

Storage Foundation and High Availability Solutions

7.4.1 HA and DR Solutions

Guide for Microsoft SQL Server - Windows

Last updated: 2019-02-01

Legal Notice

Copyright © 2019 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:

infoscaledocs@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	Getting started with Storage Foundation and High Availability Solutions for SQL Server	13
Chapter 1	Introducing SFW HA and the VCS agents for SQL Server	14
	About the Veritas InfoScale solutions for monitoring SQL Server	15
	How application availability is achieved in a physical environment	15
	How is application availability achieved in a VMware virtual environment	16
	How the VMwareDisks agent communicates with the vCenter Server instead of the ESX/ESXi host	18
	Typical VCS cluster configuration in a virtual environment	19
	Managing storage using VMware virtual disks	20
	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
	What must be protected in an SQL Server environment	25
	About the VCS agents for SQL Server	26
	About the VCS agent for SQL Server Database Engine	27
	About the VCS agent for SQL Server FILESTREAM	33
	About the VCS GenericService agent for SQL Server Agent service and Analysis service	34
	About the agent for MSDTC service	34
	About the monitoring options	36

Typical SQL Server configuration in a VCS cluster	37
Typical SQL Server disaster recovery configuration	39
SQL Server sample dependency graph	40
MSDTC sample dependency graph	41
Chapter 2 Deployment scenarios for SQL Server	42
Workflows in the Solutions Configuration Center	43
Reviewing the active-passive HA configuration	43
Sample Active-Passive configuration	44
Reviewing the prerequisites for a standalone SQL Server	45
Reviewing a standalone SQL Server configuration	46
Sample standalone SQL Server configuration	47
Reviewing the MSDTC configuration	48
VCS campus cluster configuration	50
Reviewing the campus cluster configuration	52
Campus cluster failover using the ForceImport attribute	53
Reinstating faulted hardware in a campus cluster	55
VCS Replicated Data Cluster configuration	57
Reviewing the Replicated Data Cluster configuration	60
Sample replicated data cluster configuration	60
About setting up a Replicated Data Cluster configuration	61
About setting up replication	62
About configuring and migrating the service group	62
Disaster recovery configuration	65
DR configuration tasks: Primary site	65
DR configuration tasks: Secondary site	67
Supported disaster recovery configurations for service group dependencies	69
Reviewing the disaster recovery configuration	70
Sample disaster recovery configuration	72
Notes and recommendations for cluster and application configuration	73
IPv6 support	76
IP address requirements for an Active-Passive configuration	77
IP address requirements for a disaster recovery configuration	77
Configuring the storage hardware and network	78
Configuring disk groups and volumes for SQL Server	79
About disk groups and volumes	80
Prerequisites for configuring disk groups and volumes	81
Considerations for a fast failover configuration	82

Considerations for converting existing shared storage to cluster disk groups and volumes	83
Considerations when creating disks and volumes for campus clusters	83
Considerations for volumes for a Volume Replicator configuration	85
Considerations for disk groups and volumes for multiple instances	85
Sample disk group and volume configuration	86
MSDTC sample disk group and volume configuration	87
Viewing the available disk storage	87
Creating a dynamic disk group	88
Adding disks to campus cluster sites	89
Creating volumes for high availability clusters	90
Creating volumes for campus clusters	94
About managing disk groups and volumes	98
Importing a disk group and mounting a volume	99
Unmounting a volume and deporting a disk group	100
Adding drive letters to mount the volumes	100
Configuring the cluster using the Cluster Configuration Wizard	102
Configuring notification	111
Configuring Wide-Area Connector process for global clusters	114
Adding nodes to a cluster	116

Chapter 3	Installing SQL Server	120
	About installing and configuring SQL Server	120
	About installing multiple SQL Server instances	122
	Verifying that the SQL Server databases and logs are moved to shared storage	122
	About installing SQL Server for high availability configuration	123
	About installing SQL Server on the first system	125
	About installing SQL Server on additional systems	126
	Creating a SQL Server user-defined database	128
	Completing configuration steps in SQL Server	129
	Moving the tempdb database if using Volume Replicator for disaster recovery	129
	Assigning ports for multiple SQL Server instances	129
	Enabling IPv6 support for the SQL Server Analysis Service	130

Section 2	Configuring SQL Server in a physical environment	131
Chapter 4	Configuring SQL Server for failover	132
	Tasks for configuring a new server for high availability	133
	Tasks for configuring an existing server for high availability	135
	About configuring the SQL Server service group	137
	Service group requirements for Active-Active configurations	138
	Prerequisites for configuring the SQL Server service group	138
	Creating the SQL Server service group	140
	Configuring the service group in a non-shared storage environment	148
	Assigning privileges to the existing SQL Server databases and logs	151
	Enabling fast failover for disk groups (optional)	152
	Verifying the SQL Server cluster configuration	153
	About the modifications required for tagged VLAN or teamed network	154
	Tasks for configuring MSDTC for high availability	155
	Configuring an MSDTC Server service group	156
	Prerequisites for MSDTC configuration	156
	Creating an MSDTC Server service group	157
	About configuring the MSDTC client for SQL Server	160
	About the VCS Application Manager utility	161
	Viewing DTC transaction information	161
	Modifying a SQL Server service group to add VMDg and MountV resources	163
	Determining additional steps needed	164
Chapter 5	Configuring campus clusters for SQL Server	166
	Tasks for configuring campus clusters	166
	Modifying the IP resource in the SQL Server service group	167
	Verifying the campus cluster: Switching the service group	168
	Setting the ForceImport attribute to 1 after a site failure	169
Chapter 6	Configuring Replicated Data Clusters for SQL Server	170
	Tasks for configuring Replicated Data Clusters	170
	Creating the primary system zone for the application service group	172

Creating a parallel environment in the secondary zone	173
Setting up security for Volume Replicator	174
Setting up the Replicated Data Sets (RDS)	177
Prerequisites for setting up the RDS for the primary and secondary zones	177
Creating the Replicated Data Sets with the wizard	178
Configuring a RVG service group for replication	188
Creating the RVG service group	189
Configuring the resources in the RVG service group for RDC replication	191
Configuring the RVG Primary resources	206
Configuring the primary system zone for the RVG service group	209
Setting a dependency between the service groups	209
Adding the nodes from the secondary zone to the RDC	211
Adding the nodes from the secondary zone to the RVG service group	211
Configuring secondary zone nodes in the RVG service group	213
Configuring the RVG service group NIC resource for fail over (VMNSDg only)	213
Configuring the RVG service group IP resource for failover	215
Configuring the RVG service group VMNSDg resources for fail over	217
Adding nodes from the secondary zone to the SQL Server service group	218
Configuring the zones in the SQL Server service group	220
Configuring the application service group IP resource for fail over (VMNSDg only)	221
Configuring the application service group NIC resource for fail over (VMNSDg only)	222
Verifying the RDC configuration	223
Bringing the service group online	223
Switching online nodes	223
Additional instructions for GCO disaster recovery	224

Chapter 7

Configuring disaster recovery for SQL Server

.....	226
Tasks for configuring disaster recovery for SQL Server	227
Tasks for setting up DR in a non-shared storage environment	230
Guidelines for installing Veritas InfoScale Enterprise and configuring the cluster on the secondary site	233

Verifying your primary site configuration	234
Setting up your replication environment	235
Requirements for EMC SRDF array-based hardware replication	235
Requirements for Hitachi TrueCopy array-based hardware replication	237
Assigning user privileges (secure clusters only)	239
About configuring disaster recovery with the DR wizard	240
Configuring disaster recovery with the DR wizard	242
Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)	245
Creating temporary storage on the secondary site using the DR wizard (array-based replication)	249
Installing and configuring SQL Server on the secondary site	254
Cloning the service group configuration from the primary site to the secondary site	254
Configuring the SQL Server service group in a non-shared storage environment	258
Configuring replication and global clustering	259
Configuring Volume Replicator replication and global clustering	259
Configuring EMC SRDF replication and global clustering	267
Configuring Hitachi TrueCopy replication and global clustering	270
Configuring global clustering only	274
Creating the replicated data sets (RDS) for Volume Replicator replication	277
Creating the Volume Replicator RVG service group for replication	277
Configuring the global cluster option for wide-area failover	278
Linking clusters: Adding a remote cluster to a local cluster	279
Converting a local service group to a global service group	280
Bringing a global service group online	282
Verifying the disaster recovery configuration	283
Establishing secure communication within the global cluster (optional)	285
Adding multiple DR sites (optional)	287
Recovery procedures for service group dependencies	287
Configuring DR manually without the DR wizard	289

Chapter 8	Testing fault readiness by running a fire drill	291
	About disaster recovery fire drills	291
	About the Fire Drill Wizard	292
	About Fire Drill Wizard general operations	292
	About Fire Drill Wizard operations in a Volume Replicator environment	293
	About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment	296
	About post-fire drill scripts	298
	Tasks for configuring and running fire drills	299
	Prerequisites for a fire drill	301
	Prerequisites for a fire drill in a Volume Replicator environment	302
	Prerequisites for a fire drill in a Hitachi TrueCopy environment	303
	Prerequisites for a fire drill in an EMC SRDF environment	303
	Preparing the fire drill configuration	304
	System Selection panel details	307
	Service Group Selection panel details	308
	Secondary System Selection panel details	308
	Fire Drill Service Group Settings panel details	308
	Disk Selection panel details	308
	Hitachi TrueCopy Path Information panel details	309
	HTCSnap Resource Configuration panel details	310
	SRDFSnap Resource Configuration panel details	310
	Fire Drill Preparation panel details	311
	Running a fire drill	312
	Re-creating a fire drill configuration that has changed	314
	Restoring the fire drill system to a prepared state	317
	Deleting the fire drill configuration	318
	Fire Drill Deletion panel details	319
	Considerations for switching over fire drill service groups	320

Section 3	Configuring SQL Server in a VMware environment	321
Chapter 9	Configuring application monitoring using the Veritas High Availability solution	322
	Deploying the Veritas High Availability solution for configuring application monitoring	322
	Notes and recommendations	324
	Assigning privileges for non-administrator ESX/ESXi user account	326
	Configuring application monitoring	329
	Configuring the VCS cluster	330
	Configuring the application	333
	Modifying the ESXDetails attribute	338
Chapter 10	Administering application monitoring	340
	About the various interfaces available for performing application monitoring tasks	340
	Administering application monitoring using the Veritas High Availability tab	341
	Understanding the Veritas High Availability tab work area	341
	To view the status of configured applications	344
	To configure or unconfigure application monitoring	344
	To start or stop applications	345
	To suspend or resume application monitoring	347
	To switch an application to another system	347
	To add or remove a failover system	348
	To clear Fault state	352
	To resolve a held-up operation	353
	To determine application state	353
	To remove all monitoring configurations	353
	To remove VCS cluster configurations	354
	Administering application monitoring settings	354
	Administering application availability using Veritas High Availability dashboard	355
	Understanding the dashboard work area	356
	Accessing the dashboard	360
	Monitoring applications across a data center	361
	Monitoring applications across an ESX cluster	361
	Searching for application instances by using filters	361

Selecting multiple applications for batch operations	362
Starting an application using the dashboard	362
Stopping an application by using the dashboard	363
Entering an application into maintenance mode	363
Bringing an application out of maintenance mode	364
Switching an application	365

Appendix A Using Veritas AppProtect for vSphere 366

About Just In Time Availability	367
Prerequisites	371
Setting up a plan	373
Deleting a plan	375
Managing a plan	375
Viewing the history tab	377
Limitations of Just In Time Availability	377
Getting started with Just In Time Availability	378
Supported operating systems and configurations	380
Viewing the properties	381
Log files	381
Plan states	382
Troubleshooting Just In Time Availability	384

Getting started with Storage Foundation and High Availability Solutions for SQL Server

- [Chapter 1. Introducing SFW HA and the VCS agents for SQL Server](#)
- [Chapter 2. Deployment scenarios for SQL Server](#)
- [Chapter 3. Installing SQL Server](#)

Introducing SFW HA and the VCS agents for SQL Server

This chapter includes the following topics:

- [About the Veritas InfoScale solutions for monitoring SQL Server](#)
- [How application availability is achieved in a physical environment](#)
- [How is application availability achieved in a VMware virtual environment](#)
- [Managing storage using VMware virtual disks](#)
- [How VCS monitors storage components](#)
- [What must be protected in an SQL Server environment](#)
- [About the VCS agents for SQL Server](#)
- [About the VCS agent for SQL Server Database Engine](#)
- [About the VCS agent for SQL Server FILESTREAM](#)
- [About the VCS GenericService agent for SQL Server Agent service and Analysis service](#)
- [About the agent for MSDTC service](#)
- [About the monitoring options](#)
- [Typical SQL Server configuration in a VCS cluster](#)
- [Typical SQL Server disaster recovery configuration](#)

- [SQL Server sample dependency graph](#)
- [MSDTC sample dependency graph](#)

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.

How application availability is 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.

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 18.

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 338.

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 1-1 Typical cluster configuration in a VMware virtual environment



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*.

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).

What must be protected in an 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.

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	<p>Monitors the status of SQL Server services. It queries the Windows Service Control Manager (SCM) for the status of the services.</p> <p>If detail monitoring is configured, then depending on the configuration, the agent performs a database health check or executes a monitoring script.</p> <p>See “About the monitoring options” on page 36.</p>
Clean	<p>Forcibly stops the Database Engine service.</p> <p>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.</p>

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 i18nstr IMFRegList[] = { Instance }
    static i18nstr 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 (<code>MSSQLSERVER</code>).
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
Name: SQLOfflineTimeout Type: Integer Dimension: Scalar	Number of seconds that may elapse before the Offline function aborts. Default value: 90

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

Optional attribute	Description
<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 36.</p>

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

Optional attribute	Description
Name: FaultOnDMFailure Type: Boolean Dimension: Scalar	<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>
Name: SQLDetailMonitorTimeout Type: Integer Dimension: Scalar	<p>Number of seconds that may elapse before the detail monitoring routine aborts.</p> <p>Default value: 30</p>
Name: Username Type: String Dimension: Scalar	<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 i18nstr ArgList[] = { "LanmanResName:VirtualName",  
    "MountResName:MountPath", LogPath }  
    str LanmanResName  
    str MountResName  
    i18nstr 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

Table 1-4 MSDTC agent required attributes (*continued*)

Required Attributes	Definition
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	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. <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.

Table 1-5 Methods of configuring detail monitoring (continued)

Method	Description
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> <p><code>%VCS_HOME%\bin\SQLServer\sample_script.sql</code></p> <p>where, <code>%VCS_HOME%</code> points to the default installation directory, typically: <code>C:\Program Files\Veritas\Cluster Server</code>.</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.

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.


Typical SQL Server configuration in a VCS cluster

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

log files. The cluster nodes access the shared storage. The shared storage can be managed using SFW.

The following figure illustrates a two-node cluster hosting a SQL Server service group with the different services configured. MSDTC resource is also configured on the same nodes.

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




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

Typical SQL Server disaster recovery configuration

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

The following figure illustrates a typical SQL Server DR configuration.

Figure 1-3 Typical DR configuration in a VCS cluster



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

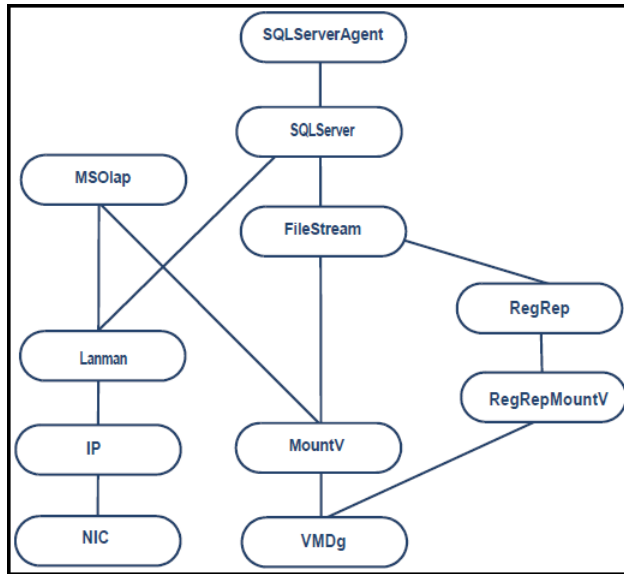
Data is replicated from the primary site to the secondary site. Replication between the storage is set up using a replication software. If the Microsoft SQL Server on System1 fails, SQL Server comes online on node System2 and begins servicing requests. From the user's perspective there might be a small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

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

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 Storage Foundation environment.

Figure 1-4 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:

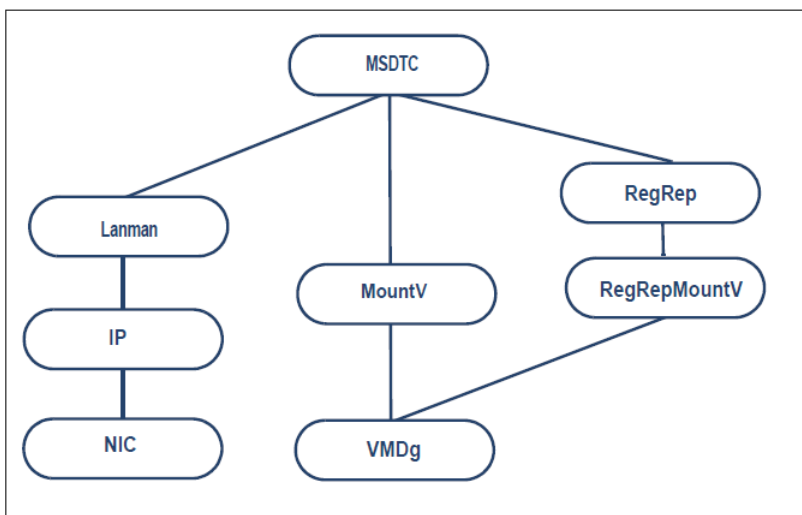
- The virtual name for the SQL Server is configured using the Lanman resource, and the service group IP address is configured using the IP and NIC resources.
- The FileStream resource monitors the Windows FILESTREAM configuration settings for the SQL Server instance.
- The storage mount point is created using the MountV resource, and the shared disk group is configured using the VMDg resource.
- The SQL Server registry is replicated using the RegRep and RegRepMount 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 1-5 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.
- The storage mount point is created using the MountV resource, and the shared disk group is configured using the VMDg resource.
- The MSDTC registry is replicated using the RegRep and RegRepMountV resources.

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

Deployment scenarios for SQL Server

This chapter includes the following topics:

- [Workflows in the Solutions Configuration Center](#)
- [Reviewing the active-passive HA configuration](#)
- [Reviewing the prerequisites for a standalone SQL Server](#)
- [Reviewing a standalone SQL Server configuration](#)
- [Reviewing the MSDTC configuration](#)
- [VCS campus cluster configuration](#)
- [Reviewing the campus cluster configuration](#)
- [VCS Replicated Data Cluster configuration](#)
- [Reviewing the Replicated Data Cluster configuration](#)
- [About setting up a Replicated Data Cluster configuration](#)
- [Disaster recovery configuration](#)
- [Reviewing the disaster recovery configuration](#)
- [Notes and recommendations for cluster and application configuration](#)
- [Configuring the storage hardware and network](#)
- [Configuring disk groups and volumes for SQL Server](#)
- [About managing disk groups and volumes](#)

- [Configuring the cluster using the Cluster Configuration Wizard](#)


Workflows in the Solutions Configuration Center

During the multi-step High Availability Configuration workflow, you may find it helpful to run an SFW HA client on another system and leave the Solutions Configuration Center open on that system. In this way, you can see what step comes next, drill down to the information about that step, and access the online help if needed. You can also print the online help topics and the documentation in PDF format.

When setting up a site for disaster recovery, you first follow the steps under High Availability (HA) Configuration and then continue with the steps under Disaster Recovery Configuration.

The following figure shows the high-level overview of the workflow steps for configuring high availability for SQL Server from the Solutions Configuration Center.


Figure 2-1 Workflow for configuring high availability for SQL Server



Reviewing the active-passive HA configuration

In a typical example of a high availability cluster, you create a virtual SQL Server in an Active-Passive 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 2-2 Active-Passive configuration



SQL Server is installed on both SYSTEM1 and SYSTEM2 and configured as a virtual server (INST1-VS) with a virtual IP address. The SQL Server databases are configured on the shared storage on volumes contained in cluster disk groups. The virtual SQL Server is configured to come online on SYSTEM1 first. If SYSTEM1 fails, SYSTEM2 becomes the active node and the virtual SQL Server comes online on SYSTEM2.

The virtual SQL Server is online on SYSTEM1, serving client requests. The shared 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 Active-Passive configuration

A sample setup is used to illustrate the installation and configuration tasks for an Active-Passive configuration.

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

Table 2-1 Active-Passive configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	servers
INST1_DG	cluster disk group
INST1_DATA_FILES	volume for SQL Server system data files
INST1_DB1_VOL	volume for SQL Server user-defined database
INST1_DB1_LOG	volume for 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_FS_VOL	volume that contains FILESTREAM enabled data objects
SQL_CLUS1	SQL Server cluster
INST1	SQL Server instance
INST1-VS	SQL Server virtual server
INST1_SG	SQL Server service group

Reviewing the prerequisites for a standalone SQL Server

This is applicable if you are configuring an existing standalone SQL Server for high availability.

Review the following requirements before you begin the process of installing the product and creating a clustered environment:

- Create a backup of the data on the existing standalone SQL Server.
- Set all SQL Server services to manual start, except for the SQL Server Browser service. Ensure that the SQL Server Browser service is set to automatic.

Refer to the SQL Server documentation for instructions.


Reviewing a standalone SQL Server configuration

You can incorporate an existing standalone SQL Server into a high availability environment in order to ensure that the mission critical SQL Server resource is always available.

This section describes the tasks necessary to create a virtual server in an Active-Passive SQL Server 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.

The following figure shows how the environment will look at the end of the configuration process.

Figure 2-3 Active-Passive configuration



The virtual SQL Server is online on SYSTEM1, serving client requests. The shared 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 standalone SQL Server configuration

A sample setup is used to illustrate the installation and configuration tasks for creating a high availability environment for a standalone SQL Server.

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

- Virtual SQL Server
The IP address should be the same on all nodes.
- Cluster IP Address

For an IPv4 network, you should obtain all required IP addresses before beginning configuration. For an IPv6 network, IP addresses are generated during configuration.

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

Table 2-2 Standalone SQL Server configuration objects

Object Name	Description
SYSTEM1 & SYSTEM2	server names; SYSTEM1 is the existing standalone SQL Server
INST1_SG	Microsoft SQL Server service group
SQL_CLUS1	virtual SQL Server cluster
INST1_DG	cluster disk group
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_FS_VOL	volume that contains FILESTREAM enabled data objects
INST1	SQL Server Instance Name
INST1-VS	name of the virtual SQL Server

Reviewing the MSDTC configuration

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

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. The MSDTC client and the MSDTC server must not run on the same cluster node.


For example, consider a SQL Server configuration in a VCS cluster that spans four nodes and two sets of shared storage. The shared storage is managed using Storage Foundation (SFW).

The following configurations are possible:

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

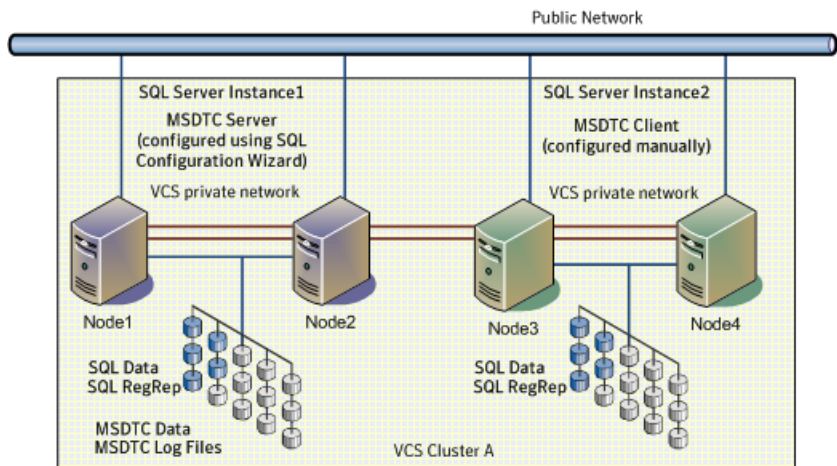
The following figure shows the configuration for SQL Server and MSDTC Server on different nodes.

Figure 2-4 MSDTC Server and SQL Server configured on different nodes



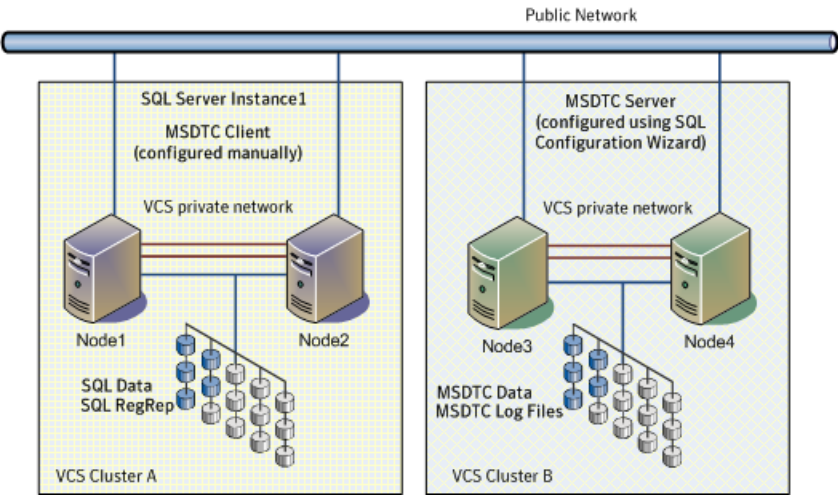
The following figure shows the configuration for SQL Server configured on the same node as the MSDTC server.

Figure 2-5 MSDTC Server and SQL Server configured on same nodes



The following figure shows the configuration where SQL Server and MSDTC Server are configured on nodes belonging to different clusters.

Figure 2-6 MSDTC Server and SQL Server configured in different clusters



VCS campus cluster configuration

You can configure a new Storage Foundation and High Availability Solutions environment for your application in a campus cluster configuration. A campus cluster environment provides high availability and disaster recovery that extends beyond local clustering and mirroring at a single site, but is not as complex as SFW HA DR solution with replication.

Veritas recommends using the Solutions Configuration Center as a guide for installing InfoScale Enterprise and configuring SFW HA for your application.

See [“Workflows in the Solutions Configuration Center”](#) on page 43.

The following table outlines the high-level tasks to complete each objective for a campus cluster configuration for your application.

Table 2-3 Task list: SQL Server campus cluster configuration

Action	Description
Review the campus cluster configuration	<ul style="list-style-type: none">Understand active-passive configurationReview the sample configuration See “Reviewing the campus cluster configuration” on page 52.

Table 2-3 Task list: SQL Server campus cluster configuration (*continued*)

Action	Description
Configure storage hardware and network	<ul style="list-style-type: none">■ Set up the network and storage for a cluster environment■ Verify the DNS entries for the systems on which the application will be installed <p>See “Configuring the storage hardware and network” on page 78.</p>
Review pre-requisites and install InfoScale Enterprise	<ul style="list-style-type: none">■ Install InfoScale Enterprise on all the systems where you want to configure SQL for high availability. <p>Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p>
Review application-specific requirements	<p>See “Notes and recommendations for cluster and application configuration” on page 73.</p>
Configure disk groups and volumes for SQL Server	<ul style="list-style-type: none">■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA)■ Add disks to campus cluster sites to enable site allocation.■ Create dynamic volumes for the SQL system database, user databases and transaction logs using the VEA <p>See “Configuring disk groups and volumes for SQL Server” on page 79.</p>
Configure the VCS cluster	<ul style="list-style-type: none">■ Verify static IP addresses and name resolution configured for each node■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 102.</p>
Install and configure SQL Server on the first cluster node	<ul style="list-style-type: none">■ Ensure that the disk group and volumes for the system database are mounted on the first node■ Follow the guidelines for installing SQL Server in the SFW HA environment <p>See “About installing and configuring SQL Server” on page 120.</p>
Install and configure SQL Server on the second or additional failover nodes	<ul style="list-style-type: none">■ Stop SQL Server services on the first node■ Follow the guidelines for installing SQL Server on failover nodes in the SFW HA environment <p>See “About installing SQL Server on additional systems” on page 126.</p>
Create a SQL user defined database	<ul style="list-style-type: none">■ Create volumes, if not created already, for a user-defined database and transaction log■ Create a user-defined database in SQL Server <p>See “Creating a SQL Server user-defined database” on page 128.</p>

Table 2-3Task list: SQL Server campus cluster configuration (*continued*)


Action	Description
Create a SQL service group	<ul style="list-style-type: none">■ Create a SQL Server service group using the SQL Server Configuration Wizard■ Bring the service group online on the node where you ran the wizard to configure the service group. This is first cluster node where you installed SQL Server. <p>See “About configuring the SQL Server service group” on page 137.</p>
Modify the Address and SubNetMask attributes if the sites are in different subnets	<p>Modify the Address and SubNetMask attributes of the IP resource if the sites are in different subnets.</p> <p>See “Modifying the IP resource in the SQL Server service group” on page 167.</p>
Set the ForceImport attribute of the VMDg resource as per the requirement	<p>If a site failure occurs, set the ForceImport attribute of the VMDg resource to 1 to ensure proper failover</p> <p>See “Setting the ForceImport attribute to 1 after a site failure” on page 169.</p>

Reviewing the campus cluster configuration

A sample campus cluster configuration is a two-node campus cluster with each node in a separate site (Site A or Site B). In this example, each node has its own storage array with the same number of disks and contains mirrored data of the storage on the other array.

The following figure illustrates an active-passive configuration with one to one failover capabilities.

Figure 2-7 Campus cluster: Active-Passive configuration



In an Active-Passive configuration, the active node of the cluster hosts the virtual SQL 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. In this case, the virtual SQL Server can fail over from SYSTEM1 to SYSTEM2 and vice versa.

The two nodes can be located miles apart and are connected via a single subnet and Fibre Channel SAN. Each node has its own storage array with an equal number of disks and contains mirrored data of the storage on the other array. The example describes a generic database application.

Plan for an equal number and size of disks on the two sites, because each disk group should contain the same number of disks on each site for the mirrored volumes.

Campus cluster failover using the ForceImport attribute

Automated recovery is handled differently in a VCS campus cluster than with a VCS local cluster. The outcomes of failure situations depend on the settings for the

ForceImport attribute of the VMDg resource. To ensure proper failover in a VCS campus cluster, you must verify the value of the ForceImport attribute.

You can set this attribute as follows:

- ForceImport set to 1 automatically forces the import of the disk groups to the other node
- ForceImport set to 0 does not force the import

The advantage of automatic failover in the event of site failure comes at the cost of potential data loss if all storage and network communication paths between the sites are severed. Choose an option that is suitable given your cluster infrastructure, uptime requirements, and administrative capabilities.

You can use the VCS Java Console or command line to modify the ForceImport attribute. For more information on modifying ForceImport:

See [“Setting the ForceImport attribute to 1 after a site failure”](#) on page 169.

The following table lists failure situations and the outcomes depending on the settings for the ForceImport attribute of the VMDg resource.

Table 2-4 Failure situations in a VCS campus cluster

Failure Situation	ForceImport set to 0 (import not forced)	ForceImport set to 1 (automatic forced import)
1) Application fault May mean the services stopped for an application, a NIC failed, or a database table went offline.	Application automatically moves to another node.	Service Group failover is automatic on the standby or preferred system or node. Service Group failover is automatic on the standby or preferred system or node.
2) Server failure May mean a power cord became unplugged or a failure caused the system to stop responding.	Application automatically moves to other node. 100% of the disks are still available.	Service Group failover is automatic on the standby or preferred system or node. 100% of the mirrored disks are still available.
3) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.	No interruption of service. Remaining disks in mirror are still accessible from the other node.	The Service Group does not failover. 50% of the mirrored disk is still available at remaining site.

Table 2-4 Failure situations in a VCS campus cluster (*continued*)

Failure Situation	Forcelmport set to 0 (import not forced)	Forcelmport set to 1 (automatic forced import)
<p>4) Zone failure</p> <p>Complete Site failure, all accessibility to the servers and storage is lost.</p>	<p>Manual intervention required to online the Service Group at remaining site. Can not automatically import 50% of mirrored disk.</p>	<p>Automatic failover of Service Group to online site. Force Import must be set to True before site failure to ensure VCS can import 50% of mirrored disk.</p>
<p>5) Split-brain (loss of both heartbeats)</p> <p>If the public network link serves as a low-priority heartbeat, the assumption is made that the link is also lost.</p>	<p>No interruption of service. Can't import disks because the original node still has the SCSI reservation.</p>	<p>No interruption of service. Failover does not occur due to Service Group resources remaining online on the original nodes. Example: Online node has SCSI reservation to own disk.</p>
<p>6) Storage interconnect lost</p> <p>Fibre interconnect severed.</p>	<p>No interruption of service. Disks on the same node are functioning. Mirroring is not working.</p>	<p>No interruption of service. Service Group resources remain online, but 50% of the mirror disk becomes detached.</p>
<p>7) Split-brain and storage interconnect lost</p> <p>If a single pipe is used between buildings for the Ethernet and storage, this situation can occur.</p>	<p>No interruption of service. Cannot import with only 50% of disks available. Disks on the same node are functioning. Mirroring is not working.</p>	<p>Automatically imports 50% of mirrored disk to the alternate node.</p> <p>Disks online for a short period in both locations but offlined again due to IP and other resources being online on original node. No interruption of service.</p>

Reinstating faulted hardware in a campus cluster

Once a failure occurs and an application is migrated to another node or site, it is important to know what will happen when the original hardware is reinstated.

The following table lists the behavior when various hardware components affecting the configuration (array or disks, site hardware, networking cards or cabling, storage interconnect, etc.) are reinstated after failure.

Table 2-5 Behavior exhibited when hardware is reinstated

Failure Situation, before Reinstating the Configuration	ForcelImport set to 0 (import not forced)	ForcelImport set to 1 (automatic force import)
3) Failure of disk array or all disks Remaining disks in mirror are still accessible from the other site.	No interruption of service. Resync the mirror from the remote site.	Same behavior.
4) Site failure All access to the server and storage is lost.	Inter-node heartbeat communication is restored and the original cluster node becomes aware that the application is online at the remote site. Resync the mirror from the remote site	Same behavior.
5) Split-brain situation (loss of both heartbeats)	No interruption of service.	Same behavior.
6) Storage interconnect lost Fibre interconnect severed.	No interruption of service. Resync the mirror from the original site.	Same behavior.
7) Split-brain situation and storage interconnect lost	No interruption of service. Resync the mirror from the original site.	VCS alerts administrator that volumes are online at both sites. Resync the mirror from the copy with the latest data.

The numbers 3 through 7 in the previous table refer to the scenarios in [Campus cluster failover using the ForcelImport attribute](#).

Situations 1 and 2 have no effect when reinstated. Keep in mind that the cluster has already responded to the initial failure.

While the outcomes of using both settings of the ForcelImport attribute for most scenarios are the same, the ForcelImport option provides automatic failover in the event of site failure. This advantage comes at the cost of potential data loss if all storage and network communication paths between the sites are severed. Choose an option that is suitable given your cluster infrastructure, uptime requirements, and administrative capabilities.

VCS Replicated Data Cluster configuration

The configuration process for a Replicated Data Cluster configuration includes the following main stages:

- Configure the SFW HA and the application components for high availability on the primary zone nodes.
- Install InfoScale Enterprise and configure SFW HA and the application components on the secondary zone.
- Configure the Volume Replicator components for both zones.
Refer to the *Volume Replicator Administrator's Guide* for additional details on Volume Replicator.

The following table outlines the high-level tasks to complete each objective for a Replicated Data Cluster configuration for your application.

Table 2-6 Process for deploying a Replicated Data Cluster

Action	Description
Understand the configuration	<ul style="list-style-type: none">■ Understand active-passive configuration and zone failover in a RDC environment■ Review the sample configuration See “Reviewing the Replicated Data Cluster configuration” on page 60.
Configure the storage hardware and network	For all nodes in the cluster: <ul style="list-style-type: none">■ Set up the storage hardware for a cluster environment■ Verify the DNS entries for the systems on which the application will be installed See “Configuring the storage hardware and network” on page 78.
Review pre-requisites and install InfoScale Enterprise	<ul style="list-style-type: none">■ Install InfoScale Enterprise on all the systems where you want to configure SQL for high availability. Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.
Review application-specific requirements	See “Notes and recommendations for cluster and application configuration” on page 73.
Configure cluster disk groups and volumes for SQL Server	<ul style="list-style-type: none">■ Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) (if using a shared storage configuration)■ Create dynamic disk groups using VEA (for a single-node configuration using non-shared storage)■ Create dynamic volumes for the SQL system database, registry replication, user databases and transaction logs using the VEA See “Configuring disk groups and volumes for SQL Server” on page 79.

Table 2-6 Process for deploying a Replicated Data Cluster *(continued)*

Action	Description
Configure the cluster	<ul style="list-style-type: none"> ■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication in the cluster <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 102.</p>
Install and configure SQL Server on the first cluster node	<ul style="list-style-type: none"> ■ Ensure that the disk group and volumes for the system database are mounted on the first node ■ Follow the guidelines for installing SQL Server in the SFW HA environment <p>See “About installing and configuring SQL Server” on page 120.</p>
Install and configure SQL Server on the failover node(s) of the primary zone	<ul style="list-style-type: none"> ■ Stop SQL Server services on the first node ■ Follow the guidelines for installing SQL Server on failover nodes in the SFW HA environment <p>See “About installing SQL Server on additional systems” on page 126.</p>
Create a SQL Server user-defined database	<ul style="list-style-type: none"> ■ If not done earlier, create volumes for a user-defined database and transaction log ■ Create a new user-defined database in SQL Server <p>See “Creating a SQL Server user-defined database” on page 128.</p>
Create a SQL Server service group	<ul style="list-style-type: none"> ■ Ensure that you have met the prerequisites ■ Create the application service group manually using templates from the Cluster Manager (Java Console) (if using a non-shared storage configuration) ■ Create a SQL Server service group using the SQL Server Configuration Wizard ■ Bring the service group online on the node where you ran the wizard to configure the service group. This is first cluster node where you installed SQL Server. <p>See “About configuring the SQL Server service group” on page 137.</p>
Create the primary system zone	<ul style="list-style-type: none"> ■ Create the primary system zone ■ Add the nodes to the primary zone <p>See “Creating the primary system zone for the application service group” on page 172.</p>
Verify failover within the primary zone	<p>Test failover between the nodes in the primary zone</p> <p>See “Verifying the SQL Server cluster configuration” on page 153.</p>

Table 2-6 Process for deploying a Replicated Data Cluster *(continued)*

Action	Description
Create a parallel environment in the secondary zone	<ul style="list-style-type: none"> ■ Install InfoScale Enterprise on the systems in the secondary zone ■ Configure disk groups and volumes using the same names as on the primary zone ■ Install your application following the prerequisites and guidelines for installing on the second zone. <p>See “Creating a parallel environment in the secondary zone” on page 173.</p>
Add the secondary zone systems to the cluster	Add the secondary zone systems to the cluster.
Set up security for Volume Replicator on all cluster nodes	<p>Set up security for Volume Replicator on all nodes in both zones.</p> <p>This step can be done at any time after installing InfoScale Enterprise on all cluster nodes, but must be done before configuring Volume Replicator replication.</p> <p>See “Setting up security for Volume Replicator” on page 174.</p>
Set up the Replicated Data Set	<p>Use the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones</p> <p>See “Setting up the Replicated Data Sets (RDS)” on page 177.</p>
Configure a RVG service group	<ul style="list-style-type: none"> ■ Create a Replicated Volume Group (RVG) service group ■ Configure the RVG service group <p>See “Configuring a RVG service group for replication” on page 188.</p>
Set a dependency between the service groups	<ul style="list-style-type: none"> ■ Set up a dependency from the Volume Replicator RVG service group to the SQL Server service group <p>See “Setting a dependency between the service groups” on page 209.</p>
Add the nodes from the secondary zone to the RDC	<ul style="list-style-type: none"> ■ Add the nodes from the secondary zone to the RVG service group ■ Configure the IP resources for failover ■ Add the nodes from the secondary zone to the SQL Server service group <p>See “Adding the nodes from the secondary zone to the RDC” on page 211.</p>
Verify the RDC configuration	<p>Verify that failover occurs first within zones and then from the primary to the secondary zone</p> <p>See “Verifying the RDC configuration” on page 223.</p>

Reviewing the Replicated Data Cluster configuration

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

- Virtual SQL Server
The IP address should be the same on all nodes at the primary and secondary zones.
- Replication IP address for the primary zone
- Replication IP address for the secondary zone

You should have these IP addresses available before you start deploying the RDC environment (for an IPv4 network, you will need to specify the addresses; for an IPv6 network, they are generated).

Sample replicated data cluster configuration

The sample setup for a Replicated Data Cluster has four servers, two for the primary zone and two for the secondary zone. The nodes will form two separate clusters, one at the primary zone and one at the secondary zone.

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

Table 2-7 Replicated Data Cluster configuration objects

Object Name	Description
Primary zone	
SYSTEM1 & SYSTEM2	First and second nodes of the primary zone
INST1_SG	Microsoft SQL Server service group
INST1-VS	Virtual SQL Server cluster
INST1	SQL Server instance name
INST1_DG	cluster disk group for SQL Server system database and files
INST1_DATA_FILES	volume for Microsoft SQL Server system data files
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_FS_VOL	volume that contains FILESTREAM enabled data objects

Table 2-7 Replicated Data Cluster configuration objects (*continued*)

Object Name	Description
INST1_REPLOG	Replicator log volume required by Volume Replicator
INST1_DB1_DG	Cluster disk group for SQL Server user-defined database and 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_DB1_REPLOG	Replicator log volume required by Volume Replicator for SQL Server user-defined database
Secondary zone	
SYSTEM3 & SYSTEM4	First and second nodes of the secondary zone
	All the other parameters are the same as on the primary zone.
RDS and VVR Components	
INST1_RDS	RDS name for SQL Server system database and files
INST1_RVG	RVG name for SQL Server system database and files
INST1_RVG_SG	Replication service group for SQL Server system database and files
INST1_DB1_RDS	RDS name for SQL Server user-defined database and files
INST1_DB1_RVG	RVG name for SQL Server user-defined database and files
INST1_DB1_RVG_SG	Replication service group for SQL Server user-defined database and files

About setting up a Replicated Data Cluster configuration

In the example, SQL Server is configured as a VCS service group in a four-node cluster, with two nodes in the primary RDC zone and two in the secondary RDC zone. In the event of a failure on the primary node, VCS can fail over the SQL Server instance to the second node in the primary zone.

The process involves the steps described in the following topics:

- See [“About setting up replication”](#) on page 62.
- See [“About configuring and migrating the service group”](#) on page 62.

About setting up replication

Use Volume Replicator to set up replication between the disk groups in the RDC primary and secondary zones.

Note the following:


- Use Volume Replicator to group the data volumes into a Replicated Volume Group (RVG). Create a primary RVG on the hosts in the first zone (primary zone) and create a secondary RVG on hosts in the second zone (secondary zone).
- If using shared storage, the primary RVG consists of volumes shared between the cluster nodes in the primary zone and the secondary RVG consists of volumes shared between the cluster nodes in the secondary zone.
- If using non-shared storage, the primary RVG consists of volumes created on the local disks on the node in the primary zone and the secondary RVG consists of the volumes created on the local disks on the node in the secondary zone.
- Create a Replicated Data Set (RDS) with the primary RVG and the secondary RVG.
- Use the same disk group name and RVG name in both the zones so that the MountV resources are able to mount the same block devices.

About configuring and migrating the service group

For a successful wide-area failover, the mount points and applications must fail over to the secondary RDC zone. Additionally, the Volume Replicator secondary disk group and RVG must be imported and started on the secondary RDC zone.


The following figure shows a screen from the VCS Cluster Manager (Java Console) that depicts a typical SQL Server service group RDC configuration.

Figure 2-8 Typical SQL Server service group RDC configuration



The following figure shows a screen from the VCS Cluster Manager (Java Console) that depicts a typical SQL Server replication service group (RVG) configuration.

Figure 2-9 Typical SQL Server replication service group (RVG) configuration



In the RDC configuration, consider a case where the primary RDC zone suffers a total failure of the storage. In this situation, none of the nodes in the primary RDC zone see any device. The service group cannot fail over locally within the primary RDC zone, because the volumes cannot be mounted on any node. The service group must therefore fail over to a node in the current secondary RDC zone.

The RVGPrimary agent ensures that Volume Replicator volumes at the secondary RDC zone are made writable. The application can be started at the secondary RDC zone and run there until the problem with the local storage is corrected. If the storage problem is corrected, you can switch the application back to the primary zone using VCS.

Before switching the application back to the original primary RDC zone, you must resynchronize any changed data from the active secondary RDC zone since the failover. Once the resynchronization completes, switch the service group to the primary zone.

Disaster recovery configuration

You begin by configuring the primary site for high availability. After setting up an SFW HA high availability environment for your application on a primary site, you can create a secondary or “failover” site for disaster recovery.

The Disaster Recovery (DR) wizard helps you to clone the storage and service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

The DR wizard also helps you set up replication and the global clustering (GCO option). You can choose to configure replication using Volume Replicator (Volume Replicator) or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Veritas recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

See [“Workflows in the Solutions Configuration Center”](#) on page 43.

To follow the workflow in the Solutions Configuration Center, the disaster recovery workflow has been split into two tables, one covering the steps for configuring high availability at the primary site, and the other covering the steps for completing the disaster recovery configuration at the secondary site.

DR configuration tasks: Primary site

The following table outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the primary site.

Table 2-8 Configuring the primary site for disaster recovery

Action	Description
Understand the configuration	Understand the DR configuration See “Reviewing the disaster recovery configuration” on page 70.
Configure the storage hardware and network	For all nodes in the cluster: <ul style="list-style-type: none">■ Set up the storage hardware for a cluster environment■ Verify the DNS entries for the systems on which the application will be installed See “Configuring the storage hardware and network” on page 78.

Table 2-8 Configuring the primary site for disaster recovery (*continued*)

Action	Description
Review pre-requisites and install InfoScale Enterprise	<ul style="list-style-type: none"> Install InfoScale Enterprise on all the systems where you want to configure SQL for high availability. <p>Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p>
Review application-specific requirements	See “Notes and recommendations for cluster and application configuration” on page 73.
Configure the VCS cluster	<ul style="list-style-type: none"> Verify static IP addresses and name resolution configured for each node Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 102.</p>
Configure cluster disk groups and volumes for SQL Server	<ul style="list-style-type: none"> Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) (if using a shared storage configuration) Create dynamic disk groups using VEA (for a single-node configuration using non-shared storage) Create dynamic volumes for the SQL system database, user databases and transaction logs using the VEA <p>See “Configuring disk groups and volumes for SQL Server” on page 79.</p>
Install and configure SQL Server on the first cluster node	<ul style="list-style-type: none"> Ensure that the disk group and volumes for the system database are mounted on the first node Follow the guidelines for installing SQL Server in the SFW HA environment <p>See “About installing and configuring SQL Server” on page 120.</p>
Install and configure SQL Server on the failover node(s)	<ul style="list-style-type: none"> Stop SQL Server services on the first node Follow the guidelines for installing SQL Server on failover nodes in the SFW HA environment <p>See “About installing SQL Server on additional systems” on page 126.</p>
Create a SQL Server user-defined database	<ul style="list-style-type: none"> If not done earlier, create volumes for a user-defined database and transaction log Create a new user-defined database in SQL Server <p>See “Creating a SQL Server user-defined database” on page 128.</p>

Table 2-8 Configuring the primary site for disaster recovery (*continued*)

Action	Description
Create a SQL Server service group	<ul style="list-style-type: none">■ Ensure that you have met the prerequisites■ Create the application service group manually using templates from the Cluster Manager (Java Console) (if using a non-shared storage configuration)■ Create a SQL Server service group using the SQL Server Configuration Wizard■ Bring the service group online on the node where you ran the wizard to configure the service group. This is first cluster node where you installed SQL Server. <p>See “About configuring the SQL Server service group” on page 137.</p>
Verify the primary site configuration	<p>Test failover between nodes on the primary site</p> <p>See “Verifying the SQL Server cluster configuration” on page 153.</p>

DR configuration tasks: Secondary site

The following table outlines the high-level objectives and the tasks to complete each objective for a DR configuration at the secondary site.

Table 2-9 Configuring the secondary site for disaster recovery

Action	Description
Review pre-requisites, install InfoScale Enterprise, configure the cluster on the secondary site	<p>Warning: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.</p> <ul style="list-style-type: none">■ Install InfoScale Enterprise on all the systems where you want to configure SQL for high availability. <p>Refer to <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p>
Verify that SQL Server has been configured for high availability at the primary site	<p>Verify that your application has been configured for high availability at the primary site and that the service groups are online</p> <p>See “Verifying your primary site configuration” on page 234.</p>

Table 2-9 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
Set up the replication prerequisites	<p>Ensure that replication prerequisites for your selected method of replication are met before running the DR wizard</p> <p>See “Setting up security for Volume Replicator” on page 174.</p> <p>See “Requirements for EMC SRDF array-based hardware replication” on page 235.</p> <p>See “Configuring Hitachi TrueCopy replication and global clustering” on page 270.</p>
(Secure cluster only) Assign user privileges	<p>For a secure cluster only, assign user privileges</p> <p>See “Assigning user privileges (secure clusters only)” on page 239.</p>
Start running the DR wizard	<ul style="list-style-type: none"> ■ Review prerequisites for the DR wizard ■ Start the DR wizard and make the initial selections required for each task: selecting a primary site system, the service group, the secondary site system, and the replication method <p>See “Configuring disaster recovery with the DR wizard” on page 242.</p>
Clone the storage configuration (Volume Replicator replication only)	<p>(Volume Replicator replication option)</p> <p>Clone the storage configuration on the secondary site using the DR wizard</p> <p>See “Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)” on page 245.</p>
Create temporary storage for application installation (other replication methods)	<p>(EMC SRDF, Hitachi TrueCopy, or GCO only replication option)</p> <p>Use the DR wizard to create temporary storage for installation on the secondary site</p> <p>See “Creating temporary storage on the secondary site using the DR wizard (array-based replication)” on page 249.</p>
Install and configure SQL Server on the first cluster node	<ul style="list-style-type: none"> ■ Ensure that the disk group and volumes for the system database are mounted on the first node ■ Follow the guidelines for installing SQL Server in the SFW HA environment <p>See “About installing and configuring SQL Server” on page 120.</p>
Install and configure SQL Server on the failover node(s)	<ul style="list-style-type: none"> ■ Stop SQL Server services on the first node ■ Follow the guidelines for installing SQL Server on failover nodes in the SFW HA environment <p>See “About installing SQL Server on additional systems” on page 126.</p>

Table 2-9 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
Clone the service group configuration	<p>Clone the service group configuration from the primary to the secondary site using the DR wizard</p> <p>See “Cloning the service group configuration from the primary site to the secondary site” on page 254.</p>
Configure replication and global clustering, or configure global clustering only	<ul style="list-style-type: none"> ■ (Volume Replicator replication) Use the wizard to configure replication and global clustering ■ (EMC SRDF replication) Set up replication and then use the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Set up replication and then use the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Use the wizard to configure global clustering, and then set up replication <p>See “Configuring replication and global clustering” on page 259.</p>
Verify the disaster recover configuration	<p>Verify that the secondary site has been fully configured for disaster recovery</p> <p>See “Verifying the disaster recovery configuration” on page 283.</p>
(Optional) Add secure communication	<p>Add secure communication between local clusters within the global cluster (optional task)</p> <p>See “Establishing secure communication within the global cluster (optional)” on page 285.</p>
(Optional) Add additional DR sites	<p>Optionally, add additional DR sites to a Volume Replicator environment</p> <p>See “Adding multiple DR sites (optional)” on page 287.</p>
Handling service group dependencies after failover	<p>If your environment includes dependent service groups, review the considerations for bringing the service groups online after failover to the secondary site</p> <p>See “Recovery procedures for service group dependencies” on page 287.</p>

Supported disaster recovery configurations for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

Service group dependency configurations are described in detail in the VCS documentation.

See the *Cluster Server Administrator’s Guide*.

For disaster recovery only certain dependent service group configurations are supported:

- Online local soft
- Online local firm
- Online local hard

If the service group has an unsupported type of dependency and you select it in the DR wizard, you receive an error notification when you attempt to move to the next wizard page.

In a hardware replication environment, the Disaster Recovery wizard supports one level of dependency (one child). If you need to configure more levels, you will need to add the service group and the dependency link manually on the secondary site after you finish running the DR wizard.

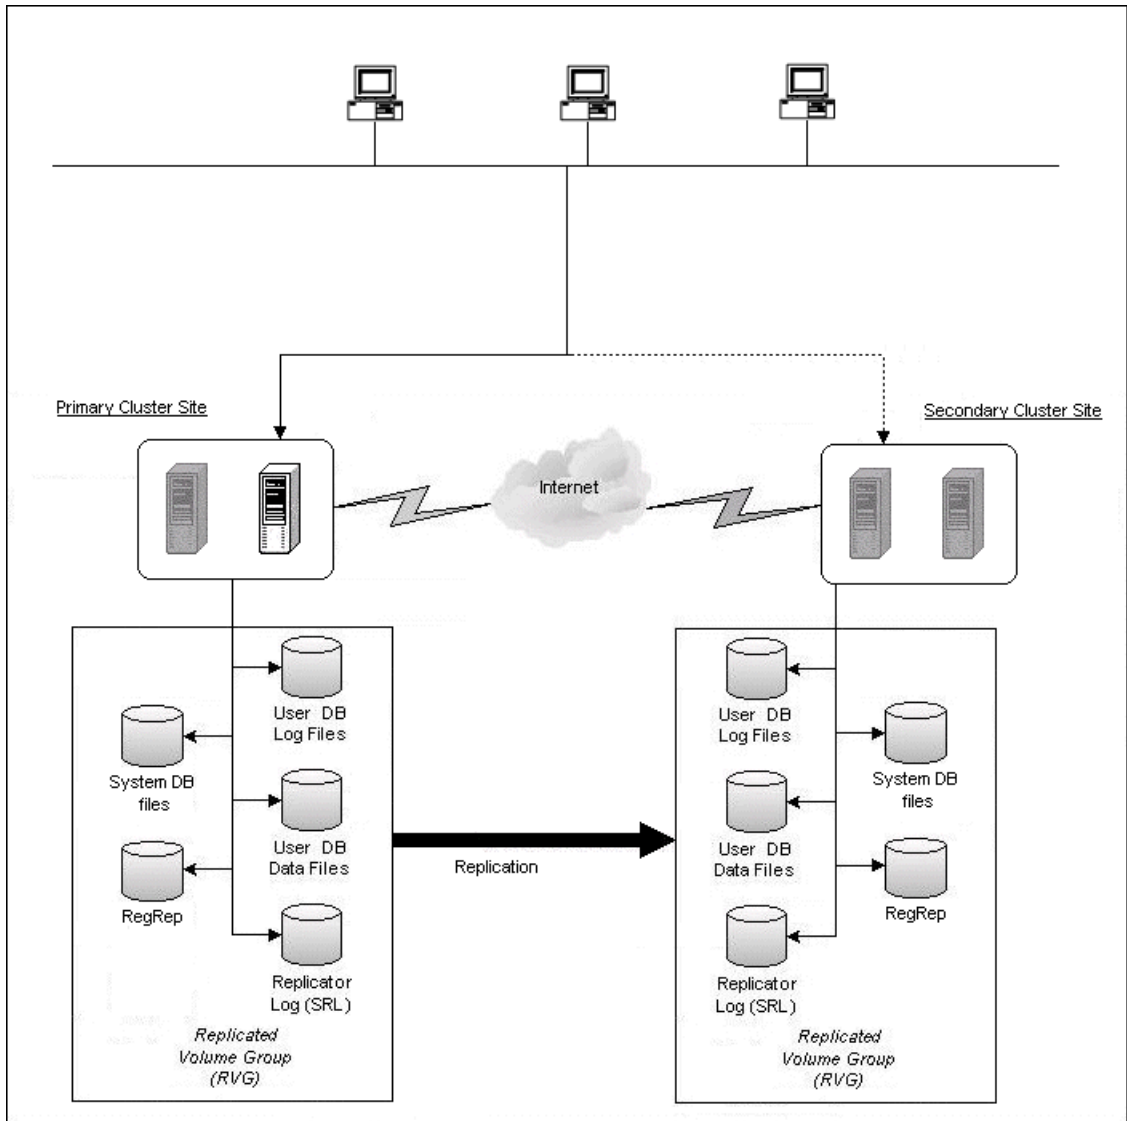
In a Volume Replicator environment, the wizard cannot configure DR for a service group that has a child and you will need to configure the secondary site manually. For more information on configuring Volume Replicator, see the *Volume Replicator Administrator's Guide*. For more information on configuring GCO, see the *Cluster Server Administrator's Guide*.

Reviewing the disaster recovery configuration

You may be preparing to configure both a primary site and a secondary site for disaster recovery.

The following figure illustrates a typical Active-Passive disaster recovery configuration using Volume Replicator (Volume Replicator).

Figure 2-10 Typical Volume Replicator configuration



In the example illustration, the primary site consists of two nodes, SYSTEM1 and SYSTEM2. Similarly the secondary setup consists of two nodes, SYSTEM3 and SYSTEM4. Each site has a clustered setup with the nodes set up appropriately for failover within the site. The cluster on the primary site has a shared disk group that is used to create the volumes required by Volume Replicator for setting up the

Replicated Volume Group (RVG). The Microsoft SQL Server application data is stored on the volumes that are under the control of the RVG.

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. If there is a disaster at the primary site, SYSTEM3 at the secondary site takes over.

You can choose to configure replication using Volume Replicator or an agent-supported array-based hardware replication. You can use the DR wizard to configure Volume Replicator replication or required options for the VCS agents for EMC SRDF or Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

Sample disaster recovery configuration

The sample setup has four servers, two for the primary site and two for the secondary site. The nodes will form two separate clusters, one at the primary site and one at the secondary site.

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

Table 2-10 Sample Disaster Recovery configuration objects

Object Name	Description
Primary site	
SYSTEM1 & SYSTEM2	first and second nodes of the primary site
INST1_SG	Microsoft SQL Server service group
SQL_CLUS1	virtual SQL Server cluster
INST1-VS	virtual server name
INST1_DG	cluster disk group
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

Table 2-10 Sample Disaster Recovery configuration objects (*continued*)

Object Name	Description
INST1_REGREP_VOL	volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server
INST1_FS_VOL	Volume that contains the FILESTREAM enabled data objects for the SQL Server instance
INST1_REPLOG	(Volume Replicator only) replicator log volume required by Volume Replicator
INST1	SQL Server Instance Name
Secondary site	
SYSTEM3 & SYSTEM4	First and second nodes of the secondary site
	All the other parameters are the same as on the primary site.
DR Components (Volume Replicator only)	
INST1_DB1_RDS	RDS Name
INST1_DB1_RVG	RVG Name
INST1_DB1_RVG_SG	Replication service group

Notes and recommendations for cluster and application configuration

- Review the Hardware compatibility list (HCL) and Software Compatibility List (SCL) at:
<https://sort.veritas.com/documents>

Note: Solutions wizards cannot be used to perform Disaster Recovery, Fire Drill, or Quick Recovery remotely on Windows Server Core systems.

The DR, FD, and QR wizards require that the .NET Framework is present on the system where these operations are to be performed. As the .NET Framework is not supported on the Windows Server Core systems, the wizards cannot be used to perform DR, FD, or QR on these systems.

Refer to the following Microsoft knowledge database article for more details:

<http://technet.microsoft.com/en-us/library/dd184075.aspx>

- Refer to the Microsoft documentation for SQL Server memory requirements.
- Shared disks to support applications that migrate between nodes in the cluster. Campus clusters require more than one array for mirroring. Disaster recovery configurations require one array for each site. Replicated data clusters with no shared storage are also supported.

If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA).

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

- SCSI, Fibre Channel, iSCSI host bus adapters (HBAs), or iSCSI Initiator supported NICs to access shared storage.
- A minimum of two NICs is required. One NIC will be used exclusively for private network communication between the nodes of the cluster. The second NIC will be used for both private cluster communications and for public access to the cluster. Veritas recommends three NICs.
- NIC teaming is not supported for the VCS private network.
- Static IP addresses are required for certain purposes when configuring high availability or disaster recovery solutions. For IPv4 networks, ensure that you have the addresses available to enter. For IPv6 networks, ensure that the network advertises the prefix so that addresses are autogenerated.

Static IP addresses are required for the following purposes:

- One static IP address per site for each SQL Virtual Server.
- A minimum of one static IP address for each physical node in the cluster.
- One static IP address per cluster used when configuring Notification or the Global Cluster Option. The same IP address may be used for all options.

- For Volume Replicator replication in a disaster recovery configuration, a minimum of one static IP address per site for each application instance running in the cluster.
- For Volume Replicator replication in a Replicated Data Cluster configuration, a minimum of one static IP address per zone for each application instance running in the cluster.
- Configure name resolution for each node.
- Verify the availability of DNS Services. AD-integrated DNS or BIND 8.2 or higher are supported.
 Make sure a reverse lookup zone exists in the DNS. Refer to the application documentation for instructions on creating a reverse lookup zone.
- DNS scavenging affects virtual servers configured in SFW HA because the Lanman agent uses Dynamic DNS (DDNS) to map virtual names with IP addresses. If you use scavenging, then you must set the DNSRefreshInterval attribute for the Lanman agent. This enables the Lanman agent to refresh the resource records on the DNS servers.
 See the *Cluster Server Bundled Agents Reference Guide*.
- 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 Network Basic Input/Output System (NetBIOS) is disabled over the TCP/IP, then you must set the Lanman agent's DNSUpdateRequired attribute to 1 (True).
- You must have write permissions for the Active Directory objects corresponding to all the nodes.
- If you plan to create a new user account for the VCS Helper service, you must have Domain Administrator privileges or belong to the Account Operators group. If you plan to use an existing user account context for the VCS Helper service, you must know the password for the user account.
- If User Access Control (UAC) is enabled on Windows systems, then you cannot log on to VEA GUI with an account that is not a member of the Administrators group, such as a guest user. This happens because such user does not have the "Write" permission for the "Veritas" folder in the installation directory (typically, C:\Program Files\Veritas). As a workaround, an OS administrator user can set "Write" permission for the guest user using the Security tab of the "Veritas" folder's properties.
- For a Replicated Data Cluster, install only in a single domain.

- Configure Microsoft Exchange Server and Microsoft SQL Server on separate failover nodes within a cluster.
- Route each private NIC through a separate hub or switch to avoid single points of failure.
- NIC teaming is not supported for the VCS private network.
- Verify that your DNS server is configured for secure dynamic updates. For the Forward and Reverse Lookup Zones, set the Dynamic updates option to "Secure only". (DNS > Zone Properties > General tab)
- This is applicable for a Replicated Data Cluster configuration.
This is applicable for a Replicated Data Cluster configuration. You can configure single node clusters as the primary and secondary zones. However, if using a shared storage configuration, you must create the disk groups as clustered disk groups. If you cannot create a clustered disk group due to the unavailability of disks on a shared bus, use the `vxclus UseSystemBus ON` command.
- To configure a RDC cluster, you need to create virtual IP addresses for the following:
 - Application virtual server; this IP address should be the same on all nodes at the primary and secondary zones
 - Replication IP address for the primary zone
 - Replication IP address for the secondary zone
 Before you start deploying your environment, you should have these IP addresses available.

IPv6 support

For IPv6 networks, the following is supported:

Types of addresses	<p>The following types of IPv6 addresses are supported:</p> <ul style="list-style-type: none"> ■ Unicast addresses: Only Global Unicast and Unique Local Unicast addresses are supported. ■ Automatic configuration: Only Stateless IPv6 address configuration is supported. In stateless mode, the IP address is configured automatically based on router advertisements. The prefix must be advertised.
LLT over UDP	<p>LLT over UDP is supported on both IPv4 and IPv6.</p> <p>You can use the Cluster Configuration Wizard (VCW) to configure LLT over UDP over IPv6.</p>

VCS agents, wizards, and other components	<p>VCS agents that require an IP address attribute and wizards that configure or discover IP addresses now support IPv6 addresses (of the type described above).</p> <p>The Veritas High Availability Engine (HAD) and the Global Cluster resource (WAC) also support IPv6 addresses.</p>
---	---

Note: Pure IPv4, pure IPv6, and dual-stack (IPv4 and IPv6 on the same system) configurations are supported.

IP address requirements for an Active-Passive configuration

In addition to preparing the names you want to assign the Active-Passive configuration objects, for an IPv4 network, you should obtain all required IP addresses before beginning configuration. For an IPv6 network, IP addresses are generated during configuration.

See [“Sample Active-Passive configuration”](#) on page 44.

Each SQL Server virtual server requires its own virtual IP address. In the sample configuration there is one SQL Server virtual server. Therefore you would need one virtual server IP address. If you want to use the notification service, you require a cluster IP address. The cluster IP address is also used by the Global Cluster Option for disaster recovery.

IP address requirements for a disaster recovery configuration

In addition to preparing the names you want to assign configuration objects, for an IPv4 network, you should obtain all required IP addresses before beginning configuration. For an IPv6 network, IP addresses are generated during configuration.

See [“Sample disaster recovery configuration”](#) on page 72.

You specify the following addresses during the DR configuration process:

Virtual SQL Server IP address	For a disaster recovery configuration, the virtual IP address for the virtual SQL Server at the primary and disaster recovery site can be the same if both sites can exist on the same network segment. Otherwise, you need to allocate one IP address for the virtual server at the primary site and a different IP address for the virtual server at the disaster recovery site.
Cluster IP address	You need one for the primary site cluster and one for the secondary site cluster.

Replication IP address You need two IP addresses per application instance, one for the primary site and one for the secondary site.

Configuring the storage hardware and network

Use the following procedures to configure the hardware and verify DNS settings. Repeat this procedure for every node in the cluster.

To configure the hardware

- 1 Install the required network adapters, and SCSI controllers or Fibre Channel HBA.
- 2 Connect the network adapters on 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 autonegotiation options on the private network adapters. Contact the NIC manufacturer for details on this process.
 - Veritas recommends removing TCP/IP from private NICs to lower system overhead.
- 3 Use independent hubs or switches for each VCS communication network (GAB and LLT). You can use cross-over Ethernet cables for two-node clusters. LLT supports hub-based or switch network paths, or two-system clusters with direct network links.
- 4 Verify that each system can access the storage devices. Verify that each system recognizes the attached shared disk and that the attached shared disks are visible.

To verify the DNS settings and binding order for all systems

- 1 Open the Control Panel by clicking **Start > Control Panel**.
- 2 Click **Network and Internet**, and then click **Network and Sharing Center**.
- 3 In the Network and Sharing Center window, on the left side of the screen under Tasks, click **Adapter settings**.
- 4 Ensure the public network adapter is the first bound adapter by following these steps sequentially:
 - In the Network Connections window, click **Advanced > Advanced Settings**.
 - In the Adapters and Bindings tab, verify the public adapter is the first adapter in the **Connections** list. If necessary, use the arrow button to move the adapter to the top of the list.
 - Click **OK**.

- 5 Open the Public status dialog box by doing one of the following in the Network Connections window:
 - Double-click the adapter for the public network.
 - Right-click the adapter for the public network and click **Status**.
 - Select the adapter for the public network and click **View status of this connection** in the toolbar.

When enabling DNS name resolution, make sure that you use the public network adapters, and not those configured for the VCS private network.
- 6 In the Public Status dialog box, on the General tab, click **Properties**.
- 7 In the Public Properties dialog box, on the General tab, select the appropriate IP version and then click **Properties**.
- 8 Select the **Use the following DNS server addresses** option.
- 9 Verify the correct value for the IP address of the DNS server.
- 10 Click **Advanced**.
- 11 In the DNS tab, make sure that the **Register this connection's address in DNS** check box is selected.
- 12 Make sure that the correct domain suffix is entered in the **DNS suffix for this connection** field.
- 13 Click **OK**.

Configuring disk groups and volumes for SQL Server

Before installing SQL Server, you must create disk groups and volumes using the Veritas Enterprise Administrator (VEA) console installed with SFW.

You create cluster disk groups if you are using a shared storage environment and dynamic disk groups in case of a non-shared storage environment. Planning disk groups and volumes is covered in the following topics:

Planning disk groups and volumes is covered in the following topics:

See [“About disk groups and volumes”](#) on page 80.

See [“Prerequisites for configuring disk groups and volumes”](#) on page 81.

See [“Considerations for a fast failover configuration”](#) on page 82.

See [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 83.

See [“Considerations when creating disks and volumes for campus clusters”](#) on page 83.

See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 85.

See [“Considerations for disk groups and volumes for multiple instances”](#) on page 85.

See [“Sample disk group and volume configuration”](#) on page 86.

See [“MSDTC sample disk group and volume configuration”](#) on page 87.

Configuring disk groups and volumes is covered in the following topics:

See [“Viewing the available disk storage”](#) on page 87.

See [“Creating a dynamic disk group”](#) on page 88.

See [“Adding disks to campus cluster sites”](#) on page 89.

See [“Creating volumes for high availability clusters”](#) on page 90.

See [“Creating volumes for campus clusters”](#) on page 94.

About disk groups and volumes

SFW uses disk groups to organize disks or LUNs for management purposes. A dynamic disk group is a collection of disks that is imported or deported as a single unit. A cluster disk group is a special type of dynamic disk group that is created on shared storage and is designed to be moved or to failover between hosts. In order to prevent data corruption a cluster disk group uses SCSI reservations to protect the shared disks and limits access to a single host at a time.

Volumes are logical entities that are comprised of portions of one or more physical disks and are accessed by a drive letter or mount point. Volumes can be configured for performance and high availability.

If using a shared storage, you create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in a high-availability cluster by first deporting the cluster disk group from the current node and then importing it on the desired node. In a campus cluster, the volumes are mirrored across the storage arrays.

If using non-shared storage, you create dynamic disk groups and volumes on the locally attached storage on each node separately. Replication is set up between the volumes to ensure data concurrency.

Note: If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR) in a shared storage environment, you must enable SCSI-3 support using the Veritas Enterprise Administrator (VEA - Control Panel - System Settings). See the *Storage Foundation Administrator's Guide* for more information.

Prerequisites for configuring disk groups and volumes

Before you create a disk group (cluster disk group or dynamic disk group), consider the following items:

- The type of volume configurations that are required
- The number of volumes or LUNs required for the disk group
- The implications of backup and restore operations on the disk group setup
- The size of databases and logs that depend on the traffic load
- The disk groups and number of disks on each site
- Types of volumes required and location of the plex of each volume in the storage array

Complete the following tasks before you create the disk group and volumes for the SQL Server instance:

- Determine the layout or configuration for each volume and the total number of disks needed. Veritas recommends that you place SQL Server user database files and log files on separate volumes.
- Determine the initial size necessary for the volumes. You may increase the volume size at a later time using the Expand Volume command but you can not decrease the size.
- If using shared storage, verify that the disks you plan to include in the cluster disk group are shared and are available from all nodes. If new disks are installed, you must rescan, and if necessary, use the Write Signature command in order to identify the disks to the operating system.
- Verify that the drive letters that will be assigned to the volumes are available on all nodes so that the volumes can be accessed from any node.

For more information on disk group and volume requirements for specific configurations, see the following topics:

- For service groups with many disk groups, you may want to implement the fast failover feature.

See [“Considerations for a fast failover configuration”](#) on page 82.

- You may be configuring new shared storage for the high availability environment, or the existing standalone SQL Server databases and logs may already be on shared storage. If the existing databases and logs are already on shared storage: See [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 83.
- For more information on disk groups and volumes for campus clusters: See [“Considerations when creating disks and volumes for campus clusters”](#) on page 83.
- For a Replicated Data Cluster configuration or a disaster recovery configuration using Volume Replicator: See [“Considerations when creating disks and volumes for campus clusters”](#) on page 83.

Considerations for a fast failover configuration

For VCS service groups that contain many disk groups, you can greatly reduce failover time by implementing fast failover.

Fast failover speeds up the failover of storage resources in several ways:

- Fast failover provides a "read-only deported" mode for disk groups on inactive nodes. This mode speeds up the process of importing a disk group.
- Fast failover maintains the current disk group configuration in memory on the inactive nodes. Any changes are automatically synchronized so that all nodes maintain an identical disk group configuration.

For more details about fast failover, refer to the *Storage Foundation Administrator's Guide*.

Take the following storage-related requirements into account if you are planning to implement fast failover:

- Fast failover is currently not supported for the following:
 - RAID-5 volumes
 - SCSI-2
 - Active/Passive (A/P) arrays for DMPW
- In synchronous mode of replication, if fast failover is set, then the RVG cannot be stopped and started when a disk group fails over to another node. If the RLINK is in hard synchronous mode, it may not be connected when the volume arrives, and the I/Os may fail. In such case, the Event Viewer displays NTFS and ReFS errors and file system reports the volume as RAW. Therefore, fast failover is not supported if the RLINK is in hard synchronous mode.

- The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click Properties. Disk group version upgrade is required after upgrading SFW HA on the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

Considerations for converting existing shared storage to cluster disk groups and volumes

The databases and logs for your existing standalone SQL Server may already be on shared storage. In this case, when you create cluster disk groups, you specify the disks that contain the existing databases and logs.

Creating a disk group converts the disks from basic disks to dynamic disks. Partitions on the disks are automatically converted to volumes on the dynamic disks.

Therefore, if your existing disk layout contains databases and logs in the same partition, they become part of the same volume in the cluster disk group. If the disk contains multiple partitions, each containing a user database, each partition becomes a separate volume, but all will become part of the same cluster disk group. If this configuration does not meet your requirements, you may want to modify your disk layout before creating the cluster disk group.

For additional information on converting basic to dynamic disks, see *Storage Foundation Administrator's Guide*.

Veritas recommends creating a separate 100 MB RegRep volume that contains the list of registry keys that must be replicated among cluster systems for the SQL Server service. However, if no additional disks are available on the shared storage, you can specify an existing volume as the registry replication path during service group creation.

For a disaster recovery configuration using Volume Replicator, you need to allow additional disk space for a Storage Replicator Log volume.

See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 85.

Considerations when creating disks and volumes for campus clusters

When you create the disk groups for a campus cluster, ensure that each disk group has the same number of disks on each physical site. You create each volume as a mirrored volume with one plex of the volume on Site A's storage array and the other plex of the volume on Site B's storage array.

Veritas recommends using the SFW site-aware allocation feature for campus cluster storage. Site-aware allocation can ensure that site boundary limits are maintained for operations like volume grow, subdisk move, and disk relocation.

Enabling site-aware allocation for campus clusters requires the following steps in the VEA:

- After creating the disk groups, you tag the disks with site names to enable site-aware allocation. This is a separate operation, referred to in the VEA as adding disks to a site.
As an example, say you had a disk group with four disks. Disk1 and Disk2 are physically located on Site A. Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to "site_a" and add Disk3 and Disk4 to "site_b".
- During volume creation, you specify the volume site type as Site Separated. This ensures that the volume is restricted to the disks on the selected site.

Note: The hot relocation operation does not adhere to site boundary restrictions. If hot relocation causes the site boundary to be crossed, then the Site Separated property of the volumes is changed to Siteless. This is done so as not to disable hot relocation. To restore site boundaries later, you can relocate the data that crossed the site boundary back to a disk on the original site and then change back the properties of the affected volumes.

For more information on site-aware allocation, refer to the *Storage Foundation Administrator's Guide*.

When you create the volumes for a campus cluster, consider the following:

- During disk selection, configure the volume as "Site Separated" and select the two sites of the campus cluster from the site list.
- For volume attributes, select the "mirrored" and "mirrored across enclosures" options.
- Veritas recommends using either simple mirrored (concatenated) or striped mirrored options for the new volumes. Striped mirrored gives you better performance compared to concatenated.
When selecting striped mirrored, select two columns in order to stripe one enclosure that is mirrored to the second enclosure.
- During the volume creation procedure for Site Separated volumes, you can only create as many mirrors as there are sites. However, once volume creation is complete, you can add additional mirrors if desired.
- Choosing "Mirrored" and the "mirrored across" option without having two enclosures that meet requirements causes new volume creation to fail.
- You cannot selecting RAID-5 for mirroring.

- Selecting "stripe across enclosures" is not recommended because then you need four enclosures, instead of two.
- Logging can slow performance.

Considerations for volumes for a Volume Replicator configuration

For a configuration using Volume Replicator, either a disaster recovery configuration on a secondary site or a Replicated Data Cluster, note the following:

- Volume Replicator does not support the following types of volumes:
 - SFW (software) RAID 5 volumes
 - Volumes with the Dirty Region Log (DRL)
 - Volumes with commas in the names
 - Data Change Object (DCO)
- A configuration with Volume Replicator requires a Storage Replicator Log (SRL) volume for each disk group that contains volumes that are replicated. You can create the SRL volume when configuring the other volumes for the application or you can create it later when you set up replication. If you create it later, ensure that you allow sufficient disk space for this volume. For more about Volume Replicator planning, see the *Volume Replicator Administrator's Guide*.
- Do not assign a drive letter to the Storage Replicator Log volume. This will limit access to that volume and avoid potential data corruption.
- In a disaster recovery configuration, Veritas recommends that for replication considerations, you create a separate volume for tempdb, for example, INST1_TEMPDB, within the system database disk group. When you later configure replication for disaster recovery, you replicate that disk group but exclude the tempdb volume from the replication.

It would waste bandwidth to replicate tempdb because the data is transitory and is not needed for DR site recovery.

You can create the volume now and later, after the SQL Server installation is complete and before configuring replication, move tempdb to the volume.

See ["Moving the tempdb database if using Volume Replicator for disaster recovery"](#) on page 129.

Considerations for disk groups and volumes for multiple instances

For an Active-Active configuration or other cases where you are setting up multiple SQL Server instances in the cluster, you create a separate set of disk groups and volumes for each instance.

For example, if you have a Billing instance and a Payroll instance, you could create the following disk groups and volumes.

For the Billing instance, create the following:

BILLING_DG

BILLING_DATA_FILES Volume for the SQL Server system data files

BILLING_REGREP_VOL Volume for the list of registry keys replicated among cluster nodes for the Billing instance

BILLING_DB1_VOL Volume for the user database files

BILLING_DB1_LOG Volume for the user database log files

PAYROLL_DG

PAYROLL_DATA_FILES Volume for the SQL Server system data files

PAYROLL_REGREP_VOL Volume for the list of registry keys replicated among cluster nodes for the Payroll instance

PAYROLL_DB1_VOL Volume for the user database files

PAYROLL_DB1_LOG Volume for the user database log files

You can choose either of the following:

- Set up disk groups and volumes for all instances at one time.
- Set up disk groups and volumes for the current instance only and complete all configuration steps for this instance. Then return to this step for the next instance.

Sample disk group and volume configuration

For an SFW HA solution, you first create a disk group (INST1_DG) (cluster disk group in case of shared storage and dynamic disk group in case of non-shared storage) and then create the following volumes:

- 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 volume for this purpose.
- INST1_FS_VOL contains the FILESTREAM enabled database objects for the SQL Server instance

- `INST1_REPLOG` contains the Volume Replicator Storage Replicator Log. This is required only for a configuration that uses Volume Replicator replication, either a Replicated Data Cluster or a disaster recovery configuration using Volume Replicator. You can create this volume later while setting up replication.

Optionally place user database files in a separate cluster disk group from the system database files, for example, by creating `INST1_SHARED_DG` for system files and `INST1_USER_DG` for user database files.

As a best practice, place SQL Server user database files and log files on separate volumes.

The following disk group and volumes may be created now or later in the configuration process:

- `INST1_DB1_DG` (optional) is a separate disk group for the SQL Server user-defined database and files
- `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 objects for the user database
- `INST1_DB1_REPLOG` contains the Volume Replicator Storage Replicator Log (required only for a configuration using Volume Replicator replication).

This configuration is a simple example. The recommended practice for disk groups and volume layout is dependent on your environment.

MSDTC sample disk group and volume configuration

For an MSDTC configuration, you will first need to create a disk group (`MSDTC_DG`) (cluster disk group in case of shared storage and dynamic disk group in case of non-shared storage) and then create the following volumes:

`MSDTC_LOG` contains the MSDTC log files

`MSDTC_REGREP` contains the list of registry keys that must be replicated among cluster systems for the MSDTC service group

Create a 100 MB volume for this purpose.

Viewing the available disk storage

Before creating disk groups and volumes you may want to view available disk storage.

To view the available disk storage

- 1 Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator**, or, on Windows 2012 operating systems, from the **Apps** menu, and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.

To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

- 4 In the VEA configuration tree, expand **hostname > StorageAgent** and then click **Disks**.

The internal names for the disks that the current system can access for available storage are displayed, with names Harddisk1, Harddisk2, etc. The list includes both disks internal to the local system and any external storage that is available.

Creating a dynamic disk group

Create a cluster disk group and volumes on only one node of a cluster. The volumes can be accessed by other nodes in the cluster by first deporting the cluster disk group from the current node and then importing it on the desired node.

Note: Setting up the cluster creates physical disk resources for all the basic disks on the shared bus. To use those disks for the SFW cluster disk groups, you must remove the physical disk resources from the cluster. Otherwise, a reservation conflict occurs.

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

Follow the steps in this section to create one or more disk groups for your application.

To create a dynamic disk group

- 1 Open the VEA console by clicking **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** (or launch the VEA from the Solutions Configuration Center) or, on Windows 2012 operating systems, from the **Apps** menu, and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.

- 3 In the Connect dialog box, select the host name from the pull-down menu and click **Connect**.

To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
 - 4 To start the New Dynamic Disk Group wizard, expand the tree view under the host node, right click the **Disk Groups** icon, and select **New Dynamic Disk Group** from the context menu.
 - 5 In the Welcome screen of the New Dynamic Disk Group wizard, click **Next**.
 - 6 Provide information about the cluster disk group as follows:
 - In the **Group name** field, enter a name for the disk group (for example, **INST1_DG**).
 - Check the **Create cluster group** check box if you wish to create cluster dynamic disk groups that are used in a shared storage environment.
 - Select the appropriate disks in the **Available disks** list, and use the **Add** button to move them to the **Selected disks** list.
Optionally, check the **Disk names prefix** checkbox and enter a disk name prefix to give the disks in the disk group a specific identifier.
For example, entering **TestGroup** as the prefix for a disk group that contains three disks creates TestGroup1, TestGroup2, and TestGroup3 as internal names for the disks in the disk group.
 - Click **Next**.
-
- Note:** Windows Disk Management Compatible Dynamic Disk Group creates a type of disk group that is created by Windows Disk Management (LDM).
-
- 7 Click **Next** to accept the confirmation screen with the selected disks.
 - 8 Click **Finish** to create the dynamic disk group.

Adding disks to campus cluster sites

For campus cluster storage, Veritas recommends using Storage Foundation (SFW) site-aware allocation. To enable site-aware allocation, you assign a site name to disks after they are added to a disk group. In the VEA assigning a site name is referred to as adding disks to a site.

For example, Disk1 and Disk2 are physically located on Site A and Disk3 and Disk4 are physically located on Site B. Therefore, you add Disk1 and Disk2 to site_a and add Disk3 and Disk4 to site_b.

To add disks to a site

- 1 From the VEA console, right-click a disk that needs to be added to a site and select **Add Disk to Site**.

Disks must be part of a dynamic disk group in order to add them to a site.
- 2 In the Add Disk to a Site screen, choose one of the following:
 - Choose **Select a new site** and specify a new site name.
The site name can include any alphanumeric value and valid characters like the period (.), dash (-), and underscore (_). It cannot exceed 31 characters. Site names are case insensitive; all names are converted to lowercase.
 - Choose **Available Sites** and select a site from the list.
- 3 From the **Available Disks** column, select the disk or disks to add to the specified site.
- 4 Click **OK**.

Creating volumes for high availability clusters

This procedure will guide you through the process of creating a volume on a dynamic disk group. Repeat the procedure to create additional volumes.

You can use this procedure for volumes in a high availability cluster. For volumes in a campus cluster, see the following:

- See [“Creating volumes for campus clusters”](#) on page 94.

Before you begin, make sure to review the following topics if they are applicable to your environment:

- See [“Considerations for converting existing shared storage to cluster disk groups and volumes”](#) on page 83.
- See [“Considerations for volumes for a Volume Replicator configuration”](#) on page 85.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 If the VEA console is not already open, click **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu, and select a profile if prompted.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name from the pull-down menu and click **Connect**.


To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.

- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.

You can right-click the disk group you just created.

For example, **INST1_DG**.

- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume.



- Make sure that the appropriate disk group name appears that in the **Group name** drop-down list.

- For Site Preference, leave the setting as **Siteless** (the default).
- Automatic disk selection is the default setting. To manually select the disks, click **Manually select disks** and use the **Add** and **Remove** buttons to move the appropriate disks to the **Selected disks** list. Manual selection of disks is recommended.
- You may also check **Disable Track Alignment** to disable track alignment for the volume. Disabling track alignment means that the volume does not store blocks of data in alignment with the boundaries of the physical track of the disk.
- Click **Next**.

7 Specify the volume attributes.

- Enter a volume name. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Provide a size for the volume. If you click the **Max Size** button, a size appears in the **Size** box that represents the maximum possible volume size for that layout in the dynamic disk group.
- Select a volume layout type. To select mirrored striped, select both the **Mirrored** checkbox and the **Striped** radio button.
- If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided.

In the Mirror Info area, select the appropriate mirroring options.

- Verify that **Enable logging** is not selected.
- Click **Next**.

8 Assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.

- To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
- To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the disk.
- If creating a Replicator Log volume for Volume Replicator, select **Do not assign a drive letter**.
- Click **Next**.

9 Create an NTFS file system.

- Make sure that the Format this volume checkbox is checked and click **NTFS**.
- For a Volume Replicator configuration, for the Replicator Log volume only, clear the **Format this volume** check box.
- Select an allocation size or accept the default.
- The file system label is optional. SFW makes the volume name the file system label.
- Select **Perform a quick format** if you want to save time.
- Select **Enable file and folder compression** to save disk space. Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.

Note: If you plan to use this volume to install SQL Server, do not select the **Enable file and folder compression** checkbox. The SQL Server installation cannot copy files on a compressed or encrypted folder.

- Click **Next**.

10 Click **Finish** to create the new volume.

11 Repeat these steps to create additional volumes.

Create the cluster disk group and volumes on the first node of the cluster only.

Creating volumes for campus clusters

This section will guide you through the process of creating a volume on a dynamic disk group for a campus cluster.

For creating volumes for other types of clusters:

- See [“Creating volumes for high availability clusters”](#) on page 90.

Before you begin, review the following topics:

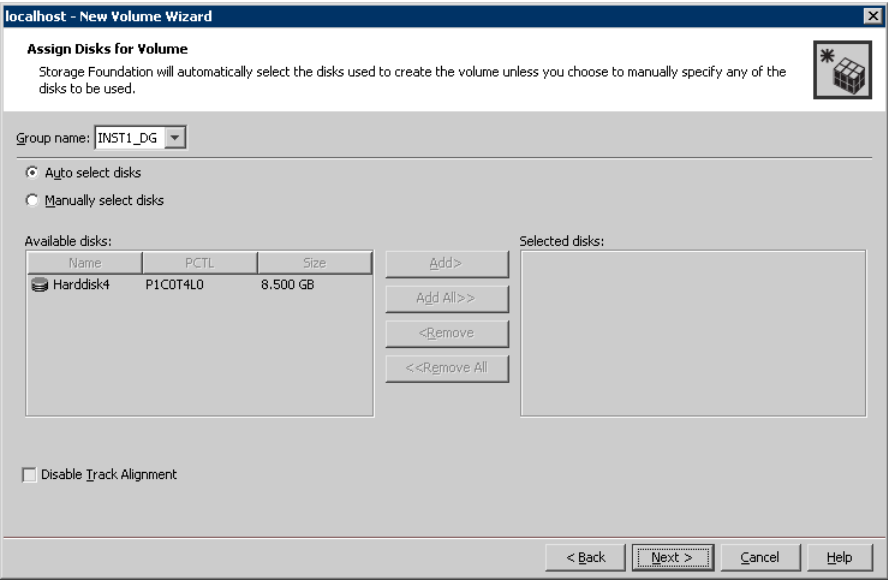
- See [“Considerations when creating disks and volumes for campus clusters”](#) on page 83.
- See [“Adding disks to campus cluster sites”](#) on page 89.

Use the following procedure to create dynamic volumes for a campus cluster.

Note: When assigning drive letters to volumes, ensure that the drive letters that you assign are available on all nodes.

To create dynamic volumes

- 1 Launch the VEA console from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator** or, on Windows 2012 operating systems, from the **Apps** menu.
- 2 Click **Connect to a Host or Domain**.
- 3 In the Connect dialog box select the host name and click **Connect**.
To connect to the local system, select **localhost**. Provide the user name, password, and domain if prompted.
- 4 To start the New Volume wizard, expand the tree view under the host node to display all the disk groups. Right click a disk group and select **New Volume** from the context menu.
You can right-click the disk group you have just created.
For example, **INST1_DG**.
- 5 At the New Volume wizard opening screen, click **Next**.
- 6 Select the disks for the volume as follows:



Group name	Make sure the appropriate disk group is selected.
Site preference	Select the Site Separated option.
Select site from	Select the campus cluster sites. Press CTRL to select multiple sites. Note: If no sites are listed, the disks have not yet been added to a site.
Auto select disks	Automatic disk selection is recommended for campus clusters. SFW automatically selects the disks based on the following criteria: <ul style="list-style-type: none">■ Their port assignment (disks with two different ports are selected): Note that in the list of available disks, the entry after each disk name starts with the port number. For example, the “P3” in the entry P3C0T2L1 refers to port 3.■ Amount of available space on the disks: SFW picks two disks (one from each array) with the most space.
Manually select disks	If you manually select disks, use the Add and Remove buttons to move the appropriate disks to the Selected disks list.

Disable Track
Alignment

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

Click **Next**.

7 Specify the volume attributes as follows:

localhost - New Volume Wizard

New Volume Wizard

Select the attributes for this volume.

Volume name: INST1_DB1_VOL

Size: 10 GB Max Size

Layout

☒ Concatenated

Columns: 2

☐ Striped

Stripe unit size (Sectors): 128

☐ RAID-5

☐ Stripe across: Port

Mirror Info

☐ Mirrored

Total mirrors: 2

☐ Mirror across: Port

☐ Enable logging

Concatenated: A simple volume with a single copy of data on one or more disks.

< Back

Next >

Cancel

Help

- Volume name

Specify a name for the volume. The name is limited to 18 ASCII characters and cannot contain spaces or forward or backward slashes.
- Size

Specify a size for the volume. If you click **Max Size**, the **Size** box shows the maximum possible volume size for that layout in the dynamic disk group.
- Layout

Ensure that the **Mirrored** checkbox is selected.

Select either the **Concatenated** or **Striped** layout type.

If you are creating a striped volume, the **Columns** and **Stripe unit size** boxes need to have entries. Defaults are provided. In addition, click the **Stripe across** checkbox and select **Ports** from the drop-down list.
- Mirror Info

Click **Mirror across** and select **Enclosures** from the drop-down list.

When creating a site separated volume, as required for campus clusters, the number of mirrors must correspond to the number of sites. If needed, you can add more mirrors after creating the volume.
- Enable logging

Verify that this option is not selected.

Click **Next**.

- 8 In the Add Drive Letter and Path dialog box, assign a drive letter or mount point to the volume. You must use the same drive letter or mount point on all systems in the cluster. Make sure to verify the availability of the drive letter before assigning it.
 - To assign a drive letter, select **Assign a Drive Letter**, and choose a drive letter.
 - To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.

Click **Next**.

- 9 Create an NTFS file system.
 - Make sure the **Format this volume** checkbox is checked and select **NTFS**.
 - Select an allocation size or accept the default.
 - The file system label is optional. SFW makes the volume name the file system label.
 - Select **Perform a quick format** if you want to save time.
 - Select **Enable file and folder compression** to save disk space.
Note that compression consumes system resources and performs encryption and decryption, which may result in reduced system performance.
 - Click **Next**.
- 10 Click **Finish** to create the new volume.
- 11 Repeat these steps to create additional volumes as needed.

Note: Create the cluster disk group and volumes on the first node of the cluster only.

About managing disk groups and volumes

During the process of setting up an SFW environment, refer to these general procedures for managing disk groups and volumes:

- When a disk group is initially created, it is imported on the node where it is created.
- A disk group can be imported on only one node at a time.

- To move a cluster dynamic disk group from one node to another, unmount the volumes in the disk group, deport the disk group from its current node, import it to a new node and mount the volumes.

Managing disk groups and volumes involves the following:

- See [“Importing a disk group and mounting a volume”](#) on page 99.
- See [“Unmounting a volume and deporting a disk group”](#) on page 100.

Note: (Disaster recovery configurations only) If your storage devices are SCSI-3 compliant, and you wish to use SCSI-3 Persistent Group Reservations (PGR), you must enable SCSI-3 support using the Veritas Enterprise Administrator (**VEA > Control Panel > System Settings**). See the *Storage Foundation Administrator's Guide* for more information.

Importing a disk group and mounting a volume

Use the VEA Console to import a disk group and mount a volume.

To import a disk group

- 1 From the VEA Console, right-click a disk name in a disk group or the group name in the Groups tab or tree view.
- 2 From the menu, click **Import Dynamic Disk Group**.

To mount a volume

- 1 If the disk group is not imported, import it.
- 2 To verify if a disk group is imported, from the VEA Console, click the Disks tab and check if the status is imported.
- 3 Right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 4 Select one of the following options in the Drive Letter and Paths dialog box depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter, select **Assign a Drive Letter**, and select a drive letter.
 - To mount the volume as a folder, select **Mount as an empty NTFS folder**, and click **Browse** to locate an empty folder on the shared disk.
- 5 Click **OK**.

Unmounting a volume and deporting a disk group

Use the VEA Console to unmount a volume and deport a disk group.

To unmount a volume and deport the dynamic disk group

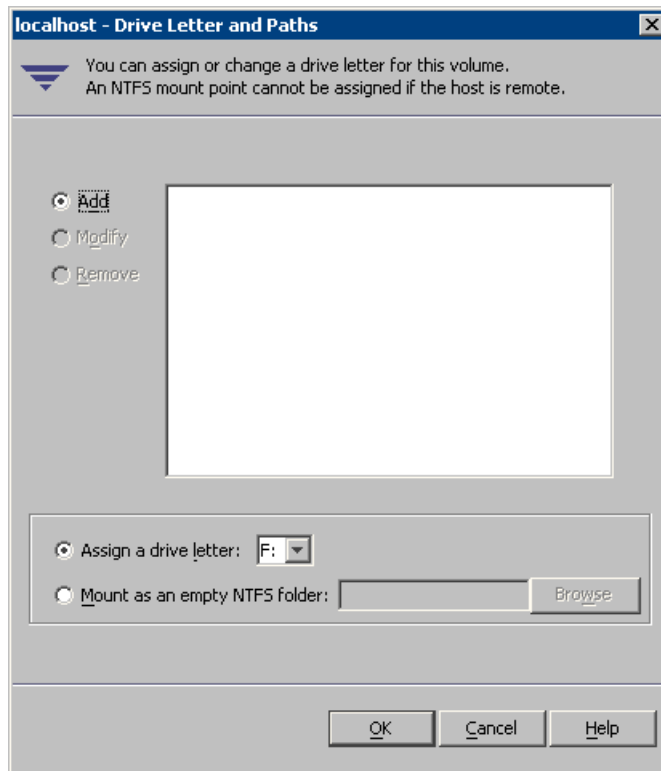
- 1** From the VEA tree view, right-click the volume, click **File System**, and click **Change Drive Letter and Path**.
- 2** In the Drive Letter and Paths dialog box, click **Remove**.
Click **OK** to continue.
- 3** Click **Yes** to confirm.
- 4** From the VEA tree view, right-click the disk group, and click **Deport Dynamic Disk Group**.
- 5** Click **Yes**.

Adding drive letters to mount the volumes

Occasionally, when a disk group is imported a drive letter may not be associated with an existing volume. If this occurs, use the VEA console to add a drive letter and mount the volume so that it can be seen by the operating system. You can also mount the volume as a folder. Verify that all volumes are mounted.

To add a drive letter or path to a volume

- 1 Navigate to the `Volumes` folder.
- 2 Right-click the volume, click **File System** and click **Change Drive Letter and Path**.



- 3 In the Drive Letter and Paths dialog box, click **Add**.
- 4 Select one of the following options depending on whether you want to assign a drive letter to the volume or mount it as a folder.
 - To assign a drive letter, select **Assign a Drive Letter** and select a drive letter from the drop-down list.
 - To mount the volume as a folder, select **Mount as an empty NTFS folder** and click **Browse** to locate an empty folder on the shared disk.

Note: Assign the same drive letter or mount path that was assigned when the volume was created.

- 5 Click **OK**.

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 (NICs), and SCSI controllers are installed and connected to each system.
Veritas recommends the following actions for network adapters:
 - Disable the ethernet auto-negotiation options on the private NICs to prevent:
 - Loss of heartbeats on the private networks
 - VCS from mistakenly declaring a system as offlineContact the NIC manufacturer for details on this process.
 - Remove TCP/IP from the 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, VCW 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

Back Next Cancel

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 111.
- 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 114.

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' window of the Veritas Cluster Server Configuration Wizard. The window has a blue title bar and a sidebar on the left with navigation links: Domain Selection, Create Cluster, Select Components, Configure, Summary, and Finish. The main area is titled 'Notifier SNMP Configuration' with the subtitle 'Specify information about SNMP console.' Below this, it says 'Enter the name or the IP address of the SNMP console and then select the desired severity level.' There is a table with two columns: 'SNMP Console' and 'Severity'. The 'SNMP Console' column has a placeholder text 'Click here to change the text..'. The 'Severity' column has a placeholder text 'Information'. Below the table, there are instructions: 'Click on '+' button to add more consoles.' and 'Click '-' to remove a console.' with corresponding '+' and '-' buttons. There is also 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 a prompt says 'Click 'Next' to continue.' At the bottom right, there are 'Back', 'Next', and 'Cancel' buttons. The 'Next' button is highlighted with a mouse cursor.

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

Click on '+' button to add more consoles. + -

Click '-' to remove a console.

SNMP Trap Port:

Note: SNMP console must be MIB 2.0 compliant.

Click 'Next' to continue.

Back Next Cancel

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**.

Notifier SMTP Configuration
Specify information about SMTP recipients.

Domain Selection
 Create Cluster
 Select Components
 Configure
 Summary
 Finish

SMTP Server Name / IP

Enter SMTP recipients and select a severity level for each recipient.

Recipients	Severity
Click here to change the text..	Information

Click '+' to add a recipient.
 Click '-' to remove a recipient.

Click 'Next' to continue.

VERITAS

Back Next Cancel

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.

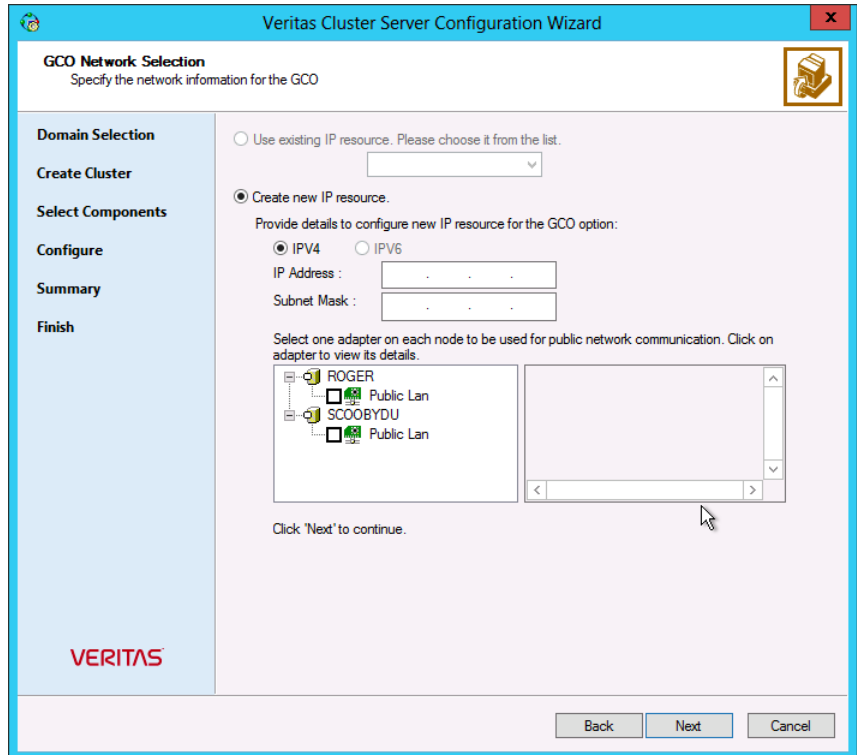
Configuring Wide-Area Connector process for global clusters

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.

Adding nodes to a cluster

To add a node to a VCS cluster

- 1 Start the VCS Cluster Configuration wizard.

Click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Cluster Configuration Wizard**.

Run the wizard from the node to be added or from a node in the cluster. The node that is being added should be part of the domain to which the cluster belongs.

- 2 Read the information on the Welcome panel and click **Next**.
- 3 On the Configuration Options panel, click **Cluster Operations** and click **Next**.
- 4 In 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 the systems and users in the domain, do the following:

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

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

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

- 5 On the System Selection panel, complete the following and click **Next**:
 - Type the name of an existing node in the cluster and click **Add**.
 - Type the name of the system to be added to the cluster and click **Add**.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

Proceed to step 8.

- 6 On the System Selection panel, specify the systems to be added and the nodes for the cluster to which you are adding the systems.

Enter the system name and click **Add** to add the system to the **Selected Systems** list. Alternatively, you can select the systems from the **Domain Systems** list and click the right-arrow icon.

If you specify only one node of an existing cluster, the wizard discovers all nodes for that cluster. To add a node to an existing cluster, you must specify a minimum of two nodes; one that is already a part of a cluster and the other that is to be added to the cluster.

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

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

- The system does not respond to a ping request.
- WMI access is disabled on the system.
- The wizard is unable to retrieve information about the system's architecture or operating system.
- VCS is either not installed on the system or the version of VCS is different from what is installed on the system on which you are running the wizard.

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

Click **Next** to proceed.

- 8 On the Cluster Configuration Options panel, click **Edit Existing Cluster** and click **Next**.
- 9 On the Cluster Selection panel, select the cluster to be edited and click **Next**.

If you chose to specify the systems manually in step 4, only the clusters configured with the specified systems are displayed.

- 10** On the Edit Cluster Options panel, click **Add Nodes** and click **Next**.

In the Cluster User Information dialog box, type the user name and password for a user with administrative privileges to the cluster and click **OK**.

The Cluster User Information dialog box appears only when you add a node to a cluster with VCS user privileges (a cluster that is not a secure cluster).

- 11** On the Cluster Details panel, check the check boxes next to the systems to be added to the cluster and click **Next**.

The right pane lists nodes that are part of the cluster. The left pane lists systems that can be added to the cluster.

- 12** The wizard validates the selected systems for cluster membership. After the nodes have been validated, click **Next**.

If a node does not get validated, review the message associated with the failure and restart the wizard after rectifying the problem.

- 13** On the Private Network Configuration panel, configure the VCS private network communication on each system being added and then click **Next**. How you configure the VCS private network communication depends on how it is configured in the cluster. If LLT is configured over Ethernet, you have to use the same on the nodes being added. Similarly, if LLT is configured over UDP in the cluster, you have use the same on the nodes being added.

Do one of the following:

- To configure the VCS private network over Ethernet, do the following:
 - Select the check boxes next to the two NICs to be assigned to the private network.
Veritas recommends reserving two NICs exclusively for the private network. However, you could lower the priority of one NIC and use the low-priority NIC for public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication.
To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - If your configuration contains teamed NICs, the wizard groups them as "NIC Group #N" where "N" is a number assigned to the teamed NIC. A teamed NIC is a logical NIC, formed by grouping several physical NICs together. All NICs in a team have an identical MAC address. 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, do the following:
 - Select the check boxes next to the two NICs to be assigned to the private network. You can assign maximum eight network links. Veritas recommends reserving at least two NICs exclusively for the VCS private network. You could lower the priority of one NIC and use the low-priority NIC for both public and private communication.
 - If you have only two NICs on a selected system, it is recommended that you lower the priority of at least one NIC that will be used for private as well as public network communication. To lower the priority of a NIC, right-click the NIC and select **Low Priority** from the pop-up menu.
 - Specify a unique UDP port for each of the link. Click **Edit Ports** if you wish to edit the UDP ports for the links. You can use ports in the range 49152 to 65535. The default ports numbers are 50000 and 50001 respectively. Click **OK**.
 - For each selected NIC, verify the displayed IP address. If a selected NIC has multiple IP addresses assigned, double-click the field and choose the desired IP address from the drop-down list. 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.
 - 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.
- 14** On the Public Network Communication panel, select a NIC for public network communication, for each system that is being added, and then click **Next**.

This step is applicable only if you have configured the ClusterService service group, and the system being added has multiple adapters. If the system has only one adapter for public network communication, the wizard configures that adapter automatically.
- 15** Specify the credentials for the user in whose context the VCS Helper service runs.
- 16** Review the summary information and click **Add**.
- 17** The wizard starts running commands to add the node. After all commands have been successfully run, click **Finish**.

Installing SQL Server

This chapter includes the following topics:

- [About installing and configuring SQL Server](#)
- [About installing multiple SQL Server instances](#)
- [Verifying that the SQL Server databases and logs are moved to shared storage](#)
- [About installing SQL Server for high availability configuration](#)
- [About installing SQL Server on the first system](#)
- [About installing SQL Server on additional systems](#)
- [Creating a SQL Server user-defined database](#)
- [Completing configuration steps in SQL Server](#)

About installing and configuring SQL Server

At a high level, installing and configuring SQL Server in an InfoScale Enterprise environment involves the following tasks:

Task	Physical environment	VMware virtual environment
1 Install SQL Server on the first system.	<ul style="list-style-type: none"> ■ Install the SQL Server instance on the local disk. ■ In case of shared storage, install the SQL Server database files and 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 locally-attached non-shared storage (non-system drive). 	<p>Install the SQL Server instance at the default path. Ensure that you configure the data folders and log folders on a VMDK.</p>
2 Install SQL Server on additional systems.	<ul style="list-style-type: none"> ■ Install the SQL Server instances, database files, and analysis service files on the local system disk. ■ In case of shared storage, you <i>do not</i> need to install the database files and the analysis service files on shared storage. ■ In case of non-shared storage, the installation of SQL Server on additional systems is not applicable. <p>The advantage of this method is that while you install SQL Server on the first system, you can run parallel installations on the remaining systems.</p>	<p>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.</p>

Task	Physical environment	VMware virtual environment
3 Configure the SQL Server service group using the SQL Server Configuration Wizard and bring it online on the first system.	<ul style="list-style-type: none"> ■ In case of shared storage, run the wizard from the first system on which you installed the SQL Server database and analysis service files. Doing so is necessary, because the wizard configures the resources for the SQL Server database and the registry information that is installed on the shared storage. Then, it propagates this information to the remaining systems that are part of the SQL Server service group. ■ In case of non-shared storage (dynamic disk groups that are configured on local disks), you must configure the service group manually. To do so, use either the Cluster Manager (Java Console) or the command line. Currently, the wizard cannot configure the resources (VMNSDg agent) that are required to monitor non-shared storage. 	Configure the service group manually using either the Cluster Manager or the command line as you would do in a physical environment with non-shared storage.

About installing multiple SQL Server instances

If you are installing multiple instances of SQL Server on the same system, as in an active-active cluster configuration, some additional requirements apply.

The following summary is provided for your review to assist you in planning the installation:

- 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.
- Assign a unique port number for each instance.

Verifying that the SQL Server databases and logs are moved to shared storage

This task is applicable only if you are configuring an existing standalone SQL Server for high availability in an InfoScale Enterprise environment.

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 cluster disk groups and volumes on the shared storage. Doing so is necessary for proper failover operations in the cluster.

To move the database and logs to shared storage

- 1 Stop the SQL Server service.
- 2 Verify that you have backed up your existing data.
- 3 Create the required disk group and volumes for the SQL Server database.
See [“Creating a dynamic disk group”](#) on page 88.
See [“Adding disks to campus cluster sites”](#) on page 89.
- 4 Ensure that the dynamic disk group is imported on the system where the original database files are located on the local drives.
See [“Importing a disk group and mounting a volume”](#) on page 99.
- 5 Mount the volumes.
See [“Adding drive letters to mount the volumes”](#) on page 100.
- 6 Move the SQL Server data file and user database locations.
Refer to the Microsoft SQL Server documentation for instructions.
- 7 Restart SQL Server.

About installing SQL Server for high availability configuration

Before installation, note the following prerequisites:

- Ensure that you have installed the product and configured SFW HA on all the systems on which you wish to configure SQL Server.
For information about installing the Veritas InfoScale products using the installation wizard or the CLI, see the *Veritas InfoScale Installation and Upgrade Guide*.
- If you are using Windows LDM storage, ensure that the disks are accessible from all the systems where 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
- Installing SQL Server on the first system

In case of shared storage, you must install the SQL Server instance on the local disk and install the SQL Server database files and analysis service files on the shared disk. The shared disks must be accessible from all the systems where you wish to install and configure SQL Server.

In case of non-shared storage, you must install the database files on the disks residing on a datastore that is accessible from all the systems where you want to install and configure SQL Server. These disks are deported and imported during a failover.
- Installing SQL Server on additional systems

Irrespective of how the storage is configured (whether shared or non-shared), install the SQL Server instances, database files and analysis service files on the local system disk.
- Make sure you do not install SQL Server on the systems which belong to the Exchange service group SystemList attribute.
- 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.
 - Ensure that the [NT AUTHORITY\SYSTEM] account is granted the sysadmin server role (from SQL Management Studio Console) on each system.
- 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 systems where you want to install SQL Server.
- Ensure that the disk group is imported and the volumes are mounted to the first system for this SQL Server instance.

See [“About disk groups and volumes”](#) on page 80.

See [“Importing a disk group and mounting a volume”](#) on page 99.

See [“Adding drive letters to mount the volumes”](#) on page 100.

About installing SQL Server on the first system

Note: To install a standalone SQL Server instance, refer to the topic about installing SQL Server on additional systems.

See [“About installing SQL Server on additional systems”](#) on page 126.

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

- Ensure that you have installed the product and configured InfoScale Enterprise on all the systems on which you plan to configure SQL Server.
See the *Veritas InfoScale Installation and Upgrade Guide*.
- 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 name and 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 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 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 125.

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

- Ensure that you have installed the product and configured InfoScale Enterprise on all the systems on which you plan to configure SQL Server.
See the *Veritas InfoScale Installation and Upgrade Guide*.
- 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.
- After installing SQL Server, use the SQL Management Studio Console to grant the `sysadmin` server role to the [NT AUTHORITY\SYSTEM] account on each system.

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 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.

Creating a SQL Server user-defined database

You can use SFW HA to manage SQL Server user-defined databases. For making the user-defined databases highly available, create the user-defined databases and then configure them in the SQL Server service group. Refer to the Microsoft SQL Server documentation for instructions on how to create databases.

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

- If you have not already created volumes for a user-defined SQL Server database and its transaction log, create them first.
See [“Configuring disk groups and volumes for SQL Server”](#) on page 79.
- Create the SQL Server database for the desired SQL Server virtual server instance, and point the database files and transaction log to the new volumes created for them.
Refer to the Microsoft SQL Server documentation for instructions.
- After creating the database, you may have additional steps to complete in the SQL Server configuration. Perform the desired steps depending on your configuration plans.
See [“Completing configuration steps in SQL Server”](#) on page 129.
- If you have already configured the SQL Server service group, run the SQL Server Configuration Wizard again to modify the SQL Server service group. This allows the wizard to add storage agent resources for the new database, to the existing service group.
See [“Modifying a SQL Server service group to add VMDg and MountV resources”](#) on page 163.
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.

Note: If you have configured the SQL Server service group in a non-shared storage configuration (dynamic disk groups configured on local disks), you have to modify the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

Completing configuration steps in SQL Server

Depending on your configuration, you may have additional steps to complete in SQL Server.

If you plan to implement a disaster recovery configuration using Volume Replicator (Volume Replicator), Veritas recommends that you exclude the tempdb database from replication. To do this, you need to first move it to a separate volume within the system database disk group.

See [“Moving the tempdb database if using Volume Replicator for disaster recovery”](#) on page 129.

If you are running multiple SQL Server instances, you must assign a different port to each SQL Server instance.

See [“Assigning ports for multiple SQL Server instances”](#) on page 129.

Moving the tempdb database if using Volume Replicator for disaster recovery

If you plan to implement a disaster recovery configuration using Volume Replicator, Veritas recommends that you move tempdb to a separate volume within the system database disk group in order to be able to exclude it from replication.

If you have not yet created the volume for tempdb, you can do that now.

See [“Creating volumes for high availability clusters”](#) on page 90.

Then, refer to the Microsoft SQL Server documentation for the instructions on moving the tempdb database.

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 instructions on how to assign ports.

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

- Bring the SQL Server service group online or partially online (up to the registry replication resource) on a cluster node.
- On the system on which the SQL Server service group is online or partially online, change the port assigned to the SQL Server instance. Refer to the Microsoft SQL Server documentation for instructions.
- Take the SQL Server service group offline on the system, 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 Server 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 Server 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.

Configuring SQL Server in a physical environment

- [Chapter 4. Configuring SQL Server for failover](#)
- [Chapter 5. Configuring campus clusters for SQL Server](#)
- [Chapter 6. Configuring Replicated Data Clusters for SQL Server](#)
- [Chapter 7. Configuring disaster recovery for SQL Server](#)
- [Chapter 8. Testing fault readiness by running a fire drill](#)

Configuring SQL Server for failover

This chapter includes the following topics:

- [Tasks for configuring a new server for high availability](#)
- [Tasks for configuring an existing server for high availability](#)
- [About configuring the SQL Server service group](#)
- [Configuring the service group in a non-shared storage environment](#)
- [Verifying the SQL Server cluster configuration](#)
- [About the modifications required for tagged VLAN or teamed network](#)
- [Tasks for configuring MSDTC for high availability](#)
- [Configuring an MSDTC Server service group](#)
- [About configuring the MSDTC client for SQL Server](#)
- [About the VCS Application Manager utility](#)
- [Viewing DTC transaction information](#)
- [Modifying a SQL Server service group to add VMDg and MountV resources](#)
- [Determining additional steps needed](#)

Tasks for configuring a new server for high availability

The following table outlines the high-level objectives and the tasks to complete each objective.

Table 4-1 SQL Server: Active-Passive configuration tasks

Action	Description
Review the HA configuration	<ul style="list-style-type: none"> ■ Understand active-passive and active-active configuration ■ Review the sample configuration <p>See “Reviewing the active-passive HA configuration” on page 43.</p>
Configure the storage hardware and network	<ul style="list-style-type: none"> ■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which the application will be installed <p>See “Configuring the storage hardware and network” on page 78.</p>
Review pre-requisites and install InfoScale Enterprise	<p>Install InfoScale Enterprise on all the systems where you want to configure SQL Server for high availability.</p> <p>For more information, see the <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p>
Review application-specific requirements	<p>See “Notes and recommendations for cluster and application configuration” on page 73.</p>
Configure disk groups and volumes for the SQL Server	<ul style="list-style-type: none"> ■ Create dynamic cluster disk groups using the Veritas Enterprise Administrator (VEA) (if using a shared storage configuration) ■ Create dynamic disk groups using VEA (for a single-node configuration using non-shared storage) ■ Create dynamic volumes for the SQL system database, user databases, transaction logs, and replicated registry keys for each instance <p>See “Configuring disk groups and volumes for SQL Server” on page 79.</p>
Configure the VCS cluster	<ul style="list-style-type: none"> ■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster <p>See “Configuring the cluster using the Cluster Configuration Wizard” on page 102.</p>
Review considerations before installing and configuring multiple instances of SQL Server	<p>If you are installing multiple instances of SQL Server, refer to the following topic.</p> <p>See “About installing multiple SQL Server instances” on page 122.</p>

Table 4-1 SQL Server: Active-Passive configuration tasks *(continued)*

Action	Description
Install and configure SQL Server on the first cluster node	<ul style="list-style-type: none"> ■ Ensure that the disk group and volumes for the system database are mounted on the first node ■ Follow the guidelines for installing SQL Server in the SFW HA environment <p>See “About installing and configuring SQL Server” on page 120.</p>
Install and configure SQL Server on the second or additional failover nodes	<ul style="list-style-type: none"> ■ Stop SQL Server services on the first node ■ Follow the guidelines for installing SQL Server on failover nodes in the SFW HA environment <p>See “About installing SQL Server on additional systems” on page 126.</p>
Create a SQL Server user-defined database	<ul style="list-style-type: none"> ■ Create volumes, if not created already, for a user-defined database and transaction log ■ Create a user-defined database in SQL Server <p>See “Creating a SQL Server user-defined database” on page 128.</p>
Create a SQL Server service group	<ul style="list-style-type: none"> ■ Create the application service group manually using templates from the Cluster Manager (Java Console) (if using a non-shared storage configuration) ■ Create a SQL Server service group using the SQL Server Configuration Wizard ■ Bring the service group online on the node where you ran the wizard to configure the service group. This is the first cluster node where you installed SQL Server. ■ For an active-active SQL Server configuration, ensure that the priority order of the systems in the service group for each instance is set up in reverse order <p>See “About configuring the SQL Server service group” on page 137.</p>
Configure fast failover for disk groups (optional)	<ul style="list-style-type: none"> ■ Ensure that you have installed the Fast Failover option and met the prerequisites for storage ■ Use the Java Console to enable the FastFailover attribute for VMDg resources. <p>See “Enabling fast failover for disk groups (optional)” on page 152.</p>
Perform additional configuration steps for multiple instances or disaster recovery configuration	<p>See “Completing configuration steps in SQL Server” on page 129.</p>
Verify the HA configuration	<p>Test failover between nodes</p> <p>See “Verifying the SQL Server cluster configuration” on page 153.</p>

Table 4-1 SQL Server: Active-Passive configuration tasks *(continued)*

Action	Description
In case of an active-active configuration, repeat the installation and configuration steps for the next SQL instance, or proceed to the additional steps depending on the desired HA configuration.	See “Determining additional steps needed” on page 164.

Tasks for configuring an existing server for high availability

You can convert an existing standalone SQL Server into a “clustered” SQL Server in a new Storage Foundation and High Availability Solutions environment. This environment involves an active-passive configuration with one to one failover capabilities.

Note: Some installation and configuration options are identified as required “for a disaster recovery configuration.” These options apply only if you intend to set up a secondary site for disaster recovery using Volume Replicator (Volume Replicator).

The following table outlines the high-level objectives and the tasks to complete each objective for converting an existing standalone application server for high availability.

Table 4-2 Task list: Standalone SQL Server HA configuration tasks

Action	Description
Review the HA configuration	<ul style="list-style-type: none"> ■ Understand active-passive configuration ■ Review the sample configuration <p>See “Reviewing a standalone SQL Server configuration” on page 46.</p>
Configure the storage hardware and network	<ul style="list-style-type: none"> ■ Set up the storage hardware for a cluster environment ■ Verify the DNS entries for the systems on which the application will be installed <p>See “Configuring the storage hardware and network” on page 78.</p>

Table 4-2 Task list: Standalone SQL Server HA configuration tasks
(continued)

Action	Description
Review pre-requisites and install InfoScale Enterprise	<p>Install InfoScale Enterprise on all the systems where you want to configure SQL for high availability.</p> <p>See the <i>Veritas InfoScale Installation and Upgrade Guide</i>.</p>
Review application-specific requirements	<p>See "Notes and recommendations for cluster and application configuration" on page 73.</p>
Configure disk groups and volumes for SQL Server	<ul style="list-style-type: none"> ■ Plan the storage layout ■ Create a dynamic cluster disk group using the Veritas Enterprise Administrator (VEA) ■ For a new shared storage configuration, create dynamic volumes for the SQL system database, user databases and transaction logs using VEA <p>See "Configuring disk groups and volumes for SQL Server" on page 79.</p>
Configure the VCS cluster	<ul style="list-style-type: none"> ■ Verify static IP addresses and name resolution configured for each node ■ Run the VCS Cluster Configuration Wizard (VCW) to configure cluster components and set up secure communication for the cluster ■ If the cluster is already configured, run VCW to add the application server systems to the cluster <p>See "Configuring the cluster using the Cluster Configuration Wizard" on page 102.</p>
Move SQL Server database and log files to shared storage	<ul style="list-style-type: none"> ■ Back up existing SQL data ■ Set SQL Server services to manual start ■ Stop SQL Server service ■ Modify data file and user database locations <p>See "Verifying that the SQL Server databases and logs are moved to shared storage" on page 122.</p>
Install and configure SQL Server on the second or additional failover nodes	<ul style="list-style-type: none"> ■ Stop SQL Server services on the first node ■ Optionally, rename system data files ■ Follow the guidelines for installing SQL Server on failover nodes in the SFW HA environment <p>See "About installing SQL Server on additional systems" on page 126.</p>

Table 4-2 Task list: Standalone SQL Server HA configuration tasks
(continued)

Action	Description
Create a SQL user defined database	<ul style="list-style-type: none"> ■ Create volumes, if not created already, for a user-defined database and transaction log ■ Create a new user-defined database in SQL Server ■ Add resources for a user-defined database in VCS <p>See “Creating a SQL Server user-defined database” on page 128.</p>
Create a SQL Server service group	<ul style="list-style-type: none"> ■ Create a SQL Server service group using the SQL Server Configuration Wizard ■ Bring the service group online on the node where you ran the wizard to configure the service group. This is cluster node where the shared drives containing the SQL Server database and log files are mounted. ■ For an active-active SQL Server configuration, ensure that the priority order of the systems in the service group for each instance is set up in reverse order <p>See “About configuring the SQL Server service group” on page 137.</p>
Configure fast failover for disk groups (optional)	<ul style="list-style-type: none"> ■ Ensure that you have installed the Fast Failover option and met the prerequisites for storage ■ Use the Java Console to enable the FastFailover attribute for VMDg resources. <p>See “Enabling fast failover for disk groups (optional)” on page 152.</p>
Verify the HA configuration	<p>Test fail over between nodes</p> <p>See “Verifying the SQL Server cluster configuration” on page 153.</p>

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, you specify the priority of the failover node while configuring the service group.

You use the VCS SQL Server Agent Configuration Wizard to configure a SQL Server 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 following topics describe requirements and procedures for creating the SQL Server service group:

- See [“Service group requirements for Active-Active configurations”](#) on page 138.

- See [“Prerequisites for configuring the SQL Server service group”](#) on page 138.
- See [“Creating the SQL Server service group”](#) on page 140.
- See [“Assigning privileges to the existing SQL Server databases and logs”](#) on page 151.
- See [“Enabling fast failover for disk groups \(optional\)”](#) on page 152.

Service group requirements for Active-Active configurations

Note the following requirements for Active-Active configurations:

- For an Active-Active configuration, you must create a separate service group for each instance.
- Each service group that you create must have a unique service group name and virtual IP address.
- For an Active-Active configuration, when you specify the priority order of systems, reverse the order for each service group so that the active system and failover system are opposite for each instance. For example, if you have two instances and two systems, you would set the priority order as follows:

INSTANCE 1	Priority order:
	SYSTEM1
	SYSTEM2
INSTANCE2	Priority order:
	SYSTEM2
	SYSTEM1

Prerequisites for configuring the SQL Server service group

Note the following prerequisites before configuring the SQL Server service group for a high availability cluster, campus cluster, or a Replicated Data Cluster:

- Verify that you have completed the steps in the high availability, campus cluster, or RDC workflows up through the step of installing SQL Server on all nodes.
 See the following topics as appropriate:
 See [“Tasks for configuring a new server for high availability”](#) on page 133.
 See [“Tasks for configuring an existing server for high availability”](#) on page 135.
 See [“VCS campus cluster configuration”](#) on page 50.
 See [“VCS Replicated Data Cluster configuration”](#) on page 57.

- Verify that you have VCS Cluster Administrator privileges. This user classification is required to configure service groups.
- You must be a local Administrator on the node where you run the wizard. If you wish to configure detail monitoring, you must have the permission to log on to the respective SQL Server instance.
- Verify that the SQL Server instance is installed identically on all nodes that will participate in the service group.
- For configuring a SQL Server service group, you must run the SQL Server Configuration Wizard from the first cluster node where you installed the SQL Server data directories on shared storage.
 This is required as the wizard configures the resources for the SQL Server database and registry information installed on the shared storage and propagates this information to the remaining nodes that are part of the SQL Server service group.
 Do not run the wizard from the additional nodes.

Note: If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage.

- Verify that the drive containing the SQL Server system data files, registry replication information, and FILESTREAM enabled data objects is mounted on the node on which you are configuring the service group. For creating a service group, this must be the first cluster node where you installed SQL Server.
 See [“About managing disk groups and volumes”](#) on page 98.
- If you wish to configure high availability for FILESTREAM, ensure that FILESTREAM is configured for the SQL Server instance, and is enabled for the SQL Server instance on the node on which you run the wizard and disabled on all the remaining nodes. You can use the SQL Server Configuration Manager to enable FILESTREAM.
 Refer to the Microsoft SQL Server documentation for instructions.
- Assign a unique virtual IP address for the SQL Server instance. You specify this IP address when configuring the service group.
- If you wish to use a script for detail monitoring, for example, to create a table and write data to it, note the location(s) of the script to use. Locate the script file.

In case of shared storage, ensure that the same file exists in the same location on all the cluster nodes.

See “[About the monitoring options](#)” on page 36.

A sample script is supplied in the following directory:

```
C:\Program Files\Veritas\Cluster
Server\bin\SQLServer\sample_script.sql
```

- 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 Server instances that you wish to configure in the service group.
 - If you have configured a Firewall, ensure that your firewall settings allow access to ports used by the product.
- For a detailed list of services and ports, refer to the *Veritas InfoScale Installation and Upgrade Guide*.

Creating the SQL Server service group

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

To create a SQL Server service group on the cluster

- 1 Ensure that you run the wizard from the first cluster node where you installed SQL Server.
- 2 Ensure that you have stopped the SQL Server service for the instances that you wish to configure.

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.

- 3 Start the SQL Server Configuration Wizard from the Solutions Configuration Center or from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** or, on Windows Server 2012 systems, from the **Apps** menu in the Start screen.
- 4 Review the prerequisites on the **Welcome** panel and then click **Next**.

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

For an active/active configuration, ensure that the active and failover systems are set differently for each instance. For example, if the system priority for the first instance is SYSTEM1, then SYSTEM2, the system priority for the second instance should be SYSTEM2, then SYSTEM1.
 - 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**.
- 7 On the **SQL Server Instance Selection** panel, 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 Server service groups
 - Select the SQL Server instance(s) that you wish to configure in the service group.

- If required, select the other services that you wish to make highly available. These options are available for selection only if the corresponding services are installed.

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

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

- Select **SQLFILESTREAM** if you wish to configure high availability for FILESTREAM enabled database objects. The wizard configures a resource only if FILESTREAM is enabled for the instance on the current node.

Note that FILESTREAM option will not appear for selection if it is not enabled on the node.

- 8 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.

- 9 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 Server 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.

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

The SQL Server 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:

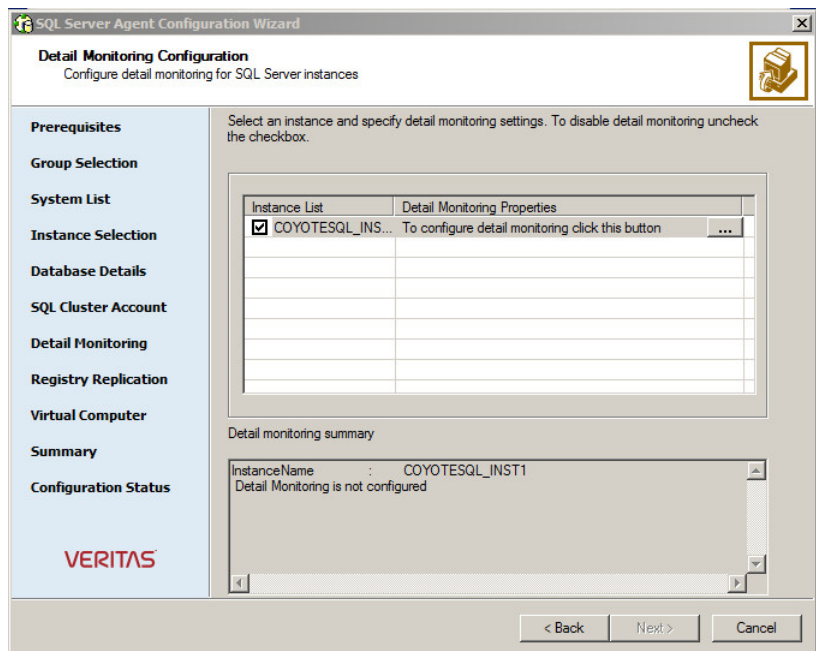
- Select a SQL Server instance in the **Instance Name** box.
- Check the **Configure SQL Server Cluster Account** check box.

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

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

The SQL Server 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.

- 11 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 Server instance, and then click the button from the Detail Monitoring Properties column to specify the detail monitoring settings.
 Clear the check box to disable detail monitoring for the instance.
- On the Detail Monitor configuration dialog box, specify the monitoring interval in the Detail monitoring interval field.
 This sets the value for the DetailMonitoringInterval attribute of the SQL 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 Server agent. In this mode of detail monitoring the agent monitors the health of the databases by connecting to those databases. The agent monitors only the databases specified in the DBList attribute.
- Select **SQLFile Detail Monitoring** if you wish to use a script to monitor SQL Server databases. In this mode of detail monitoring, the agent executes the script that you specify for detail monitoring.
- Specify the fully qualified user name and the password for connecting to the SQL Server database. Make sure that the user has SQL Server logon permissions.
- Select Global or Per System depending on whether the monitoring script location is the same for all the nodes or is unique for each cluster node, and then specify the path of the script appropriately.
- Check Fail over service group if detail monitoring fails check box, if not already checked. This allows the SQL Server agent to fail over the service group to another node if the detail monitoring fails.
- On the Detail Monitor configuration dialog box, specify the monitoring frequency in the Detail monitoring frequency field.
 This sets the value for the LevelTwoMonitorFreq 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 Server 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.
- 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 SQL script-based detail monitoring.

- Select **Global** or **Per System** depending on whether the monitoring script location is the same for all the nodes or is unique for each cluster node, and then specify the path of the script appropriately.
- Check **Fail over service group if detail monitoring fails** check box, if not already checked. This allows the VCS agent for SQL Server to fail over the service group to another node if the detail monitoring fails.
- Click **Apply**.
- Repeat these steps for each SQL Server instance that you wish to configure detail monitoring for, and then click **Next**.

- 12** 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 Server data be in separate volumes.

- 13** On the **Virtual Server Configuration** panel, specify the virtual server and network information and then click **Next**.

SQL Server Agent Configuration Wizard

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

Prerequisites

Group Selection

System List

Instance Selection

Database Details

SQL Cluster Account

Detail Monitoring

Registry Replication

Virtual Computer

Summary

Configuration Status

VERITAS

☒ IPv4 ☐ IPv6

Virtual Server Name: coyotesql_vname

Virtual IP Address: 10 . 209 . 104 . 251

Subnet Mask: 255 . 255 . 255 . 0

Specify the adapter to be used on each system.

System Name	Adapter Display Name
COYOTE	coyote_publicnic

Advanced Settings...

< Back Next > Cancel

Complete the following steps:

- 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 virtual 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.
- In the Virtual Server Name field 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. Select the **Adapter Display Name** field to view the adapters associated with a system. The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that

you select the adapters to be assigned to the public network, and not those assigned to the private network.

- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** check box, specify the desired Organizational Unit in the domain and then click **OK**. The user account configured for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

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

You can type the OU details in the format:

CN=Computers,DC=domainname,DC=com.

To search for the OU, click the ellipsis button and specify the search criteria in the Windows Find Organizational Units dialog box.

By default, the Lanman resource adds the virtual server name to the default container "Computers."

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 154.

14 In the **Service Group Summary** panel, review the service group configuration.

- The Resources box lists the configured resources. The wizard assigns unique names to resources based on their respective name rules. Click a resource to view its attributes and their configured values in the Attributes box.

Optionally, if desired, change the names of the resources as follows:

- To edit a resource name, click the resource name or press the F2 key. Press the **Enter** key after editing each resource name.
- To cancel editing a resource name, press the **Esc** key.
- To enable all the VMDg resources in the service group for fast failover, select the **Enable FastFailOver attribute for all the VMDg resources in the service group** checkbox.

For information about the FastFailOver attribute, see the *Cluster Server Administrator's Guide*.

- Click **Next**.

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

16 Check the **Bring the service group online** check box and then click **Finish**. This brings the service group online on the current node.

You must bring the SQL Server 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 VMDg and MountV resources to the service group by running the SQL Server Agent Configuration Wizard.

See [“Modifying a SQL Server service group to add VMDg and MountV resources”](#) on page 163.

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

Configuring the 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.

Assigning privileges to the existing SQL Server databases and logs

Note: The following steps are required only if you have configured the SQL Server 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 Server services and databases may not be accessible after a service group failover. For such a case, you configure the SQL Server cluster account while creating the SQL Server service group.

The SQL Server 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 Server 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 Server 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 Server cluster account with full control privileges to the following directories:
 - DATA
 - Log
- 3 Navigate inside the DATA folder and then assign the SQL Server 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 Server databases are accessible after a service group failover.

Enabling fast failover for disk groups (optional)

For service groups that contain many disk groups, you can greatly reduce failover time by implementing the SFW fast failover feature for disk groups.

More information is available about fast failover benefits and requirements.

See [“Considerations for a fast failover configuration”](#) on page 82.

For implementing the fast failover feature, VCS provides a new attribute, FastFailOver, for the Volume Manager Diskgroup (VMDg) resource. This attribute determines whether or not a disk group is enabled for fast failover.

Note: The disk group version must be 60 or later for fast failover to work. To verify the disk group version, from the VEA console, right-click the disk group and click **Properties**. Disk group version upgrade is required after upgrading SFW HA on the cluster nodes. Refer to the *Veritas InfoScale Installation and Upgrade Guide* for more information.

You can enable fast failover for all the VMDg resources while configuring the service group using the configuration wizard. The service group configuration wizard provides a checkbox to enable fast failover.

Perform these steps if you did not enable fast failover using the wizard or if you have configured the service group manually.

The following procedure describes how to enable the FastFailOver attribute using the VCS Java Console.

To enable the FastFailover attribute for a VMDg resource

- 1 In Cluster Manager (Java Console), select a service group with a VMDg resource configured for it.

Select the Properties tab from the right pane.
- 2 Scroll down to choose the **FastFailOver** attribute and click to edit the attribute value.
- 3 In the Edit Attribute dialog box, check the **FastFailOver** check box and then click **OK**.
- 4 Repeat these steps for every VMDg resource for which you want to enable fast failover.

Verifying the SQL Server cluster configuration

Simulating a failover is an important part of configuration testing. After completing the configuration, verify that failover occurs as desired.

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.

- 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 that you selected to switch to in the first step.
- 3 To move all the resources back to the original node, repeat the first step of this procedure 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, perform these steps sequentially:
 - Restart the node that you shut down in the first step.
 - 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.

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.

Tasks for configuring MSDTC for high availability

You can configure high availability for MSDTC either before or after configuring high availability for Microsoft SQL Server. The MSDTC agent comprises two parts, MSDTC Server and MSDTC client.

To configure high availability for MSDTC in a SQL server environment, you first use the MSDTC Configuration Wizard to create a service group for the MSDTC Server and then configure the MSDTC client manually.

Note: You cannot use the MSDTC Configuration Wizard to configure the MSDTC clients for SQL Server.

The following table outlines the high-level objectives and the tasks to complete each objective for an MSDTC configuration.

Table 4-3 Tasks for configuring MSDTC for high availability

Action	Description
Verifying hardware and software prerequisites	Review the compatibility lists for your specific Veritas InfoScale product and release on SORT at: https://sort.veritas.com/documents
Review the MSDTC configuration	<ul style="list-style-type: none">■ Understand MSDTC service group configuration■ Review the sample configuration See “ Reviewing the MSDTC configuration ” on page 48.
Configure disk groups and volumes for MSDTC	Configure cluster disk groups and volumes for an MSDTC Server service group See “ Configuring disk groups and volumes for SQL Server ” on page 79.
Create an MSDTC Server service group	Create an MSDTC Server service group See “ Configuring an MSDTC Server service group ” on page 156.
Configure the MSDTC client	Configure the MSDTC client See “ About configuring the MSDTC client for SQL Server ” on page 160.
Viewing DTC transactions	View DTC transaction lists and statistics See “ Viewing DTC transaction information ” on page 161.

Configuring an MSDTC Server service group

MSDTC is a global resource and can be accessed by more than one SQL Server service group. Veritas recommends that you configure only one MSDTC service group in a VCS cluster.

To configure high availability for MSDTC Server, you first use the MSDTC Configuration Wizard to create a service group for the MSDTC Server and then configure the MSDTC client manually.

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

Prerequisites for MSDTC configuration

Review the following prerequisites before configuring the MSDTC service group:

- Verify that the MSDTC service is installed on all nodes that will participate in the MSDTC Server 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 drives for the MSDTC logs and registry replication information are mounted on the node on which you are configuring the service group and unmounted on all other nodes.
- Verify that the Microsoft Distributed Transaction Coordinator (MSDTC) service is stopped.
- Keep the following information ready with you; the wizard prompts you for these details:
 - A unique virtual server name for the MSDTC Server. This is the name that is used by MSDTC clients to connect to the MSDTC Server. The DTC service runs under this name.
 - A unique virtual IP address for the MSDTC Server
 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 on the network. The wizard uses the network prefix that is advertised by the router on the IPv6 network.

Creating an MSDTC Server service group

Use the MSDTC Configuration Wizard (not the SQL Server Configuration Wizard) to configure a service group for the MSDTC Server. After configuring the service group, proceed to configuring the MSDTC client.

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

To configure an MSDTC Server service group

- 1 Start the MSDTC Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > MSDTC Configuration Wizard** or on Windows 2012 operating systems, from the **Apps** menu on the Start screen.
- 2 On the **SQL Configuration Option** panel, click **MSDTC Server - Service Group Configuration**, click **Create** and then click **Next**.

- 3 Review and verify that you have met the prerequisites and then click **Next**.
- 4 On the **Service Group Configuration** panel, specify the service group name and the system list, as follows:
 - Enter a name for MSDTC service group.
 - In the Available Cluster Systems box, select the systems on which to configure the service group and click the right-arrow to move the systems to the service group's system list. Make sure you select the systems that are not in the SystemList attribute for an Exchange service group that may be configured in the cluster.
 - To change a system's priority, in the Systems in Priority Order list, select the system and click the up and down arrows. 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.
 - 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 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 virtual server and network details and then click **Next**.
Complete the following steps:
 - Select **IPv4** to configure an IPv4 address for the MSDTC virtual server.
 - In the Virtual IP Address field, type a unique virtual IPv4 address for the virtual 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.
 - In the Virtual Server Name field enter a virtual server name for the node on which the DTC service is running. Ensure that the virtual server name you

enter is unique in the cluster. This is the name that is used by MSDTC clients to connect to the MSDTC Server.

- For each system in the cluster, select the public network adapter name. Click the **Adapter Display Name** field to view the adapters associated with a system. The wizard displays all TCP/IP enabled adapters on a system, including the private network adapters, if they are TCP/IP enabled. Make sure that you select the adapters to be assigned to the public network, and not those assigned to the private network.
- If you require a computer object to be created in the Active Directory (AD), click **Advanced Settings**, check the **Active Directory Update Required** check box, specify the desired Organizational Unit in the domain and then click **OK**. The user account configured for VCS Helper service must have adequate privileges on the specified container to create and update computer accounts.

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

You can type the OU details in the format

CN=Computers,DC=domainname,DC=com.

To search for the OU, click the ellipsis button and specify the search criteria in the Windows Find Organizational Units dialog box.

By default, the Lanman resource adds the virtual server name to the default container "Computers."

- 6** On the **Specify Data Path** panel, specify the volumes for the MSDTC log and the replication directory and then click **Next**.

Veritas recommends using different paths for these directories. If the directory does not exist, the wizard creates it.

- 7** On the **Service Group Summary** panel, review the service group configuration.

- Change the resource names if desired, as follows:
 - The Resources box lists the configured resources. Click on a resource to view its attributes and their configured values in the Attributes box.
 - The wizard assigns unique names to resources. Change names of the resources, if desired. To edit a resource name, select the resource name and either click it or press the **F2** key. Press **Enter** after editing each resource name. To cancel editing a resource name, press the **Esc** key.
- To enable all the VMDg resources in the service group for fast failover, select the **Enable FastFailOver attribute for all the VMDg resources in the service group** checkbox.

For information about the FastFailOver attribute, see the *Cluster Server Administrator's Guide*.

- Click **Next**.
- 8 Click **Yes** on the message that prompts you that the wizard will run commands to modify the service group configuration.
 Various messages indicate the status of these commands.
- 9 In the **Configuration Complete** panel, check **Bring the service group online** to bring the configured service group online and then click Finish to exit the wizard.
 To bring the service group online later, uncheck the option.

About configuring the MSDTC client for SQL Server

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

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

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

To configure an MSDTC client

- 1 Ensure that the MSDTC service group is online.
- 2 Launch the Windows Component Services Administrative tool.
 Click **Start > All Programs > Administrative Tools > Component Services**
 or
 Click **Start > Run**, type `dcomcnfg` and click **OK**.
- 3 In the console tree of the Component Services administrative tool, expand **Component Services > Computers**, right-click **My Computer** and then click Properties.
- 4 On the MSDTC tab, perform the following steps:

- Clear the **Use local coordinator** check box.
- In the Remote Host field, specify the virtual server name that you specified while creating the MSDTC Server service group.
If you are unsure of the exact name, click **Select** to search from a list of all computers on the network and select the virtual computer name from the list.
- Click **Apply** and then click **OK**.

Note: If you are using a non-shared storage configuration (dynamic disk groups configured on local disks), you have to configure the service group manually either using the Cluster Manager (Java Console) or the command line. The wizard currently cannot configure resources (VMNSDg agent) required for monitoring non-shared storage. See [“Configuring the service group in a non-shared storage environment”](#) on page 148.

About the VCS Application Manager utility

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 utility, the VCS Application Manager (VAM), 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.

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

or

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

or

On Windows 2012 operating systems, launch the wizard from the **Apps** menu in the Start screen.

The VCS Application Manager displays a list of supported application service groups configured in the cluster. For each service group it also displays the state of the service group, the name of the virtual server resource (Lanman resource) and the corresponding management tools used for that application.

- 2 Select **MSDTC** from the Select the resource type drop-down list.
- 3 Select the MSDTC resource that is online and then click **Manage**, or double-click the MSDTC resource name.

VAM launches the Component Services snap-in in the virtual server context.

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

You can use transaction statistics to get an overview of DTC performance.

Refer to the Microsoft documentation for further information.

To view transactions from any node in the domain

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

You can use transaction statistics to get an overview of DTC performance.

Refer to the Microsoft documentation for further information.

Modifying a SQL Server service group to add VMDg and MountV resources

If you create a new SQL Server database after you have created the SQL Server service group, you must rerun the SQL Server Configuration Wizard to modify the service group. This allows the wizard to add VMDg and MountV resources for the new databases, to the existing SQL Server service group.

You must run the wizard in the modify mode even if you have added or changed volumes in your existing configuration. This allows the wizard to make the necessary changes to the SQL Server service group.

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

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 add VMDg and MountV resources using the SQL Server Configuration Wizard

- 1** Start the SQL Server Configuration Wizard from the Solutions Configuration Center or click **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard** or, on Windows 2012 operating system, from the **Apps** menu in the Start screen.
- 2** Review the Prerequisites page and click **Next**.
- 3** On the Wizard Options panel, click Modify service group, select the service group, and then click **Next**.
- 4** Click **Yes** on the message 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**.

Databases that are marked with a red cross will not contain MountV resources.
- 8** If a database is not configured correctly, a 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 message indicates the configuration will be modified.
- 11** Click **Finish** to exit the wizard.

Determining additional steps needed

This completes the high availability configuration steps. Depending on the configuration being deployed, there are additional steps that you must perform to set up and complete the configuration.

The following table contains a list of references to the chapters that describe configuration specific tasks in detail. Proceed to the desired chapter depending on the desired configuration.

You must perform the configuration specific tasks only after you complete the high availability steps mentioned in this and the earlier chapters.

Table 4-4 Additional SQL Server configuration steps

Tasks	Refer to
Setting up a campus cluster configuration for SQL Server	See “Tasks for configuring campus clusters” on page 166.
Setting up a replicated data cluster configuration for SQL Server	See “Tasks for configuring Replicated Data Clusters” on page 170.
Setting up a disaster recovery configuration for SQL Server	See “Tasks for configuring disaster recovery for SQL Server” on page 227.
Configuring and running a fire drill for SQL Server configuration	See “About disaster recovery fire drills” on page 291.

Configuring campus clusters for SQL Server

This chapter includes the following topics:

- [Tasks for configuring campus clusters](#)
- [Modifying the IP resource in the SQL Server service group](#)
- [Verifying the campus cluster: Switching the service group](#)
- [Setting the ForceImport attribute to 1 after a site failure](#)

Tasks for configuring campus clusters

In campus clusters you begin by configuring a high availability cluster and then continue with the steps specific to the campus cluster configuration.

Refer to the campus cluster configuration workflow table for a complete list of configuration steps.

See [“VCS campus cluster configuration”](#) on page 50.

The following table shows the steps specific to the campus cluster configuration that are done after configuring high availability on the nodes.

Table 5-1 Completing campus cluster configuration

Action	Description
Modify the IP resource in the SQL Server service group	See “Modifying the IP resource in the SQL Server service group” on page 167.

Table 5-1 Completing campus cluster configuration (*continued*)

Action	Description
Verify the campus cluster configuration	Verify that failover occurs between the nodes. See “Verifying the campus cluster: Switching the service group” on page 168.
Set the ForceImport attribute	In case of a site failure, you may have to set the ForceImport attribute to ensure proper failover. See “Setting the ForceImport attribute to 1 after a site failure” on page 169.

Modifying the IP resource in the SQL Server service group

Note: This procedure is only applicable to a campus cluster with sites in different subnets.

Use the Java Console to modify the IP resource in the application service group. For IPv6 networks, modify the IPv6 resource.

Choose the appropriate procedure below depending on whether you are working with an IP resource or IPv6 resource.

To modify the IP resource

- 1** From the Cluster Explorer configuration tree, select the IP resource in the application service group.
- 2** In the **Properties View**, click the **Edit** icon for the **Address** attribute.
- 3** In the **Edit Attribute** dialog box, make the following selections:
 - Select the **Per System** option.
 - Select the system at Site B.
 - Enter the virtual IP address at Site B.
 - Click **OK**.
- 4** In the Properties View, click the **Edit** icon for the **SubNetMask** attribute.
- 5** In the Edit Attribute dialog box, make the following selections:
 - Select the **Per System** option.

- Select the system at Site B.
 - Enter the subnet mask at Site B.
 - Click **OK**.
- 6** From the **File** menu of Cluster Explorer, click **Close Configuration**.

To modify the IPv6 resource

- 1** From the Cluster Explorer configuration tree, select the IPv6 resource in the application service group.
- 2** In the Properties View, click the **Edit** icon for the **Address** attribute.
- 3** In the Edit Attribute dialog box, make the following selections:
 - Select the **Per System** option.
 - Select the system at Site B.
 - Enter the virtual IPv6 address at Site B.
 - Click **OK**.
- 4** In the Properties View, click the **Edit** icon for the **Prefix** attribute.
- 5** In the Edit Attribute dialog box, make the following selections:
 - Select the **Per System** option.
 - Select the system at Site B.
 - Enter the prefix at Site B.
 - Click **OK**.
- 6** From the File menu of Cluster Explorer, click **Close Configuration**.

Verifying the campus cluster: Switching the service group

Failover simulation is an important part of configuration testing.

To verify the campus cluster is functioning properly

- 1** Bring the service group online on one node as follows:
 - In the Cluster Explorer configuration tree, right-click the service group.
 - Click **Online**, and click the appropriate system from the menu.
- 2** Switch the service group to the other node as follows:
 - In the Cluster Explorer configuration tree, right-click the service group.

- Click **Switch To**, and click the appropriate system from the menu.

Setting the ForceImport attribute to 1 after a site failure

ForceImport is a flag that defines whether the agent forcibly imports the disk group when exactly half the disks are available. The value 1 indicates the agent imports the configured disk group when half the disks are available. The value 0 indicates it does not. Default is 0. This means that the disk group will be imported only when SFW acquires control over the majority of the disks.

Warning: Set this attribute to 1 only after verifying the integrity of your data. If due caution is not exercised before setting this attribute to 1, you risk potential data loss.

You must set the ForceImport attribute for the VMDg resource to 1 after a site failure to ensure proper failover.

To set the ForceImport attribute to 1 from the Java Console

- 1 From the Cluster Explorer configuration tree, select the VMDg resource in the application service group.
- 2 In the Properties View, click the **Edit** icon for the **ForceImport** attribute.
- 3 In the Edit Attribute dialog box, make the following selections:
 - Select the **Per System** option.
 - Select the system in Site B.
 - Select the **ForceImport** check box.
 - Click **OK**.
- 4 From the **File** menu of Cluster Explorer, click **Close Configuration**.
- 5 After the failover takes place, revert the ForceImport attribute to its original value.

To set the ForceImport attribute to 1 from the command line

- Use the following command for implementing the force import setting in VCS:

```
hares -modify vmdgResourceName ForceImport 1|0
```

Example:

```
hares -modify vmdg_Dg1 ForceImport 1
```

Import is forced on **vmdg_Dg1**.

Configuring Replicated Data Clusters for SQL Server

This chapter includes the following topics:

- [Tasks for configuring Replicated Data Clusters](#)
- [Creating the primary system zone for the application service group](#)
- [Creating a parallel environment in the secondary zone](#)
- [Setting up security for Volume Replicator](#)
- [Setting up the Replicated Data Sets \(RDS\)](#)
- [Configuring a RVG service group for replication](#)
- [Setting a dependency between the service groups](#)
- [Adding the nodes from the secondary zone to the RDC](#)
- [Verifying the RDC configuration](#)
- [Additional instructions for GCO disaster recovery](#)

Tasks for configuring Replicated Data Clusters

For a Replicated Data Cluster (RDC) you begin by configuring a high availability cluster on the primary zone systems.

You then continue with the steps specific to the RDC configuration.

The following table shows the steps specific to the RDC configuration that are done after configuring high availability on the primary zone.

Table 6-1 Completing the configuration of a Replicated Data Cluster

Action	Description
Create the primary system zone and then verify failover within the primary zone	<ul style="list-style-type: none"> ■ Create the primary system zone ■ Add the nodes to the primary zone <p>See “Creating the primary system zone for the application service group” on page 172.</p>
Create a parallel environment in the secondary zone	<ul style="list-style-type: none"> ■ Install the product on the systems in the secondary zone ■ Configure disk groups and volumes using the same names as on the primary zone ■ Install SQL Server following the prerequisites and guidelines for installing on the second zone. <p>See “Creating a parallel environment in the secondary zone” on page 173.</p>
Add the secondary zone systems to the cluster	<ul style="list-style-type: none"> ■ Add the secondary zone systems to the cluster. <p>See “Adding the nodes from the secondary zone to the RDC” on page 211.</p>
Set up security for Volume Replicator on all cluster nodes	<p>Set up security for Volume Replicator on all nodes in both zones.</p> <p>This step must be done before configuring Volume Replicator replication.</p> <p>See “Setting up security for Volume Replicator” on page 174.</p>
Set up the Replicated Data Set	<ul style="list-style-type: none"> ■ Use the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary zones <p>See “Setting up the Replicated Data Sets (RDS)” on page 177.</p>
Configure a RVG service group	<ul style="list-style-type: none"> ■ Create a Replicated Volume Group (RVG) service group ■ Configure the RVG service group <p>See “Configuring a RVG service group for replication” on page 188.</p>

Table 6-1 Completing the configuration of a Replicated Data Cluster
(continued)

Action	Description
Set a dependency between the service groups	<ul style="list-style-type: none"> ■ Set up a dependency from the RVG service group to the SQL Server service group <p>See “Setting a dependency between the service groups” on page 209.</p>
Add the nodes from the secondary zone to the RDC	<ul style="list-style-type: none"> ■ Add the nodes from the secondary zone to the RVG service group ■ Configure the IP resources for failover ■ Add the nodes from the secondary zone to the SQL Server service group <p>See “Adding the nodes from the secondary zone to the RDC” on page 211.</p>
Verify the RDC configuration	<ul style="list-style-type: none"> ■ Verify that failover occurs first within zones and then from the primary to the secondary zone <p>See “Verifying the RDC configuration” on page 223.</p>

Creating the primary system zone for the application service group


In the service group, set up systems in the primary zone (zone 0) to specify that initial failover occurs to systems within the primary zone.

To set up the primary system zone

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 Select the SQL Server service group (INST1_SG) in the left pane and the **Properties** tab in the right pane.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.

- 6 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 0) for the primary zone. Make sure you specify the systems in uppercase.

In case of a non-shared storage configuration, add only the single node to the primary zone.



- 7 Click **OK**.
- 8 After setting up the primary system zone, you can verify the service group failover on systems within the primary zone.

See [“Verifying the RDC configuration”](#) on page 223.

Creating a parallel environment in the secondary zone

After setting up a SFW HA environment in the primary zone, you set up a parallel environment in the secondary zone (zone 1).

Before you begin to configure the secondary zone, do the following:

- Offline the following resources in the SQL Server service group in the primary zone:
 - SQL Server resource (*sql/ServiceGroupName* - SQLServer)
 - SQL Virtual Server name resource (*sql/ServiceGroupName* - Lanman)

- SQL Virtual IP resource (sqlServiceGroupName - IP)

The remaining resources should be online, including the storage resources.

- In VEA, make sure to remove all the drive letters from the configured volumes, to avoid conflicts when configuring the zones.

Then complete the following tasks to configure the secondary zone, using the guidelines shown:

- See [“Configuring the storage hardware and network”](#) on page 78.
- Installing the Veritas InfoScale product.
For more information, see the *Veritas InfoScale Installation and Upgrade Guide*.
- See [“Configuring disk groups and volumes for SQL Server”](#) on page 79.
During the creation of disk groups and volumes for the secondary zone, make sure the following is exactly the same as the cluster at the primary zone:
 - Disk group name
 - Volume sizes
 - Volume names
 - Drive letters
- Installing and configuring SQL Server on the first cluster node
See [“About installing SQL Server on the first system”](#) on page 125.
When installing SQL Server make sure that you select the same installation options as you did for the primary zone. The instance name must be the same in the primary zone and secondary zone
- See [“About installing SQL Server on additional systems”](#) on page 126.
After you install SQL Server on the nodes in the secondary zone, make sure to use VEA to remove all the drive letters from the configured volumes.
- Adding a system to the existing cluster
You do not create another cluster in the secondary zone. Instead you add the systems to the existing cluster.

You do not create another SQL Server service group in the secondary zone. You continue with the remaining Volume Replicator configuration tasks, during which the secondary zone nodes will be added to the SQL Server service group.

Setting up security for Volume Replicator

If you use Volume Replicator for replication, you must configure the Veritas Volume Replicator Security Service (VxSAS) on all the cluster nodes.

For details on this required service, see the *Volume Replicator Administrator's Guide*.

After you install InfoScale Storage or InfoScale Enterprise, launch the Veritas Volume Replicator Security Service Configuration Wizard. This wizard lets you complete the Volume Replicator security service configuration.

Prerequisites for configuring VxSAS

- The wizard requires you to be logged on with administrative privileges.
- The account that you specify must have administrative and log-on as service privileges on all the specified hosts.
- Avoid specifying blank passwords. In a Windows Server environment, accounts with blank passwords are not supported for log-on service privileges.
- The systems on which you want to configure VxSAS must be accessible from the local system.

To configure VxSAS

- 1 Launch the Veritas Volume replicator Security Service Configuration Wizard from **Start > All Programs > Veritas > Veritas Storage Foundation > Configuration Wizards > VVR Security Service Configuration Wizard** or, on Windows 2012 operating systems, from the **Apps** menu in the **Start** screen.

Optionally, run `vxsascfg.exe` from the command prompt of the required machine.

- 2 Read the information provided on the Welcome page and click **Next**.
- 3 Complete the Account Information panel as follows:

Account name (domain\account)	Enter the administrative account name.
----------------------------------	--

Password	Specify a password
----------	--------------------

If you have already configured VxSAS for one host that is intended to be a part of the RDS, make sure you specify the same username and password when configuring VxSAS on the other hosts.

Click **Next**.

- 4 On the Domain Selection panel, select the domain to which the hosts that you want to configure belong:

Selecting domains	<p>The Available domains pane lists all the domains that are present in the Windows network neighborhood.</p> <p>Move the appropriate name from the Available domains list to the Selected domains list, either by double-clicking it or using the arrow button.</p>
Adding a domain	<p>If the domain name that you require is not displayed, click Add domain. This displays a dialog that lets you specify the domain name. Click Add to add the name to the Selected domains list.</p>

Click **Next**.

- 5 On the Host Selection panel, select the required hosts:

Selecting hosts	<p>The Available hosts pane lists the hosts that are present in the specified domain.</p> <p>Move the appropriate host from the Available hosts list to the Selected hosts list, either by double-clicking it or using the arrow button. Use the Shift key with the up or down arrow keys to select multiple hosts.</p>
Adding a host	<p>If the host name you require is not displayed, click Add host. In the Add Host dialog specify the required host name or IP in the Host Name field. Click Add to add the name to the Selected hosts list.</p>

After you have selected a host name, the **Configure** button is enabled. Click **Configure** to proceed with configuring VxSAS.

- 6 After the configuration completes, the Configuration Results page displays whether or not the operation was successful. If the operation was not successful, the page displays the details on why the account update failed, along with the possible reasons for failure and recommendations on getting over the failure.

When configuring VxSAS in a firewall setup, the VxSAS wizard may not be able to configure the machines that are across the firewall, although the Host Selection dialog may list these nodes. In this case, configure VxSAS locally on the machines that are across the firewall.

Click **Back** to change any information you had provided earlier.

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

Setting up the Replicated Data Sets (RDS)

Set up the Replicated Data Sets (RDS) in the primary zone (zone0) and secondary zone (zone1). You can configure an RDS for both zones using the Setup Replicated Data Set Wizard.

Prerequisites for setting up the RDS for the primary and secondary zones

Before you run the Setup Replicated Data Set Wizard, verify the following:

- Verify that the intended Primary host is connected to VEA, if you are configuring the RDS from a remote client or from a host that is not the Primary.
- Verify whether the IP version preference is set before you configure replication. If you specify host names when you configure replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP version Volume Replicator uses to resolve the host names.

Use one of the following methods to set the IP preference:

- Veritas Enterprise Administrator (VEA) GUI—select the appropriate options on the Control Panel > VVR Configuration > IP Settings tab.
- Run the `vxtune ip_mode [ipv4 | ipv6]` command at the primary site as well as the secondary site.
- Verify that the data volumes are not of the following types as Volume Replicator does not support these types of volumes:
 - Storage Foundation (software) RAID 5 volumes
 - Volumes with a Dirty Region Log (DRL)
 - Volumes that are already part of another RVG
For the Replicator Log volume, in addition to the above types, make sure that the volume does not have a DCM
 - Volumes names containing a comma
 - Secondary volume of a size smaller or greater than that on the Primary
- Verify that the disk group is imported and the volumes are mounted in the primary and secondary zone.
- Verify that you have configured security for Volume Replicator.
Verify that the VxSAS account has been configured with the same username and password for all the hosts, which are intended to be a part of the same RDS.

See [“Setting up security for Volume Replicator”](#) on page 174.

Creating the Replicated Data Sets with the wizard

To create the Replicated Data Set

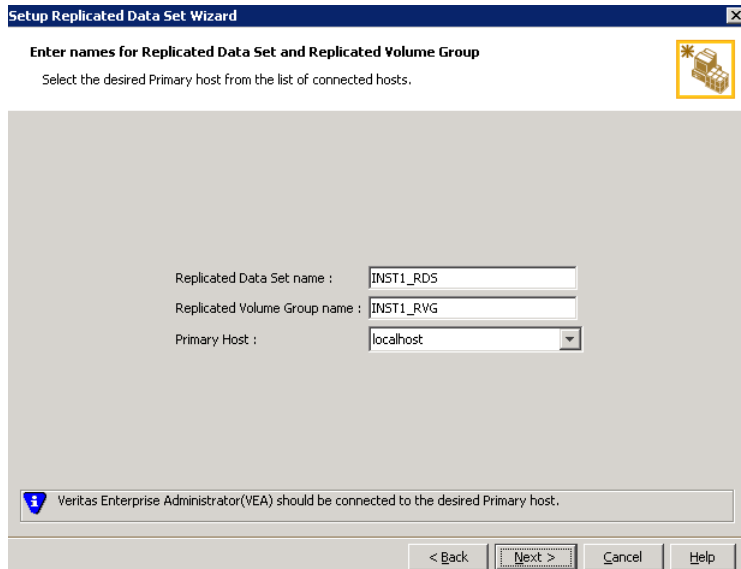
- 1 Use the Veritas Enterprise Administrator (VEA) console to launch the Setup Replicated Data Set Wizard from the cluster node on the Primary where the cluster disk group is imported.

Start VEA from **Start > All Programs > Veritas > Veritas Storage Foundation > Veritas Enterprise Administrator**.

On Windows 2012 operating systems, from the **Apps** menu in the Start screen.

From the VEA console, click **View > Connection > Replication Network**.

- 2 Right-click **Replication Network** and select **Setup Replicated Data Set**.
- 3 Read the information on the Welcome page and then click **Next**.
- 4 Specify names for the Replicated Data Set (RDS) and Replicated Volume Group (RVG) and then click **Next**.



Setup Replicated Data Set Wizard

Enter names for Replicated Data Set and Replicated Volume Group

Select the desired Primary host from the list of connected hosts.

Replicated Data Set name : INST1_RD5

Replicated Volume Group name : INST1_RVG

Primary Host : localhost

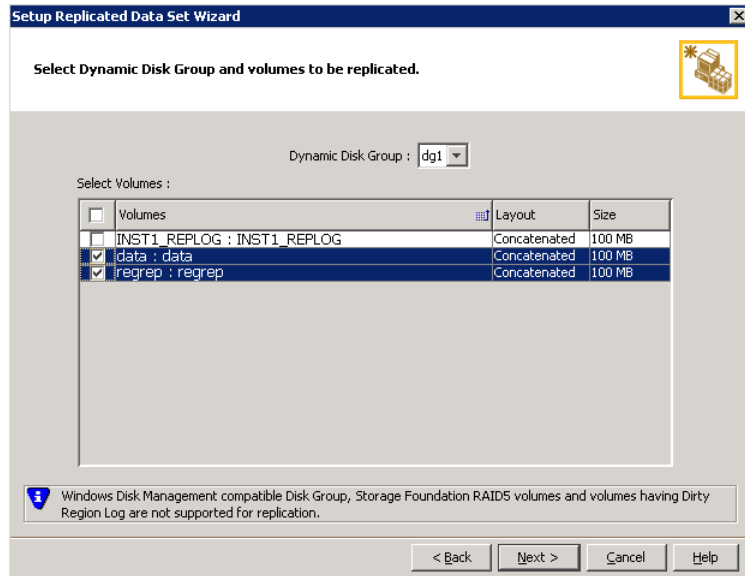
Veritas Enterprise Administrator(VEA) should be connected to the desired Primary host.

< Back Next > Cancel Help

By default, the local host is selected as the Primary Host. To specify a different host name, make sure the required host is connected to the VEA console and select it in the Primary Host list.

If the required primary host is not connected to the VEA console, it does not appear in the drop-down list of the Primary Host field. Use the VEA console to connect to the host.


- 5 Select from the table the dynamic disk group and data volumes that will undergo replication and then click **Next**.



To select multiple volumes, press the Shift or Control key while using the up or down arrow keys.

By default, a mirrored DCM log is automatically added for all selected volumes. If disk space is inadequate to create a DCM log with two plexes, a single plex is created.

- 6 Complete the Select or create a volume for Replicator Log page as follows:



To select an existing volume

- Select the volume for the Replicator Log in the table (INST1_REPLOG). If the volume does not appear in the table, click Back and verify that the Replicator Log volume was not selected on the previous page.
- Click **Next**.

To create a new volume

- Click Create Volume and enter the following information in the dialog box that appears:

Name	Enter the name for the volume in the Name field.
Size	Enter a size for the volume in the Size field.
Layout	Select the desired volume layout.

Disk Selection Enables you to specify the disk selection method.

- Enable the **Thin Provisioned Disks Only** check box to ensure that the Replicator Log volume is created only on Thin Provisioned (TP) disks.

Note: The check box will remain disabled if the diskgroup does not have any TP disk.

If this option is selected along with the **Select disks automatically** option, then the Replicator Log volume will be created only on TP disks. However, if you enable this check box along with **Select disks manually** option, then the user can select only TP disks from Available Disks.

For more information on Thin Provisioning, refer to the *Storage Foundation Administrator's Guide*.

- Choose the **Select disks automatically** option if you want Volume Replicator to select the disks.
- Choose the **Select disks manually** option to use specific disks from the Available disks pane for creating the volume. Either double-click on it or select **Add** to move the disks into the Selected disks pane.

- Click **OK** to create the Replicator Log volume.
- Click **Next** in the Select or create a volume for Replicator Log dialog box.

7 Review the information on the summary page and click **Create Primary RVG**.

8 After the Primary RVG has been created successfully, Volume Replicator displays the following message:

```
RDS with Primary RVG has been created successfully.
```

```
Do you want to add Secondary host to this RDS for replication now?
```

Click **No** to exit the Setup Replicated Data Set wizard without adding the Secondary host. To add the Secondary host later, use the **Add Secondary** option from the RDS right-click menu.

Click **Yes** to add the Secondary host to the Primary RDS now. The Specify Secondary host for replication page appears.

- 9 On the Specify Secondary host for replication page, enter the name or IP address of the Secondary host in the **Secondary Host** field and then click **Next**.

If the Secondary host is not connected to VEA, the wizard tries to connect it when you click **Next**. This wizard allows you to specify only one Secondary host. Additional Secondary hosts can be added using the **Add Secondary** option from the RDS right-click menu.

Wait till the connection process is complete and then click **Next** again.

- 10 If only a disk group without any data volumes or Replicator Log, as on the Primary host exists on the Secondary, then Volume Replicator displays a message. Read the message carefully.

The option to automatically create volumes on the Secondary host is available only if the disks that are part of the disk group have:

- The same or larger amount of space as that on the Primary
- Enough space to create volumes with the same layout as on the Primary
Otherwise, the RDS setup wizard enables you to create the required volumes manually.
- Click **Yes** to automatically create the Secondary data volumes and the Replicator Log.
- Click **No** to create the Secondary data volumes and the Replicator Log manually, using the Volume Information on the connected hosts page.

- 11 The Volume Information on connected hosts page appears. This page displays information on the availability of volumes on the Secondary nodes, if the Primary and Secondary hosts are connected to VEA.

This page does not appear if all the required volumes that are available on the Primary host are also available on the Secondary hosts.

- If the required data volumes and the Replicator Log have not been created on the Secondary host, then the page displays the appropriate message against the volume name on the Secondary.
- If an error occurs or a volume needs to be created, a volume displays with a red icon and a description of the situation. To address the error, or to create a new Replicator Log volume on the secondary site, click the volume on the secondary site, click the available task button and follow the wizard. Depending on the discrepancies between the volumes on the primary site and the secondary site, you may have to create a new volume, recreate or resize a volume (change attributes), or remove either a DRL or DCM log. When all the replicated volumes meet the replication requirements and display a green check mark, click **Next**.

- If all the data volumes to be replicated meet the requirements, this screen does not occur.
- 12** Complete the Edit replication settings page to specify the basic and advanced replication settings for a Secondary host as follows:

Setup Replicated Data Set Wizard

Edit replication settings
 Edit replication settings or click next.

Primary side IP: 10.217.53.214

Secondary side IP: 10.217.53.215

Replication Mode: Synchronous Override

Replicator Log Protection: AutoDCM

Primary RLINK Name: Pri_RLINK

Secondary RLINK Name: Sec_RLINK

Advanced

DHCP addresses are not supported by VWR.

< Back Next > Cancel Help

- To modify each of the default values listed on this page, select the required value from the drop-down list for each property. If you do not wish to modify basic properties then replication can be started with the default values when you click **Next**.

Primary side IP Enter the virtual IP address for the Primary IP resource that will be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Secondary side IP Enter the virtual IP address on the Secondary that is to be used for replication. If there is more than one IP address available for replication, you can choose the one that you want to use from the drop-down list. If the required IP address is not displayed in the list then edit the field to add the IP address.

Replication Mode Select the required mode of replication:

- **Synchronous Override** (default) enables synchronous updates under typical operating conditions. If the Secondary site is disconnected from the Primary site, and write operations occur on the Primary site, the mode of replication temporarily switches to Asynchronous.
- **Synchronous** determines updates from the application on the Primary site are completed only after the Secondary site successfully receives the updates.
- **Asynchronous** determines updates from the application on the Primary site are completed after Volume Replicator updates in the Replicator Log. From there, Volume Replicator writes the data to the data volume and replicates the updates to the secondary site asynchronously.

If the Secondary is set to the synchronous mode of replication and is disconnected, the Primary data volumes with NTFS file systems may be displayed with the status as missing.

Replicator Log
Protection

- **AutoDCM** is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.
- The **DCM** option enables the Replicator Log protection for the Secondary host when the Replicator Log overflows, and the connection between the Primary and Secondary is lost. This option is available only if all the data volumes under the Primary RVG have a DCM Log associated with them.
- The **Off** option disables Replicator Log Overflow protection. In the case of the Bunker node. Replicator Log protection is set to **Off**, by default. Thus, if the Primary RLINK overflows due to the Bunker RLINK, then this RLINK is detached.
- The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.
 If the Secondary becomes inactive due to disconnection or administrative action then Replicator Log protection is disabled, and the Replicator Log overflows.
- The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between Primary and Secondary RVG is broken, then, any new writes to the Primary RVG are failed.


Primary RLINK
Name

This option enables you to specify a Primary RLINK name of your choice. If you do not specify any name then Volume Replicator assigns a default name.

Secondary RLINK
Name

This option enables you to specify a Secondary RLINK name of your choice. If you do not specify any name then Volume Replicator assigns a default name.

- If you want to specify advanced replication settings, click **Advanced**. Edit the replication settings for a secondary host as needed.



The image shows a Windows-style dialog box titled "Advanced Replication Settings". It contains several configuration options:

- Latency Protection:** A dropdown menu currently set to "Off".
- High Mark Value:** An empty text input field.
- Low Mark Value:** An empty text input field.
- Protocol:** A dropdown menu currently set to "UDP/IP".
- Packet Size(Bytes):** A dropdown menu currently set to "1400".
- Bandwidth:** A dropdown menu set to "Maximum Available" next to a text input field with "Mbps" as a unit.
- Enable Compression:** An unchecked checkbox.
- Buttons:** "OK", "Cancel", and "Help" buttons at the bottom.

Caution: When determining the high mark and low mark values for latency protection, select a range that is sufficient but not too large to prevent long durations of throttling for write operations.

Latency protection Determines the extent of stalling write operations on the primary site to allow the secondary site to “catch up” with the updates before new write operations can occur.

Off is the default option and disables latency protection.

Fail enables latency protection. If the number of outstanding write operations reaches the **High Mark Value** (described below), and the secondary site is connected, Volume Replicator stalls the subsequent write operations until the number of outstanding write operations is lowered to the **Low Mark Value** (described below). If the secondary site is disconnected, the subsequent write operations fail.

Override enables latency protection. This option resembles the Off option when the secondary site is disconnected, and the Fail option when the secondary site is connected.

Throttling of write operations affects application performance on the primary site; use this protection only when necessary according to replication throughput and application write patterns.

High Mark Value	Is enabled only when either the Override or Fail latency protection option is selected. This value triggers the stalling of write operations and specifies the maximum number of pending updates on the Replicator Log waiting for replication to the secondary site. The default value is 10000, the maximum number of updates allowed in a Replicator Log.
Low Mark Value	Is enabled only when either the Override or Fail latency protection options is selected. After reaching the High Mark Value, write operations on the Replicator Log are stalled until the number of pending updates drops to an acceptable point at which the secondary site can "catch up" to the activity on the primary site; this acceptable point is determined by the Low Mark Value. The default value is 9950.
Protocol	UDP/IP is the default protocol for replication.
Packet Size	Updates to the host on the secondary site are sent in packets; the default size 1400 bytes. The option to select the packet size is enabled only when UDP/IP protocol is selected.
Bandwidth	By default, Volume Replicator uses the maximum available bandwidth. To control the bandwidth used, specify the bandwidth limit in Mbps.
Enable Compression	Enable this checkbox if you want to enable Compression for the secondary host.

Click **OK** to close the dialog box and then click **Next**.

- 13** On the Start Replication page, choose the appropriate option as follows:
- To add the Secondary and start replication immediately, select **Start Replication** with one of the following options:

Synchronize
Automatically

If virtual IPs have been created, select the **Synchronize Automatically** option, which is the default recommended for initial setup to start synchronization of Secondary and start replication immediately.

If the virtual IPs for replication are not yet created, automatic synchronization remains paused and resumes after the Replication Service Group is created and brought online.

When this option is selected, Volume Replicator by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

Note: Intelligent synchronization is applicable only to volumes with the NTFS and ReFS file systems and not to raw volumes or volumes with FAT/FAT32 file systems.

Synchronize from
Checkpoint

If you want to use this method, then you must first create a checkpoint.

If you have considerable amount of data on the Primary data volumes, then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the Synchronize from Checkpoint option to start replication from checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

For information on synchronizing from checkpoints, refer *Volume Replicator Administrator's Guide*.

- To add the secondary without starting replication, deselect the **Start Replication** option. You can start replication later by using the **Start Replication** option from the Secondary RVG right-click menu.
- Click **Next** to display the Summary page.

14 Review the information.

Click **Back** to change any information you had specified.

Otherwise, click **Finish** to add the secondary host to the RDS and exit the wizard.

Configuring a RVG service group for replication

If you are setting up a RDC configuration, create and configure a hybrid Replicated Volume Group (RVG) service group for replication. The RVG service group is hybrid

because it behaves as a failover service group within a zone and as a parallel service group between zones.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG service group.

For additional information about service group types, see the *Cluster Server Administrator's Guide*.

Configure the RVG service group's resources manually by copying and modifying components of the SQL Server service group. Then, create new RVG resources and bring them online.

The following table shows the resources in the RVG service group for replication.

Table 6-2 Replication service group resources

Resource	Description
IP	IP address for replication
NIC	Associated NIC for this IP
VMDg (shared storage) or VMNSDg (non-shared storage) for the system files disk group	Disk group with SQL Server system files
VvrRvg for the system files disk group	VvrRvg for the system files disk group Replicated volume group with SQL Server system files
VMDg (shared storage) or VMNSDg (non-shared storage) for the user-defined database disk group	Disk group with SQL Server user-defined files
VvrRvg for the user-defined database disk group	Replicated volume group with SQL Server user-defined files

Creating the RVG service group

Create a hybrid replicated volume (RVG) service group, to contain the resources for replication.

Note: If you are creating a DR configuration manually in a non-shared storage environment, create a failover type of RVG service group.

To create a RVG service group

- 1 From VCS Cluster Manager (Java Console), log on to the cluster.
- 2 In the VCS Cluster Explorer window, right-click the cluster in the left pane and select **Add Service Group**.
- 3 In the **Add Service Group** window, specify the following:

Add Service Group

Service Group name:

Available Systems		Systems for Service Group		
		System name	Startup	Priority
SYSTEM1	▶ ◀	SYSTEM1	<input type="checkbox"/>	0
SYSTEM2		SYSTEM2	<input type="checkbox"/>	1

Service Group Type: ☐ Failover ☐ Parallel ☒ Hybrid

Selected Template: None Templates ...

Show Command OK Cancel

- Enter a name for the service group. Make sure the service group name is in uppercase.
For example, enter `INST1_RVG_SG`.
- Select the systems in the primary zone (zone 0) and click the right arrow to add them to the service group.
- Select **Hybrid**.
If you creating the RVG service group for a DR configuration in a non-shared storage environment, select **Failover**.
- Click **OK**.

Note: If you are setting up replication in a non-shared storage environment, you can use the replication service group template, **VvrRvgVMNSRVGGroup**, available in the Java Console. For an RDC configuration, ensure that you select the service group type as **Hybrid** while creating the service group using the Configuration Wizard from Java Console.

Configuring the resources in the RVG service group for RDC replication

Configure the RVG service group's resources manually for RVG by completing the following tasks:

- Copy IP and NIC resources of the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
See [“Configuring the IP and NIC resources”](#) on page 191.
- Copy the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resources for the disk groups in the SQL Server service group (INST1_SG), paste and modify them for the RVG service group (INST1_RVG_SG).
See [“Configuring the VMDg or VMNSDg resources for the disk groups”](#) on page 194.
See [“Configuring the VMDg or VMNSDg resources for the disk group for the user-defined database”](#) on page 197.
- Create the Volume Replicator RVG resources for the disk groups and enter the attributes for each of the disk groups and the replication IP address.
See [“Adding the Volume Replicator RVG resources for the disk groups”](#) on page 199.
- Link the Volume Replicator RVG resources to establish the dependencies between the VMDg or VMNSDg resources, the IP resource for replication, and the Volume Replicator RVG resources for the disk groups. Configure the RVG service group's VMDg or VMNSDg resources to point to the disk groups that contain the RVGs.
See [“Linking the Volume Replicator RVG resources to establish dependencies”](#) on page 203.
- Delete the VMDg or VMNSDg resources from the SQL Server service group, because they depend on the replication and were configured in the RVG service group.
See [“Deleting the VMDg or VMNSDg resource from the SQL Server service group”](#) on page 205.

Configuring the IP and NIC resources

Configure the following resources and attributes for the IP and NIC agents.

The following table shows the resource attributes to modify.

Table 6-3 IP and NIC resources


Resource	Attributes to modify
IP	Address
NIC	(none)

Note: In a non-shared storage environment, if you use the Java Console template, **VvrRvgVMNSRVGGroup**, to create the RVG service group, then do not recreate these resources; you modify the attributes of the existing IP and NIC resources in the service group.

To create the IP resource and NIC resource

- 1 In the **VCS Cluster Explorer** window, select the **SQL Server service group (INST1_SG)** in the left pane.
- 2 On the **Resources** tab, right-click the **IP resource (INST1_SG-IP)**, and click **Copy > Self and Child Nodes**.
- 3 In the left pane, select the **RVG service group (INST1_RVG_SG)**.
- 4 On the **Resources** tab, right-click in the blank resource display area and click **Paste**.

- 5 In the **Name Clashes** window, change the names of the IP and NIC resources for the RVG service group.



- 6 Click **OK**.

To modify the IP resource and NIC

- 1 In the Resources tab display area, right-click the IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.
- 2 In the Properties View window, for the **Address** attribute, click **Edit**.
- 3 In the Edit Attribute window, enter the Volume Replicator IP address for the Primary Zone as the scalar value. This is the IP address you specified as the Primary side IP address while configuring the Replicated Data Set (RDS) earlier using the RDS wizard.
- 4 Close the Properties View window.

To enable the IP resource and NIC

- 1 In the Resources tab display area, right-click the IP resource (INST1_RVG_SG-IP) and select **Enabled**.
- 2 In the Resources tab display area, right-click the NIC resource (INST1_RVG_SG-NIC) and select **Enabled**.

Configuring the VMDg or VMNSDg resources for the disk groups

Configuration involves the following tasks:

- Create the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resource in the RVG service group by copying it from the SQL Server service group and renaming it.
- Modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service group to ensure the desired failover behavior.
- Modify the attributes of the MountV resources in the SQL Server service group for the new VMDg or VMNSDg in the RVG service group.
- Repeat these procedures for any additional VMDg or VMNSDg resources that you want to create for replication.
- If you are creating a DR configuration in a non-shared storage environment, modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service groups at both the sites (primary site and the disaster recovery site) separately.

Note: The MountV resources correspond to the volumes that you are configuring for replication. The table shows an example configuration. You may have additional volumes you want to include for replication, such as a FILESTREAM volume for SQL Server.

The following table shows the MountV resources and attributes to configure for the example configuration.

Table 6-4 MountV resources and attributes to modify

Resource	Attributes to modify
MountV (for the SQL Server system volume)	VMDg Resource Name Volume Name
MountV (for the registry volume)	VMDg Resource Name Volume Name
MountV (for the SQL Server user-defined database log)	VMDg Resource Name Volume Name

Table 6-4 MountV resources and attributes to modify (continued)

Resource	Attributes to modify
MountV	VMDg Resource Name
(for the SQL Server user-defined database)	Volume Name

To create the VMDg or VMNSDg resource in the RVG service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 On the Resources tab, right-click the VMDg or VMNSDg resource for the disk group that you want to configure for the RVG and click **Copy > Self**.
For example, right-click INST1_SG-VMDg or INST1_SG-VMNSDg.
- 3 In the left pane, select the RVG service group (INST1_RVG_SG).
- 4 On the Resources tab, right-click in the blank resource display area and click **Paste**.
- 5 In the Name Clashes window, change the name of the VMDg or VMNSDg resource for the RVG service group.
For example change INST1_SG-VMDg to INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg.
- 6 Click **OK**.

Note: Modify the DGGuid attribute of the new VMDg or VMNSDg resource before you perform this next procedure.

See [“Modifying the DGGuid attribute for the new disk group resource in the RVG service group”](#) on page 196.

To modify the MountV resources in the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) in the left pane.
- 2 In the Resources tab display area, right-click the MountV resource for the SQL Server system data files (INST1_SG-MountV) and select **View > Properties View**.
- 3 In the Properties View window, verify that the **Volume Name** attribute is the SQL Server system data files (INST1_DATA_FILES).
- 4 In the same Properties View window, for the **VMDg Resource Name** attribute, click **Edit**.

- 5 In the Edit Attribute window, modify the **VMDGResName** scalar value to be the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resource that was just created in the RVG service group.

For example, INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg.
- 6 Close the Properties View window.
- 7 In the Resources tab display area, right-click the MountV resource for the registry volume (INST1_SG-MountV-1) and select **View > Properties View**.
- 8 In the Properties View window, verify that the **Volume Name** attribute is the registry volume (INST1_REGREP_VOL).
- 9 In the same Properties View window, for the VMDg Resource Name attribute, click **Edit**.
- 10 In the Edit Attribute window, modify the **VMDGResName** scalar value to be the VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resource that was just created.

For example INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg .
- 11 Close the Properties View window.

To enable the VMDg or VMNSDg resource in the RVG service group

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the Resources tab display area, right-click the VMDg or VMNSDg resource (INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg) and select **Enabled**.

Modifying the DGGuid attribute for the new disk group resource in the RVG service group

To modify the DGGuid attribute for the new VMDg or VMNSDg resource in the RVG service group

- 1 From the VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the Resources tab.
- 3 In the Resources tab display area, right-click the new VMDg or VMNSDg resource and click **View > Properties View**.
- 4 In the Properties View window, locate the DGGuid attribute and click the edit icon.
- 5 In the Edit Attribute window, edit the attribute by performing the following actions sequentially:
 - Select **Per System**.
 - From the dropdown list select the first node in the primary zone (Zone 0).

- In the **Scalar Value** field specify the GUID of the disk group that is imported on the node.
 Run the `VMGetDrive` utility at the command prompt to retrieve the GUID.
 - Repeat the previous two steps, and select a different node from the dropdown list each time. You must specify the GUID separately for each node displayed in the dropdown list.
 In case of a shared storage environment (VMDg resource), if there are multiple nodes in the primary zone, then the disk group GUID will be the same for all systems within the zone. However, the GUID will always be different across zones.
- 6** In the Properties View window, verify that all nodes in the RDC primary zone have DGGuid values specified.

Note: If you are creating a DR configuration manually for a non-shared storage environment, you have to modify the DGGuid attribute of the VMNSDg resource in the RVG service groups at both the sites (primary site and the disaster recovery site) separately.

- 7** Close the Properties View window.

Configuring the VMDg or VMNSDg resources for the disk group for the user-defined database

Repeat the VMDg or VMNSDg and MountV configuration for any additional disk group you may have created for a user-defined database.

This is an example configuration. You should modify these steps as necessary to match the disk groups and volumes you want to include in replication.

To create the VMDg or VMNSDg resource for the disk group for the user-defined database

- 1** In the **VCS Cluster Explorer** window, select the **SQL Server service group (INST1_SG)** in the left pane.
- 2** On the **Resources** tab, right-click the VMDg or VMNSDg resource for the disk group, with SQL Server user-defined files (`INST1_SG-VMDg-1` or `INST1_SG-VMNSDg-1`), and click **Copy > Self**.
- 3** In the left pane, select the **RVG service group (INST1_RVG_SG)**.
- 4** On the **Resources** tab, right-click in the blank resource display area and click **Paste**.

- 5 In the **Name Clashes** window, change the name of the **VMDg** or **VMNSDg** resource for the RVG service group.

For example, change it to `INST1_RVG_SG-VMDg-1` or `INST1_RVG_SG-VMNSDg-1`.

- 6 Click **OK**.

To modify the MountV resources in the SQL Server service group

- 1 In the **VCS Cluster Explorer** window, select the **SQL Server service group (INST1_SG)** in the left pane.
- 2 In the **Resources** tab display area, right-click the **MountV resource for the SQL Server user-defined log (INST1_SG-MountV-2)** and select **View > Properties View**.

- 3 In the **Properties View** window, verify that the Volume Name attribute is the SQL Server user-defined log (`INST1_DB1_LOG`).
- 4 In the same **Properties View** window, for the VMDg Resource Name attribute, click **Edit**.

- 5 In the **Edit Attribute** window, modify the **VMDGResName** scalar value to be the VMDg or VMNSDg resource that was just created.

For example `INST1_RVG_SG-VMDg-1` or `INST1_RVG_SG-VMNSDg-1`.

- 6 Close the **Properties View** window.
- 7 In the **Resources** tab display area, right-click the MountV resource for the SQL Server user-defined database (`INST1_SG-MountV-3`) and select **View > Properties View**.
- 8 In the **Properties View** window, verify that the Volume Name attribute is the SQL Server user-defined database (`INST1_DB1_VOL`).
- 9 In the same **Properties View** window, for the VMDg Resource Name attribute, click **Edit**.
- 10 In the **Edit Attribute** window, modify the VMDGResName scalar value to be the VMDg or VMNSDg resource that was just created.

For example `INST1_RVG_SG-VMDg` or `INST1_RVG_SG-VMNSDg`.

- 11 Close the **Properties View** window.

To enable the VMDg or VMNSDg resource

- 1 In the left pane, select the RVG service group (`INST1_RVG_SG`).
- 2 In the **Resources** tab display area, right-click the VMDg or VMNSDg resource (`INST1_RVG_SG-VMDg-1` or `INST1_RVG_SG-VMNSDg-1`) and select **Enabled**.

Adding the Volume Replicator RVG resources for the disk groups

Add a Volume Replicator RVG resource for each disk group that you want to replicate.

For the example configuration, you add a Volume Replicator RVG resource for the disk group for the SQL Server system files. You then add a Volume Replicator RVG resource for the disk group for the SQL Server user-defined database files.


The following table lists the attributes that you must configure in the RVG service group for the VvrRvg resource.

Table 6-5 VvrRvg resource and attributes to modify

Resource	Attributes to Modify
VvrRvg	VMDgResName IPResName

To create the Volume Replicator RVG resource for a disk group containing the system files

- 1 In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2 In the Add Resource window, specify the following:



- Enter a resource name for the Volume Replicator RVG resource. For example, enter INST1-VvrRvg.

- In the Resource Type list, select **VvrRvg**.
- 3** In the Add Resource window the attributes appear. For the **RVG** attribute, click **Edit**.
- 4** In the Edit Attribute window, enter the name of the RVG group that is being managed.

 For example, enter INST1_RVG.

 The RVG name is the name you specified when you created the Replicated Data Set (RDS) earlier using the RDS wizard. You can retrieve the RVG name by running the command `vxprint -VPL`.
- 5** Click **OK**.
- 6** In the Add Resource window, for the **VMDGResName** attribute, click **Edit**.
- 7** In the Edit Attribute window, enter the name of disk group containing the RVG.

 For example, for the system files disk group, enter INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg.
- 8** Click **OK**.
- 9** In the Add Resource window, for the **IPResName** attribute, click **Edit**.
- 10** In the Edit Attribute window, enter the name of the IP resource managing the IP address for replication.

 For example, enter INST1_RVG_SG-IP.
- 11** Click **OK**.

- 12** In the Add Resource window, verify that the attributes have been modified:

Add Resource

Group name: INST1_RVG_SG

Resource name: INST1_VvrRvg

Resource Type: VvrRvg

Attribute name	Type	Dimension	Value	Edit
RVG	String	Scalar	INST1_RVG	
VMDgResName	String	Scalar	INST1_RVG_SG-...	
IPResName	String	Scalar	INST1_RVG_SG-IP	
SRL	String	Scalar		
RLinks	String	Vector		

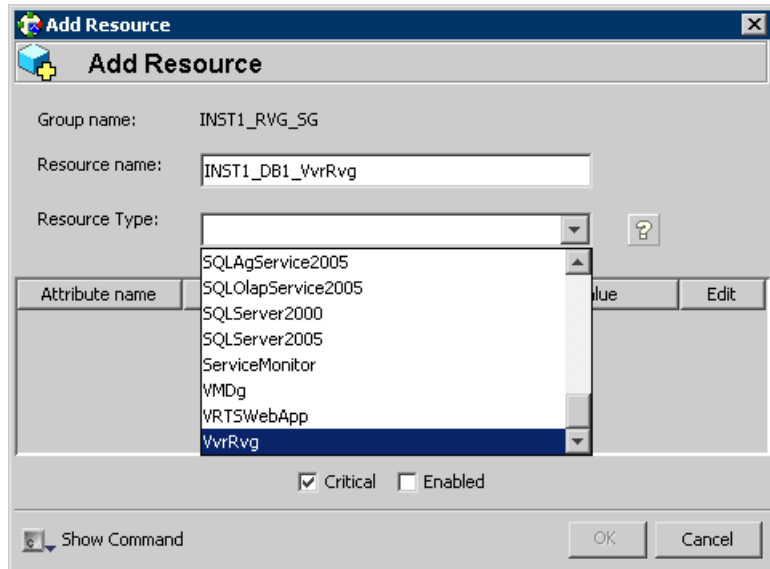
☒ Critical ☒ Enabled

Show Command OK Cancel

- 13** Click **OK**.

To create the Volume Replicator RVG resource for the disk group containing the user-defined database files

- 1** In the left pane, select the RVG service group (INST1_RVG_SG). Right-click it and select **Add Resource**.
- 2** In the Add Resource window, specify the following:



- Enter a resource name for the Volume Replicator RVG resource.
For example, enter INST1-DB1-VvrRvg.
 - In the Resource Type list, select **VvrRvg**.
- 3** In the Add Resource window the attributes appear. For the **RVG** attribute, click **Edit**.
 - 4** In the Edit Attribute window, enter the name of the RVG group that is being managed.
For example INST1_DB1_RVG.

The RVG name is the name you specified when you created the Replicated Data Set (RDS) earlier using the RDS wizard. You can retrieve the RVG name by running the command vxprint -VPI.
 - 5** Click **OK**.
 - 6** In the Add Resource window, for the **VMDGResName** attribute, click **Edit**.
 - 7** In the Edit Attribute window, enter the name of disk group containing the RVG.
For example INST1_RVG_SG-VMDg-1 or INST1_RVG_SG-VMNSDg-1.
 - 8** Click **OK**.
 - 9** In the Add Resource window, for the **IPResName** attribute, click **Edit**.

- 10
- In the Edit Attribute window, enter the name of the IP resource managing the IP address for replication.
- For example, enter INST1_RVG_SG-IP.
- In this example both disk groups are using the same IP resource for replication.
- 11
- Click **OK**.
- 12
- In the Add Resource window, verify that the attributes have been modified:

Add Resource

Add Resource

Group name:INST1_RVG_SG

Resource name:INST1_DB1_VvrRvg

Resource Type:VvrRvg

Attribute name	Type	Dimension	Value	Edit
RVG	String	Scalar	INST1_DB1_RVG	
VMDgResName	String	Scalar	INST1_RVG_SG-...	
IPResName	String	Scalar	INST1_RVG_SG-...	
SRL	String	Scalar		
RLinks	String	Vector		

☒ Critical ☒ Enabled

Show Command

OKCancel

- 13
- Click **OK**.

Linking the Volume Replicator RVG resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the Volume Replicator RVG service group to establish the dependencies between the resources.

You start from the top parent and link the parent and child resources as shown in the following table\.

Table 6-6 Dependencies for Volume Replicator RVG resources for RDC

Parent	Child
INST1_ VvrRvg	The IP for replication, for example INST1_RVG_SG-IP.


Table 6-6 Dependencies for Volume Replicator RVG resources for RDC
(continued)

Parent	Child
INST1_VvrRvg	The VMDg or VMNSDg for the SQL Server system files. For example, INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg.
INST1_DB1_VvrRvg	The IP for replication, for example INST1_RVG_SG-IP.
INST1_DB1_VvrRvg	The VMDg or VMNSDg for the SQL Server user-defined database files. For example, INST1_RVG_SG-VMDg-1 or INST1_RVG_SG-VMNSDg-1.

To link the Volume Replicator RVG resources

- 1 In the left pane, select the RVG service group (INST1_RVG_SG).
- 2 Click the **Link** button in the right pane.
- 3 To link the VvrRvg resource to the IP resource, click the parent resource, for example INST1_DB1_VvrRvg, and then click the child resource, for example INST1_RVG_SG-IP.
- 4 When prompted to confirm, click **OK**.
- 5 To link the VvrRvg resource to the VMDg or VMNSDg resource, click the parent resource, for example INST1_DB1_VvrRvg, and then click the child resource, for example INST1_RVG_SG-VMDg or INST1_RVG_SG-VMNSDg.

- 6 When prompted to confirm, click **OK**.
- 7 Repeat these steps to link all the RVG resources:



Notice that when you enable a resource and the state of the entity which it is monitoring is online, the corresponding VCS agent reports status for that resource as online. You do not have to bring the resource online manually.

Deleting the VMDg or VMNSDg resource from the SQL Server service group

The VMDg (in case of shared storage) or VMNSDg (in case of non-shared storage) resources must now be manually deleted from the SQL Server service group, because they depend on replication and were configured in the RVG service group.

To delete the VMDg or VMNSDg Resources from the SQL Server service group

- 1 In the VCS Cluster Explorer window, select the SQL Server service group (INST1_SG) from the left pane.
- 2 In the Resources tab display area, right-click the VMDg or VMNSDg resource for the system files disk group (INST1_SG-VMDg or INST1_SG-VMNSDg) and select **Delete**.

- 3 Click **Yes** to confirm that you want to delete it (even if it is online).
- 4 In the Resources tab display area, right-click the VMDg or VMNSDg resource for the user-defined database disk group (INST1_SG-VMDg-1 or INST1_SG-VMNSDg-1) and select **Delete**.
- 5 Click **Yes** to confirm that you want to delete it (even if it is online).

Configuring the RVG Primary resources

Add resources of type RVGPrimary to the SQL Server service group for each of the SQL Server disk groups (system and user-defined) and configure the attributes.

Set the value of the RvgResourceName attribute to the name of the RVG resource for the RVGPrimary agent. This is the name of the VvrRvg resource in the RVG replication service group.

The following table lists the RVG Primary resources and attributes that you must configure in the SQL Server service group for the sample configuration.

Table 6-7 RVG Primary resources and attributes to modify

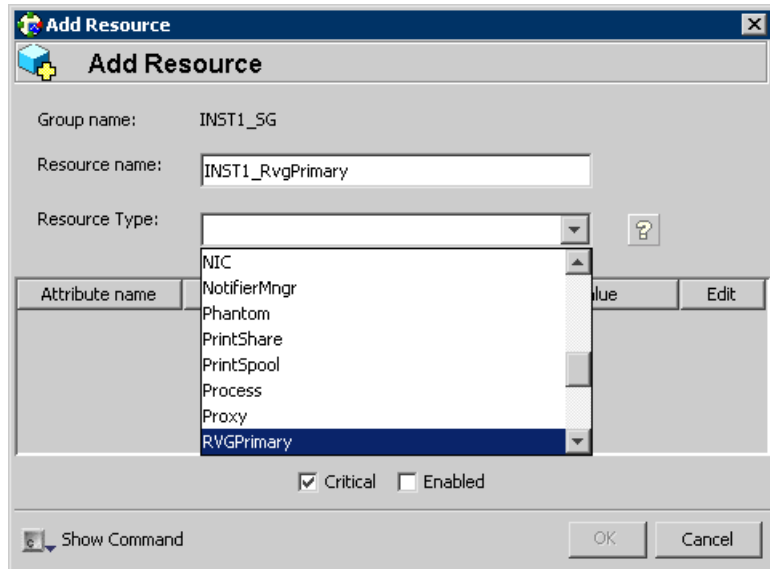
Resource	Attributes to Modify
RVGPrimary (for the disk group for the SQL Server system files)	RvgResourceName
RVGPrimary (for the disk group for the SQL Server user-defined database files)	RvgResourceName

Creating the RVG Primary resources

For each disk group created for the SQL Server system and user-defined databases, create a separate RVG Primary resource for replication.

To create the RVG Primary resource for the SQL Server system disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the **Add Resource** window, specify the following:



- Enter a resource name for the RVG Primary resource for the SQL Server system files disk group. For example, enter INST1_RvgPrimary.
 - Select the **Resource Type** of RVGPrimary.
- 3 In the Add Resource window the attributes appear. For the **RvgResourceName** attribute, click **Edit**.
 - 4 In the Edit Attribute window, enter the name of the Volume Replicator RVG resource, for example INST1_VvrRvg and click **OK**. This is the name of the VvrRvg resource in the RVG replication service group.
 - 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults. See the *Volume Replicator Administrator's Guide* for more information about the RVG Primary agent.
 - 6 Verify that **Critical** and **Enabled** are both checked.
 - 7 Click **OK**.

To create the RVG Primary resource for the SQL Server user-defined database disk group

- 1 In the VCS Cluster Explorer window, right-click the SQL Server service group (INST1_SG) in the left pane, and select **Add Resource**.
- 2 In the Add Resource window, specify the following:

- Enter a resource name for the RVG Primary resource for the SQL Server user-defined database disk group. For example, enter INST1_DB1_RvgPrimary.
 - Select **RVGPrimary** as the Resource Type.
- 3 In the Add Resource window the attributes appear. For the RvgResourceName attribute, click Edit.
 - 4 In the Edit Attribute window, enter the name of the Volume Replicator RVG resource, for example INST1_DB1_VvrRvg and click **OK**.
 - 5 If desired, set the AutoTakeover and AutoResync attributes from their defaults.
 - 6 Verify that **Critical** and **Enabled** are both checked.
 - 7 Click **OK**.

Linking the RVG Primary resources to establish dependencies

In the VCS Cluster Explorer window, link the resources in the SQL Server service group (INST1_SG) to establish the dependencies between the resources for replication.

You start from the top parent and link the parent and child resources as shown in the following table.

Table 6-8 Dependencies for the RVG Primary resources for RDC

Parent	Child
INST1_SG-MountV	INST1_RvgPrimary
INST1_SG-MountV-1	INST1_RvgPrimary
INST1_SG-MountV-2	INST1_DB1_RvgPrimary
INST1_SG-MountV-3	INST1_DB1_RvgPrimary

To link the RVG Primary resources

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 Click the **Link** button in the right pane.
- 3 Click the parent resource.
For example INST1_SG-MountV.
- 4 Click the child resource.
For example INST1_RvgPrimary.

- 5 When prompted to confirm, click **OK**.
- 6 Repeat these steps to link all the RVG Primary resources.

Bringing the RVG Primary resources online

In the VCS Cluster Explorer window, bring the RVG Primary resources in the SQL Server service group (INST1_SG) online on the first node in the primary zone.

To bring the RVG Primary resources online

- 1 In the left pane, select the SQL Server service group (INST1_SG).
- 2 In the right pane on the Resources tab, right-click the first RVG Primary resource and select **Online > SYSTEM1**.
- 3 In the right pane on the Resources tab, right click the second RVG Primary resource and select **Online > SYSTEM1**.

Configuring the primary system zone for the RVG service group

In the RVG service group, set up systems in the primary zone (Zone 0) to specify that initial failover occurs to systems within the primary zone for the RVG service group.

To configure the primary system zone for the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the Properties tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the SystemZones attribute.
- 6 In the **Edit Attribute** dialog box, click the plus sign and enter the systems and the zone number (type 0 for Zone 0) for the primary zone.

In case of a non-shared storage configuration, add only the single node to the primary zone.

- 7 Click **OK**.

Setting a dependency between the service groups


The RVG service group must be online on both the primary and secondary zones. However, if a failover occurs from one node to another within the same zone, the RVG service group must fail over along with the application service group.

To ensure that the SQL Server service group and the RVG service group fail over and switch together, set up an online local hard dependency from the RVG service group to the SQL Server service group.

The SQL Server service group (for example, INST1_SG) is dependent on the replication service group (for example, INST1_RVG_GRP).

To set up an online local hard dependency

- 1 From VCS Cluster Explorer, in the left pane, select the cluster (MYCLUSTER).
- 2 In the right pane, select the **Service Groups** tab.
- 3 Click the **Link** button to create a dependency between service groups.
- 4 Click the application service group (parent service group).
For example, click the SQL Server service group INST1_SG.
- 5 Click the RVG service group (the child resource).
For example, click the RVG service group INST1_RVG_SG.
- 6 In the **Link Service Groups** window, specify the following:



- Select the Relationship of **online local**.
- Select the Dependency Type of **hard**.
- Click **OK**.

Adding the nodes from the secondary zone to the RDC

Configuration of the systems in the Primary Zone (zone 0) is complete. The nodes in the Secondary Zone (zone 1) can now be added to the RDC configuration.

See the following topics:

- See [“Configuring secondary zone nodes in the RVG service group”](#) on page 213.
- See [“Configuring the RVG service group IP resource for failover”](#) on page 215.
- See [“Adding nodes from the secondary zone to the SQL Server service group”](#) on page 218.

Adding the nodes from the secondary zone to the RVG service group

Use the Volume Replicator Agent Configuration Wizard to add the nodes from the secondary zone to the RVG.

Note: In case of a non-shared storage environment, perform this task manually using the Java Console. You cannot use the wizard to modify the RVG service group.

Use the following procedure if the RVG service group contains a VMDg resource (shared storage environment).

To add the nodes from the secondary zone to the RVG

- 1 From the active node of the cluster in the primary zone, launch the configuration wizard from **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > Volume Replicator Agent Configuration Wizard**, or, on Windows 2012 operating systems, from the **Apps** menu.
- 2 Read and verify the requirements on the Welcome page, and click Next.
- 3 In the Wizard Options panel, specify the following:
 - Click **Modify an existing replication service group**. The existing replication service group is selected, by default.
 - Click **Next**.
- 4 If a VCS notice message appears, asking if you want to continue, click **Yes**.
- 5 Specify the system priority list as follows:

- In the Available Cluster Systems box, click the nodes in the secondary zone to add to the service group, and click the right-arrow icon to move the nodes to the service group's system list.
 - To remove a node from the service group's system list, click the node in the Systems in Priority Order box, and click the left arrow icon.
 - To change the priority of a node in the system list, click the node in the Systems in Priority Order box, then click the up and down arrow icons. The node at the top of the list has the highest 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**.
- 6 If a message appears, indicating that the configuration will be changed from Read Only to Read/Write, click **Yes** to continue.
 - 7 Review the Disk Group and Replicated Volume Group Configuration and click **Next**.
 - 8 In the IP Resource Options panel, select **Modify IP resource** and click **Next**.
 - 9 If a VCS error appears, click **OK**.
 - 10 In the Network Configuration panel, verify that the selected adapters are correct and click **Next**.
 - 11 Review the summary of the service group configuration as follows:
 - The Resources box lists the configured resources. Click a resource to view its attributes and their configured values in the Attributes box.
 - Click **Next** to modify the replication service group.
 - 12 When prompted, click **Yes** to modify the service group.
 - 13 Click **Finish**.

Use the following procedure if the RVG service group contains a VMNSDg resource (non-shared storage environment).

To add nodes from the secondary zone to the RVG service group using Java Console

- 1 From VCS Cluster Explorer, in the left pane, right-click the RVG service group and select **View > Properties View**.
- 2 In the Attributes window, click **Show all attributes**.

- 3 From the attributes list, select the attribute **SystemList** and click the edit icon.
- 4 In the Edit Attribute window, edit the **SystemList** attribute as follows:
 - Click the **+** button to add an empty row.
 - In the **System** field type the cluster node name from the secondary zone.
 - In the **Priority** field type **1**.
 - Click **OK**.
- 5 Close the Attributes window.

Configuring secondary zone nodes in the RVG service group

Specify zone 1 for the nodes in the secondary zone.

To specify the secondary zone for the nodes in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the **Properties** tab.
- 3 In the Properties pane, click the button **Show All Attributes**.
- 4 In the Attributes View, scroll down and select the **SystemZones** attribute.
- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.

 In case of a non-shared storage configuration, add only the single node to the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

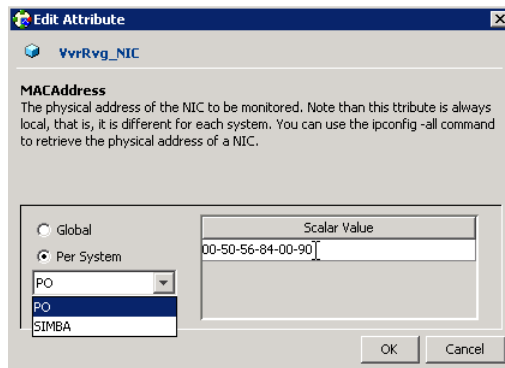
Configuring the RVG service group NIC resource for fail over (VMNSDg only)

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment where the RVG service group contains a VMNSDg resource.

Modify the **MACAddress** attribute of the NIC resource in the RVG service group to ensure desired fail over behavior in the RDC.

To modify the NIC resource in the RVG service group

- 1 From the VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the Resources tab.
- 3 In the Resources tab display area, right-click the NIC resource and click **View > Properties View**.
- 4 In the Properties View window, locate the MACAddress attribute and click the edit icon.
- 5 In the Edit Attribute window, edit the attribute by performing the following actions sequentially:



- **Select Per System.**
 - From the dropdown list, select the node in the RDC primary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - Click **OK**.
- 6 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
 - 7 Close the Properties View window.

Configuring the RVG service group IP resource for failover

Modify the IP resources in the RVG service group to ensure the desired failover behavior in the RDC.


In the event of a system or SQL Server campus cluster failure, VCS attempts to fail over the SQL Server service group to another system within the same RDC system zone. However, in the event that VCS fails to find a failover target node within the primary zone, VCS switches the service group to a node in the current secondary system zone.

Use the following procedure to modify the IP resources.

Note: For IPv6 networks, modify the IPv6 resources.

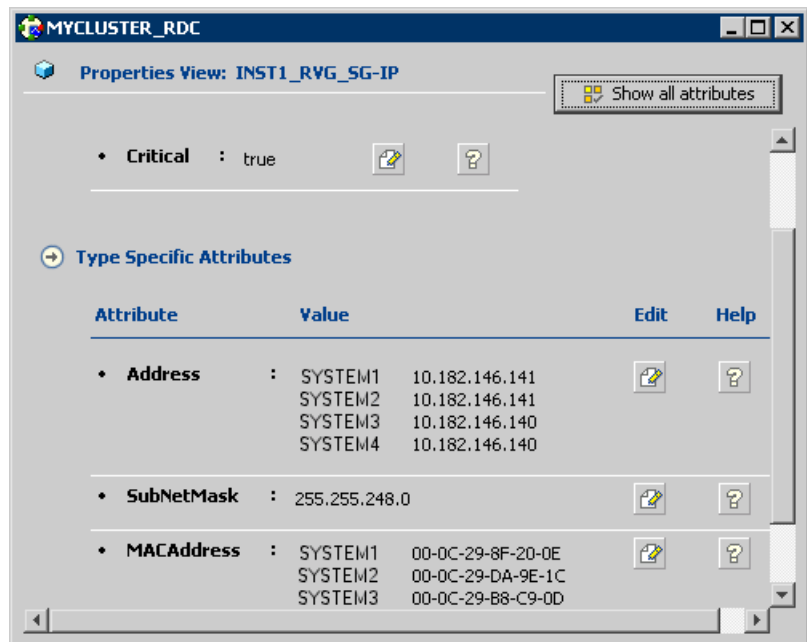
To modify the IP resources in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group (INST1_RVG_SG).
- 2 In the right pane, select the **Resources** tab.
- 3 Right-click the RVG IP resource (INST1_RVG_SG-IP) and select **View > Properties View**.
- 4 In the Edit Attributes window, edit the Address attribute.



- Select **Per System**.

- Select the first node in the primary zone and enter the virtual IP address for the primary zone.
 - Select the second node in the primary zone and enter the virtual IP address for the primary zone (the same IP address as the first node).
 - Repeat for all nodes in the primary zone.
 - Select the first node in the secondary zone (SYSTEM3) and enter the virtual IP address for the secondary zone.
 - Select the second node in the secondary zone and enter the virtual IP address for the secondary zone (the same IP address as the first node in the secondary zone).
 - Repeat for all nodes in the secondary zone.
 - Click **OK**.
- 5** In the Properties View window, verify that all nodes in the primary zone have the same IP address. Also verify that all nodes in the secondary zone have the same IP address. The IP address at the primary zone and the secondary zone should be different



- 6** This step is applicable only if you are using a non-shared storage environment (VMNSDg agent).

In the Edit Attributes window, edit the MACAddress attribute as follows:

- **Select Per System.**
 - From the dropdown list, select the node in the RDC primary zone.
 - In the Scalar Value field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the ipconfig -all command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.
 - In the Scalar Value field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the ipconfig -all command to retrieve the physical address.
 - Click **OK**.
- 7 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
 - 8 Close the Properties View window.

Since this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.

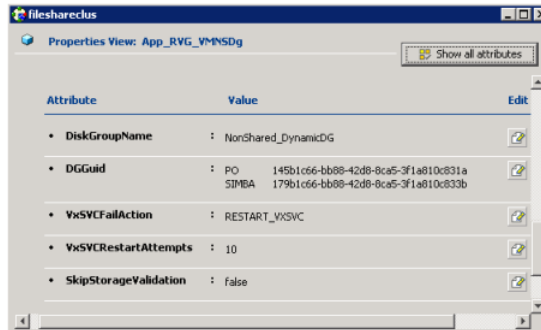
Configuring the RVG service group VMNSDg resources for fail over

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment where the RVG service group contains a VMNSDg resource. Modify the DGGuid attribute of the VMNSDg resources in the RVG service group to ensure the desired failover behavior in the RDC.

To modify the VMNSDg resources in the RVG service group

- 1 From VCS Cluster Explorer, in the left pane, select the RVG service group.
- 2 In the right pane, select the Resources tab.
- 3 Right-click the RVG VMNSDg resource and select **View > Properties View**.
- 4 In the Edit Attributes window, edit the DGGuid attribute by performing the following actions sequentially:
 - **Select Per System.**
 - Select the node in the RDC secondary zone.
 - In the Scalar Value field, enter the GUID of the dynamic disk group that is imported on the single node in the RDC secondary zone.

- You can retrieve the disk group details running the VMGetDrive utility from the command prompt.
 - Click **OK**.
- 5 In the Properties View window, verify that the DGGuid for the nodes in the primary and secondary zone are different.



- 6 Close the Properties View window.

As this is the final task in configuring the RVG service group for the primary and secondary zones, you can now bring the RVG service group online in both the primary and secondary zones.

Adding nodes from the secondary zone to the SQL Server service group

Use the SQL Server Configuration Wizard to add the nodes from the secondary zone to the SQL Server service group.

Note: In case of a non-shared storage environment, perform this task manually using the Java Console. You cannot use the wizard to modify the application service group.

Use the following procedure if the service group contains a VMDg resource (shared storage environment).

To add the nodes from the secondary zone to the SQL Server service group

- 1 Select **Start > All Programs > Veritas > Veritas Cluster Server > Configuration Tools > SQL Server Configuration Wizard**.
- 2 Verify that you have met the prerequisites listed and click **Next**.

- 3 In the Wizard Options panel, click **Modify service group**, select the SQL Server service group to be modified (INST1_SG) and click **Next**. If a VCS notice message appears indicating that resources are online, click **Yes** to continue.
- 4 On the Service Group Configuration panel, select the nodes in the secondary zone, use the arrow button to move them from Available Cluster Systems to **Systems in Priority Order** and then click **Next**.

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 in as failover targets for the group. The server that needs to come online first must be at the top of the list followed by the next one that will be brought online.

This set of nodes selected for the SQL Server service group must be the same as the nodes selected for the RVG service group. Ensure that the nodes are also in the same priority order.

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*.

- 5 On the SQL Server Instance Selection page, click **Next**.
- 6 On the User Databases List panel, click **Next**. This panel summarizes the databases for this instance of SQL.
- 7 On the Detail Monitoring Configuration panel, clear the box in the SQL Server instance list to disable monitoring, if required, and then click **Next**. Detailed monitoring is not necessary.
- 8 On the Registry Replication Path panel, click **Next**.
- 9 On the Virtual Server Configuration panel, verify that the public adapter is used on each system and click **Next**.

- 10** In the Service Group Summary, review the service group configuration.

To enable all the VMDg resources in the service group for fast failover, select the **Enable FastFailOver attribute for all the VMDg resources in the service group** checkbox. For information about the FastFailOver attribute, see the *Storage Foundation Administrator's Guide*.

Click **Next**.

A message appears if the configuration is currently in the Read Only mode. Click **Yes** to make the configuration read and write enabled. The wizard validates the configuration and modifies it.

- 11** Click **Finish**.

Use the following procedure if the service group contains a VMNSDg resource (non-shared storage environment).

To add nodes from the secondary zone to the application service group using Java Console

- 1** From VCS Cluster Explorer, in the left pane, right-click the SQL Server service group (INST1_SG) and select **View > Properties View**.
- 2** In the Attributes window, click **Show all attributes**.
- 3** From the attributes list, select the attribute **SystemList** and click the edit icon.
- 4** In the Edit Attribute window, edit the SystemList attribute as follows:
 - Click the + button to add an empty row.
 - In the System field type the cluster node name from the secondary zone.
 - In the Priority field type 1.
 - Click OK.
- 5** Close the Attributes windows.

Configuring the zones in the SQL Server service group

Specify zone 1 for the SQL Server service group nodes in the secondary zone.

To specify the secondary zone for the nodes in the SQL Server service group

- 1** From VCS Cluster Explorer, in the left pane, select the SQL Server service group (INST1_SG).
- 2** In the right pane, select the **Properties** tab.
- 3** In the Properties pane, click the button **Show All Attributes**.
- 4** In the Attributes View, scroll down and select the SystemZones attribute.

- 5 Click the **Edit** icon for the **SystemZones** attribute.
- 6 If a message appears indicating that the configuration be changed to read/write, click **Yes**.
- 7 In the Edit Attribute dialog box, click the plus sign and enter the systems and the zone number (zone 1) for the secondary zone.

In case of a non-shared storage configuration, add only the single node to the secondary zone.
- 8 Click **OK**.
- 9 Close the Attributes View window.

Configuring the application service group IP resource for fail over (VMNSDg only)

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment (VMNSDg agent).

Modify the IP resource in the application service group to ensure the desired failover behavior in the RDC.

Note: For IPv6 networks, modify the IPv6 resources.

To modify the IP resource in the application service group

- 1 From VCS Cluster Explorer, in the left pane, select the application service group.
- 2 In the right pane, select the Resources tab.
- 3 Right-click the IP resource and select **View > Properties View**.
- 4 In the Edit Attributes window, edit the MACAddress attribute by performing these actions sequentially:
 - Select **Per System**.
 - From the dropdown list, select the node in the RDC primary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system.

Use the `ipconfig -all` command to retrieve the physical address.

- Click **OK**.
- 5 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
- 6 Close the Properties View window.

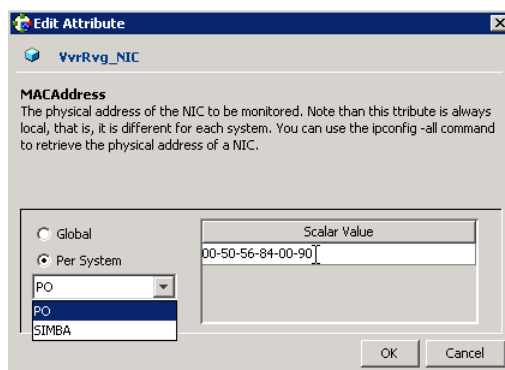
Configuring the application service group NIC resource for fail over (VMNSDg only)

This procedure is applicable only if you setting up RDC configuration in a non-shared storage environment (VMNSDg agent).

Modify the MACAddress attribute of the NIC resource in the application service group to ensure desired fail over behavior in the RDC.

To modify the NIC resource in the application service group

- 1 From the VCS Cluster Explorer, in the left pane, select the application service group.
- 2 In the right pane, select the Resources tab.
- 3 In the Resources tab display area, right-click the NIC resource and click **View > Properties View**.
- 4 In the Properties View window, locate the MACAddress attribute and click the edit icon.
- 5 In the Edit Attribute window, edit the attribute by performing the following actions sequentially:



- Select **Per System**.
- From the dropdown list, select the node in the RDC primary zone.

- In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the primary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - From the dropdown list, select the node in the RDC secondary zone.
 - In the **Scalar Value** field, enter the physical address (MAC address) of the network interface card (NIC) to be monitored on the secondary zone system. Use the `ipconfig -all` command to retrieve the physical address.
 - Click **OK**.
- 6 In the Properties View window, verify that the MACAddress attribute for the nodes in the primary and secondary zone are different.
 - 7 Close the Properties View window.

Verifying the RDC configuration

After completing all the configuration tasks for the primary and secondary zones, you can bring the service group online, then verify the configuration.

Perform the following tasks:

- See [“Bringing the service group online”](#) on page 223.
- See [“Switching online nodes”](#) on page 223.

Bringing the service group online

After completing all configuration, ensure that the RVG service group is online in both the primary and secondary zone. Then you can bring the application service group online in the primary zone.

To bring the application service group online

- 1 From VCS Cluster Explorer, in the left pane, right-click the application service group.
- 2 Click **Online**.

Switching online nodes

Failover simulation is an important part of configuration testing. Test the failover by switching online nodes.

The RVG service group is online in both the primary and secondary zone. However, within a zone, if more than node is configured, the RVG service group should fail over with the application service group.

Note: This should never be tested on systems with live data. A reliable and tested backup should be available. A tested backup means that it has been tested successfully by a restore.

Switch the application service group between nodes using Veritas Cluster Manager (Java Console). When you complete the procedure, you will see the online system role shift from one system to another.

If you enter the system name manually from the Java Console, specify the name in upper case.

To switch online nodes

- 1 In the Veritas Cluster Manager (Java Console), click the cluster in the configuration tree, and click the Service Groups tab.
- 2 Switch the service group as follows:
 - Right-click the service group icon in the view panel.
 - Click **Switch To**, and click the appropriate node from the menu.
 - In the dialog box, click **Yes**. The service group you selected is taken offline on the original node and brought online on the node you selected.

If there is more than one service group, you must repeat this step until all the service groups are switched.

- 3 Verify that the service group is online on the node you selected.
- 4 To move all the resources back to the original node, repeat step 2 for each of the service groups.

Additional instructions for GCO disaster recovery

After completing the tasks for setting up a replicated data cluster, you can optionally create a secondary site for wide area disaster recovery using the SFW HA Global Cluster option (GCO).

With this option, if a disaster affects a local or metropolitan area, data and critical services are failed over to a site hundreds or thousands of miles away.

To configure disaster recovery using a secondary site, you must install the SFW HA Global Cluster Option on all nodes on the primary (replicated data cluster) site cluster, as well as the secondary (DR) site cluster. GCO configuration also requires a static IP address available for each site.

You can use the Disaster Recovery (DR) wizard when setting up the secondary site. The secondary site is not configured as a replicated data cluster. There can be only one replicated data cluster in the DR environment.

The DR wizard does the following tasks:

- Clones the storage
- Clones the application service group
- Sets up Volume Replicator replication for the secondary site
- Configures the primary and secondary site clusters as global clusters
See [“Disaster recovery configuration”](#) on page 65.

When cloning the service group, the wizard does not clone the settings that specify primary and secondary zones, because the secondary site cluster is not divided into zones.

Configuring disaster recovery for SQL Server

This chapter includes the following topics:

- [Tasks for configuring disaster recovery for SQL Server](#)
- [Tasks for setting up DR in a non-shared storage environment](#)
- [Guidelines for installing Veritas InfoScale Enterprise and configuring the cluster on the secondary site](#)
- [Verifying your primary site configuration](#)
- [Setting up your replication environment](#)
- [Assigning user privileges \(secure clusters only\)](#)
- [About configuring disaster recovery with the DR wizard](#)
- [Cloning the storage on the secondary site using the DR wizard \(Volume Replicator replication option\)](#)
- [Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)](#)
- [Installing and configuring SQL Server on the secondary site](#)
- [Cloning the service group configuration from the primary site to the secondary site](#)
- [Configuring the SQL Server service group in a non-shared storage environment](#)
- [Configuring replication and global clustering](#)
- [Creating the replicated data sets \(RDS\) for Volume Replicator replication](#)

- [Creating the Volume Replicator RVG service group for replication](#)
- [Configuring the global cluster option for wide-area failover](#)
- [Verifying the disaster recovery configuration](#)
- [Establishing secure communication within the global cluster \(optional\)](#)
- [Adding multiple DR sites \(optional\)](#)
- [Recovery procedures for service group dependencies](#)
- [Configuring DR manually without the DR wizard](#)

Tasks for configuring disaster recovery for SQL Server

After setting up an SFW HA high availability environment for SQL Server on a primary site, you can create a secondary or “failover” site for disaster recovery.

The Disaster Recovery (DR) wizard helps you to clone the storage and service group configuration from the primary site to the secondary site. You can install the application on the secondary site during the DR wizard workflow.

The DR wizard also helps you set up replication and the global clustering option (GCO). You can choose to configure replication using Volume Replicator (Volume Replicator) or an agent-supported array-based hardware replication. The DR wizard can configure required options for the VCS agents for EMC SRDF and for Hitachi TrueCopy. To use the wizard with any other agent-supported array-based replication, you must complete configuring global clustering with the wizard before configuring replication on the array.

The DR wizard is available from the Solutions Configuration Center. Veritas recommends using the Solutions Configuration Center as a guide for installing and configuring disaster recovery.

Note: If you are using non-shared storage (dynamic disk groups monitored using VMNSDg agent), you cannot use the DR wizard to configure disaster recovery. You have to set up DR manually. Refer to the separate workflow table (See Table 10-2) available for configuring DR manually.

The following table shows the steps specific to the DR configuration that are done after configuring high availability on the primary zone.

Table 7-1 Configuring the secondary site for disaster recovery

Action	Description
Install the product and configure the cluster on the secondary site	<p>Warning: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.</p> <p>See “Guidelines for installing Veritas InfoScale Enterprise and configuring the cluster on the secondary site” on page 233.</p>
Verify that SQL Server has been configured for high availability at the primary site	<p>Verify that SQL Server has been configured for high availability at the primary site and that the service groups are online</p> <p>See “Verifying your primary site configuration” on page 234.</p>
Set up the replication prerequisites	<p>Ensure that replication prerequisites for your selected method of replication are met before running the DR wizard</p> <p>See “Setting up security for Volume Replicator” on page 174.</p> <p>See “Requirements for EMC SRDF array-based hardware replication” on page 235.</p> <p>See “Requirements for Hitachi TrueCopy array-based hardware replication” on page 237.</p>
(Secure cluster only) Assign user privileges	<p>For a secure cluster only, assign user privileges</p> <p>See “Assigning user privileges (secure clusters only)” on page 239.</p>
Start running the DR wizard	<ul style="list-style-type: none"> Review prerequisites for the DR wizard Start the DR wizard and make the initial selections required for each task: selecting a primary site system, the service group, the secondary site system, and the replication method <p>See “Configuring disaster recovery with the DR wizard” on page 242.</p>
Clone the storage configuration (Volume Replicator replication only)	<p>(Volume Replicator replication option)</p> <p>Clone the storage configuration on the secondary site using the DR wizard</p> <p>See “Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)” on page 245.</p>

Table 7-1 Configuring the secondary site for disaster recovery (*continued*)

Action	Description
Create temporary storage for application installation (other replication methods)	<p>(EMC SRDF, Hitachi TrueCopy, or GCO only replication option)</p> <p>Use the DR wizard to create temporary storage for application installation on the secondary site</p> <p>See “Creating temporary storage on the secondary site using the DR wizard (array-based replication)” on page 249.</p>
Install and configure SQL Server on the first cluster node	<ul style="list-style-type: none"> ■ Ensure that the disk group and volumes for the system database are mounted on the first node ■ Follow the guidelines for installing SQL Server in the SFW HA environment <p>See “About installing SQL Server on the first system” on page 125.</p>
Install and configure SQL Server on the failover node(s)	<ul style="list-style-type: none"> ■ Stop SQL Server services on the first node ■ Follow the guidelines for installing SQL Server on failover nodes in the SFW HA environment <p>See “About installing SQL Server on additional systems” on page 126.</p>
Clone the service group configuration	<p>Clone the service group configuration from the primary to the secondary site using the DR wizard</p> <p>See “Cloning the service group configuration from the primary site to the secondary site” on page 254.</p>
Configure replication and global clustering, or configure global clustering only	<ul style="list-style-type: none"> ■ (Volume Replicator replication) Use the wizard to configure replication and global clustering ■ (EMC SRDF replication) Set up replication and then use the wizard to configure the SRDF resource and global clustering ■ (Hitachi TrueCopy) Set up replication and then use the wizard to configure the HTC resource and global clustering ■ (Other array-based replication) Use the wizard to configure global clustering, and then set up replication <p>See “Configuring replication and global clustering” on page 259.</p>

Table 7-1 Configuring the secondary site for disaster recovery *(continued)*

Action	Description
Verify the disaster recover configuration	<p>Verify that the secondary site has been fully configured for disaster recovery</p> <p>See “Verifying the disaster recovery configuration” on page 283.</p>
(Optional) Add secure communication	<p>Add secure communication between local clusters within the global cluster (optional task)</p> <p>See “Establishing secure communication within the global cluster (optional)” on page 285.</p>
(Optional) Add additional DR sites	<p>Optionally, add additional DR sites to a Volume Replicator environment</p> <p>See “Adding multiple DR sites (optional)” on page 287.</p>
Handling service group dependencies after failover	<p>If your environment includes dependent service groups, review the considerations for bringing the service groups online after failover to the secondary site</p> <p>See “Recovery procedures for service group dependencies” on page 287.</p>

Tasks for setting up DR in a non-shared storage environment

The following table outlines the high-level objectives and tasks for a creating a single-node DR configuration at the secondary site. Refer to this table if you are setting up DR in a non-shared storage environment (dynamic disk groups configured using VMNSDg agent).

You cannot use the DR wizard to configure disaster recovery in a non-shared storage environment. You have to configure DR manually.

Note: Some procedures (for example, configuring Volume Replicator replication) are common if you are setting up a DR or an RDC configuration. To avoid duplication, the topics referenced in this table point to the procedures described in the RDC chapter covered earlier.

Table 7-2 Non-shared storage: Configuring Disaster Recovery

Action	Description
Install the product and configure the cluster on the secondary site	<ul style="list-style-type: none"> ■ Verify the software and hardware prerequisites ■ Set up the network and storage ■ Install the product ■ Configure the cluster at the secondary site ■ Ensure that you select the GCO option to configure the Global Cluster Option resource for the cluster <p>Warning: Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.</p> <ul style="list-style-type: none"> ■ Configure disk groups and volumes. <p>See “Guidelines for installing Veritas InfoScale Enterprise and configuring the cluster on the secondary site” on page 233.</p>
Verify that SQL Server has been configured for high availability at the primary site	<p>Verify that SQL Server has been configured for high availability at the primary site and that the service groups are online.</p> <p>See “Verifying your primary site configuration” on page 234.</p>
Set up the replication prerequisites	<ul style="list-style-type: none"> ■ Ensure that Volume Replicator replication prerequisites are met ■ Configure the VxSAS service for Volume Replicator, specifying the cluster nodes at both primary and secondary sites <p>See “Setting up security for Volume Replicator” on page 174.</p>
(Secure cluster only) Assign user privileges	<p>For a secure cluster only, assign user privileges.</p> <p>See “Assigning user privileges (secure clusters only)” on page 239.</p>
Install and configure SQL Server on the first cluster node	<ul style="list-style-type: none"> ■ Ensure that the disk group and volumes for the system database are mounted on the first node ■ Follow the pre-installation, installation, and post-installation procedures for the first node on the secondary site <p>See “About installing SQL Server on the first system” on page 125.</p>

Table 7-2 Non-shared storage: Configuring Disaster Recovery (*continued*)

Action	Description
Configure the SQL Server service group for VCS (secondary site)	<ul style="list-style-type: none"> ■ Configure the application service group manually using the Cluster Manager (Java Console) ■ Ensure that the name of the service group is the same as that on the primary site <p>See “Configuring the SQL Server service group in a non-shared storage environment” on page 258.</p>
Set up the replicated data sets (RDS) for Volume Replicator replication	<ul style="list-style-type: none"> ■ Use the Setup Replicated Data Set Wizard to create RDS and start replication for the primary and secondary sites ■ Use the Setup Replicated Data Set Wizard to create Replicator Log volumes for the primary and secondary sites <p>See “Creating the replicated data sets (RDS) for Volume Replicator replication” on page 277.</p>
Create the Volume Replicator RVG service group (repeat steps on primary and secondary site separately)	<ul style="list-style-type: none"> ■ Create the Volume Replicator RVG service group for the replicated volume group ■ Use the Cluster Manager (Java Console) to manually create the service group ■ Create the RVG service group at the primary site and the secondary site separately ■ Bring the RVG service group online on the primary site <p>See “Creating the Volume Replicator RVG service group for replication” on page 277.</p>
Configure the global cluster option for wide-area failover	<ul style="list-style-type: none"> ■ Link clusters (adding a remote cluster to a local cluster) ■ Converting the application service group that is common to all the clusters to a global service group ■ Converting the local service group to a global group ■ Bringing the global service group online <p>See “Configuring the global cluster option for wide-area failover” on page 278.</p>
Verify the disaster recover configuration	<p>Verify that the secondary site has been fully configured for disaster recovery.</p> <p>See “Verifying the disaster recovery configuration” on page 283.</p>

Table 7-2 Non-shared storage: Configuring Disaster Recovery *(continued)*

Action	Description
(Optional) Add secure communication	Add secure communication between local clusters within the global cluster (optional task) See “Establishing secure communication within the global cluster (optional)” on page 285.
(Optional) Add additional DR sites	Optionally, add additional DR sites to a Volume Replicator environment. See “Adding multiple DR sites (optional)” on page 287.
Handling service group dependencies after failover	If your environment includes dependent service groups, review the considerations for bringing the service groups online after failover to the secondary site. See “Recovery procedures for service group dependencies” on page 287.

Guidelines for installing Veritas InfoScale Enterprise and configuring the cluster on the secondary site

Use the following guidelines for installing Veritas InfoScale Enterprise and configuring the cluster on the secondary site.

- Ensure that you have set up the components required to run a cluster.
See [“Configuring the storage hardware and network”](#) on page 78.
- Ensure that when installing Veritas InfoScale Enterprise you install the appropriate disaster recovery options at both the primary and secondary sites. Be sure to select the following installation options as appropriate for your environment:

Global Cluster Option	Required for a disaster recovery configuration only.
Volume Replicator	If you plan to use Volume Replicator for replication, select the option to install Volume Replicator.

For more information see the *Veritas InfoScale Installation and Upgrade Guide*.

- Configure the cluster with the VCS Cluster Configuration Wizard (VCW). Ensure that the name you assign to the secondary site cluster is different from the name assigned to the primary site cluster.

See [“Configuring the cluster using the Cluster Configuration Wizard”](#) on page 102.

Note: You do not need to configure the GCO option while configuring the cluster. This is done later using the Disaster Recovery wizard. However, in case you are configuring DR manually in a non-shared storage environment, you must configure GCO using VCW.

- The storage configuration will be handled by the DR wizard (in case of shared storage only). If you are configuring a DR setup that uses non-shared storage (VMNSDg agent), you have to configure the storage using VEA console. While creating disk groups and volumes for the secondary site, make sure to use the same names of volumes as those on the primary site. The size of the volumes on the secondary site must be equal to or larger than the size of the volumes on the primary site.
Ensure that you allow sufficient disk space to create a volume for the Volume Replicator Storage Replicator Log for each storage group. You can create the volume now, or later, when running the wizard to create replicated data sets (RDS).

See [“Configuring disk groups and volumes for SQL Server”](#) on page 79.

Verifying your primary site configuration

Before you begin configuring disaster recovery, make sure that SQL Server has been configured for high availability at the primary site. If you have not yet configured SQL Server for high availability at the primary site, go to High Availability (HA) Configuration in the Solutions Configuration Center and follow the steps in the order shown.

See [“Tasks for configuring a new server for high availability”](#) on page 133.

See [“Tasks for configuring an existing server for high availability”](#) on page 135.

To verify the configuration, use the Cluster Manager (Java console) on the primary site and check the status of the service group in the tree view. Verify that all the resources are online.

Note: If you are setting up a replicated data cluster at the primary site, use the replicated data cluster instructions rather than the high availability configuration steps in the Solutions Configuration Center.

See [“VCS Replicated Data Cluster configuration”](#) on page 57.

Setting up your replication environment

The DR wizard can assist you with setting up replication for the following methods of replication:

- Volume Replicator
- EMC SRDF
- Hitachi TrueCopy

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

For array-based hardware replication, you can use any replication agent supported by Cluster Server. The DR wizard can help with configuring the methods listed above. If you choose a different replication method, you must run the wizard first to complete configuring global clustering; then afterwards, you configure replication separately.

See [“Configuring global clustering only”](#) on page 274.

Before configuring replication with the wizard, ensure that you set up the replication environment prerequisites.

Choose from the following topics, depending on which replication method you are using:

- Volume Replicator
See [“Setting up security for Volume Replicator”](#) on page 174.
- EMC SRDF
See [“Requirements for EMC SRDF array-based hardware replication”](#) on page 235.
- Hitachi TrueCopy
See [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 237.

Requirements for EMC SRDF array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for EMC SRDF. The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also configures the Symm heartbeat. Optional resource settings are left in the default state.

For more information about the EMC SRDF agent functions and the configuration options, see *Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide*.

Before using the DR wizard, review the following topics:

- See [“Software requirements for configuring EMC SRDF”](#) on page 236.
- See [“Replication requirements for EMC SRDF”](#) on page 236.

Software requirements for configuring EMC SRDF

The EMC SRDF agent supports SYMCLI versions that EMC recommends for the firmware on the array. The agent supports SRDF on all microcode levels on all Symmetrix arrays, provided that the host/HBA/array combination is in EMC's hardware compatibility list.

To use the DR wizard to configure the required agent settings for EMC SRDF, ensure that the following software requirements are met:

- The EMC Solutions Enabler is installed on all cluster nodes.
- The SYMCLI version that is installed supports the generation of XML output.
- The SYMCLI version and the microcode level support dynamic swapping.
- The VCS EMC SRDF agent is installed on all cluster nodes.

Replication requirements for EMC SRDF

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that no devices are RDF2.
- On the secondary site, the wizard verifies that no devices are RDF1.

Otherwise, the wizard displays an invalid configuration message and is unable to proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All disks in SFW disk groups must belong to the same device group.
- The device group must not span more than one array (no composite device groups).
- A device group can contain one or more disk groups.
- Dynamic swap must be enabled on both sites.

- On the primary site:
 - All devices must be RDF1 and part of an RDF1 device group.
 - Devices must have write access.
- On the secondary site:
 - All devices must be RDF2 and part of an RDF2 device group.
 - Write access must be disabled.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the SRDF resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the SRDF resource, not to the array configuration. However, the SRDF resource will be unable to come online in the service group until replication has been configured correctly.

Note: In addition, the agent requires that the device group configuration must be the same on all nodes of the cluster.

Requirements for Hitachi TrueCopy array-based hardware replication

The DR wizard configures the settings required for the VCS hardware replication agent for Hitachi TrueCopy. The wizard configures the required settings for the HTC resource in the VCS application service group. Optional settings are left in the default state.

For more information about the Hitachi TrueCopy agent functions and the configuration options, see *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*.

Before using the DR wizard, review the following topics:

- See [“Software requirements for Hitachi TrueCopy”](#) on page 237.
- See [“Replication requirements for Hitachi TrueCopy”](#) on page 238.

Software requirements for Hitachi TrueCopy

The Hitachi TrueCopy agent supports all versions of Hitachi RAID Manager.

To use the DR wizard to configure the required agent settings for Hitachi TrueCopy, ensure that the following requirements are met:

- RAID Manager is installed in the same location on all nodes on a site.
- Enter the primary and secondary site file paths for the `horcm` files on the Hitachi TrueCopy Path Information panel in the wizard. The default location is:

`systemDriver\Windows`

- The `horcm` files are named `horcmnn.conf` (where `nn` is a positive number without a leading zero, for example, **horcm1.conf**, but not **horcm01.conf**).

Replication requirements for Hitachi TrueCopy

Before it performs any tasks, the wizard validates the array configuration as follows:

- On the primary site, the wizard verifies that all devices are the same type, but not S-SWS or SSUS.
- On the secondary site, the wizard verifies that all devices are the same type, but not P-VOL or PSUS.

Otherwise, the wizard displays an invalid configuration message and does not proceed.

The DR wizard does not start or stop replication. Array replication configuration is not a prerequisite for the wizard to perform storage cloning or service group cloning.

After the service group cloning task is complete, the DR wizard displays a screen describing the following replication requirements:

- All configured instances are running.
- No disks in the SFW disk group span across the Device Group.
- A device group can contain one or more disk groups.
- The device group does not span more than one array.
- At the primary site, all devices are of the type P-VOL.
- At the secondary site, all devices are of the type S-VOL.
- All device groups at the primary site are paired to an IP address which must be online on the secondary node.
- Device group and device names include only alphanumeric characters or the underscore character.

It is recommended that you ensure that these requirements are met before proceeding with the wizard. The wizard then validates the array replication configuration.

If replication is configured correctly, the wizard populates the resource configuration screen with the required replication settings for the HTC resource.

If the replication configuration does not meet the requirements, the wizard leaves the fields on the resource configuration screen blank. You can optionally enter the resource configuration information in the wizard and configure the array replication requirements later. The information you enter is applied only to the HTC resource, not to the array configuration. However, the HTC resource will be unable to come online in the service group until replication has been configured correctly.

Assigning user privileges (secure clusters only)

In order to enable remote cluster operations you must configure a VCS user with the same name and privileges in each cluster.

When assigning privileges in secure clusters, you must specify fully-qualified user names, in the format `username@domain`. You cannot assign or change passwords for users when VCS is running in secure mode.

You must assign service group rights to the SQL Server service group as well as any dependent service groups except for the RVG service group.

See the *Cluster Server Administrator's Guide*.

To assign user privileges at the primary site

- 1 Set the configuration to read/write mode:

```
haconf -makerw
```

- 2 Add the user. Specify the name in the format `username@domain`.

```
hauser -add user [-priv <Administrator|Operator>]
```

- 3 Modify the attribute of the service group to add the user. Specify the SQL Server service group and any dependent service groups except for the RVG service group.

```
hauser -add user [-priv <Administrator|Operator>  
[-group service_groups]]
```

- 4 Reset the configuration to read-only:

```
haconf -dump -makero
```

To assign user privileges at the secondary site

- 1 Set the configuration to read/write mode:

```
haconf -makerw
```

- 2 Add the user. Specify the name in the format `username@domain`.

```
hauser -add user [-priv <Administrator|Operator>]
```

- 3 Reset the configuration to read-only:

```
haconf -dump -makero
```

About configuring disaster recovery with the DR wizard

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

The Disaster Recovery Configuration Wizard (DR wizard) assists you to perform the following tasks for the selected service group:

- Clone the storage configuration (Volume Replicator replication) or prepare a temporary storage configuration for application installation (array-based hardware replication)
- Clone the service group
- Optionally, configure Volume Replicator replication, or configure the VCS hardware replication agent settings for EMC SRDF or Hitachi TrueCopy
- Configure global clustering

Warning: To use the Disaster Recovery Configuration Wizard in an array-based hardware replication environment that is not configured by the wizard, you must first run the wizard to configure global clustering before configuring replication.

The wizard allows you to exit after the logical completion of each task. Each time you re-start the wizard, you specify the primary site system, service group, secondary site system, and replication method, as described in the following procedure. Clicking

Next then takes you to the start page of the process following the one that you had last completed.

The DR Wizard list of service groups shows only those that contain a MountV resource. For a dependent service group to be listed, the parent service group must also contain a MountV resource.

Warning: Once you have completed configuring replication and global clustering with the DR wizard, you cannot use the wizard to change the method of replication.

Before running the DR wizard to configure disaster recovery, ensure that you meet the following prerequisites:

- InfoScale Enterprise is installed and a cluster is configured at the secondary site. Ensure that the name assigned to the secondary site cluster is different than the name assigned to the primary site cluster.
- Your application or server role is configured for HA at the primary site and all required services are running at the primary site.
- The clusters taking part in the DR configuration should have distinct names.
- Enough free disk space is available at the secondary site to duplicate the storage configuration at the primary site.
- Global Cluster Option (GCO) is installed at the primary and secondary site, and one static IP address is available at each site for configuring GCO.
- For IPv4 networks, static IP addresses are available to enter for the following (for IPv6, they are generated during configuration):
 - One static IP address per application service group to be cloned.
 - One static IP address at each site for configuring GCO.
 - If using Volume Replicator for replication, a minimum of one static IP address per site for each application instance running in the cluster.
- The service group to be cloned can use either IPv4 IP addresses or IPv6 addresses but not a mixture of both.
- To configure IPv6 settings, the wizard must be launched from a system on which the IPv6 stack is installed.
- For Volume Replicator replication, the service group to be cloned cannot contain a child service group.
- A VCS user is configured with the same name and privileges in each cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall is set to allow both ingoing and outgoing TCP requests on port 7419.

Note: The DR wizard does not support Volume Replicator configurations that include a Bunker secondary site.

In addition, see the following replication prerequisites, depending on the replication method you are using:

- See [“Setting up security for Volume Replicator”](#) on page 174.
- See [“Requirements for EMC SRDF array-based hardware replication”](#) on page 235.
- See [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 237.

Configuring disaster recovery with the DR wizard

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

This procedure describes how to configure disaster recovery using the wizard.

To start configuring disaster recovery with the DR wizard

- 1 Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center**, or, on Windows 2012 operating systems, from the **Apps** menu.
- 2 Expand the Solutions for Microsoft SQL Server tab.
 Expand the Solutions for Additional Applications tab.
 Click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

Note: By design, the DR wizard requires specific settings for the Lanman attributes on the primary and secondary sites. Before beginning the DR configuration, the wizard checks for these values, and if they are not set as required, the wizard will automatically proceed with setting these values, both at the primary and secondary sites.

- 3 In the Welcome panel, review the prerequisites to ensure that they are met and click **Next**.

- 4 In the System Selection panel, provide information in the **System Name** field:

Enter the IP address or Fully Qualified Host Name (FQHN) of the primary system where the SQL Server instance is online.

If you have launched the wizard on the system where the instance is online at the primary site, you can also specify `localhost` to connect to the system.

Click **Next**.

- 5 In the Service Group Selection panel, select the service group that you want to clone to the secondary site.

For a hardware replication environment, you can choose to clone only the parent service group by not selecting the dependent service group. Only online and local dependencies are supported, in soft, firm, or hard configurations. The wizard can configure only one level of dependency.

In a Volume Replicator environment, the DR wizard does not support configuring DR for a service group that has a child. If you select a service group that has a child, you will receive an error message when you select the Volume Replicator replication method later in the wizard.

The panel lists only service groups that contain a MountV resource.

Click **Next**.

- 6 In the Secondary System Selection panel, enter the Fully Qualified Host Name (FQHN) or the IP address of the secondary system for which you want to configure disaster recovery.

Click **Next**.

- 7 In the Replication Options panel, select the replication method. Although you must select the replication method now, configuring replication and the global cluster option is done later, after service group cloning.

Configure Volume Replicator (Volume Replicator) and the Global Cluster Option (GCO)	<p>Select this option if you want to configure Volume Replicator replication.</p> <p>Select this option even if you plan to configure Volume Replicator replication or the GCO option manually. This option is required for the wizard to configure the storage cloning correctly for a Volume Replicator environment.</p> <p>The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p> <p>You cannot mix replication methods. That is, if your primary site is using array-based replication, and you select the Volume Replicator option, the wizard will warn you that you cannot use Volume Replicator replication for the disaster recovery site.</p>
Configure EMC SRDF and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS EMC SRDF agent. All disks used for the service group on the primary site must belong to an EMC SRDF array.</p> <p>Select this option even if you plan to configure EMC SRDF replication or the GCO option manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>
Configure Hitachi TrueCopy and the Global Cluster Option (GCO)	<p>Select this replication option if you want to configure the settings for the VCS Hitachi TrueCopy agent. All disks used for the service group on the primary site must belong to a Hitachi TrueCopy array.</p> <p>Select this option even if you configure GCO manually. The wizard verifies each configuration task and recognizes if a task has been completed successfully.</p>

Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)

Configure the Global Cluster Option (GCO) only

If you select this option, the DR wizard does not configure any replication settings. It configures the global cluster option.

Select this option if you want to use the wizard in an array-based replication environment that is not supported by this wizard. You must configure replication manually after you finish the wizard.

If you select the GCO only option, the DR wizard sets up the storage and service group configuration on the secondary site for an array-based hardware replication environment.

Therefore, you cannot use this option to clone the storage and service group for a Volume Replicator replication environment.

Click **Next**.

8 Continue with the next DR configuration task.

For Volume Replicator replication:

See [“Cloning the storage on the secondary site using the DR wizard \(Volume Replicator replication option\)”](#) on page 245.

For array-based replication:

See [“Creating temporary storage on the secondary site using the DR wizard \(array-based replication\)”](#) on page 249.

Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)

The DR wizard enables you to clone the storage configuration present at the primary site on to the secondary site. To do this successfully, the systems at the secondary site must have adequate free storage. If you have created the configuration but there is a mismatch in the volume sizes, the wizard can correct this and then complete the configuration.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you have not yet started the wizard, refer to the following topic before continuing with the storage cloning procedure:

To clone the storage configuration from the primary site to the secondary site (Volume Replicator replication method)

- 1
- If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system. In the Replication Options panel, select the Volume Replicator replication method and click **Next**.
- 2
- Review the information in the Storage Validation Results panel. This panel compares the configuration at the secondary site with that on the primary. If the storage is already configured identically on both sites, the panel shows that results are identical. Otherwise, the panel shows the differences and recommended actions. You can toggle between a summary and detailed view of information about the differences.

The detailed view shows the following:

Disk Group	Displays the disk group name that needs to be created on the secondary site.
Volume	Displays the list of volumes, if necessary, that need to be created at the secondary site.
Size	Displays the size of the volume that needs to be created on the secondary site.
Mount	Displays the mount to be assigned the volume on the secondary site.
Recommended Action	<div>Indicates the action that needs to be taken at the secondary to make the configuration similar to that on the primary.<ul style="list-style-type: none">■ If the volume does not exist, a new volume will be created.■ If the volume exists but is of a smaller size than that on the primary, the volume will be expanded to the required size.■ If the volume is of a greater size than that on the primary, the volume will be recreated using the appropriate size.■ If the volume is the same as that on the primary, the message indicates that the volumes are identical and no action is required.</div>

The summary view shows the following:

Disk groups that do not exist	Displays the names of any disk groups that exist on the primary but do not exist on the secondary.
Existing disk groups that need modification	Displays the names of any disk groups on the secondary that need to be modified to match the primary.

Cloning the storage on the secondary site using the DR wizard (Volume Replicator replication option)

Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.
---------------------------------	---

If the panel displays a message indicating that the available disks are inadequate to clone the primary site configuration on the secondary, you can free some disks on the secondary or add more storage. Then, click **Refresh/Validate** to have the wizard update its information about the secondary storage configuration.

You continue with the wizard to provide information for the recommended actions. Before proceeding to the service group configuration, the wizard ensures that the configuration of the disk groups and volumes for the service group is the same at the primary and secondary site.

Click **Next**.

- 3** In the Disk Selection for Storage Cloning panel, for each of the disk groups that does not exist or is not same as the corresponding disk group at the primary site, select disks that the wizard can use to create the respective disk groups at the secondary site.

Selecting Disks	For each of the disk groups that needs to be created, select the required disks from the Available Disks pane. Either double-click on the host name or the >> button to move the hosts into the Selected disks pane.
-----------------	--

Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures.

Click **Next**.

- 4 In the Volume Layout for Secondary Site Storage panel, complete the requested information:

Disk Group	Displays the disk group name to which the volume belongs.
Volume (Volume Size)	Displays the name and the size of the volume, corresponding to that on the primary, that needs to be created on the secondary.
Available Disks	<p>Select the disks on which you want the wizard to create the volumes. From the Available Disks pane, either double-click on the disk name or the >> button to move the disks into the Selected Disks pane. For each disk group the Available disks pane displays the list of disks that are part of the disk group.</p> <p>Select disks for each unavailable volume that you want to clone on to the secondary.</p>
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been moved in from the Available Disks pane.
View Primary Layout	Displays the volume layout at the primary site. Use this information as a reference to specify the details for the Secondary layout.

Click **Next**.

- 5 In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the storage configuration at the secondary site.

Note: On the VEA GUI of the secondary site, a Windows dialog box might appear prompting you to format a disk. Click **Cancel** to close the dialog.

The appearance of this dialog box has no impact on the operations being performed by the DR wizard. You can safely ignore it.

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

- 6 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 7 In the Storage Cloning Configuration Result screen, view the results and click **Next**.
- 8 In the SQL Server Installation panel, review the information. If the application is already installed on the secondary site nodes, click **Next** to proceed with service group cloning.

Otherwise, proceed with installation on the required nodes as follows:

- Before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.
- The system may get restarted when the SQL installation is complete. Therefore, if you are running the DR wizard from a system where you need to install SQL, click **Finish** to exit the wizard before proceeding with installation. Afterwards, restart the Disaster Recovery wizard and continue through the wizard from the Welcome panel.
- If the DR Wizard is run from a remote node, you can keep the wizard running on that node while you install the application locally on each of the required nodes. Once application installation is complete, click **Next** to proceed with service group cloning.

Creating temporary storage on the secondary site using the DR wizard (array-based replication)

To enable you to install applications, the DR wizard can create a temporary disk group, DR_APP_INSTALL_DG, which contains the volumes and mount points for use in application installation. The temporary configuration uses 500 MB volumes or the volume size at the primary site, depending on which is smaller. The wizard deletes the temporary configuration after application installation.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you have already installed the application on all nodes, you can skip this storage cloning step by unchecking the Perform storage cloning check box on the Storage Cloning panel.

If you are starting the wizard for the first time, refer to the following topic before continuing with the storage cloning procedure:

See [“About configuring disaster recovery with the DR wizard”](#) on page 240.

To create temporary storage for application installation (array-based replication)

- 1 If you have not yet done so, start the Disaster Recovery Configuration Wizard and specify the information for the primary site system, the service group, and the secondary site system.
- 2 In the Replication Options panel, select the array-based replication method you plan to use and click **Next**:
 - EMC SRDF
 - Hitachi TrueCopy
 - Global Cluster Option only (select if you are using another agent-supported array-based replication method)
- 3 If you selected Hitachi TrueCopy replication, the Hitachi TrueCopy File Paths panel is displayed. The wizard populates the fields if it locates the files in the default location. Otherwise, fill in the file path information for both the primary and secondary sites as follows:

RAID Manager bin path	Path to the RAID Manager Command Line interface The default path is C:\HORCM\etc, where C is the system drive.
HORCM files location	Path to the horcm configuration files (horcmnn.conf) The default path is: C:\Windows, where C is the system drive. The horcm configuration file is required by the RAID Manager on all nodes; however, the wizard does not validate its presence.

- 4 In the Storage Cloning panel, you can choose whether or not to perform storage cloning, which creates a temporary storage disk group and volumes for application installation. The wizard will delete the temporary storage once you confirm application installation is complete.

Choose one of the following:

- If you have not yet installed the application on all nodes, leave **Perform storage cloning** checked and click **Next**. Continue with the next step in this procedure.
 - If you have already installed the application on all nodes, uncheck **Perform storage cloning** and click **Next**. Continue with the procedure for service group cloning.
- 5 The Storage Validation Results panel shows the temporary storage configuration that the wizard will configure at the secondary site. You can click **Show Summary** to toggle to a summary view and toggle back to a detailed view by clicking **Show Details**.

The detailed view shows the following:

Disk Group	Displays the name of the single disk group required on the secondary site for temporary storage: DR_APP_INSTALL__DG
Volume	Displays the list of volumes required at the secondary site.
Size	Displays the size of the volumes required on the secondary site.
Mount	Displays the mounts required at the secondary site.
Recommended Action	Indicates the action that the wizard will take at the secondary site.

The summary view shows the following:

Existing configuration	Displays the existing secondary configuration.
Free disks present on secondary	Displays the list of free disks that exist on the secondary along with details about the free space and total disk space information.

If the panel displays a message indicating that the available disks on the secondary are inadequate, you can free some disks on the secondary or add more storage. Then, click **Refresh/Validate** so that the wizard can update its information about the secondary storage configuration. Click **Next**.

- 6
- In the Disk Selection for Storage Cloning panel, a default disk selection is shown for the temporary storage at the secondary site. You can change the selection by moving disks to and from the Available Disks and Selected Disks pane. Under the Available Disks label, a drop-down list allows you to filter available disks by disk enclosure name. The default is All, which displays all free disks available on all enclosures. Click **Next**.
- 7
- The Volume Layout for Secondary Site Storage panel shows a default volume layout for the temporary storage based on the primary site volume layout. Optionally, you can change the default disk assignment and layout for any volume:

Disk Group	Displays the DR_APP_INSTALL__DG disk group.
Volume (Volume Size)	Displays the name and the size of the volume to be created on the secondary.
Available Disks	Displays the disks that are available for the volumes. To select a disk, either double-click on the host name or click the >> button to move the hosts into the Selected Disks pane.
Layout	By default, the same layout as the one specified for the primary volume is selected. Click Edit to change the layout to suit your specific requirements.
Selected Disks	Displays the list of disks that have been selected for the volume. To remove a disk from the list, select it and click the << button.
View Primary Layout	Displays the volume layout at the primary site.

Click **Next**.

- 8
- In the Storage Configuration Cloning Summary panel, review the displayed information. If you want to change any selection, click **Back**. Otherwise, click **Next** to allow the wizard to implement the temporary storage configuration at the secondary site.

- 9 In the Implementation panel, wait until the status for all the completed tasks is marked with a check symbol, indicating successful completion. Wait until the wizard completes cloning the storage. The progress bar indicates the status of the tasks. If some task could not be completed successfully, then the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure.

Click **Next**.

Note: If SCSI-3 support is enabled for using Persistent Group Reservations (PGR), and if one of the selected disks is not SCSI-3 compliant, the following error is displayed: “Unable to reserve a majority of dynamic disk group members. Failed to start SCSI reservation thread.”

Recommended action: Click **Finish** to exit the wizard. Either replace the non-compliant disk with a SCSI-3 compliant disk, or enable SCSI-2 support, and then run the wizard again.

- 10 In the Storage Configuration Cloning Result screen, view the results and click **Next**.
- 11 In the SQL Server Installation panel, review the information and do one of the following:
 - Before you begin installation, ensure that your disk groups are imported and volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the wizard does not mount the volumes on the secondary site. You must manually format the volumes and assign the drive path to the volumes using Veritas Enterprise Administrator. Use the same letters and folder names that were assigned on the primary site.
 - The system may get restarted when the SQL Server installation is complete. Therefore, if you are running the DR wizard from a system where you need to install SQL Server, click **Finish** to exit the wizard and proceed with installing the application on the required nodes. Afterwards, restart the Disaster Recovery wizard and continue through the wizard from the Welcome panel.
 - If the DR Wizard is run from a remote node, you can keep the wizard running on that node while you install the application locally on each of the required nodes. Once application installation is complete, click **Next** to proceed with service group cloning.

Once the application is installed, the temporary cloned storage is no longer needed. Before beginning service group cloning, the wizard displays the

Temporary Storage Deletion panel to confirm the deletion of the temporary storage.

Installing and configuring SQL Server on the secondary site

Use the same installation and configuration procedures for SQL Server as on the primary site but note the following considerations when installing SQL Server on the secondary site.

- Before installing Microsoft SQL Server, verify that the cluster disk group is imported to the first node and the volumes are mounted. If volumes were mounted as drive paths (folder mount) on the primary site, the DR Wizard does not mount the volumes on the secondary site and you must format the volumes and mount them manually.

See [“About managing disk groups and volumes”](#) on page 98.

- During installation, use the same instance name as on the primary site.

Cloning the service group configuration from the primary site to the secondary site

Before cloning a service group on the secondary site, verify if the application is installed on the secondary site.

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

Before cloning the service group on the secondary site, verify that you have installed the application on the secondary site on all nodes for this SQL Server instance.

Ensure that the SQL Server Full-Text Search service on the secondary site is configured to start in the manual mode and is initially in the stopped state.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you are launching the wizard for the first time, refer to the following topic for additional information:

See [“About configuring disaster recovery with the DR wizard”](#) on page 240.

Note: Although you can view the cloning progress in the VCS Java Console, do not save and close the configuration while cloning is in progress. Otherwise, the cloning fails and you have to delete the service group on the secondary site and run the wizard again.

To clone the service group configuration from the primary site to the secondary site

- 1 At the primary site, verify that you have brought the application service group online.
- 2 Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.

Expand the Solutions for Microsoft SQL Server tab.

Click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

- 3 In the Welcome panel, click **Next** and continue through the wizard, providing the requested information for the primary site system, the service group, the secondary site system, and the replication method.

If you selected the Volume Replicator replication method, the wizard proceeds to the storage cloning task and notifies you if it detects that the storage is identical. Click **Next** until you reach the Service Group Analysis panel.

If you selected an array-based replication method (EMC SRDF, HTC, or GCO only), the temporary storage is no longer needed once the application is installed and the wizard confirms whether or not to delete it.

- 4 (Array-based replication method only) In the Temporary Storage Deletion panel, confirm whether or not to delete the cloned storage:
 - If the application is already installed on the required nodes, leave **Delete cloned storage** checked and click **Next**. When the wizard prompts you to confirm deleting the shared storage, click **Yes**.
 - If you want to delete the cloned storage manually later, uncheck **Delete cloned storage** and click **Next**.
- 5 (Array-based replication method only) If you selected to delete the cloned storage, the wizard shows the progress of the tasks in the Implementation panel. If the storage deletion fails, the wizard will show a failure summary page. Otherwise, when it shows the tasks are complete, click **Next**.

Cloning the service group configuration from the primary site to the secondary site

- 6** Review the following information displayed in the Service Group Analysis panel and click Next to continue with service group cloning.

Service Group Name Displays the list of application-related service groups present on the cluster at the primary site.

Service Group Details on the Primary Cluster Displays the resource attributes for the service group at the primary site.

The NIC resource consists of the MAC address.

The IP resource consists of the IP address and subnet mask.

Service Group Details on the Secondary Cluster Displays a message to indicate whether the service group or the corresponding attributes have been configured at the secondary site.

- 7** In the Service Group Cloning panel, specify the requested system information for the secondary site.

Service Group Name Depending on the application service group already created at the primary site, and subsequently selected on the Service Group Selection page, the wizard displays the names of the service groups that will be cloned at the secondary site.

Available Systems Displays a list of available systems on the secondary cluster that are not yet selected for service group cloning.

Select any additional secondary systems on which you want the wizard to clone the application service group configuration.

Either double-click on the system name or use the > option to move the hosts into the Selected Systems pane.

Note: If you want to add systems to a service group after you finish cloning the service group configuration with the DR wizard, you cannot do so by running the DR wizard again. Instead, run the VCS configuration wizard and edit the system list of the existing service group.

Selected Systems Displays the list of selected systems. The secondary system that you selected earlier in the wizard is listed by default.

Click **Next**.

- 8 In the Service Group Attribute Selection panel, complete the requested information to create the required resources on the secondary site. The panel also displays the service group resource name and the attribute information at the primary site.

Resource Name	Displays the list of resources that exist on the primary cluster.
Attribute Name	<p>Displays the attribute name associated with each of the resources displayed in the Resource Name column.</p> <p>If you need to edit additional attributes that are not shown, you must edit them manually on the secondary site service group once service group cloning is complete.</p>
Primary Cluster	Displays the primary attribute values for each of the displayed attributes.
Secondary Cluster	<p>For an IPv4 network, the default is the same as the primary cluster; the same virtual IP address can be used if both sites exist on the same network segment. You can specify different attributes depending on your environment.</p> <p>For IPv6, select the network from the dropdown list. If you select the same subnet as the primary site, the primary site IP address will be used. Otherwise the IP address will be generated from the network.</p> <p>For the MACAddress attribute select the appropriate public NIC from the drop-down list.</p> <p>For IPv6 available NICs are those belonging to the selected IPv6 network.</p>

Click **Next**.

- 9 In the Service Group Summary, review the attribute information that will be cloned on to the secondary cluster. Click **Back** to change any of the secondary service group attributes. Otherwise, click **Next** to proceed with cloning the service group configuration on the secondary site.

- 10 In the Implementation panel, wait until all the tasks are completed. The progress bar indicates the status of the tasks. Successful tasks are marked with a check symbol. If some task could not be completed successfully, the task is marked with an (x) symbol. The Information column displays details about the reasons for task failure. Click **Next**.
- 11 If the cloning failed, review the troubleshooting information. Otherwise, click **Next** to continue with the replication and GCO configuration, or with GCO only, depending on which option you selected. Optionally, you can exit the wizard at this point and launch the wizard again later. When you launch the wizard again, continue through the wizard, specifying the primary site system, the service group, the secondary site system, and the replication method. Click **Next** to continue to the replication and/or GCO configuration task.

To configure an MSDTC service group:

See [“Tasks for configuring MSDTC for high availability”](#) on page 155.

Configuring the SQL Server service group in a non-shared storage environment

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to clone the application service group created at the primary site if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure the service group manually using the Cluster Manager (Java Console).

Note the following before configuring the service group at the secondary site:

- Ensure that the application agent resources, the Lanman resource (if configured), and the IP resource is offline in the service group on the primary site. The remaining resources, including the storage resources, must be online.
- Ensure that the name of the service group is the same as that on the primary site.
- After configuring the service group do not bring it online on the secondary site at this time. You can bring it online later after completing all the DR configuration steps.

See [“Configuring the service group in a non-shared storage environment”](#) on page 148.

Configuring replication and global clustering

After creating the identical service group configuration on both sites, the DR wizard helps you set up replication and global clustering (GCO option). You can choose to configure replication using Volume Replicator or an agent-supported array-based hardware replication.

Note: The DR wizard cannot be used if you are setting up DR in a non-shared storage environment.

If you are using an array-based replication that is not supported by the wizard, you configure global clustering only. In this case, you must complete configuring global clustering before configuring replication.

The following topics cover the steps required for each replication method:

- See [“Configuring Volume Replicator replication and global clustering”](#) on page 259.
- See [“Configuring EMC SRDF replication and global clustering”](#) on page 267.
- See [“Configuring Hitachi TrueCopy replication and global clustering”](#) on page 270.
- See [“Configuring global clustering only”](#) on page 274.

Configuring Volume Replicator replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure Volume Replicator replication and global clustering.

Note: The DR configuration wizard clubs the data and log volumes of one database in one RVG. It clubs the system database volumes and the RegRep volume in one RVG, and uses a separate RVG for each user-created database. For information about setting up Volume Replicator replication with VEA, see the *Volume Replicator Administrator's Guide*

Before you begin, ensure that you have met the following prerequisites:

- Ensure that Volume Replicator Security Service (VxSAS) is configured at the primary and secondary site.
See [“Setting up security for Volume Replicator”](#) on page 174.
- Verify whether the IP version preference is set before you configure replication. If you specify host names when you configure replication, Volume Replicator resolves the host names with the IP addresses associated with them. This setting determines which IP version Volume Replicator uses to resolve the host names.

Use one of the following methods to set the IP preference:

- Veritas Enterprise Administrator (VEA) GUI—select the appropriate options on the Control Panel > VVR Configuration > IP Settings tab.
- Run the `vxtune ip_mode [ipv4 | ipv6]` command at the primary site as well as the secondary site.
- Ensure that a minimum of one static IP address per site is available for each application instance running in the cluster.
- Ensure that, for remote cluster configuration, you configure a VCS user with the same name and privileges in each cluster.

Use the following procedure to configure Volume Replicator replication and global clustering with the DR wizard.

To configure Volume Replicator replication and GCO

- 1 Verify that the application server service group is online at the primary site and the appropriate disk groups are imported at the secondary site.
- 2 If the DR wizard is still open after the previous wizard task, continue with the Replication Setup panel.

Otherwise, launch the wizard and proceed to the Replication Setup panel as follows:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
 - Expand the Solutions for Microsoft SQL Server tab and click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- 3 On the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
 - 4 On the Replication Methods panel, click **Configure Volume Replicator and the Global Cluster Option (GCO)**. Click **Next**.
 - 5 In the Internet Protocol panel, select IPv4 or IPv6 depending on which type of network you are using. (You must use the same on primary and secondary sites.) Click **Next**.
 - 6 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**. If not, click **Cancel** and restart the wizard after meeting the requirements.

- 7 In the Replication Settings for Replicated Volume Group panel, specify the requested information. If you are adding a DR site to an existing DR configuration, fields that must match the existing settings, such as the RVG or RDS name, are dimmed so that you cannot change them.

Disk Group	The left column lists the disk groups. By design, an RVG is created for each disk group.
RVG Name	Displays the default RVG name. If required, change this to a name of your choice.
RDS Name	Displays the default Replicated Data Set (RDS) name. If required, change this to a name of your choice.
Available Volumes	<p>Displays the list of available volumes that have not been selected to be a part of the RVG.</p> <p>Either double-click on the volume name or use the > button to move the volumes into the Selected RVG Volumes pane.</p>
Selected RVG Volumes	<p>Displays the list of volumes that have been selected to be a part of the RVG.</p> <p>To remove a selected volume, either double-click the volume name or use the < button to move the volumes into the Available Volumes pane.</p> <p>Veritas recommends excluding tempdb from replication. If you earlier moved tempdb to a separate volume in the same disk group as the system database volumes, you can exclude tempdb from replication by removing the tempdb volume from the Selected RVG Volumes pane.</p>
Primary SRL	<p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the name, size, and disk.</p> <p>Otherwise, select the appropriate primary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>
Secondary SRL	<p>If you did not create a Replicator Log volume on the primary site, click Create New on the drop-down menu. On the New Volume dialog box, specify the same name and size as you specified for the primary SRL.</p> <p>Otherwise, select the appropriate secondary Replicator Log volume from the drop-down menu and enter an appropriate size.</p>

Start Replication after the wizard completes Select this check box to start replication automatically after the wizard completes the necessary configurations.

Once replication is configured and running, deselecting the checkbox does not stop replication.

Click **Advanced Settings** to specify some additional replication properties.

Advanced Replication Settings for RVG_TESTFS_0

Replication Mode:	Synchronous Override	Protocol:	UDP
Log Protection:	AutoDCM	Packet Size (Bytes):	1400
Primary RLINK Name:	48326630361624	Latency Protection:	Fail
Secondary RLINK Name:	48326630361623	High Mark Value:	10000
Bandwidth:	Maximum	Low Mark Value:	9950
	Mbps	Initial Synchronization:	Auto Synchronous

OK Cancel

The options on the dialog box are described column-wise, from left to right:

- Replication Mode Select the required mode of replication; **Synchronous**, **Asynchronous**, or **Synchronous Override** (default).
- Log Protection Select the appropriate log protection from the list:
- **AutoDCM** is the default selected mode for the Replicator Log overflow protection when all the volumes in the Primary RVG have a DCM log. The DCM is enabled when the Replicator Log overflows.
 - The **Off** option disables Replicator Log Overflow protection.
 - The **Override** option enables log protection. If the Secondary node is still connected and the Replicator Log is about to overflow then the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log.
If the Secondary becomes inactive due to disconnection or administrative action then Replicator log protection is disabled, and the Replicator Log overflows.
 - The **Fail** option enables log protection. If the log is about to overflow the writes are stalled until a predetermined amount of space, that is, 5% or 20 MB (whichever is lesser) becomes available in the Replicator Log. If the connection between primary and secondary RVG is broken, then, any new writes to the primary RVG are failed.

Primary RLINK Name	Enter a name of your choice for the primary RLINK. If you do not specify any name then the wizard assigns a default name.
Secondary RLINK Name	Enter a name of your choice for the Secondary RLINK. If you do not specify any name then the wizard assigns a default name.
Bandwidth	<p>By default, Volume Replicator replication uses the maximum available bandwidth. You can select Specify to specify a bandwidth limit.</p> <p>The default unit is Megabits per second (Mbps) and the minimum allowed value is 1 Mbps.</p>
Protocol	Choose TCP or UDP. UDP/IP is the default replication protocol.
Packet Size (Bytes)	Default is 1400 Bytes. From the drop-down list, choose the required packet size for data transfer. The default unit for the packet size is Bytes. You can set the packet size only if the protocol is UDP/IP.
Latency Protection	<p>By default, latency protection is set to Off.</p> <p>When this option is selected the High Mark Value and the Low Mark Value are disabled. Select the Fail or Override option to enable Latency protection.</p> <p>This Override option behaves like the Off option when the Secondary is disconnected and behaves like the Fail option when the Secondary is connected.</p>
High Mark Value	<p>This option is enabled only when Latency Protection is set to Override or Fail. It specifies the maximum number of pending updates by which the secondary site can be behind the primary site. The default value is 10000.</p> <p>To ensure that latency protection is most effective the difference between the high and low mark values must not be very large.</p>
Low Mark Value	<p>This option is enabled only when Latency Protection is set to Override or Fail. When the updates in the Replicator log reach the High Mark Value, then the writes to the system at the primary site continues to be stalled until the number of pending updates on the Replicator log falls back to the Low Mark Value. The default is 9950.</p>

Initial Synchronization If you are doing an initial setup, then use the **Auto Synchronous** option to synchronize the secondary site and start replication. This is the default.

When this option is selected, Volume Replicator by default performs intelligent synchronization to replicate only those blocks on a volume that are being used by the file system. If required, you can disable intelligent synchronization.

If you want to use the **Synchronize from Checkpoint** method then you must first create a checkpoint.

If you have a considerable amount of data on the primary data volumes then you may first want to synchronize the secondary for existing data using the backup-restore method with checkpoint. After the restore is complete, use the **Synchronize from Checkpoint** option to start replication from the checkpoint to synchronize the secondary with the writes that happened when backup-restore was in progress.

To apply changes to advanced settings, click **OK**.

For additional information on Volume Replicator replication options, refer to the *Volume Replicator Administrator's Guide*.

Click **Next**.

- 8 In the Replication Attribute Settings panel, specify required replication attribute information for the cluster at the primary and secondary site. Click the arrow icon to expand an RVG row and display the replication attribute fields. If you are configuring an additional secondary site (multiple DR sites), some fields are disabled.

Disk Group	Displays the list of disk groups that have been configured.
RVG Name	Displays the Replicated Volume Groups corresponding to the disk groups.
IP Address	<p>For IPv4 networks, enter replication IPs that will be used for replication, one for the primary site and another for the secondary site.</p> <p>For IPv6, select the network from the dropdown list. An IP address will be generated.</p>
Subnet Mask or Prefix	<p>For IPv4, enter the subnet mask for the system at the primary site and the secondary site.</p> <p>For IPv6, enter the prefix.</p>
Public NIC	<p>Select the public NIC from the drop-down list for the system at the primary and secondary site.</p> <p>For IPv6, available NICs are those belonging to the selected network.</p>
Copy	Enables you to copy the above network settings to any additional RVGs that are listed on this screen. If there is only one RVG, this option does not apply.

After specifying the replication attributes for each of the RVGs, click **Next**.

- 9 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	For IPv4, enter a virtual IP for the WAC resource. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 10 In the Settings Summary panel, review the displayed information.

Click **Back** if you want to change any of the parameters. If you have a printer installed, you can click the printer icon at the bottom of the scrollable list to print the settings.

Otherwise, click **Next** to implement the settings.

- 11 In the Implementation panel, wait till the wizard completes creating the replication configuration and the WAC resource required for global clustering. If a task could not be completed successfully, it is marked with an (x) symbol. For any critical errors, the wizard displays an error message. For less critical errors, the Information column displays a brief description about the task failure and the next screen displays additional information on what action you can take to remedy it. Click **Next**.
- 12 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Configuring EMC SRDF replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the SRDF resource in the application service group.

Ensure that you have one static address is available per site for configuring GCO.

See [“Requirements for EMC SRDF array-based hardware replication”](#) on page 235.

The wizard configures the required agent settings as well as the SYMM heartbeat. It uses defaults for optional settings.

See [“Optional settings for EMC SRDF”](#) on page 270.

To configure EMC SRDF replication and GCO

- 1 Verify that you have brought the application service group online at the primary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel.

Otherwise, launch the wizard and proceed to the Replication Setup panel by following these steps sequentially:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
- Expand the Solutions for Microsoft SQL Server tab.
Click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.

- In the Replication Methods panel, click **Configure EMC SRDF and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 3 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the SRDF resource cannot come online in the service group.

- 4 In the SRDF Resource Configuration panel, the wizard populates the required resource fields if replication has been configured. Otherwise, you must enter the required resource settings manually.

Symmetrix Array ID (SID)	Specify the array ID for the primary site and for the secondary site.
Device Group name	Specify the name of the Symmetrix device group that contains the disks of the disk group for the selected instance.
Available VMDG Resources	Select the disk groups associated with the selected application instance.

- 5 If you want to configure an additional SRDF resource for the instance, click **Add**. Otherwise, click **Next**.

- 6 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	<p>Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes.</p> <p>Once GCO is configured and running, deselecting the checkbox does not stop GCO.</p>

Click **Next**.

- 7 In the Settings Summary panel, review the displayed information.

Click **Back** if you want to change any of the parameters specified for the replication resource settings or the global cluster settings.

Otherwise, click **Next**.

- 8 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check symbol indicates successful completion of a task. An (x) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.

- 9 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.
- 10 Proceed with configuring additional optional settings for the SRDF resource if desired, and then verifying the disaster recovery configuration.

Optional settings for EMC SRDF

The wizard configures the required settings for the SRDF resource in the VCS application service group. The wizard also detects and configures the SymHome attribute.

Other settings are left in the default state. For information on configuring the optional settings, see *Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide*. If you change any settings, ensure that you edit the resource on both the primary and secondary sites.

The optional settings use the following defaults:

Option	Default setting
DevFOTime	2 seconds per device required for a device to fail over
AutoTakeover	The default is 1; the agent performs a read-write enable on partitioned devices in the write-disabled state during a failover, if devices are consistent.
SplitTakeover	The default is 1; the agent brings service groups online on the R2 side even if the devices are in the split state because they are read-write enabled.

Configuring Hitachi TrueCopy replication and global clustering

After you complete the service group configuration task in the DR wizard, you configure replication and global clustering.

The wizard helps you to configure the settings for the HTC resource in the application service group.

See [“Requirements for Hitachi TrueCopy array-based hardware replication”](#) on page 237.

Ensure that you have one static address is available per site for configuring GCO.

The wizard configures the required agent settings. It uses defaults for optional settings.

See [“Optional settings for HTC”](#) on page 274.

To configure Hitachi TrueCopy replication and GCO

- 1 Verify that you have brought the application server service group online at the primary site.
- 2 If the wizard is still open after the service group cloning task, continue with the Replication Setup panel.

Otherwise, launch the wizard and proceed to the Replication Setup panel by following these steps sequentially:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
Expand the Solutions for Microsoft SQL Server tab.
Click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.

- 3 In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
- 4 In the Replication Methods panel, click **Configure Hitachi TrueCopy and the Global Cluster Option (GCO)**. Click **Next** and continue to the Replication Setup panel.
- 5 In the Replication Setup panel, review the replication requirements. If you have met the requirements, click **Next**.

Warning: Although you can continue with the wizard even if replication requirements are not met, the wizard will warn you that the configuration is not valid. If the configuration is not valid, the HTC resource cannot come online in the service group.

- 6 In the HTC Resource Configuration panel, the wizard populates the required resource fields if the `horcm` file is configured properly. If not, you can configure the `horcm` file and click **Refresh** to populate the fields. Alternatively, enter the required resource settings manually:

Instance ID	Specify the instance number of the device group. Multiple device groups may have the same instance number.
Device Group name	Specify the name of the Hitachi device group that contains the disk group for the selected instance. The device group name must be the same on both the primary and secondary sites.
Available VMDG Resources	Select the disk groups associated with the selected application instance.
Add, Remove, Reset buttons	Click Add or Remove to display empty fields so that you can manually add or remove additional resources. Click Refresh to repopulate all fields from the current <code>horcm</code> file.

- 7 If you want to configure an additional HTC resource for the instance, click **Add**. Otherwise, click **Next**.

- 8 In the Global Cluster Settings panel, specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site; click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	Enter a virtual IP for the WAC resource.
Subnet Mask	Enter the subnet mask for the system at the primary site and the secondary site.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 9 In the Settings Summary panel, review the displayed information.

If you want to change any of the parameters specified for the replication resource settings or the global cluster settings, click **Back**.

Otherwise, click **Next**.

- 10 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check symbol indicates successful completion of a task. An (x) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.

- 11 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.
- 12 Proceed with configuring additional optional settings for the HTC resource if desired, and then verifying the disaster recovery configuration.

Optional settings for HTC

The wizard configures the required settings for the HTC resource in the VCS application service group.

Optional settings are left in the default state. For information on configuring the optional settings, see *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*

The optional settings use the following defaults

Option	Default setting
LinkMonitor	The default is 0; the agent does not periodically attempt to resynchronize the S-VOL side if the replication link is disconnected. The value 1 indicates that when the replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the pairresync command.
SplitTakeover	The default is 0; the agent does not permit a failover to S-VOL devices if the replication link is disconnected; that is, if P-VOL devices are in the PSUE state.

Configuring global clustering only

If you are using a replication method that the DR wizard does not configure, you must select the replication option to configure global clustering only.

For the GCO only option, you use the wizard to complete all DR tasks except the replication configuration task. You must complete the final wizard task of configuring global clustering before configuring replication.

Before configuring GCO:

- Ensure that you have one static address is available per site for configuring GCO.
- If you created secure clusters at the primary site and secondary site, ensure that you have configured a VCS user with the same name and privileges in each cluster, and the user must be added in the Administrator role.

The following procedure assumes that you have completed the earlier wizard tasks through the service group cloning task and are continuing with the final step of configuring global clustering.

To configure GCO only

- 1 If the wizard is still open after the service group cloning task, continue with the GCO Setup panel.

Otherwise, launch the wizard and proceed to the GCO Setup panel by following these steps sequentially:

- Start the DR Configuration Wizard from the Solutions Configuration Center by clicking **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu.
- Expand the Solutions for Microsoft SQL Server tab.
Click **Disaster Recovery Configuration > Configure Disaster Recovery > Disaster Recovery Configuration Wizard**.
- In the Welcome panel, click **Next** and continue through the wizard, providing the requested information.
- In the Replication Methods panel, click **Configure Global Cluster Option (GCO) only**. Click **Next** and continue to the GCO Setup panel.

- 2 In the GCO Setup panel, review the requirements. If you have met the requirements, click **Next**.

- 3 In the Global Cluster Settings panel specify the heartbeat information for the wide-area connector resource. You must specify this information for the primary and the secondary cluster. Any existing WAC resource information can be reused. If you are adding a DR site to an existing DR configuration, GCO is already configured at the primary site, so the primary site fields are dimmed.

Use existing settings	Allows you to use a WAC resource that already exists at either the primary or secondary site. Click Primary or Secondary, depending on the site at which the WAC resource already exists.
Resource Name	Select the existing WAC resource name from the resource name list box.
Create new settings	Select the appropriate site, primary or secondary, for which you want to create a new WAC resource.
IP Address	For IPv4, enter a virtual IP for the WAC resource. For IPv6, select the network from the dropdown list. An IP address will be generated.
Subnet Mask or Prefix	For IPv4, enter the subnet mask for the system at the primary site and the secondary site. For IPv6, enter the prefix.
Public NIC	Select the public NIC for each system from the drop-down list for the system at the primary and secondary site.
Start GCO after configuration	Select this check box to bring the cluster service group online and start GCO automatically after the wizard completes the necessary configurations. Otherwise, you must bring the service group online and start GCO manually, after the wizard completes. Once GCO is configured and running, deselecting the checkbox does not stop GCO.

- 4 In the Settings Summary panel, review the displayed information.
 If you want to change any of the parameters specified, click **Back**.
 Otherwise, click **Next**.

- 5 In the Implementation panel, wait until the wizard completes creating the replication configuration and the WAC resource required for global clustering. A check symbol indicates successful completion of a task. An (x) symbol indicates a task that could not be completed successfully. The Information column shows details about the reasons for task failure. Click **Next**.
- 6 In the Finish panel, review the displayed information. If a task did not complete successfully, the panel displays an error message, which will provide some insight into the cause for failure. Click **Finish** to exit the wizard.

Creating the replicated data sets (RDS) for Volume Replicator replication

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to configure Volume Replicator replication if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure replication using the Setup Replicated Data Set Wizard.

Configuring Volume Replicator involves setting up the replicated data sets (RDS) on the hosts for the primary and secondary sites. The Setup Replicated Data Set Wizard enables you to configure Replicated Data Sets for both sites.

See [“Setting up the Replicated Data Sets \(RDS\)”](#) on page 177.

Creating the Volume Replicator RVG service group for replication

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

You cannot use the DR wizard to configure Volume Replicator replication if the service group contains VCS resources configured for non-shared storage (VMNSDg agent resources). You must configure the replication service group manually using the Cluster Manager (Java Console).

Complete the following procedures first on the node in the primary site. Then repeat all the steps on the node in the secondary site. You must follow the order of the procedures as mentioned.

Refer to the following topics:

- See [“Configuring a RVG service group for replication”](#) on page 188.

- See [“Creating the RVG service group”](#) on page 189.
- See [“Configuring the IP and NIC resources”](#) on page 191.
- See [“Configuring the VMDg or VMNSDg resources for the disk groups”](#) on page 194.
- See [“Configuring the VMDg or VMNSDg resources for the disk group for the user-defined database”](#) on page 197.
- See [“Adding the Volume Replicator RVG resources for the disk groups”](#) on page 199.
- See [“Linking the Volume Replicator RVG resources to establish dependencies”](#) on page 203.
- See [“Deleting the VMDg or VMNSDg resource from the SQL Server service group”](#) on page 205.
- See [“Configuring the RVG Primary resources”](#) on page 206.
- See [“Creating the RVG Primary resources”](#) on page 206.
- See [“Linking the RVG Primary resources to establish dependencies”](#) on page 208.
- See [“Bringing the RVG Primary resources online”](#) on page 209.
- See [“Setting a dependency between the service groups”](#) on page 209.

Configuring the global cluster option for wide-area failover

This is applicable only if you are setting up disaster recovery in a non-shared storage environment.

The Global Cluster option is required to manage global clustering for wide-area disaster recovery. The process of creating a global cluster environment involves the following tasks:

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

Use the VCS Java Console and perform the following global cluster operations:

- See [“Linking clusters: Adding a remote cluster to a local cluster”](#) on page 279.
- See [“Converting a local service group to a global service group”](#) on page 280.
- See [“Bringing a global service group online”](#) on page 282.

Linking clusters: Adding a remote cluster to a local cluster

This is applicable only if you are setting up DR manually in a non-shared storage environment.

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

Note the following uses of the wizard:

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

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

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

Veritas 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 **Edit > Add/Delete Remote Cluster**.

or

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

- 2 Review the required information for the Remote Cluster Configuration Wizard and then click **Next**.
- 3 In the Wizard Options panel, click **Add Cluster**, then click **Next**.
- 4 In the New Cluster Details panel, enter the details of the new cluster.

If the cluster is not running in secure mode, specify 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 and the password.
- Click **Next**.

If the cluster is running in secure mode, specify 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 by entering `hahb -display` in the command window.

The state attribute in the output should show "alive". If the state is unknown, then take the ClusterService group offline and bring it online again.

Converting a local service group to a global service group

This is applicable only if you are setting up DR manually in a non-shared storage environment.

To convert a local service group to a global group

- 1 From Cluster Explorer, click **Edit > Configure Global Groups**.
 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 3.
- 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, as follows:

Cluster not in
secure mode

Follow these steps sequentially:

- 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

Follow these steps sequentially:

- 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

This is applicable only if you are setting up DR manually in a non-shared storage environment.

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, specify 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**.

Verifying the disaster recovery configuration

The steps you need to take to verify your DR configuration depend on the type of replication you are using.

After the DR wizard has completed, you can confirm the following to verify the DR configuration:

- For Volume Replicator replication, confirm that the configuration of disk groups and volumes at the DR site have been created by the DR wizard storage cloning.
- Confirm that the application VCS service group has been created in the DR cluster including the same service group name, same resources, and same dependency structure as the primary site's application VCS service group.
- Confirm that the application service group is online at the primary site. The application service group should remain offline at the DR site.
- For Volume Replicator replication:
 - Ensure Volume Replicator replication configuration. This includes ensuring that the RVGs have been created at primary and secondary with the correct volume inclusion, replication mode, Replicator Log configuration, and any specified advanced options.
 - Confirm that the replication state matches what was specified during configuration. If specified to start immediately, ensure that it is started. If specified to start later, ensure that it is stopped.
 - Ensure that the Volume Replicator RVG VCS service group is configured on the primary and secondary clusters, including the correct dependency to the application service group, the specified IP for replication, and the correct disk group and RVG objects within the RVG VCS service group.
 - Confirm that the RVG service groups are online at the primary and secondary sites.
 - Confirm that the RVG Primary resources are online in the primary cluster's application service group. If they are offline, then bring them online in the primary site's cluster's application service group. Do not bring them online in the secondary site application service group.
- For array-based replication, verify that the required array resource is created in the primary and secondary cluster's application service group and that a dependency is set between the VMDg resource and the array resource.

- For EMC SRDF replication, verify that the SRDF resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, verify that the HTC resource is online in the primary cluster's application service group. If not, bring it online.
- For Hitachi TrueCopy replication, you must perform a manual Volume Manager rescan on all the secondary nodes after setting up replication and other dependent resources, in order to bring the disk groups online. This must be performed only once, after which the failover works uninterrupted. For more information, see *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*.
- Ensure that the application service groups are configured as global.
- Check to ensure that the two clusters are communicating and that the status of communication between the two clusters has a state of Alive.
- If you are using Volume Replicator for replication and configuring an additional DR site, verify the heartbeat and replication configuration between all sites.
- If you are using Volume Replicator for replication and chose to start replication manually in the DR wizard, to avoid replicating large amounts of data over the network the first time, then you will need to start the process necessary to synchronize from checkpoint.

This process typically consists of the following tasks:

- Starting a Volume Replicator replication checkpoint
 - Performing a block level backup
 - Ending the Volume Replicator replication checkpoint
 - Restoring the block level backup at the DR site
 - Starting replication from the Volume Replicator replication checkpoint
- To learn more about the process of starting replication from a checkpoint, refer to the *Volume Replicator Administrator's Guide*.
- Do not attempt a wide area failover until data has been replicated and the state is consistent and up to date. The Solutions Configuration Center provides a Fire Drill Wizard to test wide area failover for Volume Replicator-based replication.

Establishing secure communication within the global cluster (optional)

A global cluster is created in non-secure mode by default. You may continue to allow the global cluster to run in non-secure mode or choose to establish secure communication between clusters.

The following prerequisites are required for establishing secure communication within a global cluster:

- The clusters within the global cluster must be running in secure mode.
- You must have Administrator privileges for the domain.

The following information is required for adding secure communication to a global cluster:

- The active host name or IP address of each cluster in the global configuration.
- The user name and password of the administrator for each cluster in the configuration.
- If the local clusters do not point to the same root broker, the host name and port address of each root broker.

Adding secure communication involves the following tasks:

- Taking the ClusterService-Proc (wac) resource in the ClusterService group offline on the clusters in the global environment.
- Adding the -secure option to the StartProgram attribute on each node.
- Establishing trust between root brokers if the local clusters do not point to the same root broker.
- Bringing the ClusterService-Proc (wac) resource online on the clusters in the global cluster.

To take the ClusterService-Proc (wac) resource offline on all clusters

- 1 From Cluster Monitor, log on to a cluster in the global cluster.
- 2 In the Service Groups tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the **Process** agent.
- 3 Right-click the **ClusterService-Proc** resource, click **Offline**, and click the appropriate system from the menu.
- 4 Repeat all the previous steps for the additional clusters in the global cluster.

To add the -secure option to the StartProgram resource

- 1 In the Service Groups tab of the Cluster Explorer configuration tree, right-click the **ClusterService-Proc** resource under the **Process** type in the **ClusterService** group.
- 2 Click **View > Properties view**.
- 3 Click the Edit icon to edit the **StartProgram** attribute.
- 4 In the Edit Attribute dialog box, add -secure switch to the path of the executable Scalar Value.

For example:

```
"C:\Program Files\Veritas\Cluster Server\bin\wac.exe" -secure
```

- 5 Repeat the previous step for each system in the cluster.
- 6 Click **OK** to close the Edit Attribute dialog box.
- 7 Click the Save and Close Configuration icon in the tool bar.
- 8 Repeat all the previous steps for each cluster in the global cluster.

To establish trust between root brokers if there is more than one root broker

- ◆ Establishing trust between root brokers is only required if the local clusters do not point to the same root broker.

Log on to the root broker for each cluster and set up trust to the other root brokers in the global cluster.

The complete syntax of the command is:

```
vssat setuptrust --broker host:port --securitylevel [low|medium|high]
[--hashfile fileName | --hash rootHashInHex]
```

For example, to establish trust with a low security level in a global cluster comprised of Cluster1 pointing to RB1 and Cluster2 pointing to RB2 use the following commands:

From RB1, type:

```
vssat setuptrust --broker RB2:14141 --securitylevel low
```

From RB2, type:

```
vssat setuptrust --broker RB1:14141 --securitylevel low
```

To bring the ClusterService-Proc (wac) resource online on all clusters

- 1 In the Service Groups tab of the Cluster Explorer configuration tree, expand the **ClusterService** group and the **Process** agent.
- 2 Right-click the **ClusterService-Proc** resource, click **Online**, and click the appropriate system from the menu.
- 3 Repeat all the previous steps for the additional clusters in the global cluster.

Adding multiple DR sites (optional)

In a Volume Replicator replication environment only, you can use the DR wizard to add additional secondary DR sites. Cluster Server supports up to four DR sites. In other replication environments, additional DR sites require manual configuration.

Run the DR wizard and on the Secondary System selection panel, select the new site.

Before you start the wizard on the task of configuring replication and global clustering, ensure that the cluster service group is online at the existing primary and secondary sites. This enables the wizard to configure GCO not only between the selected primary site and the new secondary site but also between the new site and the earlier configured secondary site. Otherwise, the wizard displays a warning message after the global clustering task.

When configuring the Volume Replicator replication settings with the wizard for the additional site, fields that must match existing settings are dimmed so that you cannot change them. For example, you cannot change the RVG name or RVG layout on the Replication Settings panel. Similarly, on the Global Cluster Settings panel, GCO has already been configured at the primary site, so the primary site fields are dimmed.

Recovery procedures for service group dependencies

Service group dependencies have special requirements and limitations for disaster recovery configuration and for actions to be taken in a disaster recovery scenario.

See [“Supported disaster recovery configurations for service group dependencies”](#) on page 69.

The procedure and requirements for bringing service group dependencies online at the secondary site depends on their configuration: soft, firm, or hard.

In general, if a child or parent remains online at the primary site, you take it offline before you bring the child and parent service groups online in the correct order on the secondary site.

An exception is the RVG service group, used for Volume Replicator replication, which the wizard creates with an online, local, hard dependency. The RVG group remains online at the primary site in all cases and should be left online at the primary site.

The following tables show the recovery requirements if a child or parent service group fails at the primary site and is unable to fail over on the primary site, thus requiring the secondary site to be brought online.

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, soft dependency link.

Table 7-3 Online, local, soft dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	<ul style="list-style-type: none"> ■ The parent remains online on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<ul style="list-style-type: none"> ■ Primary site: Manually take the parent service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the parent and child service groups online in the appropriate order (child first, then parent).
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ul style="list-style-type: none"> ■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, firm dependency link.

Table 7-4 Online, local, firm dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Leave the RVG group online at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ul style="list-style-type: none"> ■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Using a scenario of a parent and one child, the following table shows the expected results and necessary actions you must take for an online, local, hard dependency link.

Table 7-5 Online, local, hard dependency link

Failure condition	Result	Action required (sequentially)
The child service group fails	<ul style="list-style-type: none"> ■ The parent goes offline on the primary site. ■ An alert notification at the secondary site occurs for the child service group only. ■ The RVG group remains online. 	<p>Secondary site: Bring the service groups online in the appropriate order (child first, then parent).</p> <p>Do not take the RVG group offline at the primary site.</p>
The parent service group fails	<ul style="list-style-type: none"> ■ The child remains online on the primary site. ■ An alert notification at the secondary site occurs for the parent only. ■ The RVG group remains online. 	<ul style="list-style-type: none"> ■ Primary site: Manually take the child service group offline at the primary site. Leave the RVG group online. ■ Secondary site: Bring the service groups online in the appropriate order (child first, then parent).

Configuring DR manually without the DR wizard

The Disaster Recovery solution workflow describes how to use the DR wizard from the Solutions Configuration Center to configure disaster recovery. However, if necessary, you can configure disaster recovery manually.

To do so, you first configure the high availability for SQL Server on both sites. You then configure the disaster recovery components: Volume Replicator and Global Cluster Option. You also have the choice of using array-based hardware replication for your disaster recovery solution.

See the *Cluster Server Administrator's Guide for information on configuring the Global Cluster Option*.

See the *Volume Replicator Administrator's Guide for information on configuring Volume Replicator*.

For information on configuring array-based hardware replication, see the VCS hardware agent documentation for the particular array you want to configure.

Testing fault readiness by running a fire drill

This chapter includes the following topics:

- [About disaster recovery fire drills](#)
- [About the Fire Drill Wizard](#)
- [About post-fire drill scripts](#)
- [Tasks for configuring and running fire drills](#)
- [Prerequisites for a fire drill](#)
- [Preparing the fire drill configuration](#)
- [Running a fire drill](#)
- [Re-creating a fire drill configuration that has changed](#)
- [Restoring the fire drill system to a prepared state](#)
- [Deleting the fire drill configuration](#)
- [Considerations for switching over fire drill service groups](#)

About disaster recovery fire drills

A disaster recovery (DR) plan should include regular testing of an environment to ensure that a DR solution is effective and ready if a disaster strikes. This testing is called a fire drill.

A fire drill is performed at the secondary site using a special service group for fire drills. The fire drill service group uses a copy of the data that is used by the application service group.

About the Fire Drill Wizard

Storage Foundation and High Availability Solutions (SFW HA) provides a Fire Drill Wizard to help you set up and run a fire drill on a disaster recovery environment. You launch the Fire Drill Wizard from the Solutions Configuration Center.

The Fire Drill Wizard tests the fault readiness of a disaster recovery configuration by mimicking a failover from the primary site to the secondary site. The wizard does this without stopping the application at the primary site and disrupting user access.

The Fire Drill Wizard supports conducting a fire drill for a disaster recovery site that uses Volume Replicator (Volume Replicator) or that uses Hitachi TrueCopy or EMC SRDF hardware replication.

In the Hitachi TrueCopy or EMC SRDF environments, the Fire Drill Wizard supports only the Gold configuration. For the Silver or Bronze configuration, you must manage (create, restore, delete) the fire drill configurations and run the fire drills manually. For further information about the Gold, Silver, and Bronze configurations, refer to the following documents:

Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide

Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide

Note: After upgrading to 6.0.1 or later, the existing fire drill service groups will not be usable. In a Hitachi TrueCopy or EMC SRDF environment, you must manually edit the existing fire drill service groups. In a Volume Replicator environment, you must use the Fire Drill Wizard to re-create them. For more information, see the *Veritas InfoScale Installation and Upgrade Guide*.

About Fire Drill Wizard general operations

The Fire Drill Wizard performs the following operations:

- Prepares for the fire drill by creating a fire drill service group on the secondary site
The fire drill service group is a copy of the application service group. When creating the fire drill service group, the wizard uses the application service group name, with the suffix `_fd`. The wizard renames the fire drill service group

resources by adding a prefix FDnn and changes attribute values as necessary to refer to the FD resources.

The wizard also supports fire drill service groups created under a different naming convention by an earlier version of the wizard.

- Runs the fire drill by bringing the fire drill service group online on the secondary site

This operation demonstrates the ability of the application service group to failover and come online at the secondary site should the need arise.

Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online on the primary site.

- Restores the fire drill configuration, taking the fire drill service group offline
After you complete a fire drill, run the wizard to restore the fire drill configuration to a prepared state. Otherwise, the fire drill service group remains online.
If you run a fire drill on one service group, restore that service group before you continue with a fire drill on another service group.
You must also restore the fire drill configuration before you can delete it.

Warning: If the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting. Therefore, after completing the fire drill testing for a service group, always use the wizard to restore the fire drill configuration to a prepared state as soon as possible.

See [“Restoring the fire drill system to a prepared state”](#) on page 317.

- Deletes the fire drill configuration

The details of some Fire Drill Wizard operations are different depending on the replication environment.

See [“About Fire Drill Wizard operations in a Volume Replicator environment”](#) on page 293.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 296.

About Fire Drill Wizard operations in a Volume Replicator environment

The general operations of the Fire Drill Wizard are the same in all replication environments.

- Prepares for the fire drill, creating a fire drill service group on the secondary site

- Runs the fire drill, bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration, taking the fire drill service group offline
- Deletes the fire drill configuration

See [“About the Fire Drill Wizard”](#) on page 292.

However, the following additional Fire Drill Wizard operations are specific to a Volume Replicator environment.

Preparing the fire drill configuration

In a Volume Replicator environment, when preparing the fire drill configuration, the wizard does the following:

- Replaces the RVGPrimary resources with VVRSnap resources in the fire drill service group
- Uses the SFW HA VxSnap feature to prepare snapshot mirrors for use during the fire drill
While running the wizard, you assign one or more disks for the mirrored volumes. Mirror preparation can take some time, so you can exit the wizard after this step is started and let the preparation continue in the background.
- Sets the `offline-local-firm` dependency between the service groups, where the fire drill service group is the parent and the application service group is the child
- Configures the VVRSnap resource by setting the following attributes to the appropriate values:
 - RVG
 - AppDiskGroupName
 - DiskGroupName
- Sets the FireDrill attribute of the following resources to true:
 - IP
 - Lanman
 - RegRep
 - SQLServer
- Sets the ForFireDrill attribute of the following resources to `true` in the fire drill service group:
 - MountV

- VMDg

This indicates that the volume being monitored by the VVRSnap agent belongs to the fire drill disk group.

About running the fire drill

The Fire Drill Wizard brings the fire drill service group online. Optionally, you can also run the fire drill using the Veritas InfoScale Operations Manager console.

In a Volume Replicator environment, when running the fire drill, the VVRSnap agent does the following:

- Detaches the mirrors from the original volumes to create point-in-time snapshots of the production data
- Creates a fire drill disk group on the secondary site with a snapshot of the application data to use for testing purposes

About restoring the fire drill configuration

The Fire Drill Wizard takes the fire drill service group offline. Optionally, you can also restore the fire drill using the Veritas InfoScale Operations Manager console.

In a Volume Replicator environment, restoring the fire drill system to a prepared state, the VVRSnap agent does the following:

- Joins the fire drill disk group to the application service group disk group
- Snaps back the snapshot mirrors to reattach to the original volumes

About deleting the fire drill configuration

In a Volume Replicator environment, when deleting the fire drill configuration, the wizard does the following:

- Sets the FireDrill attribute of the following resources to false:
 - IP
 - Lanman
 - RegRep
 - SQLServer
- Unlinks the fire drill service group
- Deletes the fire drill service group and any associated registry entry
- Performs the snap abort operation on the snapshot mirrors to free up the disk space

About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment

The Fire Drill Wizard performs the following basic operations in all replication environments:

- Prepares for the fire drill by creating a fire drill service group on the secondary site
- Runs the fire drill by bringing the fire drill service group online on the secondary site
- Restores the fire drill configuration by taking the fire drill service group offline
- Deletes the fire drill service group and any associated registry entries

See [“About the Fire Drill Wizard”](#) on page 292.

In Hitachi TrueCopy or EMC SRDF replication environments, the Fire Drill Wizard performs the following additional actions during preparation, running of the fire drill, restoring the configuration, and deleting the configuration. You must configure the ShadowImage (for Hitachi) or BCV (for SRDF) pairs before running the wizard.

About preparing the fire drill configuration

When preparing the fire drill configuration, the wizard does the following:

- In the fire drill service group, the wizard creates HTCSnap or SRDFSnap resources for each HTC and SRDF resource in the application service group. It links the fire drill service group to the corresponding application service group.
- In an HTC or SRDF environment, the wizard configures the Snap resource and sets the following attributes to the value 1, which indicates:
 - **UseSnapshot**: Take a local snapshot of the target array.
 - **RequireSnapshot**: Require a successful snapshot for the Snap resource to come online.
 - **MountSnapshot**: Use the snapshot to bring the fire drill service group online.
- In an EMC SRDF environment, the wizard sets the following attribute values:
 - It sets **CopyMode** to one of the following, which indicates:
 - **Mirror**: Use the TimeFinder Mirror technology to create snapshots.
 - **Clone**: Use the TimeFinder Clone technology to create snapshots.
 - **Snap**: Use the TimeFinder Snap technology to create snapshots.
 - When the TimeFinder Clone technology is used, it sets **UseTgt** to one of the following, which indicates:

- **0:** Use BCV devices to create snapshots.
- **1:** Use STD devices to create snapshots.
- When the TimeFinder Snap technology is used, if a custom save pool area name is specified, it sets **SavePoolName** accordingly. The specified save pool area is used to create snapshots.
If no value is specified on the SRDFSnap Resource Configuration panel, the default save pool area is used.

For information about the actual procedure:

See [“Preparing the fire drill configuration”](#) on page 304.

About running the fire drill

When running the fire drill, the wizard brings the HTCSnap or SRDFSnap agent online. The HTCSnap or SRDFSnap agent manage the replication and mirroring functionality according to the attribute settings. The Snap agents take a consistent snapshot of the replicating data using the snapshot or mirroring technology provided by the array vendor. The Snap agents also import the disk group present on the snapshot devices with a different name.

In more detail, the Snap agent does the following:

- Suspends replication to get a consistent snapshot
- For HTCSnap, takes a snapshot of the replicating application data on a ShadowImage device
- For SRDFSnap, takes a snapshot of the replicating application data on a BCV, STD, or VDEV device
- Resumes replication
- Modifies the disk group name in the snapshot

For information about the actual procedure:

See [“Running a fire drill”](#) on page 312.

About restoring the fire drill configuration

When restoring the fire drill configuration to a prepared state, the wizard takes the fire drill service group offline, thereby taking the SRDF and HTC Snap agents offline.

This action reattaches the hardware mirrors to the replicating secondary devices and resynchronizes them.

For information about the actual procedure:

See [“Restoring the fire drill system to a prepared state”](#) on page 317.

About deleting the fire drill configuration

When deleting the fire drill configuration, the wizard does the following:

- Delinks the fire drill service group from the corresponding application service group.
- Deletes the fire drill service group
- Deletes any associated registry entries

If you want to remove the hardware mirrors, you must do so manually.

For information about the actual procedure:

See [“Deleting the fire drill configuration”](#) on page 318.

For more information about the Hitachi TrueCopy Snap agent functions, see *Cluster Server Hardware Replication Agent for Hitachi TrueCopy Configuration Guide*.

For more information about the EMC SRDF Snap agent functions, see *Cluster Server Hardware Replication Agent for EMC SRDF Configuration Guide*.

About post-fire drill scripts

You can specify a script for the Fire Drill Wizard to run on the secondary site at the end of the fire drill.

For example, in a SQL Server environment, if you create and populate a test table at the primary site, you could create a script to verify the replication of the data at the secondary site.

For the wizard to run the script, the script must exist on the secondary system that you are specifying for the fire drill.

Note: The wizard does not support using script commands to launch a user interface window. In such a case, the process is created but the UI window does not display.

Optionally, you can specify to run a Windows PowerShell cmdlet by creating a `.bat` file.

To run a cmdlet

- 1 Create a .bat file with the following entry:

```
%windir%\system32\WindowsPowerShell\v1.0\  
PowerShell.exe -command "$ScriptName"
```

In this entry, `$ScriptName` is either the fully qualified .ps1 script, or the cmdlet entered by user.

For example:

```
D:\WINDOWS\system32\WindowsPowerShell\v1.0\  
PowerShell.exe -command C:\myTest.ps1
```

- 2 Specify the name of the .bat file as the script to run.

Tasks for configuring and running fire drills

While running the Fire Drill Wizard, the following sequence of actions are available:

- Prepare the fire drill configuration
- Run the fire drill or delete the configuration
- Restore the fire drill configuration after running a fire drill
- Run another fire drill or delete the configuration

In addition, you have the option to re-create a fire drill configuration that has changed.

After an action is complete, the next action becomes available in the wizard. You can select the next action or exit the wizard and perform the next action later.

The following table gives more details of the process of configuring and running fire drills with the wizard.

Table 8-1 Tasks for configuring and running fire drills

Action	Description
Verify the hardware and software prerequisites	Before running the wizard, review the prerequisites and make sure that they are met. See "Prerequisites for a fire drill" on page 301.
Prepare the fire drill configuration	Use the wizard to configure the fire drill. See "Preparing the fire drill configuration" on page 304.

Table 8-1 Tasks for configuring and running fire drills (*continued*)

Action	Description
Re-create a fire drill configuration that has changed	<p>If a fire drill configuration exists for the selected service group, the wizard checks for differences between the fire drill service group and the application service group. If differences are found, the wizard can re-create the fire drill configuration before running the fire drill.</p> <p>See “Re-creating a fire drill configuration that has changed” on page 314.</p>
Run the fire drill	<p>Use the wizard to run the fire drill. Running the fire drill brings the fire drill service group online. Optionally you can specify a script to be run once the fire drill is complete.</p> <p>See “Running a fire drill” on page 312.</p> <p>Perform your own tests of the application to confirm that it is operational.</p> <p>Note: After completing the fire drill testing, run the wizard again as soon as possible to restore the configuration. Otherwise the fire drill service groups remain online. It is recommended that you restore a fire drill service group to a prepared state before running a fire drill on another service group.</p>
Restore the fire drill configuration to a prepared state	<p>Use the wizard to restore the fire drill system to a state of readiness for future fire drills or to prepare for removal of the fire drill configuration.</p> <p>This is a required action after running the fire drill.</p> <p>See “Restoring the fire drill system to a prepared state” on page 317.</p> <p>This operation takes the fire drill service group offline and reattaches snapshot mirrors.</p>

Table 8-1 Tasks for configuring and running fire drills (continued)

Action	Description
Delete the fire drill configuration	<p>If a fire drill service group is no longer needed, or if you want to free up resources, use the wizard to remove the fire drill configuration.</p> <p>See “Deleting the fire drill configuration” on page 318.</p> <p>The wizard deletes the service group on the secondary site. In a Volume Replicator environment, the wizard performs a snap abort to delete the snapshot mirrors created on the secondary site for use in the fire drill. In hardware replication environments, you can delete these manually.</p> <p>If a fire drill has been run, the wizard ensures that you first restore the fire drill configuration to a prepared state before this option becomes available. This ensures that mirrors are reattached and the fire drill service group is offline before the configuration is deleted.</p>

Prerequisites for a fire drill

Before running the Fire Drill Wizard make sure that you meet the following general requirements:

- You can run the Fire Drill Wizard from any node in the domain of the cluster, as long as the SFW HA client is installed on that node.
- If the cluster is secured, the login you use to run the Fire Drill Wizard must have the appropriate permissions to make changes in the cluster.
- If a firewall exists between the wizard and any systems it needs access to, the firewall must be set to allow both ingoing and outgoing TCP requests on port 7419.
- For an application service group using IPv4 addresses, for each IP address in the service group, an IP address must be available to use on the secondary site for the fire drill service group.
 For IPv6, the IP address will be autogenerated.
 To configure IPv6 settings, the wizard must be launched from a system on which the IPv6 stack is installed. The wizard can accept input for one IP address and Lanman resource. If the application service group has multiple IP addresses and Lanman resources, the wizard notifies you to edit the fire drill service group resources to supply these values. More information on editing service group resources is available.
 See the *Cluster Server Administrator’s Guide*.

- If you want the fire drill wizard to run a script that you supply, ensure that the script file is available on any secondary site nodes where you plan to run the fire drill.

For testing purposes, you may want to create and populate a new table from the active node at the primary site. After you run the fire drill to bring the fire drill service group online and create the fire drill snapshots, you can check that the table and its data were replicated and are available from the fire drill service group. You can automate this process with a script and when preparing to run the fire drill, specify it as a post-fire drill script.

Additional requirements apply to specific replication environments.

See [“Prerequisites for a fire drill in a Volume Replicator environment”](#) on page 302.

See [“Prerequisites for a fire drill in a Hitachi TrueCopy environment”](#) on page 303.

See [“Prerequisites for a fire drill in an EMC SRDF environment”](#) on page 303.

Prerequisites for a fire drill in a Volume Replicator environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 301.

Make sure that the following additional prerequisites are met before configuring and running a fire drill in a Volume Replicator environment:

- The primary and secondary sites must be fully configured with Volume Replicator replication and the global cluster option.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- On the secondary site, empty disks must be available with enough disk space to create snapshot mirrors of the volumes. Snapshot mirrors take up the same amount of space as the original volumes. In addition, two disk change object (DCO) volumes are created for each snapshot mirror, one for the source volume and one for the snapshot volume. The two DCO volumes must be on different disks. Allow 2 MB additional space for each DCO volume.

The empty disks must be in the same disk group that contains the RVG. If the disk group does not have empty disks available, you must use the VEA to add the disks to the disk group before you run the wizard. The secondary system must have access to the disks or LUNs.

- All disk groups in the service group must be configured for replication. The Fire Drill wizard does not support a Volume Replicator configuration in which disk

groups are excluded from replication. However, you can exclude individual volumes within a disk group from replication.

Prerequisites for a fire drill in a Hitachi TrueCopy environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 301.

Make sure that the following prerequisites are met before configuring and running a fire drill in a Hitachi TrueCopy environment:

- The primary and secondary sites must be fully configured with Hitachi TrueCopy replication and the global cluster option. Make sure that you have configured disaster recovery with Hitachi TrueCopy.
- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- Make sure that Hitachi RAID Manager/Command Control Interface (CCI) is installed.
- ShadowImage for TrueCopy must be installed and configured for each LUN on the secondary site target array. ShadowImage pairs must be created to allow for mirroring at the secondary site.
- The name of the ShadowImage device group must be the same as the replicated device group for both replicated and non-replicated LUNs that are to be snapshot. The instance number should be different.
- Make sure the HORCM instance managing the S-VOLs runs continuously; the agent does not start this instance.

Prerequisites for a fire drill in an EMC SRDF environment

Before you run the Fire Drill Wizard make sure that you meet both the general requirements and the specific requirements for your replication environment.

General requirements are covered separately.

See [“Prerequisites for a fire drill”](#) on page 301.

Make sure that the following prerequisites are met before configuring and running a fire drill in an EMC SRDF environment:

- The primary and secondary sites must be fully configured with EMC SRDF replication and the global cluster option. Make sure that you have configured disaster recovery with EMC SRDF.

- The secondary system where you plan to run the fire drill must have access to the replicated volumes.
- To take snapshots of R2 devices, appropriate additional devices must be associated with the RDF2 device group and fully established with the devices.
- The infrastructure to take snapshots at the secondary site must be properly configured between the secondary site source and target arrays. Depending on the snapshot technology in use, this process involves the following tasks:
 - Mirror: Associate Symmetric Business Continuance Volumes (BCVs) and synchronize them with the secondary site source (STD devices).
 - Clone: Make sure that no clone session is in progress.
The source and target devices must be of the exact same size.
 - Snap: Make sure that sufficient save pool area is configured and that no snap session is in progress.
The source and target devices must be of the exact same size.
- If you plan to run a fire drill on SRDF/A devices, you must have a TimeFinder/CG license. Make sure TimeFinder for SRDF is installed and configured at the target array.
- To take snapshots of non-replicated devices, create a EMC Symmetrix device group with the same name as the SFW disk group. The device group must contain the same devices as in the disk group and have the corresponding BCV, STD, or VDEV devices associated.

Preparing the fire drill configuration

Preparing the fire drill configuration creates a fire drill service group. You specify the application service group and the secondary system to use. Only one service group can be prepared for a fire drill at one time.

For a Volume Replicator environment, the preparation step also prepares snapshot mirrors of production data at the specified node on the secondary site.

Note: Preparing the snapshot mirrors takes some time to complete.

Before you prepare the fire drill configuration with the Fire Drill Wizard, make sure that you meet the prerequisites.

See [“Prerequisites for a fire drill”](#) on page 301.

To prepare the fire drill configuration

- 1 Open the Solutions Configuration Center (From **Start > All Programs > Veritas > Veritas Cluster Server > Solutions Configuration Center** or, on Windows 2012 operating systems, from the **Apps** menu).
- 2 Start the Fire Drill Wizard (expand Solutions for Microsoft SQL, expand Fire Drill, expand Configure or run a fire drill, and click Fire Drill Wizard).
- 3 In the Welcome panel, review the information and click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
See [“System Selection panel details”](#) on page 307.
- 5 In the Service Group Selection panel, select the service group that you want to use for the fire drill and click **Next**.
See [“Service Group Selection panel details”](#) on page 308.
- 6 In the Secondary System Selection panel, select the cluster and the system to be used for the fire drill at the secondary site, and then click **Next**.
See [“Secondary System Selection panel details”](#) on page 308.

- 7 If the Fire Drill Prerequisites panel is displayed, review the information and ensure that all prerequisites are met. Click **Next**.

See [“Prerequisites for a fire drill”](#) on page 301.

Otherwise, if a fire drill service group already exists on this system for the specified service group, one of the following panels is displayed:

If the Run Fire Drill option or Delete Fire Drill options are shown, a fire drill service group has already been prepared.	You can run the fire drill with no further preparation. Click Run Fire Drill and follow the procedure for running a fire drill. See “Running a fire drill” on page 312.
--	--

If the Fire Drill Restoration panel is displayed, the fire drill service group remains online from a previous fire drill.	Follow the procedure for restoring the fire drill configuration to a prepared state. This must be done before running a new fire drill. See “Restoring the fire drill system to a prepared state” on page 317.
---	---

If the Re-create Fire Drill Service Group panel is displayed, a fire drill service group has already been prepared but is not up to date.	You can choose to re-create the fire drill configuration to bring it up to date. See “Re-creating a fire drill configuration that has changed” on page 314. Or you can clear the check box to re-create the configuration and run the fire drill on the existing configuration.
---	---

- 8 If the Fire Drill Service Group Settings panel is displayed, assign the virtual IP address and virtual name (Lanman name) to be used for the fire drill service group that will be created on the secondary site.

See [“Fire Drill Service Group Settings panel details”](#) on page 308.

- 9 The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

Volume Replicator replication	Disk Selection panel See “Disk Selection panel details” on page 308.
Hitachi TrueCopy replication	Horcm Files Path Selection panel See “Hitachi TrueCopy Path Information panel details” on page 309. HTCSnap Resource Configuration panel See “HTCSnap Resource Configuration panel details” on page 310.
EMC SRDF replication	SRDFSnap Resource Configuration panel See “SRDFSnap Resource Configuration panel details” on page 310.

Click **Next**.

- 10 In the Fire Drill Preparation panel, the wizard shows the status of the preparation tasks.

See [“Fire Drill Preparation panel details”](#) on page 311.

When preparation is complete, click Next.

- 11 The Summary panel displays the message that preparation is complete.

To run the fire drill now, click **Next**. Continue with the procedure to run the fire drill.

See [“Running a fire drill”](#) on page 312.

To run the fire drill later, click **Finish**. The fire drill preparation remains in place.

System Selection panel details

Use the System Selection panel of the wizard to specify a system in the primary site cluster.

All systems containing online global service groups are available to select. The default system is the node where you launched the wizard (localhost) if a global service group is online on that system. When selecting a system you can specify either a fully qualified host name or IP address.

See [“Preparing the fire drill configuration”](#) on page 304.

Service Group Selection panel details

Use the Service Group Selection panel of the wizard to select the service group that you want to use for the fire drill. You can select only one service group at a time for a fire drill.

See [“Preparing the fire drill configuration”](#) on page 304.

Secondary System Selection panel details

Use the Secondary System Selection panel of the wizard to select the cluster and the system to be used for the fire drill at the secondary site.

The selected system must have access to the replicated data.

The system must have access to disks for the snapshots that will be created for the fire drill.

See [“Preparing the fire drill configuration”](#) on page 304.

Fire Drill Service Group Settings panel details

Use the Fire Drill Service Group Settings panel of the wizard to assign the virtual IP address and virtual name (Lanman name) to be used for the fire drill service group that will be created on the secondary site. These must be an address and name not currently in use.

For IPv4, you must manually assign the IP address. For IPv6, the IP address will be autogenerated and displayed in the Virtual IP address field.

If the service group contains more than one IP and Lanman resource, this panel does not appear. After the fire drill service group is created, the wizard notifies you to manually update the IP and Lanman resources in the fire drill service group.

See [“Preparing the fire drill configuration”](#) on page 304.

Disk Selection panel details

During fire drill preparation in a Volume Replicator replication environment, you must ensure that information is available to the wizard for creating the fire drill snapshots. Use the Disk Selection panel of the wizard to review the information on disks and volumes and make the selections for the fire drill snapshots, as follows:

Volume	<p>Select the volumes for the fire drill snapshots. By default all volumes associated with the service group are selected. If you deselect a volume that might result in the fire drill service group failing to come online, the wizard displays a warning message.</p> <p>Note: The Disk Selection panel also appears if the wizard is re-creating a fire drill service group to which volumes have been added. In that case, only the new volumes are shown for selection.</p>
Disk Group	<p>Shows the name of the disk group that contains the original volumes. This field is display only.</p>
Fire Drill DG	<p>Shows the name of the fire drill disk group that running the fire drill will create on the secondary system to contain the snapshots. This field is display only. For the fire drill disk group name, the wizard prefixes the original disk group name with FDnn.</p>
Disk	<p>Click the plus icon to the right of the Disk column and specify the disk to be used for the snapshot volume. Repeat for each row that contains a selected volume.</p> <p>If the production volumes reside on disks in the same disk group, you can store multiple snapshot volumes on a single disk. If the volumes in a disk group are configured on multiple RVG resources, provide a separate disk for each RVG.</p> <p>Note: The Fire Drill Wizard does not allow creating mirrors of multiple RVGs from a single disk group on the same disk. You must select a different disk for each RVG in a disk group.</p> <p>If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard.</p>
Mount Details	<p>Shows the mount details for the snapshot volumes on the secondary system, which match the mounts for the production volumes. This is a display-only field.</p>

Hitachi TrueCopy Path Information panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the Hitachi TrueCopy Path Information panel is displayed.

The wizard populates the path field with the customary default location, `C:\Windows`, where `c` is the system drive.

If the `horcm` configuration files are in a different location, edit the field to specify that location.

HTCSnap Resource Configuration panel details

During fire drill preparation in a Hitachi TrueCopy replication environment, the wizard discovers the HTC resources and non-replicating SFW disk groups in the application service group.

This information is used to configure the HTCSnap resources.

The wizard lists each HTCSnap resource that will be configured. You can clear the HTCSnap resource name check box if you do not want to include its dependent disk groups in the fire drill.

You must specify the ShadowImage instance.

The HTCSnap Resource Configuration panel shows the following:

Target Resource Name	The panel shows the HTC resource name in the case of a Replication Device Group or the disk group resource name in the case of a non-replicating disk group.
ShadowImage Instance ID	For every HTC resource, specify the ID of the ShadowImage instance associated with the replicating secondary devices.

See [“Preparing the fire drill configuration”](#) on page 304.

More information about HTCSnap resource configuration and operation is available.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 296.

SRDFSnap Resource Configuration panel details

Depending on the snapshot technology in use, the wizard validates the following when preparing for fire drill in an EMC SRDF replication environment:

- Mirror:
 - The number of BCV devices matches that of the STD devices.
 - The BCV devices are associated and synchronized with the STD devices.
- Clone:
 - The number of BCV devices (or STD devices in case of Targets) matches that of the STD devices.
 - No clone session is in progress.
- Snap:
 - The number of VDEV devices matches that of the STD devices.
 - No snap session is in progress.

If these criteria are not satisfied, the wizard displays a warning on this panel. The wizard does not check whether the sizes of the source and target devices match, and therefore does not display a warning. The following figure depicts such a warning.

However, you can proceed with the configuration. The wizards configures the fire drill service group, but is unable to bring the service group online.

This panel lists all the SRDFSnap resources that will be configured. If you do not want to include the dependent disk groups of a SRDFSnap resource in the fire drill, clear the check box against its name.

The name of the resource that is managing the LUNs that you want to snapshot appears as the Target Resource Name. For data being replicated from the primary site, the Target Resource Name is the name of the SRDF resource. For data that is not replicated, the Target Resource Name is the name of the disk group resource.

For example, in a typical SQL Server setup, you might replicate data files and logs, but you may choose to avoid replicating temporary tempdb. The tempdb must still exist at the DR site and may be part of its own disk group.

You can specify the TimeFinder snapshot technology to be used for configuring fire drill for the SRDFSnap resources:

- **Mirror**
BCV devices are used to create snapshots.
- **Clone**
BCV devices are used to create snapshots. Optionally, you can specify that Target devices be used. If you select the **Use Target Devices** check box, STD devices are used to create snapshots.
- **Snap**
VDEV devices are used to create snapshots. The default SavePoolArea is used. Optionally, to use a different SavePoolArea, specify its name.

To discover the most recent SRDF configuration information, click **Refresh**.

See [“Preparing the fire drill configuration”](#) on page 304.

More information about SRDFSnap resource configuration and operation is available.

See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 296.

Fire Drill Preparation panel details

After you enter the information required to prepare a fire drill configuration, the Fire Drill Preparation panel is displayed. You wait while the wizard completes the preparation tasks.

The fire drill service group is created on the secondary site (but remains offline).

In addition, for a Volume Replicator replication environment, the snapshot mirrors for the volumes are prepared; this can take some time. You may want to minimize the wizard while the task runs in the background. You can also track the mirror preparation progress in the VEA. When done, the wizard displays a message that the fire drill preparation is complete. If the wizard is completing the preparation steps as part of re-creating a fire drill configuration, snapshot mirrors are prepared only for new volumes.

See [“Re-creating a fire drill configuration that has changed”](#) on page 314.

See [“Preparing the fire drill configuration”](#) on page 304.

Running a fire drill

After you complete the initial fire drill preparation step using the Fire Drill Wizard, you can run the fire drill immediately without exiting the wizard or run the wizard later to run the fire drill.

Running the fire drill does the following:

- Creates the snapshots
- Enables the firedrill resources
- Brings the fire drill service group online
- Optionally, executes a specified command to run a script
See [“About post-fire drill scripts”](#) on page 298.

For details on the operations that occur when running a fire drill, refer to the following topics:

- See [“About Fire Drill Wizard operations in a Volume Replicator environment”](#) on page 293.
- See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 296.

Warning: After running the fire drill, the fire drill service group remains online. After you verify the fire drill results, run the wizard again to restore the system to the prepared state. Otherwise, if the fire drill service group remains online, it could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

See [“Restoring the fire drill system to a prepared state”](#) on page 317.

To run a fire drill

- 1 If you completed the initial preparation and have not exited the wizard, or if you are returning to this procedure after re-creating a fire drill service group, go to step 6.

Otherwise, if you need to restart the wizard, continue with the next step.

- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand Solutions for Microsoft SQL, expand Fire Drill, expand Configure or run a fire drill, and click Fire Drill Wizard).

- 3 In the Welcome panel, click **Next**.

- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.

- 5 In the Service Group Selection panel, select the service group and click **Next**.

- 6 In the Secondary System Selection panel, specify the system previously prepared for the fire drill at the secondary site and click **Next**.

If the fire drill configuration is in a prepared state, the wizard compares the resources of the fire drill service group with the resources of the application service group.

- 7 If the application service group changed since the fire drill configuration was prepared, the wizard displays the Re-create Fire Drill Service Group panel, showing the differences.

Choose one of the following:

- Leave the option checked to re-create the configuration before running the fire drill and click **Next**. You complete additional steps in the wizard before running the fire drill.

See [“Re-creating a fire drill configuration that has changed”](#) on page 314.

- To run the fire drill on the existing configuration, clear the option to re-create the fire drill service group and click **Next**.

- 8 In the Fire Drill Mode Selection panel, click Run Fire Drill and click **Next**.

- 9 In the Post Fire Drill Script panel, optionally specify the full path to a script for the wizard to run on the secondary system right after running the fire drill. The script must already exist on the secondary system. Click **Next**.

See [“About post-fire drill scripts”](#) on page 298.

- 10 In the Fire Drill Implementation screen, wait until all fire drill tasks are performed and click **Next**.

The Summary panel displays the message that the fire drill is complete. You can leave the wizard running while you verify the results or exit the wizard.

To exit the wizard, click **Finish**.

- 11 Run your own tests to verify the fire drill results.

Warning: You should always restore the fire drill system to a prepared state immediately after completing fire drill testing on a service group.

- 12 Restore the fire drill configuration to the prepared state.

See [“Restoring the fire drill system to a prepared state”](#) on page 317.

Re-creating a fire drill configuration that has changed

When you run the Fire Drill wizard, a fire drill service group may already exist for the selected application service group. However, the application service group may have changed since the fire drill service group was created. The wizard compares the resource names of the two service groups. If differences are found, the wizard lists them on the Re-create Fire Drill Service Group panel.

The wizard also checks the RVGs configured for disk groups. If a single RVG is configured per disk, the wizard allows you to re-create the service group; the existing snapshots are retained. If multiple RVGs are configured on a disk, the wizard only allows you to delete the service group; the existing snapshots are deleted. To create a corresponding new one, you need to launch the wizard again and perform the fire drill preparation steps.

Note: The wizard does not check for changes in volume attributes, such as the MountPath attribute. For example, if you have a MountV resource with an attribute that points to drive Y and you change that attribute to point to drive X, the wizard does not identify this change and does not give the option to re-create the fire drill service group.

You can choose whether to manually edit the fire drill service group for such changes and then run the fire drill, ignore the differences, or delete the configuration and start over.

You have the following choices from the Re-create Fire Drill Service Group panel:

- Leave the option checked to re-create the fire drill service group. Proceed with using the wizard to re-create the configuration to match the application service group. The wizard deletes the existing fire drill configuration first, before creating the new one.

For a Volume Replicator replication environment, the wizard handles existing volumes as follows: It does not delete the mirrors for volumes that still exist. When it re-creates the fire drill configuration, it prepares new mirrors only for new volumes. If volumes have been removed, the wizard displays an additional option to snap abort the obsolete snapshot volumes to free up disk space.

- Clear the option to re-create the fire drill service group. You can then proceed with using the wizard to do either of the following:
 - Run the fire drill, ignoring the differences.
 - Delete the entire fire drill configuration. Then start over with preparing the fire drill configuration.

The following procedure describes the choice of re-creating the fire drill configuration.

Note: Veritas recommends that you do not use this procedure to re-create any existing fire drill service groups after performing an upgrade. Instead, use the Fire Drill Wizard to delete the existing service groups and create corresponding new ones.

To re-create the fire drill configuration if the service group has changed

- 1 In the Re-create Fire Drill Service Group panel, leave the option checked to re-create the configuration before running the fire drill.

For a Volume Replicator replication environment, if volumes have been removed, optionally select to snap abort the volumes.

Click **Next**.

- 2 In the Fire Drill Mode Selection panel, Delete Fire Drill Configuration is selected. Click **Next**, and click **Yes** to confirm the deletion.

- 3 The Fire Drill Deletion panel shows the progress of the deletion.

For a Volume Replicator replication environment, the wizard leaves the existing fire drill snapshot volumes so that those snapshot mirrors do not have to be prepared again. If volumes were removed and you selected the option to snap abort, the wizard snap aborts the snapshots of those volumes.

Warning: If you close the wizard after deleting the fire drill configuration without continuing on to the fire drill preparation step, the information of the existing snapshot volumes is lost.

When all tasks are complete, click **Next**.

- 4 In the Fire Drill Prerequisites panel, review the information and ensure that all prerequisites are met. Click **Next**.

See [“Prerequisites for a fire drill”](#) on page 301.

- 5 The wizard selects the appropriate panel to display next, depending on the replication method. Fill in any required information on the panel that is displayed.

Volume Replicator replication	If volumes have been added, the Disk Selection panel is displayed.
	Specify the information for the added volumes.
	If there is not enough disk space, you can use the Veritas Enterprise Administrator to add disks to the disk group. Then click the Refresh button in the wizard.
	See “Disk Selection panel details” on page 308.
Hitachi TrueCopy replication	Horcm Files Path Selection panel
	See “Hitachi TrueCopy Path Information panel details” on page 309.
	HTCSnap Resource Configuration panel
	See “HTCSnap Resource Configuration panel details” on page 310.
EMC SRDF replication	SRDFSnap Resource Configuration panel
	See “SRDFSnap Resource Configuration panel details” on page 310.

Click **Next**.

- 6 The Fire Drill Preparation panel is displayed. Wait while the wizard re-creates the fire drill service group.

For Volume Replicator replication environments, wait while the wizard starts mirror preparation.

Mirror creation can take some time. You may want to minimize the wizard while the task runs in the background. You can also close the wizard and track the mirror preparation progress in the VEA.

- 7 Once preparation is complete, click **Next**. The Summary page is displayed. To continue with running the fire drill, click **Next**.

See [“Running a fire drill”](#) on page 312.

Restoring the fire drill system to a prepared state

After running a fire drill and verifying the results, use the Fire Drill Wizard as soon as possible to restore the fire drill system at the secondary site to a prepared state. A prepared state is the initial fire drill configuration created by the wizard, in which the fire drill service group has been prepared but is offline.

Restoring the fire drill system to a prepared state is required for any of the following:

- Making the secondary system available for failover of the application service group at the primary site
- Running another fire drill
- Deleting the fire drill configuration after a fire drill has been run

For details on the operations that occur when restoring a fire drill configuration, see the following topics:

- See [“About Fire Drill Wizard operations in a Volume Replicator environment”](#) on page 293.
- See [“About Fire Drill Wizard operations in a Hitachi TrueCopy or EMC SRDF environment”](#) on page 296.

To restore the fire drill system to a prepared state

- 1 If you completed running a fire drill and have not exited the wizard, go to step 6.
Otherwise, continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand Solutions for Microsoft SQL, expand Fire Drill, expand Configure or run a fire drill, and click Fire Drill Wizard).
- 3 In the Welcome panel, click **Next**.

- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.
The default system is the node where you launched the wizard.
- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 In the Fire Drill Restoration Information panel, review the requirements for restoration and click **Next**.
- 8 In the Fire Drill Restoration screen, wait until the screen shows the restoration tasks are completed and click **Next**.
- 9 In the Summary screen, click **Next** if you want to delete the fire drill configuration. Otherwise click **Finish** to exit the wizard, leaving the fire drill configuration in a prepared state.

Deleting the fire drill configuration

If you no longer need a fire drill configuration you can delete it. Deleting a fire drill configuration deletes the fire drill service group on the secondary site.

In a Volume Replicator replication environment, deleting a fire drill configuration also performs a snap abort of the snapshot mirrors created on the secondary site for use in the fire drill. It frees up the disk space used for the snapshot mirrors for other use.

In a Hitachi TrueCopy or EMC SRDF environment, you could manually remove mirrors after the deletion is complete.

To delete a fire drill configuration

- 1 If you have just used the wizard to prepare or restore a fire drill configuration and have not exited the wizard, go to step 8.
Otherwise continue with the next step.
- 2 From the Solutions Configuration Center, start the Fire Drill Wizard (expand Solutions for Microsoft SQL, expand Fire Drill, expand Configure or run a fire drill, and click Fire Drill Wizard).
- 3 In the Welcome panel, click **Next**.
- 4 In the System Selection panel, specify a system in the primary site cluster and click **Next**.

The default system is the node where you launched the wizard.

- 5 In the Service Group Selection panel, select the service group that was used for the fire drill and click **Next**.
- 6 In the Secondary System Selection panel, specify the system on which the fire drill was run at the secondary site.
- 7 If the wizard detects that the fire drill service group is different from the application service group, it displays the Re-create Fire Drill Service Group panel. Clear the option to re-create the fire drill service group and click **Next**.
- 8 If the wizard detects that the fire drill service group is still online, the Fire Drill Restoration panel is displayed. Review the requirements for restoration and click **Next**.
- 9 In the Restore Fire Drill screen, wait until the screen shows the restoration tasks are completed. Then click **Next**.
- 10 In the Fire Drill Mode Selection panel, click Delete Fire Drill Configuration and click **Next**, and click Yes to confirm the deletion.
- 11 The Fire Drill Deletion panel shows the progress of the deletion. Wait until all tasks are complete and then click **Next**.

If errors occur while deleting the fire drill configuration, the wizard will list any incomplete steps so that you can complete them manually.
- 12 The Summary panel is displayed. Click **Finish**.

Fire Drill Deletion panel details

The Fire Drill Deletion panel shows the progress of the deletion tasks. Wait until all tasks are complete before you click **Next** to proceed to the Summary panel and exit the wizard.

See [“Deleting the fire drill configuration”](#) on page 318.

The Fire Drill Deletion panel also appears if you selected the option to re-create a fire drill configuration. During a re-create operation, wait until all deletion tasks are complete and then click **Next** to continue with the fire drill preparation.

When the wizard re-creates a configuration in a Volume Replicator environment, it deletes the service group but maintains the required snapshot volumes. If you close the wizard without continuing to the preparation step, the snapshot volume information is lost and the fire drill configuration is fully deleted.

Considerations for switching over fire drill service groups

In a Volume Replicator environment, if you directly switch the fire drill service group from one node to another, the VVRSnap resource fails to come online on the target node. The fire drill service group depends on the RVG service group. To make the switch successfully, you must first switch the RVG service group to the intended node and then switch the fire drill service group.

Configuring SQL Server in a VMware environment

- [Chapter 9. Configuring application monitoring using the Veritas High Availability solution](#)
- [Chapter 10. Administering application monitoring](#)

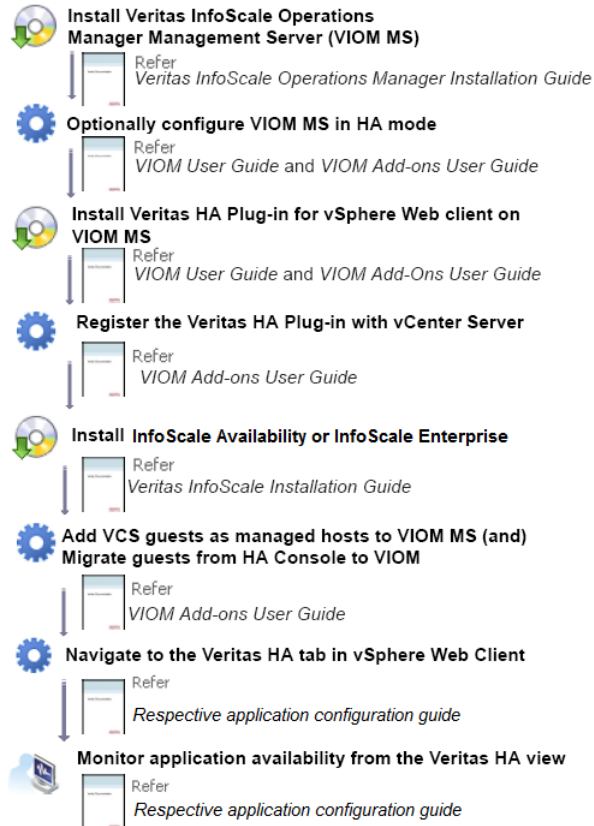
Configuring application monitoring using the Veritas High Availability solution

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 326.
- 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 327.
- Integrate with the existing authentication mechanism
 See [“Integrating with Active Directory or local authentication”](#) on page 328.
 - Add the ESX user to the created role
 See [“Adding a user to the role”](#) on page 329.

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 328.

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 329.

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 330.

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 333.

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.

- 6 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 IP version 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 version and then specify the following details for each failover system:

Note: You must select the same IP version 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 338.

- 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.
 For details on encrypting agent passwords, see the *Cluster Server Administrator's Guide*.
- 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.
 For details on encrypting agent passwords, see the *Cluster Server Administrator's Guide*.
- 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 10-1 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)7.4.1 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 356.

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>

Accessing the dashboard

You can use the Veritas High Availability dashboard to perform one of the following actions:

- Identify all instances and failover systems of one or more applications running in a data center
- Drill down to a specific application, and perform an administrative action on the application
- View alerts for faulted applications and stale application health reports

Prerequisites for accessing the dashboard

Before you access the Veritas High Availability dashboard to administer an application, ensure:

- Single sign-on is configured between the Veritas High Availability Console and the systems hosting the monitored applications
- Veritas High Availability Console is able to communicate with Veritas High Availability guest components on designated port (port 5634).
- The application that you want to administer is configured for application monitoring with Veritas High Availability

How to access the dashboard

When you install Veritas High Availability guest components, the product installation script or wizard automatically installs the required dashboard components. As a result, the Symantec High Availability dashboard appears in the **Veritas High Availability** tab of the vSphere Client.

You must, however, ensure that Veritas High Availability is successfully installed and that you have adequate user privileges to access the dashboard.

To access dashboard

Perform the following step:

- In the inventory view (left pane) of the vSphere Client, click a data center or a VMware cluster. In the right pane, to view the Veritas High Availability dashboard, click the Symantec High Availability tab.

Who can access the dashboard

To access High Availability dashboard, the VMware vCenter administrator must assign one the following roles to you:

- Guest: View application state
- Operator: View application state and perform administrative actions
- Admin: View application state and perform administrative actions. You can also configure application availability monitoring in this role, but not from the dashboard.

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.
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 361.
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 362.
- 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 361.
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 362.
- 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 361.
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 362.
- 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 361.
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 362.
- 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 361.
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 362.
- 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.

Using Veritas AppProtect for vSphere

This appendix includes the following topics:

- [About Just In Time Availability](#)
- [Prerequisites](#)
- [Setting up a plan](#)
- [Deleting a plan](#)
- [Managing a plan](#)
- [Viewing the history tab](#)
- [Limitations of Just In Time Availability](#)
- [Getting started with Just In Time Availability](#)
- [Supported operating systems and configurations](#)
- [Viewing the properties](#)
- [Log files](#)
- [Plan states](#)
- [Troubleshooting Just In Time Availability](#)

About Just In Time Availability

The Just In Time Availability solution provides increased availability to the applications on a single node InfoScale Availability cluster in VMware virtual environments.

Using the Just In Time Availability solution, you can create plans for:

- Planned Maintenance
- Unplanned Recovery

Planned Maintenance

In the event of planned maintenance, the Just In Time Availability solution enables you to clone a virtual machine, bring it online, and fail over the applications running on that virtual machine to the clone on the same ESX host. After the maintenance procedure is complete, you can fail back the applications to the original virtual machine. Besides failover and failback operations, you can delete a virtual machine clone, view the properties of the virtual machine and its clone, and so on.

Unplanned Recovery

When an application encounters an unexpected or unplanned failure on the original or primary virtual machine on the primary ESX host, the Just In Time Availability solution enables you to recover the application and bring it online using the unplanned recovery feature.

With **Unplanned Recovery Policies**, the Just In Time Availability solution enables you to set up recovery policies to mitigate unplanned failures that are encountered by an application. Just In Time Availability solution provides the following recovery policies; you may select one or all the recovery policies as per your need:

Unplanned Recovery Policies	Description
-----------------------------	-------------

Restart Application	
---------------------	--

	Just In Time Availability (JIT) solution attempts to restart the service group (SG), and bring the application online on the original virtual machine on primary ESX.
--	---

	Maximum three retry attempts are permitted under this policy.
--	---

	Note: If all the three attempts fail, application continues to remain in faulted state or continues with the next policy as selected while creating a plan.
--	--

Unplanned Recovery Policies	Description
Restart virtual machine (VM)	<p>Just In Time Availability (JIT) solution performs the following subsequent tasks:</p> <ul style="list-style-type: none">■ take the service group offline■ shut down the virtual machine■ power on the virtual machine■ bring the service group online on the original virtual machine on primary ESX <p>You are provided with Last attempt will be VM reset option to reset the virtual machine.</p> <p>By default, this checkbox is selected and the default retry attempt value is one. If you retain the default settings, then VM reset operation is performed on the virtual machine at the first attempt itself.</p> <p>Maximum three retry attempts are permitted for this operation.</p> <p>If you deselect the checkbox, then the virtual machine reset (VM Reset) operation is not performed.</p>
Restart VM on target ESX	<p>Using this policy, you can recover the faulted application on the virtual machine.</p> <p>In this policy, the original virtual machine is unregistered from the primary ESX; registered on the target ESX; and the faulted application is brought online on the target ESX.</p>
Restore VM on target ESX	<p>Using this policy, you can recover the faulted application on the virtual machine using a boot disk backup copy of the original virtual machine.</p> <p>In this policy, the original virtual machine is unregistered from the ESX and the boot disk backup copy of the original virtual machine is registered on target ESX. The faulted application is then brought online on the virtual machine.</p>

Unplanned Recovery Policies

Unplanned Failback

The **Unplanned Failback** operation lets you fail back the application from the boot disk backup copy of virtual machine on the target ESX to the original virtual machine on primary ESX.

If you have selected either **Restart VM on target ESX** or **Restore VM on target ESX** or both the recovery policies, you can perform the **Unplanned Failback** operation.

On the **Plans** tab, in the plans table list, right-click the virtual machine and click **Unplanned Failback**.

Note: **Unplanned Failback operation** operation is disabled and not available for the plans and the virtual machines which have **Restart Application** and **Restart VM** policies as the only selected options.

Based on the selected recovery policy for a plan, Just In Time Availability (JIT) solution performs the necessary operations in the sequential order.

For example, if you have selected **Restart Application** and **Restart VM** as the recovery policy, then in the event of unplanned application failure, first it performs tasks for **Restart Application** policy and if that fails, it moves to the next policy.

You may select one or all the recovery policies based on your requirement.

[Table A-1](#) lists the sequence of tasks that are performed for each Unplanned Recovery policy.

Table A-1 Tasks performed for each Unplanned Recovery policy

Unplanned Recovery Policy	Tasks Performed
Restart Application	◆ Make an attempt to restart the application.
Restart virtual machine (VM)	1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Power on the virtual machine 4 Brings the service group(s) online

Table A-1 Tasks performed for each Unplanned Recovery policy (*continued*)

Unplanned Recovery Policy	Tasks Performed
Restart VM on target ESX	<ol style="list-style-type: none"> 1 Takes the service group(s) offline 2 Shuts down the original virtual machine 3 Detaches the data disks from the original virtual machine 4 Unregisters the virtual machine from the primary ESX 5 Registers the original virtual machine on target ESX 6 Attaches the data disks back to the virtual machine 7 Power on the virtual machine 8 Brings the service group(s) online
Restore VM on target ESX	<ol style="list-style-type: none"> 1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Detaches the data disks from the virtual machine 4 Unregisters the original virtual machine from the target ESX 5 Registers the boot disk backup copy of the original virtual machine to the target ESX 6 Attaches the data disks back to the virtual machine 7 Power on the virtual machine 8 Brings the service group(s) online
Unplanned Failback	<ol style="list-style-type: none"> 1 Takes the service group(s) offline 2 Shuts down the virtual machine 3 Detaches the data disks from the virtual machine 4 Unregisters the virtual machine from the target ESX 5 Registers the virtual machine using the original boot disk backup copy to the primary ESX 6 Attaches the data disks to the virtual machine 7 Power on the virtual machine on primary ESX 8 Brings the service group(s) online on the virtual machine

Scheduler Settings

While creating a plan for unplanned recovery, with **Scheduler Settings**, you can set up a schedule for taking a back up of boot disk of all the virtual machines that are a part of the plan.

To use the Just In Time Availability solution, go to **vSphere Web Client > Home view > Veritas AppProtect**.

See [“Setting up a plan”](#) on page 373.

Prerequisites

Before getting started with Just In Time Availability, ensure that the following prerequisites are met:

- The Just In Time (JIT) solution feature cannot co-exist with VMware HA, VMware FT, and VMware DRS. This pre-requisite is applicable for **Unplanned Recovery** only.
- VIOM 7.2 version must be installed and configured using fully qualified domain name (FQDN) or IP.
- Make sure that you have the admin privileges for vCenter.
- VMware Tools must be installed and running on the guest virtual machine.
- VIOM Control Host add-on must be installed on VIOM server or machine.
- The virtual machines must be added in VIOM. The virtual machines, vSphere ESX servers, and VIOM must have the same Network Time Protocol (NTP) server configured.
- Make sure to specify VIOM Central Server FQDN or IP in the SNMP Settings of the vCenter Server.
- vCenter Server and VIOM must be configured using the same FQDN or IP address. Make sure that if FQDN is used to configure vCenter in VIOM Server that is used during the configuration.
- If raw disk mapping (RDM) disks are added to the virtual machine, then make sure that the virtual machine is in the physical compatibility mode. Veritas AppProtect does not support the virtual compatibility mode for RDM disks.
- For Microsoft Windows operating system, make sure that you have the Microsoft Windows product license key. The key is required to run the Sysprep utility, which enables customization of the Windows operating system for a clone operation.

- For RHEL7 and SUSE12 operating system, install the deployPkg plug-in file on the virtual machine.
For more information on installing the plug-in, see <https://kb.vmware.com/kb/2075048>
- Make sure that the InfoScale Availability service group is configured with one of the storage agents such as Mount, DiskGroup, LVMVolumeGroup, VMNSDg (for Windows), or DiskRes (for Windows), for the data disks. This configuration enables Veritas AppProtect to discover data disks for the applications. Also, ensure that the service group is online to determine data disk mapping.
- Virtual machines which have snapshots associated with them are not supported.
- Virtual machines with SCSI Bus Sharing are not supported.
- Make sure that the SNMP Traps are configured for the following from vCenter server to VIOM:
 - Registered virtual machine
 - Reconfigured virtual machine
 - Virtual machine which is getting cloned
- Make sure that the boot disk of VM's (vmdk) does not have spaces.
- For HA console add on upgrade from VIOM 7.1 to VIOM 7.2, refer *Veritas InfoScale Operations Manager 7.2 Add-ons User's Guide* for more details.
- Make sure to set the vSphere DRS Automation Level to manual, if you want to configure **Restart VM on target ESX** or **Restore VM on target ESX** policies for your plan.
- Ensure to update or edit the plan, when a virtual machine is migrated or if there are any modifications made to the settings of the virtual machines which are configured for that plan.
- Ensure to increase the tolerance limit of DiskRes resource to two, if you want to create a plan for unplanned recovery with **Restore VM on target ESX** as the unplanned recovery policy.

Note: This prerequisite is applicable for Windows operating system.

Setting up a plan

Plan is a template which involves a logical grouping of virtual machines so as to increase the availability of the application in the event of a planned failover and recovery of the application in the event of an unexpected application failure.

To set up a plan

- 1 Launch Veritas AppProtect from the **VMware vSphere Web Client > Home view > Veritas AppProtect** icon.
- 2 Click **Configure Plan**.
The **Plan Configuration** wizard appears.
- 3 Specify a unique **Plan Name** and **Description**, and then click **Next**.
The wizard validates the system details to ensure that all prerequisites are met.
- 4 Select the virtual machines that you want to include in the plan, review the host and operating system details, and then click **Next**.
The **Unplanned Recovery Settings** page appears.
- 5 On the **Unplanned Recovery Settings** page, you can configure the selected virtual machines for **Unplanned Recovery** as well.
Deselect the **Configure selected VMs for Unplanned Recovery as well** check box, if you do not want to include the selected virtual machines for unplanned recovery.

If you have selected the virtual machines for unplanned recovery, then set up the unplanned recovery policies as appropriate from the available options. You can set up policies to restart applications, restart virtual machines, restart virtual machine on target ESX, and restore a virtual machine on target ESX.

If you have selected **Restore VM on target ESX** as the unplanned recovery policy, then you can set up a schedule to create a boot disk back up copy of the virtual machine within the configured plan. You can set the frequency as daily, weekly, monthly, or manual as per your requirement.

After you have finished making necessary settings for Unplanned Recovery, Click **Next**.
- 6 The wizard validates the prerequisite attributes of the virtual machine and the ESX host, and adds the qualified virtual machines to the plan.
Click **Next** after the validation process completes.

- 7 In the **Disks** tab, you can view the selected application data disks. Just In Time Availability solution uses the selected data disks to perform detach-attach operation during a planned failover and unplanned recovery.

Note: If the disks are not auto-marked as selected to perform detach-attach operation, then first refresh the VIOM server and then the VCenter server in VIOM and then create a plan.

- 8 In the **Network Configuration** tab, specify the network interface configuration details for the cloned virtual machine. Make sure to specify at least one public interface and valid IP details.
- 9 In the **Unplanned Recovery Target** tab, specify the target ESX server to restore the virtual machine, and the target ESX port details.

Note: The **Unplanned Recovery Target** tab is visible only when **Restart VM on target ESX** or **Restore VM on target ESX** is selected.

- 10 In the **Windows Settings** tab, specify the domain name, Microsoft Windows product license key, domain user name, domain password, admin password, and time zone index.

Note: The **Windows Settings** tab is visible only when a Windows virtual machine is selected in the plan.

- 11 Click **Next**. The **Summary** wizard appears.
- 12 In the **Summary** wizard, review the plan details such as the plan name, unplanned recovery policies, schedule, and so on.

Deselect the **Start backup process on finish** checkbox if you do not want to initiate a backup process when the plan creation procedure is finished. This checkbox is selected by default.

Click **Create**. The plan is created and saved.

- 13 Click **Finish** to return to the plans tab and view the created plans.

See [“Managing a plan”](#) on page 375.

See [“Deleting a plan”](#) on page 375.

Deleting a plan

After you have finished performing failback operations from the clone to the primary virtual machine in case of planned maintenance and recovery operations in case of unplanned recovery, you may want to delete the plan.

To delete a plan

- 1 Launch **Veritas AppProtect** from the VMware vSphere Web Client Home view.
- 2 In the **Plans** tab, select the plan that you want to delete.
- 3 Click **Delete Plan**.

Note: The **Delete plan** icon is enabled only when the selected plan is in **Ready For Failover**, **Failed to Revert**, and **Failed to Failback** state.

Managing a plan

Planned Maintenance

After the maintenance plan is created, you can fail over the applications to the clone virtual machine and fail back the applications from the clone to the virtual machine. When the scheduled maintenance is complete, you can delete the cloned virtual machine or retain it for future use.

To perform failover, failback, revert, or delete clone operations, go to **Plans**, and select a plan. Based on the enabled operation, perform the following tasks:

To fail over the applications to the cloned virtual machine

- ◆ Click the **Failover** icon.

Just In Time Availability (JIT) performs the sequence of failover tasks, which includes taking the application offline, detaching the disks, cloning the virtual machine, attaching the disks, and so on.

To fail back the applications from the clone to the primary virtual machine

- ◆ Click the **Failback** icon.

Just In Time Availability (JIT) performs the sequence of failback tasks, which includes taking the application offline, detaching the disks, attaching the disks, and so on.

To revert a failover or a fallback operation

- ◆ Click the **Revert** icon.

If the failover or a fallback operation fails, the revert operation restores the applications on the virtual machine, and deletes the clone if created.

To delete a clone

- ◆ Click the **Delete Clone** icon.

After the fallback operation is complete, you can delete the clone. By default, the revert operation deletes the clone.

Note: Alternatively, right-click **Plan** in the **Plans** table on the **Plans** wizard to perform failover, fallback, revert, delete plan, and delete clone operations.

Unplanned Recovery

Once you have set up a plan for unplanned recovery during **Configure Plan** operation, based on the recovery policies selected for the plan, the application is recovered accordingly.

You can manage unplanned recovery policies settings by performing the following operations on the plan and its associated virtual machines.

Managing unplanned recovery settings

On the **Plans** tab, in the plans table which lists all the existing plans, navigate to the required plan and use the right-click option on the selected plan.

- **Edit:** Use this option to modify the configured plans settings such as adding or removing a virtual machine from the plan, and so on.
The same **Configuration Plan** wizard using which you had set up or configured a plan is displayed with pre-populated details.
See [“Setting up a plan”](#) on page 373.
- **Disable Unplanned Recovery:** Use this option to disable the Unplanned Recovery settings.
- **Enable Unplanned Recovery:** Use this option to enable the Unplanned Recovery settings.
- **Disable Scheduler:** Use this option to disable the scheduler settings.
- **Enable Scheduler:** Use this option to enable the scheduler settings.
- **Delete Plan:** Use this option to delete the created plan.
- **Properties:** Use this option to view the properties for unplanned recovery. It displays details such as the selected unplanned recovery policies and the

associated operations for the selected policies. It also provides information about the selected scheduler mode for performing boot disk back up operation for the selected virtual machines.

Managing virtual machines settings

On the **Plans** tab, in the plans table which lists all the existing plans and its associated virtual machines, navigate to the required virtual machine. Select the required virtual machine and use the right-click option on the selected virtual machine.

- **Remove VM From Plan:** Use this option to delete the virtual machine from the selected plan.
- **Create Clone Backup:** Use this option to create a boot disk backup copy of the virtual machine.
- **Unplanned Failback:** Use this option to fail back the application from the boot disk backup copy of the virtual machine on target ESX to the original virtual machine on primary ESX.

Note: This option is available only if you have set unplanned recovery policies as **Restart VM on target ESX** or **Restore VM on target ESX**.

- **Properties:** Use this option to view properties such as the last run time for backup operation, last successful backup attempt time and the target ESX details.

See [“Plan states”](#) on page 382.

Viewing the history tab

On the **History** tab, you can view the detailed summary of the operations that are performed on the virtual machine. The details include the plan name, virtual machine name, operation, the status of the operation, the start and the end time of the operation, and the description of the operation status.

To view the summary

- 1 Launch **Veritas AppProtect** from the VMware vSphere Web Client Home view.
- 2 Click the **History** tab.

Limitations of Just In Time Availability

The following limitations are applicable to Just In Time Availability:

- On a single ESX host only ten concurrent failover operations are supported. Across ESX hosts, twenty concurrent failover operations are supported.
- Linked mode vCenter is not supported.
- Only three backup operations per data store are active, the rest will be queued. Only five backup operations per ESX host are active, the rest will be queued.

See [“Supported operating systems and configurations”](#) on page 380.

Getting started with Just In Time Availability

You can access the Just In Time Availability solution from the **vSphere Web Client > Veritas AppProtect** interface.

The **Veritas AppProtect** is registered with Veritas InfoScale Operations Manager (VIOM), and is accessed from the **vSphere Web Client > Home** view.

[Figure A-1](#) describes the Veritas AppProtect interface in detail.

Figure A-1 Elements of the Veritas AppProtect interface




Table A-2 Elements of the Veritas AppProtect interface and the description

Label	Element	Description
1	Plans tab	<p>Enables setting up a plan for a planned failover and unplanned recovery.</p> <p>Displays the plan attributes, and the virtual machines that are added to the plan.</p> <p>Displays the status of virtual machines for unplanned recovery and schedule for virtual machine back up operation based on the criteria set while configuring or editing the plan.</p> <p>Shows the enabled or disabled failover, failback, delete clone, revert, delete plan, and properties operations icons based on the state of the selected plan for planned failover.</p>
2	History tab	Displays the status and the start and the end time of the specific operation performed on the created plans.
3	Configure Plan link	Opens the Plan Configuration wizard.
4	Plans table	Displays the attributes of the plan.
5	Failover icon	Fails over the applications from the original virtual machine to the clone.
6	Failback icon	Fails back the applications from the clone to the original virtual machine.
7	Delete Clone icon	Deletes the cloned virtual machine.
8	Revert State icon	Reverts the failed operation, restores the applications to the original virtual machines, and delete the clone virtual machines.
9	Delete Plan icon	Deletes the plan.
10	Properties icon	Displays the attributes of each virtual machine and the clone.

Table A-2 Elements of the Veritas AppProtect interface and the description
(continued)

Label	Element	Description
11	Operation-specific tabs	<p>Displays the sequence of the tasks that are performed for the selected operation.</p> <p>Based on the operation that is executed, the associate tab opens.</p> <p>For Planned Maintenance</p> <ul style="list-style-type: none"> 1 Failover 2 Failback 3 Revert 4 Delete Clone <p>For Unplanned Recovery</p> <ul style="list-style-type: none"> ◆ Unplanned Recovery Summary
12	Diagnostic information	Displays the logs that are reported for the Veritas AppProtect interface.

See “[Plan states](#)” on page 382.

Supported operating systems and configurations

Just In Time Availability supports the following operating systems:

- On Windows: Windows 2012, and Windows 2012 R2.
- On Linux: RHEL5.5, RHEL6, RHEL7, SUSE11, SUSE12.

Just In Time Availability supports the following configurations:

- Veritas Cluster Server (VCS) 6.0 or later, or InfoScale Availability 7.1 and later.
- Veritas InfoScale Operations Manager managed host (VRTSsfmh) 7.1 and 7.2 version on the virtual machines.
For more information about VRTSsfmh, see the *Veritas InfoScale Operations Manager 7.2 User Guide*.
- Veritas InfoScale Operations Manager (VIOM) 7.2 as a central or managed server.
- VMware vSphere 5.5 Update 2, Update 3, or 6.0 and 6.0 Update 1 version.

Viewing the properties

Virtual Machine Properties

The **Virtual Machine Properties** window displays information about the virtual machine and its clone such as name, operating system, cluster name, service groups, DNS server, domain, IP addresses, and data disks.

To view the properties

- 1 On the **Plans** tab, select the virtual machine.
- 2 Click the **Properties** icon or right-click the virtual machine.

The **Virtual Machine Properties** window opens and displays the attributes of the virtual machine and its clone.

Plan Properties

The **Plan Properties** window displays information about the unplanned recovery policies selected; scheduler mode set; and the time when the last backup operation was run and was successful for a virtual machine.

To view properties for the plan

- 1 In the Plan Name table, select the plan.
- 2 Right-click the selected plan. A window with a list of options is displayed.
- 3 Click **Properties**

The **Plan Properties** window opens and displays the unplanned recovery policies selected and the schedule mode for virtual machine backup operation.

Log files

The following log files are helpful for resolving the issues that you may encounter while using Veritas AppProtect:

- Console related logs:

```
/var/opt/VRTSsfmcs/logs/*
```

These log files show console messages and are useful for debugging console issues.

- Operations logs:

```
/var/opt/VRTSsfmh/logs/vm_operations.log
```

This log file shows the messages pertinent to the Veritas AppProtect interface.

- VMware vSphere 6.0 logs:

C:\ProgramData\VMware\vCenterServer\logs\vsphere-client\logs*

These log files show the messages that are reported for the VMware vSphere Web Client version 6.0.

- VMware vSphere 5.5 U2 and U3 logs:

C:\ProgramData\VMware\vsphere Web Client\serviceability\logs*

These log files show the messages that are reported for the VMware vSphere Web Client version 5.5 U2 and U3.

- Veritas AppProtect interface logs:

The log file shows the logs that are reported for the Veritas AppProtect interface. To view the log files, on the **Planned Maintenance** tab or the **History** tab > **Diagnostic Information**.

Plan states

Based on the state of the plan, the operation icons are enabled and disabled on the **Plans** tab.

Table A-3 List of plan and operation states

Plan state	Failover	Failback	Revert	Delete clone	Delete Plan	Unplanned Failback	Create Clone backup	Properties
Ready For Failover	✓	–	–	✓ Note: Enabled when the selected maintenance plan has an associate clone.	✓ Note: Enabled when the selected maintenance plan does not have an associate clone.	–	✓	✓
Failed Over	–	✓	–	–	–	–	–	✓
Failed To Failover	–	–	✓	–	–	–	–	✓

Table A-3 List of plan and operation states (*continued*)

Plan state	Failover	Failback	Revert	Delete clone	Delete Plan	Unplanned Failback	Create Clone backup	Properties
Failed To Failback	–	–	✓	–	–	–	–	✓
Failed To Revert	–	–	✓	–	✓	–	–	✓
Unknown	–	–	✓	–	–	✓	–	✓
Failed To Delete Clone	–	–	–	✓	–	–	–	✓
Failover In Progress	–	–	–	–	–	–	–	✓
Failback In Progress	–	–	–	–	–	–	–	✓
Revert In Progress	–	–	–	–	–	–	–	✓
Delete Clone In Progress	–	–	–	–	–	–	–	✓
Application Faulted	–	–	–	–	–	–	–	✓
Failed To Restart VM	–	–	–	–	–	–	–	✓
Failed To Move VM	–	–	–	–	–	✓	–	✓
Failed To Restore VM	–	–	–	–	–	✓	–	✓
Unplanned	–	–	–	–	–	✓	✓	–
Unplanned Restored VM	–	–	–	–	–	✓	–	✓
Unplanned Failed to Failback	–	–	–	–	✓	–	–	–

Troubleshooting Just In Time Availability

Table A-4 lists the issues and the recommended solutions.

Table A-4 Issues and the corresponding resolutions

Issue	Recommended Solution
When setting up a maintenance plan, the registered virtual machine is not listed on the wizard.	To troubleshoot the issue, make sure of the following: <ul style="list-style-type: none">■ ESX host on which the virtual machine resides, is connected to the vCenter.■ The virtual machine is added as a managed host to Management Server.■ On the virtual machine, at least one application is configured for monitoring, along with VCS.■ The virtual machine is registered in VIOM.■ VCS is configured on the virtual machine.■ The virtual machine runs only a supported Windows Server version.■ VCS is configured with the service groups.
When setting up a maintenance plan, the listed virtual machine is not available for selection.	To troubleshoot the issue, make sure the following: <ul style="list-style-type: none">■ The virtual machine is not configured for Global Cluster option (GCO).■ Agents that support SAN are configured.
When Veritas AppProtect executes an operation, the timeout message is reported.	To troubleshoot the issue, perform the following: <ul style="list-style-type: none">■ If the failover or the failback operation fails, then click Planned Maintenance > Revert icon. Retry the operation.■ If the delete plan or the delete clone operation fails, then retry the operation.
The revert operation failed.	Manually revert the virtual machine to its original state.