDisks resource may either fail to go offline or may report an unknown status. The resource will eventually failover and report online after the vMotion is successful and the application is online on the target system.
- VMware snapshot operations may fail if VMwareDisks agent is configured for a physical RDM type of disk. Currently only virtual RDM disks are supported.
- Non-shared disks partitioned using GUID Partition Table (GPT) are not supported. Currently only Master Boot Record (MBR) partition is supported.
- VMwareDisks agent does not support disks attached to the virtual machine using IDE controllers. The agent resource reports an unknown if IDE type of disks are configured.
- In case VMware HA is disabled and the ESX itself faults, VCS moves the application to the target failover system on another ESX host. VMwareDisks agent registers the faulted system on the new ESX host. When you try to power on the faulted system, you may see the following message in the vSphere Client:

```
This virtual machine might have been moved or copied.  
In order to configure certain management and networking features,  
VMware ESX needs to know if this virtual machine was moved or copied.  
If you don't know, answer "I copied it".
```

You must select “I moved it” (instead of the default “I copied it”) on this message prompt.

- You must not restore a snapshot on a virtual machine where an application is currently online, if the snapshot was taken when the application was offline on that virtual machine. Doing this may cause an unwanted fail over.
This also applies in the reverse scenario; you should not restore a snapshot where the application was online on a virtual machine, where the application is currently offline. This may lead to a misconfiguration where the application is online on multiple systems simultaneously.
- If you want to suspend a system on which an application is currently online, then you must first switch the application to a failover target system.

If you suspend the system without switching the application, then VCS moves the disks along with the application to another system.

Later, when you try to restore the suspended system, VMware does not allow the operation because the disks that were attached before the system was suspended are no longer with the system.

- While creating a VCS cluster in a virtual environment, you must configure one of the cluster communication link over a public adapter in addition to the link configured over a private adapter. To have less VCS cluster communication over the link using the public adapter, you may assign it low priority. This keeps the VCS cluster communication intact even if the private network adapters fail. If the cluster communication is configured over the private adapters only, the cluster systems may fail to communicate with each other in case of network failure. In this scenario, each system considers that the other system has faulted, and then try to gain access to the disks, thereby leading to an application fault.
- VMware Fault Tolerance does not support adding or removing of non-shared disks between virtual machines. During a failover, disks that contain application data cannot be moved to alternate failover systems. Applications that are being monitored thus cannot be brought online on the failover systems.
- For cluster communication, you must not select the teamed network adapter or the independently listed adapters that are a part of the teamed NIC. A teamed network adapter is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address, due to which you may experience the following issues:
 - The application monitoring configuration wizard may fail to discover the specified network adapters
 - The application monitoring configuration wizard may fail to discover/validate the specified system name

Assigning privileges for non-administrator ESX/ESXi user account

The application monitoring configuration in a VMware virtual environment using non-shared disks involves the VMwareDisks agent. In event of a failure, the VMwareDisks agent sends a disk detach request to the ESX/ESXi host and then attaches it to the new failover target system.

To enable the VMwareDisks agent to communicate with the ESX/ESXi host, we need to specify the ESX user account details during the application configuration workflow. This ESX user account must have the administrative privileges or should be a root user. If the ESX user account does not have these privileges, you must perform the following tasks:

- Create a role having the following privileges

- Low level file operations
 - Add existing disk
 - Change resource
 - Remove disk
- See [“Creating a role”](#) on page 137.
- Integrate with the existing authentication mechanism
 See [“Integrating with Active Directory or local authentication”](#) on page 138.
 - Add the ESX user to the created role
 See [“Adding a user to the role”](#) on page 139.

Note: If you do not want to add the existing user, you can create a new user and then add the same to the created role
 See [“Creating a new user”](#) on page 138.

Creating a role

Perform the following steps to create the role

- 1 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Administration > Roles**.
- 2 Click **Add Role**.
- 3 On the Add New Role panel, specify a name for the new role. For example, "ESX/ESXi User Role for Application Monitoring".
- 4 In the Privileges tree, click the following check boxes to assign the required privileges:
 - **All Privileges > Datastore > Low level file operations**
 - **All Privileges > Virtual Machine > Configuration > Adding existing disk**
 - **All Privileges > Virtual Machine > Change resource**
 - **All Privileges > Virtual Machine > Configuration > Remove disk**
- 5 Click **OK**.

Integrating with Active Directory or local authentication

To integrate with Active Directory or local authentication

- 1 Create a domain user in the Active Directory.
- 2 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**
- 3 Click the ESX host.
- 4 In the right pane, click **Configuration**.
- 5 In the Software panel, click **Authentication Services**.
- 6 Review the Directory Services Configuration.

If the Directory Service Type is not Active Directory, and you do not want to integrate with Active Directory, proceed to the section,

See [“Adding a user to the role”](#) on page 139.

If the Directory Service Type is not Active Directory, and you want to integrate with Active Directory, in the top right corner, click **Properties**.

- 7 In the Directory Service Configuration panel, from the Select Directory Service Type drop down list, select **Active Directory**.
- 8 In the Domain Settings area, specify the **Domain**, and click **Join Domain**.
Alternatively, configure vSphere Authentication proxy.
- 9 Enter the user name and password of a directory service user that has permissions to join the host to the domain, and click **OK**.

Creating a new user

You must perform this task only if you do not want to add the existing user to the created role.

Perform the following steps to create a new user

- 1 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**.
- 2 Click the ESX host.
- 3 In the right pane, click **Local Users & Groups**.

The Users list appears by default.

- 4 If the Users list is not displayed, on the View bar, click **Users**.

Alternatively, if the Users list is displayed, right-click any existing user and then click **Add**.

- 5 In the Add New User panel, specify a Login and Password to define a new user account.

To confirm the password, retype the password.

To define the new user account, you can also specify a descriptive User Name and user ID (UID). If you do not specify the UID, the vCenter server automatically assigns one.

- 6 Click **OK**.

Adding a user to the role

To add a user to the role

- 1 Using the VMware vSphere Client, log on to the ESX host, and navigate to **Home > Inventory**.
- 2 Click the ESX host.
- 3 In the right pane, click **Permissions**.
- 4 In the Permissions tab, right-click the blank space, and click **Add Permission**.
- 5 In the Assign Permissions panel, click **Add**.
- 6 In the Users and Groups frame of the Select Users and Groups panel, specify the user(s) that you want to assign the new role.

Press Ctrl and click to select multiple users, if required, and then click **Add** and click **OK**.

- 7 In the Assigned Role drop down list, click the new role and then click **OK**.

Configuring application monitoring

Configuring an application for monitoring involves the following tasks:

1. Configuring the VCS cluster

This task involves selecting the virtual machines on which you want to configure monitoring and setting up the network communication links between those virtual machines.

See [“Configuring the VCS cluster”](#) on page 140.

2. Configuring the application

This task involves configuring the application in one of the following modes:

- Start/stop mode on a single system: In the event of a failure, the application is restarted on a virtual machine for a configured number of times.

- **Failover mode on multiple systems:** In the event of a failure, the application is first restarted on a virtual machine for a configured number of times. If the application does not restart, the virtual machine fails over to another ESX host.

When you configure an application in the failover mode, you need to select the ESX hosts where the virtual machines can fail over.

See [“Configuring the application”](#) on page 143.

Configuring the VCS cluster

VCS cluster configuration involves selecting the virtual machines on which you want to configure monitoring and setting up the network communication links between the selected virtual machines.

To configure the VCS cluster

- 1** Launch the vSphere Client and connect to the vCenter Server that manages the virtual machine. On the Inventory view of the vCenter Server, in the left pane, select the virtual machine where you want to configure application monitoring. Then, in the right pane, select the Veritas High Availability tab.
- 2** On the Veritas High Availability view, click **Configure a VCS Cluster**. This launches the VCS cluster configuration wizard. Unless you configure a cluster, you cannot configure an application for monitoring.
- 3** On the Welcome panel, review the pre-requisites and then click **Next**.

- On the Configuration Inputs panel, specify the systems for the VCS cluster operations and the user account details for each system.

Note: The specified user account must have administrative privileges on the system.

The **Cluster systems** lists the systems that are included in the cluster configuration. The local system is selected by default.

To add more systems, click **Add System** and then on the Add System dialog box, specify the following details of the system that you want to add to the VCS cluster.

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
User name	Specify the user account for the system. The user name must be in the domain.com\username. Note: The specified user must be a domain user having administrative privileges on all the selected system.
Password	Specify the password for the user account mentioned.
Use the specified user account on all systems	Uses the specified user account on all the cluster systems. This option is selected by default.

The wizard validates the system details and then adds the system to VCS cluster system list.

- Skip this step if you do not want to modify the default security settings for your cluster.

To modify the security settings for the cluster, on the Configuration Inputs panel, click **Advanced Settings** . In the Advanced settings dialog box, specify the following details and then click **OK**.

Use Single Sign-on	Select to configure single sign-on using VCS Authentication Service for cluster communication. This option is enabled by default.
Use VCS user privileges	Select to configure a user with administrative privileges to the cluster. Specify the user name and password and click OK.

- On the Network Details panel, select the type of communication for the VCS cluster and then select the adapters to configure the communication links.

Depending on the network over which you want to configure the links, select:

- **Use MAC address for cluster communication (LLT over Ethernet) :**
The LLT over Ethernet communication configures the links over the non-routed network. Choose this mode only if the failover target systems reside in the same subnet.
- **Use IP address for cluster communication (LLT over UDP):** The LLT over UDP communication configures the links over the routed network. You choose this mode regardless of whether the failover target systems reside in the same subnet or in different subnets. You can select only those adapters that have an IP address.

You must select a minimum of two adapters per system.

Veritas recommends the following:

- IP addresses that are assigned to the selected adapters should be in different subnets.
- One of the network adapters must be a public adapter. You may assign low priority to the VCS cluster communication link that uses the public adapter.

Network Adapter	<p>Select a network adapter for the communication links.</p> <p>You must select a different network adapter for each communication link.</p> <p>To configure links over UDP, select the type of IP protocol and then specify the required details for each communication link.</p> <p>Note: Do not select the teamed network adapters and the independently listed adapters that are a part of the teamed NIC.</p>
IP Address	Specify the IP address for cluster communication over the specified UDP port.
Port	<p>Specify a unique port number for each link. You can use ports in the range 49152 to 65535.</p> <p>A specified port for a link is used for all the cluster systems on that link.</p>
Subnet mask	Displays the subnet masks to which the specified IP belongs.

By default, the VCS cluster communication link that uses the public adapter is configured as low-priority link. To change the priority, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

For information about adding or changing the selected network links after the configuration workflow is complete, see the *Cluster Server Administrator's Guide*.

- 7 On the Configuration Summary panel, specify a cluster name and a cluster ID of your choice. Review the VCS cluster configuration details and then click **Next** to initiate the configuration.

If the network contains multiple clusters, the wizard verifies the cluster ID with the IDs assigned to all the accessible clusters in the network. The wizard does not validate the assigned ID with the clusters that are not accessible during the validation. Veritas recommends you to validate the uniqueness of the assigned ID in the existing network.

- 8 On the Implementation panel, the wizard displays the VCS cluster configuration tasks and the status of each task. After all the tasks are complete, click **Next**.

If the configuration tasks fail, click **Diagnostic information** to check the details of the failure.

Rectify the cause of the failure and run the wizard again to configure the VCS cluster.

- 9 On the Finish panel, click **Finish** to complete the wizard workflow. This step completes the VCS cluster configuration.

The Veritas High Availability view now displays the link to configure an application for high availability.

Configuring the application

Perform the following steps to configure monitoring for SQL Server using the Veritas High Availability Configuration Wizard.

Note: Veritas High Availability does not support application monitoring for two different versions of SQL Server on the same system simultaneously.

To configure the application

- 1
- Launch the vSphere Client and connect to the vCenter Server that manages the virtual machine. On the Inventory view of the vCenter Server, in the left pane, select the virtual machine where you want to configure application monitoring. Then, in the right pane, select the Veritas High Availability tab.
- 2
- On the Veritas High Availability tab, click **Configure Application for High Availability**. This option is available only after a cluster has been configured. Unless you configure a cluster, you cannot configure an application for monitoring.
- 3
- On the Welcome panel of the Veritas High Availability Configuration wizard, review the pre-requisites and then click **Next**.
- 4
- On the Application Selection panel, select the application from the Supported Applications list and then click **Next**.

Alternatively, you can use the **Search** box to find the application and then click **Next**.

If you want to download any of the High Availability Agents, click the **Download Application Agents (SORT)** link to download the agents from the Veritas Operations Readiness Tools (SORT) site.

<https://sort.veritas.com/agents>

- 5
- On the Application Inputs panel, 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.
- 6
- On the SQL Instance Selection panel, choose the SQL Server instances and any of the following associated components that you want to monitor, and then click **Next**.

SQL Agent Service	Select this option to configure monitoring for SQL Server Agent service for the selected instance. You must select this for each selected SQL Server instance separately.
Analysis Service	Select this option to configure monitoring for SQL Server Analysis service for the selected instance. You must select this for each selected SQL Server instance separately.
FILESTREAM	Select this option to configure monitoring for FILESTREAM. You can select this option, if FILESTREAM is enabled on the selected instance.

7 To enable detail monitoring for the selected instances and associated services, select **Configure detail monitoring** and provide the following required details:

- Enter a non-zero value in the **Monitor after every ... cycles** box. This value indicates the number of online monitor cycles that the agent must wait before performing detail monitoring.

The numeric value specifies how often the monitoring check must run. 1 means, run a detail check every single monitor interval, 2 means run the detail check every second monitor interval. This interpretation may be extended to other values.

Veritas recommends that you set this value between 1 and 12. The default value is 5.

- Select one of the following modes for detail monitoring:
 - **Database monitoring**
Select this mode to enable detail monitoring by connecting to the database and choosing the desired databases from the **Databases** list.
 - **Script-based monitoring**
Select this mode to enable detail monitoring by using a user-defined SQL script.

If you choose script-based monitoring, you must enter the following details:

Script Path	Enter the script location on the virtual machine
User Name	Enter the valid user name
	Note: The user account must have rights to run the script and execute the SQL commands specified in the script.
Password	Enter the password for the user account specified in the User Name text box

- If the detail monitoring fails, the SQL Server agent can fail over the service group to another node. If you want the agent to initiate the failover in such a scenario, make sure that the **Restart the SQL instance if detail monitoring fails** check box is checked.

8 On the Registry Replication Details panel, select a location from the **Registry replication directory** drop-down list to save the registry replication data.

Veritas recommends that you store the registry replication data and the SQL Server application data at different locations.

- On the Configuration Inputs panel, select the VCS cluster systems on which you want to configure the application for high availability and move them to the **Application failover targets** list. The local system is selected by default.

Using the up-down arrow keys, you can define the priority order for the failover systems.

For each system that you assign as a failover target, you must specify the domain user account details in the appropriate fields on the Edit System dialog box. The VCS agents use these details to perform domain operations (such as Active Directory updates).

- On the Virtual Network Details panel, specify the virtual IP and the network details for the application to be configured and then click **Next**.

To specify the virtual IP and network details, select the IP protocol and then specify the following details for each failover system:

Note: You must select the same IP protocol as the one that was selected during the VCS cluster configuration.

Virtual IP address	Specify a unique virtual IP address.
Subnet mask	Specify the subnet mask to which the IP address belongs.
Virtual name	Specify a virtual name.
Network Adapter column	Select the network adapter that will host the virtual IP.

- 11** On the Storage HA Inputs panel, specify the ESX/ESXi hosts and the administrative user account details for each host, and then click **Next**.

Note: This panel appears only if you have specified a registry replication directory, or multiple application failover targets, or both for the selected SQL Server application.

To specify the ESX/ESXi hosts, click **Add ESX/ESXi Host** and on the Add ESX/ESXi Host dialogue box, specify the following details:

ESX/ESXi hostname or IP address	<p>Specify the target ESX hostname or IP address.</p> <p>The virtual machines will fail over on this ESX host during vMotion.</p> <p>The mount points configured on the ESX host where the application is currently running must be available on the target ESX host.</p>
User name	<p>Specify a user account for the ESX host.</p> <p>The user account must have administrator privileges on the specified ESX host.</p>
Password	<p>Specify the password for the user account provided in the User name text box.</p>

Note: By default, the wizard sets up a communication link with the ESX/ESXi server. You can modify the configuration to set up the communication link with vCenter Server instead. To set up a link with a vCenter Server, you must modify the ESXDetails attribute after this application monitoring configuration workflow is complete.

See [“Modifying the ESXDetails attribute”](#) on page 148.

- 12** On the Configuration Summary panel, review the application configuration details.

The wizard assigns a unique name to the application service group. Click **Rename** to rename the service group.

Click **Next** to initiate the application monitoring configuration.

- 13 On the Implementation panel, the wizard performs the application configuration tasks.

The wizard displays the status of each task. After all the tasks are complete, click **Next**.

If the configuration tasks fail, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure, and run the wizard again to configure application monitoring.

- 14 On the Finish panel, click **Finish** to complete the wizard workflow. This configures the application for high availability on the selected VCS cluster systems.

The Veritas High Availability view now displays the application status and its component dependency.

If the application status shows as not running, click **Start** to start the configured components on the system.

Note: Ensure that you start the application (bring the service group online) on the first system at least once. This is required to store the database related information and the SQL Server instances registry setting details in the VCS cluster configuration. When the application fails over to any other VCS cluster system, this information is applied to that system and the application is brought online on that system.

Modifying the ESXDetails attribute

You must modify the value of the "ESXDetails" attribute (of the VMwareDisks agent) if you want the VMwareDisks agent to communicate with the vCenter Server (instead of the ESX/ESXi host) for the disk detach and attach operations.

By default the "ESX Details" attribute of the VMwareDisks agent used the host names or IP addresses and the user account details of the ESX hosts on which the virtual machines are configured. To enable the VMwareDisks agent to communicate with the vCenter Server, you must modify the ESXDetails attribute and provide the host name or IP address and the user account details of the vCenter Server to which the virtual machines belong.

Use the Cluster Manager (Java Console) or the Command Line to modify the attribute values.

To modify the attribute from Cluster Manager

- 1 From the Cluster Manager configuration tree, select the VMwareDisks resource and then select the **Properties** tab.
- 2 On the **Properties** tab, click the Edit icon next to the ESX Details attribute.
- 3 On the Edit Attribute dialog box, select all the entries specified under the Key-Value column and press “-” to delete them.
- 4 Encrypt the password of the vCenter Server user account.
 - From the command prompt, run the following command:
`Vcsencrypt -agent`
 - Enter the vCenter Server user account password.
 - Re-enter the specified password.
 The encrypted value for the specified password is displayed.
- 5 On the Edit Attribute dialog box, click “+” to specify the values under the Key-Value column.
- 6 Under the Key column, specify the vCenter Server host name or the IP address.
- 7 Under the Value column, specify the encrypted password of the vCenter Server user account (from step 4)
- 8 Click **OK** to confirm the changes.
- 9 Repeat the steps for all VMwareDisks resources from the Cluster Manager configuration tree.
- 10 Save and close the configuration.

To modify/specify the attribute from Command Line

- 1 Change the VCS configuration to read/write mode.
`Haconf -makerw`
- 2 Delete the existing details of the ESX Server.
`hares -modify VMwareDisks ResourceName ESXDetails -delete -keys`
- 3 Encrypt the password of the vCenter Server user account.
 - From the command prompt, run the following command:
`Vcsencrypt -agent`
 - Enter the vCenter Server user account password.
 - Re-enter the specified password.

The encrypted value for the specified password is displayed.

4 Specify the vCenter Server details.

```
hares -modify <VMwareDisks ResourceName> ESXDetails  
-add <vCenter IP address or hostname> <UserName>=<encrypted password>
```

Administering application monitoring

This chapter includes the following topics:

- [About the various interfaces available for performing application monitoring tasks](#)
- [Administering application monitoring using the Veritas High Availability tab](#)
- [Administering application monitoring settings](#)
- [Administering application availability using Veritas High Availability dashboard](#)

About the various interfaces available for performing application monitoring tasks

Veritas High Availability solution enables you to configure and administer application monitoring through vSphere Web Client and directly through a browser window.

Veritas High Availability provides you with the following interfaces to perform the application monitoring tasks through vSphere Web Client:

- Veritas High Availability tab
Use the Veritas High Availability tab to configure and administer application monitoring on a virtual machine.
- Veritas High Availability dashboard
Use the Veritas High Availability dashboard to administer application monitoring for multiple applications on multiple systems in a data center.

For information about how does the Veritas High Availability solution integrates with vSphere Web Client, see:

To perform the tasks directly through a browser window, use the following URL:

`https://<virtualmachineNameorIPAddress>:5634/vcs/admin/
application_health.html?priv=ADMIN`

A prompt for the user account details will be displayed. You must enter the system user account details.

Administering application monitoring using the Veritas High Availability tab

Use the **Veritas High Availability** tab to perform the following tasks:

- To configure and unconfigure application monitoring
- To unconfigure the VCS cluster
- To start and stop configured applications
- To add and remove failover systems
- To enter and exit maintenance mode
- To switch an application
- To determine the state of an application (components)
- To resolve a held-up operation
- To modify application monitoring settings
- To view application dependency
- To view component dependency

Understanding the Veritas High Availability tab work area

The **Veritas High Availability** tab displays the consolidated health information for applications running in a Cluster Server (VCS) cluster. The cluster may include one or more systems.

When you click a system in the inventory view of the VMware vSphere Client, the **Veritas High Availability** tab displays application information for the entire VCS cluster, not just the selected system.

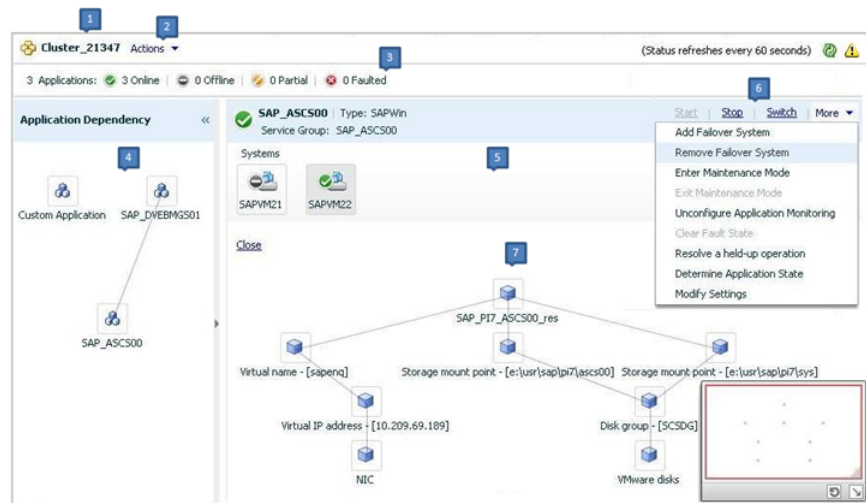
Note: If you do not configure any application for monitoring in the VCS cluster, then the **Veritas High Availability** tab displays only the following link: **Configure an application for high availability**.

The **Veritas High Availability** tab uses icons, color coding, dependency graphs, and tool tips to report the detailed status of an application.

The **Veritas High Availability** tab displays complex applications, like SAP Netweaver, in terms of multiple interdependent instances of that application. These interdependent instances represent component groups of the application. The component groups are also known as "service groups" in VCS terminology.

Each service group in turn includes several critical components of the application. The components are known as "resources" in VCS terminology.

The following figure displays two instances of SAP running in the **Veritas High Availability** tab:



1. Title bar
2. Actions menu
3. Aggregate status bar
4. Application dependency graph
5. Application table
6. Application-specific task menu
7. Component dependency graph

The **Veritas High Availability** tab graphic user interface (GUI) includes the following components:

- Title bar: Displays the name of the VCS cluster, the Actions menu, the Refresh icon, the Alert icon. Note that the Alert icon appears only if the Veritas High Availability view fails to display a virtual machine, or displays stale data

- **Actions menu:** Includes a drop-down list of operations that you can perform with effect across the cluster. These include: Configuring an application for high availability; Unconfigure all applications; and Unconfigure VCS cluster.
- **Aggregate status bar:** Displays a summary of applications running in the cluster. This summary includes the total number of applications, and the state-wise breakdown of the applications in terms of the Online, Offline, Partial, and Faulted states.
- **Application dependency graph:** Illustrates the order in which the applications or application instances, must start or stop.
If an application must start first for another application to successfully start, the former application appears at a lower level. A line connects the two applications to indicate the dependency. If no such dependency exists, all applications appear in a single horizontal line.
- **Application table:** Displays a list of all applications that are configured in the VCS cluster that is associated with the system you selected in the inventory view.
Each application is listed in a separate row. Each row displays the systems where the application is configured for monitoring.
The title bar of each row displays the following entities to identify the application or the application instance (service group):
 - Display name of the application (for example, Payroll application)
 - Type of application (for example, Custom)
 - Service group name
- **Application-specific task menu:** Appears in each application-specific row of the application table. The menu includes application-specific tasks such as Start, Stop, Switch, and a drop-down list of more tasks. The More drop-down list includes tasks such as Add a failover system, and Remove a failover system.
- **Component dependency graph:** Illustrates the order in which application components (resources) must start or stop for the related application or application instance to respectively start or stop. The component dependency graph by default does not appear in the application table. To view the component dependency graph for an application, you must click a system on which the application is running.
The track pad, at the right-bottom corner helps you navigate through complex component dependency graphs.
If you do not want to view the component dependency graph, in the top left corner of the application row, click **Close**.

To view the status of configured applications

In the application dependency graph, click the application for which you want to view the status. If the appropriate row is not already visible, the application table automatically scrolls to the appropriate row. The row displays the state of the application for each configured failover system in the cluster for that application.

If you click any system in the row, a component dependency graph appears. The graph uses symbols, color code, and tool tips to display the health of each application component. Roll the mouse over a system or component to see its health details.

The health of each application or application component on the selected system is displayed in terms of the following states:

Table 12-1 Application states

State	Description
Online	Indicates that the configured application or application components are running on the virtual machine. If the application is offline on at least one other failover system, an alert appears next to the application name.
Offline	Indicates that the configured application or its components are not running on the virtual machine.
Partial	Indicates that either the application or its components are started on the virtual machine or VCS was unable to start one or more of the configured components If the application is offline on at least one other failover system, an alert appears next to the application name.
Faulted	Indicates that the configured application or its components have unexpectedly stopped running.

To configure or unconfigure application monitoring

Use the **Veritas High Availability** tab to configure or unconfigure an application for monitoring in a cluster under Cluster Server (VCS) control.

The tab provides you with specific links to perform the following configuration tasks:

- Configure the first application for monitoring in a VCS cluster:
If you have not configured any application for monitoring in the cluster, the **Veritas High Availability** tab appears blank except for the link **Configure an application for high availability**.

Click the link to launch the **Veritas High Availability Configuration Wizard**.
Use the wizard to configure application monitoring.

- Unconfigure monitoring of an application:
In the appropriate row of the application table, click **More > Unconfigure Application Monitoring** to delete the application monitoring configuration from the VCS.
Note that this step does not remove VCS from the system or the cluster, this step only removes the monitoring configuration for that application.
Also, to unconfigure monitoring for an application, you can perform one of the following procedures: Unconfigure monitoring of all applications, or unconfigure VCS cluster.
- Unconfigure monitoring of all applications:
Click **Actions > Unconfigure all applications**. This step deletes the monitoring configuration for all the applications that are configured in the cluster.
- Unconfigure VCS cluster:
Click **Actions > Unconfigure VCS cluster**. This step stops the VCS cluster, removes VCS cluster configuration, and unconfigures application monitoring.

To start or stop applications

Use the following options on the **Veritas High Availability** tab to control the status of the configured application and the associated components or component groups (application instances).

Note that the **Start** and **Stop** links are dimmed in the following cases:

- If you have not configured any associated components or component groups (resources or service groups) for monitoring
- If the application is in maintenance mode
- If no system exists in the cluster, where the application is not already started or stopped as required.

To start an application

- 1 In the appropriate row of the application table, click **Start**.
- 2 If the application (service group) is of the failover type, on the **Start Application** panel, click **Any system**. VCS uses predefined policies to decide the system where to start the application.

If the application (service group) is of the parallel type, on the **Start Application** panel, click **All systems**. VCS starts the application on all required systems, where the service group is configured.

Note: Based on service group type, either the Any system or the All Systems link automatically appears.

To learn more about policies, and parallel and failover service groups, see the *Cluster Server Administrator's Guide*.

If you want to specify the system where you want to start the application, click **User selected system**, and then click the appropriate system.

- 3 If the application that you want to start requires other applications or component groups (service groups) to start in a specific order, then check the **Start the dependent components in order** check box, and then click **OK**.

To stop an application

- 1 In the appropriate row of the application table, click **Stop**.
- 2 If the application (service group) is of the failover type, in the Stop Application Panel, click **Any system**. VCS selects the appropriate system to stop the application.

If the application (service group) is of the parallel type, in the Stop Application Panel click **All systems**. VCS stops the application on all configured systems.

Note: Based on service group type, either the Any system or the All Systems link automatically appears.

To learn more about parallel and failover service groups, see the *Cluster Server Administrator's Guide*.

If you want to specify the system, where you want to stop the application, click **User selected system**, and then click the appropriate system.

- 3 If the application that you want to stop requires other applications or component groups (service groups) to stop in a specific order, then check the **Stop the dependent components in order** check box, and then click **OK**.

To suspend or resume application monitoring

After configuring application monitoring you may want to perform routine maintenance tasks on those applications. These tasks may or may not involve stopping the application but may temporarily affect the state of the applications and its dependent components. If there is any change to the application status, Cluster Server (VCS) may try to restore the application state. This may potentially affect the maintenance tasks that you intend to perform on those applications.

The **Enter Maintenance Mode** link is automatically dimmed if the application is already in maintenance mode. Conversely, if the application is not in maintenance mode, the **Exit Maintenance Mode** link is dimmed.

The **Veritas High Availability** tab provides the following options:

To enter maintenance mode

- 1 In the appropriate row, click **More> Enter Maintenance Mode**.
During the time the monitoring is suspended, Veritas High Availability solutions do not monitor the state of the application and its dependent components. The **Veritas High Availability** tab does not display the current status of the application. If there is any failure in the application or its components, VCS takes no action.
- 2 While in maintenance mode, if a virtual machine restarts, if you want application monitoring to remain in maintenance mode, then in the **Enter Maintenance Mode** panel, check the **Suspend the application availability even after reboot** check box, and then click **OK** to enter maintenance mode.

To exit the maintenance mode

- 1 In the appropriate row, click **More > Exit Maintenance Mode**, and then click **OK** to exit maintenance mode.
- 2 Click the Refresh icon in the top right corner of the **Veritas High Availability** tab, to confirm that the application is no longer in maintenance mode.

To switch an application to another system

If you want to gracefully stop an application on one system and start it on another system in the same cluster, you must use the Switch link. You can switch the application only to a system where it is not running.

Note that the Switch link is dimmed in the following cases:

- If you have not configured any application components for monitoring
- If you have not specified any failover system for the selected application
- If the application is in maintenance mode

- If no system exists in the cluster, where the application can be switched
- If the application is not in online or partial state on even a single system in the cluster

To switch an application

- 1 In the appropriate row of the application table, click **Switch**.
- 2 If you want VCS to decide to which system the application must switch, based on policies, then in the **Switch Application** panel, click **Any system**, and then click **OK**.

To learn more about policies, see the *Cluster Server Administrator's Guide*.

If you want to specify the system where you want to switch the application, click **User selected system**, and then click the appropriate system, and then click **OK**.

Cluster Server stops the application on the system where the application is running, and starts it on the system you specified.

To add or remove a failover system

Each row in the application table displays the status of an application on the systems that are part of a VCS cluster in a VMware environment. The displayed system/s either form a single-system Cluster Server (VCS) cluster with application restart configured as a high-availability measure, or a multi-system VCS cluster with application failover configured. In the displayed cluster, you can add a new system as a failover system for the configured application.

The system must fulfill the following conditions:

- Cluster Server (InfoScale Availability)7.2 is installed on the system.
- The system is not part of any other VCS cluster.
- The system has at least two network adapters.
- The required ports are not blocked by a firewall.
- The application is installed identically on all the systems, including the proposed new system.

To add a failover system, perform the following steps:

Note: The following procedure describes generic steps to add a failover system. The wizard automatically populates values for initially configured systems in some fields. These values are not editable.

To add a failover system

- 1 In the appropriate row of the application table, click **More > Add Failover System**.
- 2 Review the instructions on the welcome page of the Veritas High Availability Configuration Wizard, and click **Next**.
- 3 If you want to add a system from the Cluster systems list to the **Application failover targets** list, on the **Configuration Inputs** panel, select the system in the Cluster systems list. Use the Edit icon to specify an administrative user account on the virtual machine. You can then move the required system from the Cluster system list to the **Application failover targets** list. Use the up and down arrow keys to set the order of systems in which VCS agent must failover applications.

If you want to specify a failover system that is not an existing cluster node, on the **Configuration Inputs** panel, click **Add System**, and in the **Add System** dialog box, specify the following details:

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
Domain/Username	<p>Specify the user name with administrative privileges on the system.</p> <p>Specify the user name must be in the <i>domain.com\username</i> format.</p> <p>If you want to specify the same user account on all systems that you want to add, check the Use the specified user account on all systems box.</p>
Password	Specify the password for the account you specified.
Use the specified user account on all systems	This option is checked by default. You cannot modify this setting.

The wizard validates the details, and the system then appears in the **Application failover target** list.

- 4 Specify the user name and that VCS agents must use to perform domain operations such as Active Directory updates.
- 5 If you are adding a failover system from the existing VCS cluster, the Network Details panel does not appear.

If you are adding a new failover system to the existing cluster, on the **Network Details** panel, review the networking parameters used by existing failover

systems. Appropriately modify the following parameters for the new failover system.

Note: The wizard automatically populates the networking protocol (UDP or Ethernet) used by the existing failover systems for Low Latency Transport communication. You cannot modify these settings.

- To configure links over ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- To configure links over UDP, specify the required details for each communication link.

Network Adapter	<p>Select a network adapter for the communication links.</p> <p>You must select a different network adapter for each communication link.</p> <p>Veritas recommends that one of the network adapters must be a public adapter and the VCS cluster communication link using this adapter is assigned a low priority.</p> <p>Note: Do not select the teamed network adapter or the independently listed adapters that are a part of teamed NIC.</p>
IP Address	<p>Select the IP address to be used for cluster communication over the specified UDP port.</p>
Port	<p>Specify a unique port number for each link. You can use ports in the range 49152 to 65535.</p> <p>The specified port for a link is used for all the cluster systems on that link.</p>
Subnet mask	<p>Displays the subnet mask to which the specified IP belongs.</p>

- 6 If a virtual IP is not configured as part of your application monitoring configuration, the **Virtual Network** Details page is not displayed. Else, on the **Virtual Network Details** panel, review the following networking parameters that the failover system must use, and specify the NIC:

Virtual IP address	Specifies a unique virtual IP address.
Subnet mask	Specifies the subnet mask to which the IP address belongs.
Virtual name	Specifies a virtual name.
NIC	For each newly added system, specify the network adaptor that must host the specified virtual IP.

- 7 If the newly added failover system is associated with a different ESX host as compared to other systems, then on Target ESX Details page, specify the ESX host of the newly added failover system. Also specify the administrative user account details associated with the ESX host.

Note: If the application for which you are adding a failover system does not use storage attached directly to the ESX host, the wizard does not display this page.

If the new failover system runs on a different ESX host, or is configured to failover to another ESX host, specify that ESX host. To specify the ESX host, click **Add ESX Host** and on the **Add ESX Host** dialog box, specify the following details, and then click **Next**:

ESX host name or IP address	Specify the target ESX host name or IP address. The virtual machines can failover to this ESX host during vMotion. Specify an ESX host that has the same mount points as those currently used by the application.
User name	Specify a user account for the ESX host. The user account must have administrator privileges on the specified ESX host.
Password	Specify the password that is associated with the user name you specified.

The wizard validates the user account and the storage details on the specified ESX host, and uses this account to move data disks during vMotion.

- 8 On the **Configuration Summary** panel, review the VCS cluster configuration summary, and then click **Next** to proceed with the configuration.

- 9 On the **Implementation** panel, the wizard adds the specified system to the VCS cluster, if it is not already a part. It then adds the system to the list of failover targets. The wizard displays a progress report of each task.
 - If the wizard displays an error, click **View Logs** to review the error description, troubleshoot the error, and re-run the wizard from the Veritas High Availability tab.
 - Click **Next**.
- 10 On the **Finish** panel, click **Finish**. This completes the procedure for adding a failover system. You can view the system in the appropriate row of the application table.

Similarly you can also remove a system from the list of application failover targets.

Note: You cannot remove a failover system if an application is online or partially online on the system.

To remove a failover system

- 1 In the appropriate row of the application table, click **More > Remove Failover System**.
- 2 On the **Remove Failover System** panel, click the system that you want to remove from the monitoring configuration, and then click **OK**.

Note: This procedure only removes the system from the list of failover target systems, not from the VCS cluster. To remove a system from the cluster, use VCS commands. For details, see the *Cluster Server Administrator's Guide*.

To clear Fault state

When you fix an application fault on a system, you must further clear the application Faulted state on that system. Unless you clear the Faulted state, VCS cannot failover the application on that system.

You can use the Veritas High Availability tab to clear this faulted state at the level of a configured application component (resource).

The Clear Fault link is automatically dimmed if there is no faulted system in the cluster.

To clear Fault state

- 1 In the appropriate row of the application table, click **More > Clear Fault state**.
- 2 In the **Clear Fault State** panel, click the system where you want to clear the Faulted status of a component, and then click **OK**.

To resolve a held-up operation

When you try to start or stop an application, in some cases, the start or stop operation may get held-up mid course. This may be due to VCS detecting an incorrect internal state of an application component. You can resolve this issue by using the resolve a held-up operation link. When you click the link, VCS appropriately resets the internal state of any held-up application component. This process prepares the ground for you to retry the original start or stop operation, or initiate another operation.

To resolve a held-up operation

- 1 In the appropriate row of the application table, click **More > Resolve a held-up operation**.
- 2 In the **Resolve a held-up operation** panel, click the system where you want to resolve the held-up operation, and then click **OK**.

To determine application state

The **Veritas High Availability** tab displays the consolidated health information of all applications that are configured for monitoring in a VCS cluster. The tab automatically refreshes the application health information every 60 seconds.

If you do not want to wait for the automatic refresh, you can instantaneously determine the state of an application by performing the following steps:

To determine application state

- 1 In the appropriate row of the Application table, click **More > Determine Application State**.
- 2 In the **Determine Application State** panel, select a system and then click **OK**.

Note: You can also select multiple systems, and then click **OK**.

To remove all monitoring configurations

To discontinue all existing application monitoring in a VCS cluster, perform the following step:

- On the **Veritas High Availability** tab, in the Title bar, click **Actions > Unconfigure all applications**. When a confirmation message appears, click **OK**.

To remove VCS cluster configurations

If you want to create a different VCS cluster, say with new systems, a different LLT protocol, or secure communication mode, you may want to remove existing VCS cluster configurations. To remove VCS cluster configurations, perform the following steps:

Note: The following steps delete all cluster configurations, (including networking and storage configurations), as well as application-monitoring configurations.

- On the Title bar of the Veritas High Availability tab, click **Actions > Unconfigure VCS cluster**.
- In the **Unconfigure VCS Cluster** panel, review the Cluster Name and Cluster ID, and specify the User name and Password of the Cluster administrator. For non-secure clusters, specify the user name and password credentials of a domain user with local administrative privileges on each VCS cluster node, and then click **OK**.

Administering application monitoring settings

The Veritas High Availability tab lets you define and modify settings that control application monitoring with Cluster Server (VCS). You can define the settings on a per application basis. The settings apply to all systems in a VCS cluster, where that particular application is configured for monitoring.

The following settings are available:

- **App.StartStopTimeout:** When you click the **Start Application** or **Stop Application**, or **Switch Application** links in the **Veritas High Availability** tab, VCS initiates an application start or stop, respectively. This option defines the number of seconds that VCS must wait for the application to start or stop, after initiating the operation. You can set a value between 0 and 300 seconds for this attribute; the default value is 30 seconds.
If the application does not respond in the stipulated time, the tab displays an alert. The alert states that the operation may take some more time to complete and that you must check the status after some time. A delay in the application response does not indicate that the application or its dependent component has faulted. Parameters such as workload, system performance, and network

bandwidth may affect the application response. VCS continues to wait for the application response even after the timeout interval elapses.

If the application fails to start or stop, VCS takes the necessary action depending on the other configured remedial actions.

- **App.RestartAttempts:** This setting defines the number of times that VCS must try to restart a failed application. The value of App.RestartAttempts may vary between 0 and 5; the default value is 0. If an application fails to start within the specified number of attempts, VCS fails over the application to a configured failover system.
- **App.DisplayName:** This setting lets you specify an easy-to-use display name for a configured application. For example, Payroll Application. VCS may internally use a different application name to uniquely identify the application. However, the internal string, for example OraSG2, may not be intuitive to understand, or easy to recognize while navigating the application table.
Moreover, once configured, you cannot edit the application name, while you can modify the application display name as required. Note that the Veritas High Availability tab displays both the application display name and the application name.

Administering application availability using Veritas High Availability dashboard

The **Veritas High Availability** dashboard is a consolidated graphic user interface that lets you administer application monitoring on systems in a VMware vCenter administered data center.

The dashboard is fully integrated with the VMware vSphere Web Client. The dashboard appears in the **Veritas High Availability** tab of the VMware vSphere Web Client. To view the dashboard, select a data center or an ESX cluster in the inventory, and then click the Veritas High Availability tab.

On the dashboard, you can view the aggregate health statistics for monitored applications across a data center. You can also drill down to an ESX cluster and view monitored applications running in that cluster.

Note: The dashboard presents a unified view of monitored applications in a data center. It displays an application-centric view, not a product-centric view. If you have configured applications under more one Veritas High Availability product (VCS or ApplicationHA), then you cannot determine which application is under the control of which Veritas High Availability product. However, you can conclude that applications configured for failover are under VCS control. Applications configured for monitoring without a failover system may either be under VCS control or under ApplicationHA control.

To understand how to navigate across the dashboard:

See [“Understanding the dashboard work area”](#) on page 167.

You can drill down to an individual application and perform the following administrative actions:

- Start application
- Stop application
- Enter maintenance mode
- Exit maintenance mode
- Switch application (to another system)

Apart from applications on systems running Cluster Server, the Veritas High Availability dashboard also displays applications running on Symantec ApplicationHA guests (versions 5.1 SP2 and later).


For more information on monitoring applications running on Symantec ApplicationHA, refer to Symantec ApplicationHA documentation.

Understanding the dashboard work area

The Veritas High Availability dashboard displays the aggregate application health status information for a data center or an ESX cluster.

Depending on whether you click a data center or a VMware cluster in the inventory view (left pane) of the VMware vSphere Client GUI, the dashboard displays the aggregate application status information. Apart from the application table described below, the dashboard uses color code and tool tips to indicate the status of an application.

The following figure illustrates the dashboard work area. Note that the red boxes highlight the key GUI elements:



In the above figure, the labels stand for the following elements of the dashboard

- | | | | | | |
|---|----------------------|---|------------------------|---|---------------------------|
| 1 | Aggregate status bar | 2 | ESX cluster/host table | 3 | Taskbar |
| 4 | Filters menu | 5 | Application table | 6 | Systems table (drop-down) |

Aggregate status bar

The aggregate status bar of the dashboard displays the following details:

- Number of ESX clusters that have applications configured for monitoring with VCS
- Number of configured applications in the selected data center
- Number of faulted applications
- Number of applications in partial state
- Number of online applications
- Number of offline applications

ESX cluster/host table

The Veritas High Availability dashboard displays this table only if you click a data center in the inventory view of the vSphere Client, and then click the Veritas High Availability tab.

The cluster table lists the following statistics per ESX cluster (or independent ESX host) in the data center:

- Number of configured applications

- Number of faulted applications
- Number of applications in partial state
- Number of online applications
- Number of offline applications
- Overall status (percentage of healthy applications)

If you click a row in the ESX cluster/host table, the application table of the dashboard displays monitored applications running on systems hosted by the selected ESX cluster or ESX host (an ESX server that is not part of an ESX cluster).

Note: This is the only method to navigate to applications running on systems hosted by standalone ESX hosts, by using the Veritas High Availability dashboard.

Taskbar

The taskbar displays icons for various administrative tasks. A tool tip highlights the task that each icon represents.

The dashboard supports the following tasks:

- Start Application: Starts a configured application
- Stop Application: Stops a configured application
- Enter Maintenance Mode: Suspends application monitoring of the configured application. In maintenance mode, VCS does not monitor the state of the application, and its dependent components.
- Exit Maintenance Mode: Resumes application monitoring for a configured application.
- Switch Application: Switches and an application gracefully from one system to another.

Filters menu

The filters menu lets you dynamically filter the applications that are displayed in the applications table. You can filter the applications by the following parameters:

- Application name
- Application status
- Search (by a string)

Application table

If you click an ESX cluster in the ESX cluster/host table, or in the inventory view of the VMware vSphere Client, then the list of applications running in that ESX cluster appears in the application table of the dashboard.

If you click an ESX host (an ESX server that is not part of an ESX cluster) in the ESX cluster/host table, then the list of applications that are configured on systems hosted by that ESX server appears. Note that this is the only route to navigate to such applications through the dashboard

The following table lists each column in the application table and its description:

Column	Description
Applications	Indicates the application name.
Service Groups	<p>Indicates the group of critical application components that VCS uses to determine the health of a monitored application. Service group is a VCS term. The equivalent term in Veritas High Availability terminology is “component group”.</p> <p>VCS may use more than one service group to monitor a complex application. The dashboard displays each service group of such an application as a separate instance of that application.</p>
Status	<p>This column indicates the effective status of an application in a VCS cluster. It does not indicate the state of the application on per member system. For example, in a two-system cluster, if the application has faulted on one system but has failed over to another system, then this column states the state of the application as Online.</p> <p>Indicates one of the following states of an application:</p> <ul style="list-style-type: none"> ■ Online ■ Offline ■ Faulted ■ Partial <p>Note: After you perform an administrative task such as starting or stopping an application, or entering or exiting maintenance mode, it takes a few seconds for the dashboard to reflect the revised status of the configured application.</p>
Systems	Indicates the number of systems where the application is configured for monitoring. To view more information about all such systems, click the (...) icon. The System table (drop-down) appears, listing the ESX host name of each configured system, the VM name (system name), and the status of the application on each system.

Column	Description
Alerts and description	<p>Displays a triangular alert icon (!) and describes the reason for the alert. This column displays alerts in two cases: a). If the application status record is stale; b). If the application has faulted on a system.</p> <p>For stale records, the column includes the timestamp of the last received health record. In case of application fault, the column provides details of the system where the fault occurred.</p>

Monitoring applications across a data center

If you click a data center in the inventory view of the VMware vSphere Client, and then click the Veritas High Availability tab, the dashboard appears, displaying the aggregate health information of applications running inside various ESX clusters.

You can use filters to drill down from all applications running across the data center and view a single application and its various instances in the data center.

Monitoring applications across an ESX cluster

If you click an ESX cluster in the inventory view of the VMware vSphere Client, and then click the tab, the dashboard displays the consolidated information on the systems and applications running in the ESX cluster. The dashboard also displays the application health and application monitoring information.

You can use filters to drill down from all applications running in the ESX cluster, to view a single application and its various instances in the ESX cluster.

Searching for application instances by using filters

The High Availability dashboard lets you search for all instances of a particular application in the selected data center or an ESX cluster. Various filters enable you to search for the application that you want to monitor. You can use multiple filters simultaneously to search for an application.

The following table lists each field in the filter menu and its description:

Field	Description
Application	Lets you specify the name of the application that you want to filter in the application table. A drop-down list displays all the applications that are configured in the data center or ESX cluster. Click to select the name of the application that you want to filter.

Field	Description
Status	Lets you specify the status of the application by which you want to filter the application table. A drop-down list displays the following status values: Online, Offline, Faulted, and Partial.
Search	<p>Lets you search for an application by using a string or pattern of characters. Enter the string using which you want to filter applications. As you enter the string in the Search box, the dashboard dynamically filters the applications.</p> <p>Note: The dashboard searches for the specified string in the Systems column.</p>

Selecting multiple applications for batch operations

You can select one or more instances of an application for administering by using the dashboard as follows:

- To select one application instance, click inside the row of that application instance.
- To select various instances, keep the **Control** key pressed and then click inside the row of each instance.
- To select a batch of consecutive entries in the application table, keep the **Shift** key pressed, click inside the row of the first instance, and then click inside the row of the last instance. Alternatively, you can keep the **Shift** key pressed and drag the mouse to mark a block of consecutive entries.
- To select all instances in the application table, click **Select All**.

Starting an application using the dashboard

To start an application, perform the following steps in the application table of the dashboard.

To start an application

- 1 Filter the applications that you want to start.

See [“Searching for application instances by using filters”](#) on page 171.

The application table displays all the instances of the filtered applications.

- 2 If required, select multiple applications or instances to perform a batch operation.

See [“Selecting multiple applications for batch operations”](#) on page 172.

- 3 To start the application, in the taskbar, click the appropriate icon (use the tool tip to recognize the appropriate icon).
- 4 In the Start Application panel, click the systems where you want to start the application. Note that you can start the application on any of the systems displayed for each application.

Click **OK**.

Stopping an application by using the dashboard

To stop an application on one or more virtual machines, perform the following steps in the application table of the High Availability dashboard.

To stop an application

- 1 Filter the applications that you want to stop.

See [“Searching for application instances by using filters”](#) on page 171.

The application table displays all the instances of the filtered applications.
- 2 If required, select multiple applications or instances to perform a batch operation.

See [“Selecting multiple applications for batch operations”](#) on page 172.
- 3 To stop the application, in the taskbar, click the appropriate icon (use the tool tip to recognize the appropriate icon).
- 4 In the Stop Application panel, from the drop-down list, click the systems where you want to stop the application.

Click **OK**.

Entering an application into maintenance mode

You may need to intentionally take an application offline for maintenance purposes, without triggering a corrective response from Cluster Server (VCS).

To enter an application into maintenance mode, perform the following steps in the application table of the High Availability dashboard.

Note: The maintenance mode configuration is application-specific, not system-specific.

To enter maintenance mode

- 1** Filter the application that you want to gracefully take offline for maintenance.
See [“Searching for application instances by using filters”](#) on page 171.
The application table displays all the instances of the filtered applications.
- 2** If required, select multiple applications or instances to perform a batch operation.
See [“Selecting multiple applications for batch operations”](#) on page 172.
- 3** To enter maintenance mode, in the taskbar, click the appropriate icon for entering maintenance mode (use the tool tip to recognize the appropriate icon).
- 4** If a system restarts while the application is in maintenance mode, and you want the application to remain in maintenance mode, then in the Enter Maintenance Mode panel, check the **Suspend the application availability even after reboot**.
- 5** On the Enter Maintenance Mode panel, click **OK**.

Bringing an application out of maintenance mode

To bring an application out of maintenance mode on one or more systems, perform the following steps in the application table of the High Availability dashboard.

To exit maintenance mode

- 1** Filter the applications that you want to bring out of maintenance mode.
See [“Searching for application instances by using filters”](#) on page 171.
The application table displays all the instances of the filtered applications.
- 2** If required, select multiple applications or instances to bring out of maintenance mode.
See [“Selecting multiple applications for batch operations”](#) on page 172.
- 3** To bring the applications out of maintenance mode, in the taskbar, click the appropriate icon for exiting maintenance mode (use the tool tip to recognize the appropriate icon).
- 4** In the Exit Maintenance Mode panel, click **OK**.

Switching an application

To gracefully switch an application from one system to another, perform the following steps in the application table of the dashboard.

Note: You can switch an application only if the application monitoring configuration includes one or more failover systems.

To switch an application

- 1 Filter the applications that you want to switch to another node.
See [“Searching for application instances by using filters”](#) on page 171.
The application table displays all the instances of the filtered applications.
- 2 If required, select multiple applications or instances to perform a batch operation.
See [“Selecting multiple applications for batch operations”](#) on page 172.
- 3 To switch the applications, in the taskbar, click the appropriate icon (use the tool tip to recognize the appropriate icon).
- 4 In the Switch Application panel, select the systems where you want to switch the applications, and then click **OK**. Cluster Server takes the applications offline on the existing systems, and brings them online on the systems that you specified.

Appendixes

- [Appendix A. Troubleshooting](#)
- [Appendix B. Using the virtual MMC viewer](#)

Troubleshooting

This appendix includes the following topics:

- [VCS logging](#)
- [VCS Cluster Configuration Wizard \(VCW\) logs](#)
- [VCWsilent logs](#)
- [NetApp agents error messages](#)
- [Error and warning messages from VCS agent for SQL Server](#)
- [Troubleshooting application monitoring configuration issues](#)
- [Troubleshooting Veritas High Availability view issues](#)

VCS logging

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

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

The format of agent log messages is: *Timestamp (YYYY/MM/DD) | Mnemonic | Severity | UMI | Agent Type | Resource Name | Entry Point | Message Text.*

The following table describes the agent log message components and their descriptions.

Table A-1 Log message components and their description

Log message component	Description
Timestamp	Denotes the date and time when the message was logged.
Mnemonic	Denotes which Veritas product logs the message. For Cluster Server, the mnemonic is 'VCS'.
Severity	<p>Denotes the severity of the message. Severity is classified into the following types:</p> <ul style="list-style-type: none"> ■ CRITICAL indicates a critical error within a VCS process. Contact Technical Support immediately. ■ ERROR indicates failure of a cluster component, unanticipated state change, or termination or unsuccessful completion of a VCS action. ■ WARNING indicates a warning or error, but not an actual fault. ■ NOTE informs the user that VCS has initiated an action. ■ INFO informs the user of various state messages or comments. <p>Among these, CRITICAL, ERROR, and WARNING indicate actual errors. NOTE and INFO provide additional information.</p>
UMI or Unique Message ID	<p>UMI is a combination of Originator ID, Category ID, and Message ID. For example, the UMI for a message generated by the SQLServer agent would resemble: V-16-xxxxx-yy. Originator ID for all VCS products is 'V-16.' Category ID for the agents is as follows:</p> <ul style="list-style-type: none"> ■ SQL Server 2012 and 2014: 20093 ■ SQL Server FILESTREAM: 20070 ■ MSTDC: 20021 <p>Message ID is a unique number assigned to the message text.</p>
Message Text	Denotes the actual message string.

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

VCS Cluster Configuration Wizard (VCW) logs

The VCS Cluster Configuration Wizard (VCW) log is located at

`%allusersprofile%\Veritas\Cluster Server\vcw.log`.

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

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

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

A typical wizard log resembles the following:

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

VCWsilent logs

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

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

A typical VCWsilent log resembles the following:

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

NetApp agents error messages

The following table contains a list of error messages for the VCS agents for NetApp.

Table A-2 NetApp agents error messages

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

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

Message	Description
Failed to connect/disconnect virtual disk. (error_type: %s, error_code: 0x%s, error_message: %s)	<p>This could occur because one or more of the following parameters are defined incorrectly in the VCS configuration:</p> <ul style="list-style-type: none"> ■ Filer name ■ Volume name/LUN name ■ Share name ■ Storage IP <p>Verify the configuration definition of the resource. Make sure each attribute is defined correctly.</p>
Unable to create/delete online lock file %s. Error code %s,	Make sure you have write permissions on the specified directory.

Error and warning messages from VCS agent for SQL Server

The VCS agent for SQL Server returns various types of messages as part of its operations. This topic lists the messages of the `ERROR` and the `WARNING` types, further explains those messages, and provides recommended solutions wherever possible.

Messages from the VCS agent for SQL Server Database Engine

The VCS agent for SQL Server Database Engine returns the following messages when it encounters potential or actual issues with the SQL Server service group.

Table A-3 VCS agent for SQL Server 2012 and 2014 error messages

Message	Description
Failed to initialize the SQLServer agent.	The agent failed to initialize the SQLServer agent for SQL Server <i>SQLServerVersion</i> .
Failed to open the SCM handle. Error = <i>Error code</i> .	<p>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.</p> <p>Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.</p>

Table A-3 VCS agent for SQL Server 2012 and 2014 error messages
(continued)

Message	Description
Failed to open the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to open the service from the SCM. Solution: Check whether the service is present in the SCM.
Failed to query the status of the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to query the state of the service. Solution: Check whether the service is present in the SCM.
The service <i>service name</i> is not in stopped state.	The agent is trying to start the service. But the service is in an invalid state. Solution: Check the state of the service.
Failed to set the virtual computer name in the environment of the service <i>service name</i> . Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Veritas Technical Support.
Failed to start the service <i>service name</i> . Error = <i>Error code</i> .	The agent failed to start the service. Solution: Verify if you can start the service from the Windows Services console. If the service starts successfully, stop the service. If the service does not start, see the associated Windows error code for more information.
The service <i>service name</i> did not start within the specified time limit.	The agent failed to start the service within the time limit as specified in the SQLOnlineTimeout attribute. Solution: If the system is slow, you can modify the SQLOnlineTimeout attribute value to accommodate the time that the service takes to start.
Failed to wait for the service <i>service name</i> to start. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Veritas Technical Support.
Sql script failed. Enable DBG1 Logging for the Script output.	The SQL script failed to monitor the SQL Server instance. See the script output for more information.

Table A-3 VCS agent for SQL Server 2012 and 2014 error messages
(continued)

Message	Description
Failed to start the Sql script. (User = <i>user name</i> , Domain = <i>domain name</i>) Error = <i>Error code</i> .	The agent failed to execute the script for detail monitoring. See the associated Windows error code for more information.
Sql script has failed. Error : <i>Error code</i> .	The SQL script for detail monitoring failed. See the associated Windows error code for more information.
Error occurred while getting the process exit code. Error : <i>Error code</i> .	This is a VCS internal error. Solution: Contact Veritas Technical Support.
WaitForSingleObject () failed. Error : <i>Error code</i>	This is a VCS internal error. Solution: Contact Veritas Technical Support.
Failed to convert the argument list. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Veritas Technical Support.
Failed to convert the password attribute to UNICODE. Error = <i>Error code</i> .	This is a VCS internal error. Solution: Contact Veritas Technical Support.
Configuration mismatch: Detail monitoring is configured, but user credentials are not provided.	When configuring detail monitoring for the VCS agent for SQL Server, authentication is required for regardless of the monitoring option (database list or SQL script). Solution: Provide the appropriate user credentials to perform detail monitoring.
Configuration mismatch: Detail monitoring is configured, but neither a list of databases to be monitored nor a SQL monitor script is provided.	When configuring the VCS agent for SQL Server, if you specify that detail monitoring should be performed, you also need to specify the method to be used. Solution: Select either the database list-based or the SQL script-based monitoring option, and provide the appropriate user credentials.

Messages from the VCS agent for SQL Server FILESTREAM

The VCS agent for SQL Server FILESTREAM returns the following messages when it encounters potential or actual issues with the FILESTREAM service group.

Table A-4 Error and warning messages about SQL Server FILESTREAM

Message	Description
Check Filestream is enabled in MSSQL-Configuration Manager if not enable filestream with appropriate enable level [206]	<p>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.</p> <p>Solution: Enable FILESTREAM for that SQL instance and then probe the VCS Filestream resource.</p>
SQLFilestream Resource will be in UNKNOWN [Actual :Offline] State : Filestream Fileshare exists even filestream is disabled [406]	<p>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.</p> <p>Solution: Delete the FILESTREAM fileshares from the command line manually, and then probe the resource.</p>
SQLFilestream Resource is in UNKNOWN[Actual:online] : Filestream Fileshare exists even filestream is enabled for local access only [407]	<p>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.</p> <p>Solution: Delete the FILESTREAM fileshares from the command line manually, probe the resource, take the resource offline and then bring it online.</p>
Filestream will be in offline [Actual:Online] : Filestream Fileshare doesn't exists even filestream is Enabled [409]	<p>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.</p>

Messages from the VCS agent for SQL Server Agent service and Analysis service

VCS uses the GenericService agent to monitor the SQL Server Agent service and Analysis service. The GenericService agent returns the following messages when it encounters potential or actual issues with the Agent service or the Analysis service.

Table A-5 Error and warning messages about the SQL Server Agent service or Analysis service

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

SQL Server Analysis service (MSOLAP) service fails to come online with "invalid context of address" error

The SQL Server Analysis service (MSOLAP) may fail to come online and an "invalid context of address" error may be logged in the Event Viewer.

This issue occurs if there are multiple IP addresses mapped to a single virtual server in a DNS server. In a SQL Server service group, the Lanman resource comes online in the specified virtual server name context. The OLAP resource inherits the virtual server context and communicates with the DNS server to resolve the corresponding virtual IP address. If there are multiple IP addresses mapped to a single virtual server, the DNS server may provide any of these IP addresses. If the resolved IP address is different from the one that was specified during the service group configuration, the resource faults and the service group fails to come online.

Workaround: Set the value of the DNSOption attribute of the Lanman agent to **PurgeDuplicate**. Doing so removes the duplicate DNS entries from the DNS servers.

Messages from the VCS agent for MSDTC

The VCS agent for MSDTC returns the following messages when it encounters potential or actual issues with the MSDTC service group.

Table A-6 Error and warning messages about MSDTC

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

Table A-6 Error and warning messages about MSDTC (*continued*)

Message	Description
MountResName attribute has not been configured.	<p>No value specified for MountResName attribute.</p> <p>Solution: Specify a valid value for the MountResName attribute.</p>
LogPath attribute has not been configured.	<p>No value specified for LogPath attribute.</p> <p>Solution: Specify a valid value for the MountResName attribute.</p>
Failed to open the SCM handle. Error = <i>Error code</i> .	<p>The agent fails to get a handle to the Windows Service Control Manager (SCM). This situation could occur if the specified SCM database does not exist or the requested access is denied.</p> <p>Solution: Verify that SCM can be run on the host. See the associated Windows error code for more information.</p>
Failed to open the MSDTC service. Error = <i>Error code</i> .	<p>The agent failed to open the MSDTC service from the Service Control Manager (SCM).</p> <p>Solution: Check whether the service is present in the Service Control Manager.</p>
Failed to start the MSDTC service. Error = <i>Error code</i> .	<p>The agent failed to start the MSDTC service. See the associated Windows error code for more information.</p>
The MSDTC log path is ' <i>path name</i> '. Configured one is ' <i>path name</i> '.	<p>The specified path for the MSDTC logs is different from the actual path.</p> <p>Solution: Specify the correct MSDTC log path.</p>
The MSDTC service is not in the RUNNING state. Offline might be unsuccessful.	<p>The MSDTC service might be in PAUSE, PAUSE PENDING, or START PENDING state.</p> <p>Solution: Resume the service and then attempt to stop it.</p>
Failed to stop the MSDTC service. Error = <i>Error code</i> .	<p>The MSDTC service cannot be stopped. See the associated Windows error code for more information.</p>
Failed to wait for the MSDTC service to stop. Error = <i>Error code</i> .	<p>The agent cannot stop the service within the specified time limit of 20 seconds. See the associated Windows error code for more information.</p>

Troubleshooting application monitoring configuration issues

This section lists common troubleshooting scenarios that you may encounter while or after configuring application monitoring.

Running the 'hastop –all' command detaches virtual disks

The 'hastop –all' command takes offline all the components and components groups of a configured application, and then stops the VCS cluster. In the process, the command detaches the virtual disks from the VCS cluster nodes. (2920101)

Workaround: If you want to stop the VCS cluster (and not the applications running on cluster nodes), instead of the “hastop –all”, use the following command:

```
hastop -all -force
```

This command stops the cluster without affecting the virtual disks attached to the VCS cluster nodes.

Validation may fail when you add a failover system

On the Configuration Inputs panel, when you add a failover system using the **Add System** option, you may see the following error message:

```
Validation has failed for <System Name>.
```

Workaround: Verify the following:

- The Veritas High Availability guest components are installed on the system.
- The operating system running on the system is supported by this product.
- The specified system name or IP address is valid and the system is switched on and accessible over the network.
- The firewall settings allow access on port 5634 used by the Storage Foundation Messaging Service.
- If xprtd daemon is running on the system, restarting the xprtd service may resolve the issue.

Adding a failover system may fail if you configure a cluster with communication links over UDP

When you configure a cluster with communication links over UDP and add a failover system, you may see the following error message:

The same network adapter is specified for one or more links on the system. You must select a different network adapter for each communication link.

This issue occurs when you unplumb an IP address from one of the configured communication links. In this scenario, VCS may also go into jeopardy.

Workaround:

Plumb the IP address again.

Troubleshooting Veritas High Availability view issues

This section lists common troubleshooting scenarios that you may encounter when using the **Veritas High Availability** tab/dashboard.

Veritas High Availability tab not visible from a cluster node

If you click a system in the inventory view of the VMware vSphere Client GUI, then the **Veritas High Availability** tab displays the cluster view (consolidated cluster-level health information of the configured applications running on the selected system). In some multi-node cluster, the view is not visible from at least one of the cluster nodes.

This behavior occurs if connectivity of the configured LLT links fail. This may be a networking error. (2863649)

Workaround

Ensure that valid LLT links are configured for the affected cluster node, and then retry.

Veritas High Availability tab does not display the application monitoring status

The **Veritas High Availability** tab in the vSphere Client console may either display a HTTP 404 Not Found error or may not show the application health status at all.

Verify the following conditions and then refresh the **Veritas High Availability** tab in the vSphere Client console:

- Verify that the Veritas High Availability Console host is running and is accessible over the network.
- Verify that the VMware Web Service is running on the vCenter Server.

- Verify that the VMware Tools Service is running on the guest virtual machine.
- Verify that the Veritas Storage Foundation Messaging Service (xprtld process) is running on the Veritas High Availability Console and the virtual machine. If it is stopped, type the following on the command prompt:

```
net start xprtld
```
- Verify that ports 14152, 14153, and 5634 are not blocked by a firewall.
- Log out of the vSphere Client and then login again. Then, verify that the Veritas High Availability plugin is installed and enabled.

Veritas High Availability tab may freeze due to special characters in application display name

For a monitored application, if you specify a display name that contains special characters, one or both of the following symptoms may occur:

- The **Veritas High Availability** tab may freeze
- The **Veritas High Availability** tab may display an Adobe exception error message. Based on your browser settings, the Adobe exception message may or may not appear. However, in both cases the tab may freeze. (2923079)

Workaround: Reset the display name using only those characters that belong to the following list:

- any alphanumeric character
- space
- underscore

Use the following command to reset the display name:

```
hagrp -modify sg name UserAssoc -update Name modified display name  
without special characters
```

Veritas High Availability view may fail to load or refresh

The Veritas High Availability view displays health information of monitored applications in a VCS cluster. At times, the view display may fail to load or fail to refresh itself after the default interval of 60 seconds. (2932028)

Workaround:

Restarting the xprtld service may resolve the issue.

To restart the xprtld service

1 Stop the xprtld service:

```
# /etc/init.d/xprtld stop
```

2 Ensure that xprtld is stopped:

```
# ps -ef | grep xprtld
```

If the services is not stopped, terminate the process:

```
# kill -9 xprtld_pid_value
```

Where pid is the process ID of the xprtld process.

3 Start xprtld service:

```
# /etc/init.d/xprtld start
```

Operating system commands to unmount resource may fail

If a user configures a mount point as component (resource) while configuring high availability for an application, then unconfiguring the application from the Veritas High Availability view may lock the mount point. Operating system commands to unmounts the resource may fail. (3574657)

Workaround:

Ensure that you stop the application before you unconfigure the application by clicking More > Unconfigure Application Monitoring in the Veritas High Availability view.

Using the virtual MMC viewer

This appendix includes the following topics:

- [About using the virtual MMC viewer](#)
- [Viewing DTC transaction information](#)

About using the virtual MMC viewer

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

Viewing DTC transaction information

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

Prerequisites for viewing DTC transaction information are as follows:

- An MSDTC service group must be configured and online in the cluster.
- MSDTC client must be configured on the nodes on which you wish to view the transactions.

- The MSDTC service group must be online on the node where you run the VCS Application Manager utility.

To view transactions from a node where MSDTC resource is online

- 1 Start the VCS Application Manager utility.

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

The VCS Application Manager displays a list of supported application service groups configured in the cluster. For each service group it also displays the state of the service group, the name of the virtual server resource (Lanman resource) and the corresponding management tools used for that application.

- 2 Select **MSDTC** from the Select the resource type drop-down list.
- 3 Select the MSDTC resource that is online and then click **Manage**, or double-click the MSDTC resource name.

VAM launches the Component Services snap-in in the virtual server context.

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

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

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

To view transactions from any node in the domain

- 1 Launch the Windows Component Services Administrative tool.
Click **Start > All Programs > Administrative Tools > Component Services**.
- 2 In the console tree of the Component Services administrative tool, double-click **Component Services**, right-click **Computers**, click **New > Computer**.
- 3 In the Add Computer dialog box, specify the virtual server name that you specified while creating the MSDTC Server service group. If you are unsure of the exact name, click **Browse** to search from a list of all computers on the network and select the virtual computer name from the list.
- 4 Click **OK**. The virtual computer entry is added to the Computers container.

- 5** Expand the newly added virtual computer entry and double-click **Distributed Transaction Coordinator**.
- 6** Click **Transaction List** to view all transactions, their status, and their identifiers. Right-click a transaction and click **View > Properties** to list the parent transaction and its children.
- 7** Click **Transaction Statistics** to view statistical information about the transactions in which a server participated.

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