

Veritas™ Cluster Server Implementation Guide

ESX

5.1 Maintenance Pack 2



Veritas Cluster Server Implementation Guide

Copyright © 2008 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, and Veritas are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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

The Licensed Software and Documentation are deemed to be "commercial computer software" and "commercial computer software documentation" as defined in FAR Sections 12.212 and DFARS Section 227.7202.

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

Third-party legal notices

All third-party copyrights associated with this product are listed in the Third Party Copyrights document, which is included on the product disc.

Technical support

For technical assistance, visit:

http://www.symantec.com/business/support/assistance_care.jsp.

Select phone or email support. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service.

Contents

Section I Installing VCS for VMware ESX

Chapter 1 Introducing VCS for VMware ESX

Features	18
About VCS	19
Multiple nodes	20
Virtual machines and applications	20
Shared storage	20
LLT and GAB	21
Network channels for heartbeats	21
Service groups	21
About high availability using VCS for VMware ESX	22
About disaster recovery using VCS for VMware ESX	22
Replicated storage	23
Global clusters	23
Installation road map	24

Chapter 2 Requirements

VMware ESX Server software and infrastructure	28
Patches	29
Patch for VCS agents hangs and when agents return UNKNOWN states	29
Patch for ESX Server 3.0.1 freezing during rescan operations	29
Patch for ESX 3.0.x systems running VCS 5.1	29
Bind utilities update	29
Supported operating systems	30
Supported operating systems in virtual machines for high availability	30
Supported operating systems in virtual machines for application monitoring or disaster recovery	30
Supported operating systems for increasing allocated storage	31
Supported applications	32
Support for detecting intentional offline for specific applications	33
Supported hardware	34
Veritas Cluster Server hardware requirements	34

Veritas Cluster Server required disk space	35
VMware ESX components and configuration requirements	35

Chapter 3 Planning to install VCS on an ESX Server

About installing VCS on an ESX Server	38
About optional VCS features	38
Veritas Cluster Server Management Console	38
Cluster Manager (Java Console)	38
Notifications	39
Global clusters using VCS	39
Performing pre-installation tasks	39
Setting environment variables	41
Setting up the private network	41
Selecting virtual interfaces for use by the private network	42
Enabling communication between systems	43
Optimizing LLT media speed settings on private NICs	43
Setting the media speed of the LLT interconnects	43
Enabling password free SSH-communication between systems	44
Obtaining VCS license keys	46
Accessing documentation	47
Preparing your VCS installation and configuration information	47
Optional VCS RPMs	50
Service console firewall ports that installvcs opens during installation	50

Chapter 4 Installing and configuring VCS on ESX Servers

About installing and configuring VCS	52
Installing and configuring VCS 5.1 MP2	53
Pre-patching servers running ESX Server 3.5	53
VCS installation flow overview	53
VCS installation example	54
Checking the systems for installation	54
Starting the software installation	55
Specifying systems for installation	56
Licensing VCS	56
Choosing VCS RPMs	57
Installing VCS 5.1 MP2	58
Configuring VCS 5.1 MP2	59
Adding VCS users	60
Configuring cluster connector	60
Configuring Veritas Cluster Server Management Console	61

Configuring SMTP email notification	62
Configuring SNMP trap notification	63
Configuring global clusters	65
ESX Server network configuration	66
Starting VCS	66
Completing the installation	67
Enforcing compatibility between VCS and VMware features (DRS and ESX Server maintenance mode)	67
Copying the Implementation Guide to each node	68
Installing the Cluster Manager (Java Console)	68
Verifying the cluster after installation	69
Optional installvcs program actions	69
Installing VCS using installonly option	70
Configuring VCS using configure option	70
Performing VCS installation in a secure environment	70
Performing automated installations	72
Checking licensing information on the system	78
Updating product licenses using vxlicinst	79
Replacing a VCS demo license with a permanent license	79
About installvcs command options	80
About the uninstallvcs program	82
Prerequisites	82
Uninstalling VCS	83
Removing VCS RPMs	83
Running uninstallvcs from the disc	84
Uninstalling the Cluster Management Console cluster connector	84

Chapter 5

Verifying VCS on ESX Servers

About verifying the VCS installation	88
Verifying LLT and GAB configuration files	88
/etc/llthosts	88
/etc/llttab	88
/etc/gabtab	89
Verifying the main.cf file	89
Example main.cf, for clusters without the GCO option	90
Example main.cf, for clusters with the GCO option	97
Verifying LLT, GAB, and cluster operation	97
Verifying LLT	97
Verifying GAB	99
Verifying the cluster	100

Chapter 6 Adding and removing cluster nodes

About adding and removing nodes	104
Adding nodes	104
Removing nodes	104
Adding a node to a cluster	104
Setting up the hardware	105
Preparing for a manual installation	106
Installing VCS RPMs for a manual installation	106
Adding a license key	107
Configuring LLT and GAB	107
Adding the node to the existing cluster	109
Starting VCS and verifying the cluster	110
Adding the node to the VMware cluster	110
Removing a node from a cluster	110
Verify the status of nodes and service groups	111
Deleting the node from the VMware cluster	112
Deleting the departing node from VCS configuration	112
Modifying configuration files on each remaining node	114
Unloading LLT and GAB and removing VCS on the departing node	114

Chapter 7 Installing VCS on a single node

About installing VCS on a single node	118
Creating a single-node cluster using the <code>installvcs</code> program	118
Preparing for a single node installation	118
Starting the <code>installvcs</code> program for the single node cluster	119
Creating a single-node cluster manually	119
Setting the <code>PATH</code> variable	120
Installing VCS RPMs for a manual installation	121
Adding a license key	122
Renaming the LLT and GAB startup files	122
Modifying the startup files	122
Configuring VCS	123
Verifying single-node operation	124
Adding a node to a single-node cluster	125
Setting up a node to join the single-node cluster	126
Installing and configuring Ethernet cards for private network	126
Configuring the shared storage	127
Bringing up the existing node	127
Installing the VCS RPMs and license key	128
Configuring LLT and GAB	128
Starting LLT and GAB	130

	Reconfiguring VCS on the existing node	130
	Verifying configuration on both nodes	131
Chapter 8	Upgrading to VCS 5.1 MP2	
	About upgrading to VCS 5.1 MP2	134
	Freeze all the virtual machine service groups in the cluster	135
	Upgrading ESX Server	135
	Upgrading VCS on ESX Server nodes	135
	Upgrading to VCS 5.1 MP2	136
	Upgrading agents' configuration	137
	Upgrading Veritas Virtual Machine Tools in virtual machines running Windows	138
	Upgrading Veritas Virtual Machine Tools in virtual machines running Linux	139
	Unfreezing the service groups	139
	Upgrading Veritas Virtualization Manager (VVM) on clients	140
Section II	Configuring VCS for virtual machines	
Chapter 9	Installing the Veritas Virtualization Manager (VVM)	
	About Veritas Virtualization Manager (VVM)	144
	Installing the Veritas Virtualization Manager (VVM)	145
	Veritas Virtualization Manager hardware requirements	145
	Installing the Veritas Virtualization Manager	145
	Preparing keystores	146
	Keystore generation and population quickstart	147
	Finding the path information on the VirtualCenter Server	148
	Creating the certificate store for the Veritas Virtualization Manager	150
	Copying the keystore file from the VirtualCenter Server to the Veritas Virtualization Manager (VVM)	150
	Copying the keystore file from the VirtualCenter Server to each of the ESX Sever nodes in the VCS cluster	150
	Starting Veritas Virtualization Manager (VVM)	151
	Starting VVM with or with an SSL certificate	151
	Removing Veritas Virtualization Manager (VVM)	152

Chapter 10	Configuring virtual machines for high availability	
	About configuring virtual machines	154
	Configuring virtual machines for high availability using the Veritas Virtualization Manager	154
	Prerequisites for configuring virtual machines for high availability	154
	Reviewing the generated service groups	156
	Accessing the service groups	157
	Verifying virtual machine failover	157
Chapter 11	Configuring virtual machines for disaster recovery	
	About VCS global clusters	160
	VCS global clusters: The building blocks	161
	Prerequisites for global clusters	162
	Setting up a global cluster manually	163
	Configuring the ClusterService group	164
	Configuring replication	165
	Configuring the second cluster	167
	Linking clusters	167
	Creating the global service group	168
	Configuring virtual machines for disaster recovery using the Veritas Virtualization Manager	169
	Overview of tasks	170
	Prerequisites for configuring virtual machines for disaster recovery	171
	Setting up secure DNS update	172
	Using Veritas Virtualization Manager to configure virtual machines for disaster recovery	173
	Deploying VCS components on the virtual machines in the primary site	175
	Confirming service group availability	175
	Reversing the direction of replication	176
	Using VVM to configure virtual machines for disaster recovery on the secondary site	176
	Deploying VCS components on virtual machines in the secondary site	177
	Verifying the service group on the secondary site and using the Global Wizard	177
	Post-failover actions	177
	Reviewing the generated service groups	178
	Accessing the service groups	179
	Verifying virtual machine failover	179
	Disaster recovery best practices	180

	General best practices	180
	Supported configurations	180
	Planned failover across sites	184
	Failover after a disaster	184
Section III	Configuring applications in virtual machines	
Chapter 12	Configuring applications and resources in Linux virtual machines	
	About VCS components for virtual machines	
	running Linux	188
	About monitoring levels	189
	Supported software	189
	About the VCS agent for Oracle	190
	About the VCS agent for the Apache Web server	199
	About the VCS agent for SAP NetWeaver	205
	About the VCS agent for WebLogic Server	212
	About the Application agent	220
	About the Mount agent	224
	How VCS monitors applications and resources on	
	virtual machines	227
	Communication between GuestOSApp and application agents on virtual	
	machines	228
	Installing the applications	229
	Installing Veritas Virtual Machine Tools	229
	Mounting, installing, and configuring Veritas Virtual Machine	
	Tools on the virtual machine	230
	Validating the configuration of Veritas Virtual Machine Tools	232
	Configuring application and resource monitoring	
	inside of virtual machines	232
	Prerequisites	233
	Prerequisites for configuring Oracle and NetIsnr resources	233
	Configuring resources inside virtual machines	234
	Resource data types	236
	Verifying that VCS is running	236
	Applying the configuration and creating the	
	corresponding GuestOSApp resource	236
	Removing the Veritas Virtual Machine Tools	237

Chapter 13	Configuring applications and resources in Windows virtual machines	
	About VCS components on virtual machines running Windows	240
	About monitoring levels	241
	Supported software	241
	How VCS monitors applications and resources on virtual machines	242
	Overview of deploying applications in Windows virtual machines	243
	Installing the applications	245
	Installing Veritas Virtual Machine Tools	245
	Adding the tools .iso file	245
	Connecting to the virtual machine using the Remote Desktop	
	Connection program	245
	Installing Veritas Virtual Machine Tools	246
	Configuring Veritas Virtual Machine Tools	247
	Pagefile configuration in Windows Server 2008	249
	Validating the configuration of Veritas Virtual Machine Tools	250
	Prerequisites for configuring agents in Windows virtual machines	251
	Configuring the agents for SQL Server in a Windows virtual machine	252
	About the VCS agents for SQL Server	252
	Configuring the agents for SQL Server	260
	Configuring the agent for Internet Information Services in a Windows virtual machine	263
	About the VCS agent for Internet Information Services	263
	Configuring the agent for Internet Information Services	266
	Configuring the agent for Exchange 2003 in a Windows virtual machine	269
	About the VCS agent for Exchange Server 2003	269
	Configuring the agent for Exchange 2003	273
	Configuring the agent for Exchange 2007 in a Windows virtual machine	276
	About the VCS agent for Exchange Server 2007	276
	Configuring the Exchange 2007 agent	278
	Configuring WebSphere Application Server in a Windows virtual machine	284
	About WebSphere Application Server 6.0	284
	Configuring the GenericService agent to monitor WebSphere Application Server	284
	Configuring SharePoint Portal Server in a Windows virtual machine	286
	About SharePoint Portal Server 2007	286
	Configuring SharePoint Portal Server	286
	Configuring generic services in a Windows virtual machine	287
	About the VCS agent for generic services	287
	Configuring the agent for generic services	289

Troubleshooting communication between the GuestOSApp and application agents on virtual machines	291
Configuring application monitoring	292
Verifying the configuration for application monitoring	293
Applying the configuration and creating the corresponding GuestOSApp resource	293
Removing Veritas Virtual Machine Tools from the virtual machine running Windows	294

Section IV Administering VCS for VMware ESX

Chapter 14 Administration

Administering a VCS cluster	298
How VCS handles the graceful shutdown of applications outside of VCS control	298
Additional resource type attributes	299
Additional service group attribute	300
Using VMware features and commands in a VCS environment	301
Using VMotion in a VCS environment	301
Using DRS in a VCS environment	304
Using maintenance mode in VCS environment	306
Increasing allocated storage	308
Prerequisites	308
Increasing storage	308
Preserving the last-known good copy of your configuration	309
Using raw devices for the virtual machine's boot image	310
Setting up shared raw device storage under VCS	310
The service group for the shared raw device	311
Performing maintenance on virtual machines and applications in virtual machines	312
Performing maintenance on a virtual machine	312
Performing maintenance on applications inside the virtual machine	313
Troubleshooting maintenance	314
Atypical VCS configuration with the virtual machine configuration file on local storage	314

Index

Installing VCS for VMware ESX

This section contains the following chapters:

- [Chapter 1, “Introducing VCS for VMware ESX”](#) on page 17
- [Chapter 2, “Requirements”](#) on page 27
- [Chapter 3, “Planning to install VCS on an ESX Server”](#) on page 37
- [Chapter 4, “Installing and configuring VCS on ESX Servers”](#) on page 51
- [Chapter 5, “Verifying VCS on ESX Servers”](#) on page 87
- [Chapter 6, “Adding and removing cluster nodes”](#) on page 103
- [Chapter 7, “Installing VCS on a single node”](#) on page 117
- [Chapter 8, “Upgrading to VCS 5.1 MP2”](#) on page 133

Introducing VCS for VMware ESX

This chapter contains the following topics:

- [Features](#)
- [About VCS](#)
- [About high availability using VCS for VMware ESX](#)
- [About disaster recovery using VCS for VMware ESX](#)
- [Installation road map](#)

Features

The following features appear in this release of VCS.

- **High availability**
VCS provides high availability for virtual machines, the applications that run in the virtual machines, storage and networking components of the ESX Server, and the ESX Server itself.
- **Disaster recovery**
Use VCS to prepare your environments for disaster—and have confidence that your clusters can survive a disaster. Requires a supported replication infrastructure and VCS agent for replication.
- **Last-known good copy**
After testing application configuration and data integrity, you can take a snap shot of the “last known good copy” of the operating system for safe-keeping. Requires a supported VCS agent for replication.
- **Support for VMotion and Distributed Resource Scheduler (DRS)**
When VMotion or DRS moves a virtual machine, VCS correctly interprets this and does not register the movement as a failure.
- **Management options**
Manage your clusters and nodes with the Java Console, the Veritas Cluster Server Management Console, or from the command line.
- **The Veritas Virtualization Manager**
Use the Veritas Virtualization Manager for quick deployment and configuration of virtual machines to high availability and for disaster recovery.
- **Notification**
VCS can notify you of events. You have access to SMTP email notification and SNMP trap notification.
- **Virtual machine storage management**
Enables you to easily grow your application data mounts. Note that in certain configurations you can use NFS and raw devices for your virtual machine's data storage.

About VCS

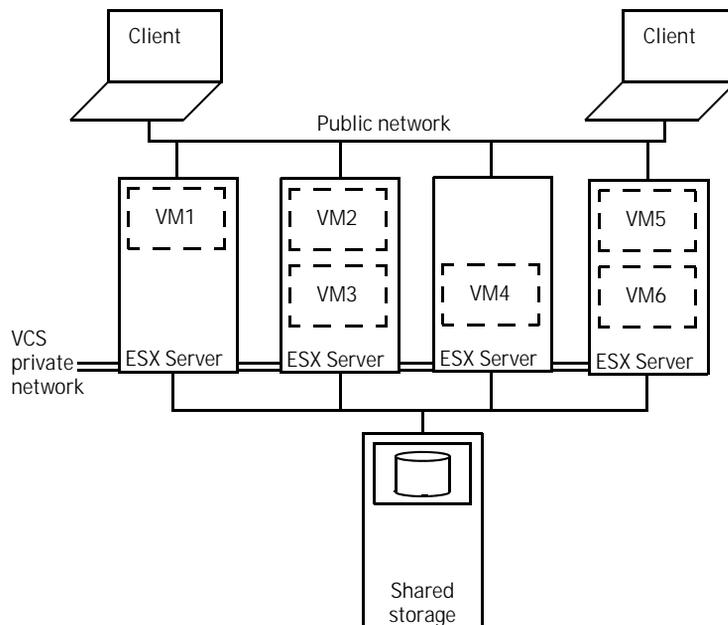
Veritas Cluster Server (VCS) monitors sites, clusters, systems, virtual machines, and applications. You can group up to 16 ESX Server systems together with VCS to form a cluster. Each ESX Server that runs VCS becomes a node in a cluster.

VCS enables you to manage groups of applications. These groups of applications are called service groups. Depending on your configuration, when hardware or software fails, VCS moves the service group to another node in the same cluster, or to a remote node in a different cluster.

VCS for VMware ESX runs the majority of the VCS components (the VCS engine, most agents, GAB, and LLT) at the ESX console operating system, and not in the individual virtual guest operating systems.

Figure 1-1 illustrates a typical four-node VCS cluster configuration connected to shared storage. Client workstations receive service over the public network from applications running on VCS nodes. VCS monitors the nodes, virtual machines, and applications on the virtual machines. VCS nodes in the cluster communicate over a private network.

Figure 1-1 Example of a four-node VCS cluster configuration



Multiple nodes

VCS runs on each node in the cluster. The private network enables the nodes to share identical state information about all resources and to recognize: active nodes, nodes that are joining or leaving the cluster, and failed nodes. The private network requires two communication channels for heartbeats.

Virtual machines and applications

You can make the virtual machines and applications that run inside the virtual machines highly available. When the virtual machine or application faults, VCS takes corrective actions. In some cases, you might want VCS to restart the application. In other situations, you might want VCS to fail over to a different node entirely.

Detailed and basic monitoring

For certain applications, VCS supports detailed monitoring and the ability to detect a graceful shut down. If confronted with the failure of an application, VCS attempts to restart the application or move the virtual machine that runs the application onto another node. If confronted with a user intentionally shutting down an application or moving the virtual machine, VCS takes no action.

In general, basic monitoring checks for running application processes. Detailed monitoring, however, performs application-specific tasks to check the application's health.

An example of this is monitoring an Apache instance in a RHEL guest operating system. With basic monitoring, VCS ensures that a specific httpd process (and pid) is in the process list. Detailed monitoring takes this a step further by attempting to connect to the Apache service and evaluate its response.

Veritas Virtual Machine Tools

Veritas Virtual Machine Tools is a package of tools that reside in the virtual machine and that provides configuration resources and wizards. You can make these tools available through the Veritas Virtualization Manager by mounting an ISO file.

Shared storage

A VCS hardware configuration usually consists of multiple nodes that are connected to shared storage through I/O channels. Shared storage provides multiple access paths to the same data, and enables VCS to restart virtual machines on alternate nodes when a node fails.

LLT and GAB

VCS uses two components, LLT and GAB, to share data over private networks among nodes.

- LLT (Low Latency Transport) provides very fast, kernel-to-kernel communications, and monitors network connections. During installation and subsequent changes to the cluster configuration, VCS creates or modifies the following text-readable files:
 - `/etc/llthosts`, which lists all the nodes in the cluster
 - `/etc/llttab`, which describes the local system's private network links to the other nodes in the cluster
- GAB (Group Membership and Atomic Broadcast) provides the global message order required to maintain a synchronized state among the nodes. It monitors disk communications such as the VCS heartbeat utility. VCS creates or modifies the gab configuration file (`/etc/gabtab`) during cluster creation or modification.

Network channels for heartbeats

For the VCS private network, heartbeats travel over two or more network channels. These network channels are also used for transmitting information.

Each cluster configuration requires at least two network channels between the systems. The requirement for two channels is to protect your cluster against the failure of a single network, and subsequent cluster corruption. For more information about network partitions refer to the following information:

Service groups

A service group is a collection of all of the items (resources) needed to provide a particular service. In the VCS for ESX paradigm, each virtual machine (and the applications it hosts) is considered a service. A service group is how you structure dependencies among resources. For example, your virtual machine must have storage to work. The virtual machine has a dependency on its storage.

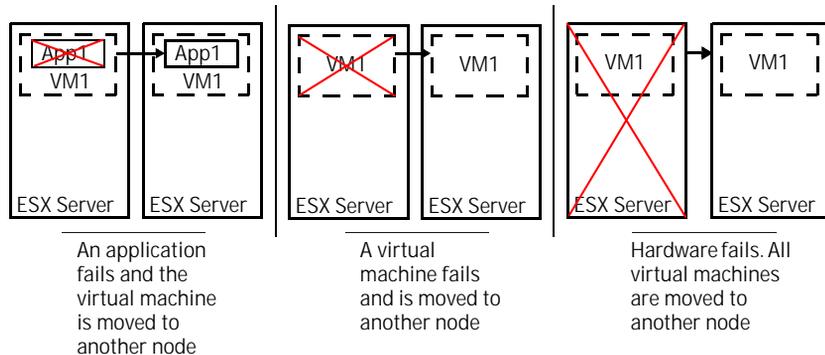
When you use the Veritas Virtualization Manager to configure virtual machines, it also creates service groups. The service group that it creates is for the virtual machine, its network, and storage. For applications, like Oracle or SQL, the service group contains additional resources for these applications. For more information on service groups:

About high availability using VCS for VMware ESX

VCS monitors applications within virtual machines. VCS does not, however, start or stop applications. For all applications inside of virtual machines, configure them to start automatically when the virtual machine starts. When you want to stop the application, stop it normally.

In all failure situations, VCS moves the virtual machine. When an ESX Server node or a configured application in a virtual machine fails, VCS fails over the virtual machine to other nodes. When a virtual machine crashes, VCS detects the failure and initiates the failover of the virtual machine.

Figure 1-2 A failed node, virtual machine, or application moved to a working system



About disaster recovery using VCS for VMware ESX

Use VCS to ensure that your applications and data remain online. Disaster recovery protects your servers from unwanted downtime due to a cluster- or site-wide event. VCS can migrate your applications to a safe, predetermined location, and with a minimum of downtime, to keep your services running.

You need to test your infrastructure and configuration to see if it can survive a disaster. VCS for VMware ESX provides for this testing with fire drills. These fire drills give your applications a full test of their functionality during an emergency.

When you prepare for disaster, you should have a last known good copy of your application and guest data available. With the last known good copy, even if a disaster strikes within a disaster (a corrupted boot image), you can recover with the last stable copy of the virtual machine.

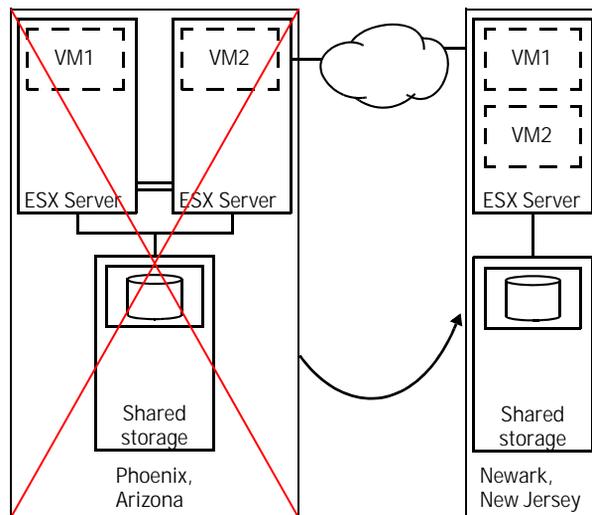
Replicated storage

VCS supports several replication technologies, and uses them for disaster recovery. VCS agents monitor, and if needed change, the replication state between primary and secondary sites. Contact your Symantec sales representative for a list of replication technologies supported with this release of VCS.

Global clusters

You can create clusters that operate in two geographically diverse locations. In the event that one cluster fails completely, the entire cluster fails over to the back-up location. As its virtual machines come back online, the network configuration is changed to re-direct clients to the new site. Applications restart in the virtual machines.

Figure 1-3 A two-node cluster with one globally clustered node



Global clustering requires a separate license. For more information, refer to the *Veritas Cluster Server User's Guide*.

Installation road map

Figure 1-4 on page 24 illustrates a VCS for VMware ESX installation.

Figure 1-4 Suggested installation flow

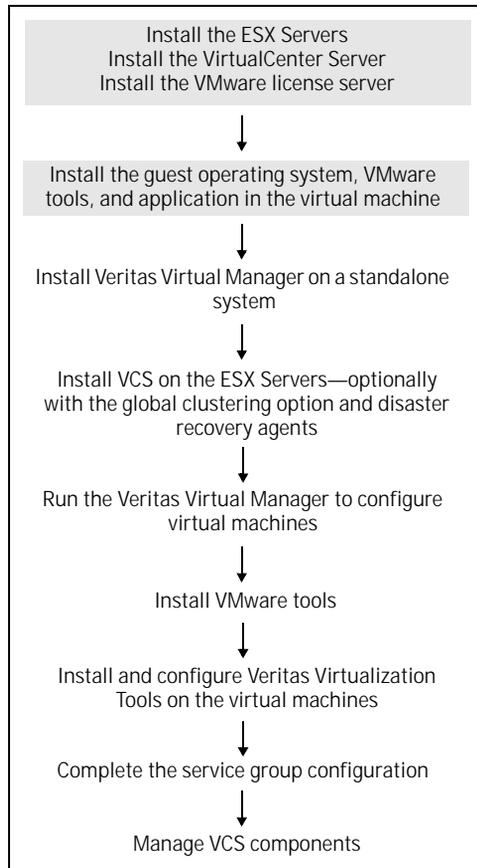


Table 1-1 on page 25 describes where to look for the pertinent road map information.

Table 1-1 Installation road map

Road map entry	Document or chapter
Install VCS on the ESX Servers. Optionally install the global clustering option and disaster recovery agents.	<ul style="list-style-type: none"> ■ See “Requirements” on page 27. ■ See “Planning to install VCS on an ESX Server” on page 37. ■ See “Installing and configuring VCS 5.1 MP2” on page 53. ■ See “Verifying VCS on ESX Servers” on page 87.
Install the Veritas Virtual Manager on a standalone system.	See “Installing the Veritas Virtualization Manager (VVM)” on page 143.
Run the Veritas Virtual Manager to configure virtual machines.	<ul style="list-style-type: none"> ■ See “Configuring virtual machines for high availability” on page 153. ■ See “Configuring virtual machines for disaster recovery” on page 159.
Install and configure Virtual Machine Tools in the virtual machine.	<ul style="list-style-type: none"> ■ See “Configuring applications and resources in Linux virtual machines” on page 187. ■ See “Configuring applications and resources in Windows virtual machines” on page 239.
Configure VCS to monitor your applications.	<ul style="list-style-type: none"> ■ See “Configuring applications and resources in Linux virtual machines” on page 187. ■ See “Configuring applications and resources in Windows virtual machines” on page 239.
Manage VCS components.	See the <i>Veritas Cluster Server User’s Guide</i> .

Requirements

This chapter contains the following topics:

- [VMware ESX Server software and infrastructure](#)
- [Patches](#)
- [Supported operating systems](#)
- [Supported applications](#)
- [Support for detecting intentional offline for specific applications](#)
- [Supported hardware](#)
- [Veritas Cluster Server hardware requirements](#)
- [Veritas Cluster Server required disk space](#)
- [VMware ESX components and configuration requirements](#)

VMware ESX Server software and infrastructure

VCS 5.1 MP2 supports the following:

- ESX Server 3.0.1, 3.0.2, 3.5, 3.5 Update 1, and later
- VirtualCenter Server 2.0, 2.5, and later
- VMotion
- Datastores on VMFS 3 (SAN-attached)
- When you do not have VMotion, you lose the following VMotion actions triggered from VCS:
 - From the command line, you cannot perform the `hagrp -migrate` command
 - From the Veritas Virtualization Manager, you cannot right-click a service group and order it to migrate
 - Other VCS Management interfaces like VCS Cluster Management Console, VCS APIs, etc.
- VCS supports the following VMware infrastructure offerings:
 - Foundation
 - Standard
 - Enterprise

Veritas products will operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility. For the ESX Server in addition to the above kernel ABI information, VCS will operate on subsequent releases provided that the Virtual Infrastructure API and ESX Host CLI compatibility is maintained.

Patches

Review the following patch recommendations, apply the patches where necessary.

Patch for VCS agents hangs and when agents return UNKNOWN states

Patches the ESX Server and the VirtualCenter Server for incomplete SOAP messages. This patch fixes situations where VCS agents hang due to monitor timeouts, or start to return UNKNOWN states.

Review the VMware knowledge base article for more information:

<http://kb.vmware.com/kb/1002415>

Find the patch number from the knowledge base article, and download it from:

http://www.vmware.com/download/vi/vi3_patches.html

Patch for ESX Server 3.0.1 freezing during rescan operations

Patches the ESX Server hosts to fix a problem where they stop responding during a rescan.

Review the VMware knowledge base articles for more information:

- <http://kb.vmware.com/kb/1000039>
- <http://kb.vmware.com/kb/10229>

Patch for ESX 3.0.x systems running VCS 5.1

On ESX 3.0.x systems that run VCS 5.1 or a subsequent VCS MP release on top of 5.1, ensure that the VMware Perl and COM Scripting API version 2.3.1 (or later) are installed. The ESXVirtualMachine agent requires this patch so that it can correctly determine virtual machine availability.

To download and install the Scripting API patch, refer to VMware documentation for details.

Bind utilities update

The disaster recovery configuration requires the latest bind utilities. The DNS agent requires bind-utils-9.2.4-16.EL4. Symantec recommends installing the latest version of bind utilities before configuring the cluster for disaster recovery.

Supported operating systems

Refer to the following information for supported operating systems for VCS for ESX:

- [“Supported operating systems in virtual machines for high availability”](#) on page 30
- [“Supported operating systems in virtual machines for application monitoring or disaster recovery”](#) on page 30
- [“Supported operating systems for increasing allocated storage”](#) on page 31

Supported operating systems in virtual machines for high availability

VCS for ESX provides high availability for all the operating systems that VMware ESX supports as virtual machine guests.

If you need application monitoring or disaster recovery refer to the following section.

Supported operating systems in virtual machines for application monitoring or disaster recovery

[Table 2-1](#) lists the architectures and operating systems that VCS for VMware supports.

Table 2-1 Supported operating systems and architectures

Guest operating systems	Architectures	File systems/ Volume managers
†Windows 2000 Server or Advanced Server with Service Pack 4	x86 (32-bit)	NTFS
†Windows Server 2003: Standard Edition or Enterprise Edition (SP1 with .NET Framework 2.0 required)	x86 (32-bit) x86 (64-bit)	NTFS
Windows Server 2008: Standard Edition or Enterprise Edition	x86 (32-bit) x86 (64-bit)	NTFS NTFS
†*Red Hat Enterprise Linux 4 (RHEL 4) Update 5	x86 (32-bit) x86 (64-bit)	ext2, ext3, reiserfs/ LVM

Table 2-1 Supported operating systems and architectures

Guest operating systems	Architectures	File systems/ Volume managers
†*SUSE Linux Enterprise Server 9 (SLES 9) with SP4	x86 (32-bit) x86 (64-bit)	ext2, ext3, reiserfs/ LVM
SUSE Linux Enterprise Server 10 (SLES 10) with SP1	x86 (32-bit) x86 (64-bit)	ext2, ext3, reiserfs/ LVM
Solaris 10	x86	

† Supports the mount .iso feature.

* This version is supported due to a known Linux file system issue. For more information, see the release notes.

On Linux-based operating systems: Veritas products will operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

On Windows-based operating systems: Veritas products will operate on subsequent service pack (SP) releases provided that the vendor maintains forward compatibility.

Note: The EMC CLARiiON series and Symmetrix series storage arrays do not support virtual machines running Solaris 10 U1 guest operating systems. See the VMware documentation for more information.

Supported operating systems for increasing allocated storage

Table 2-2 lists the guest operating systems that VCS supports for increasing allocated storage.

Table 2-2 Supported operating systems for increasing allocated storage

Guest operating systems	32-bit	64-bit	Supported file systems
Windows 2000	Yes	No	NTFS
Windows Server 2003	Yes	No	NTFS
Windows Server 2008	Yes	Yes	NTFS
*RHEL 4 Update 3	Yes	Yes	ext3

Table 2-2 Supported operating systems for increasing allocated storage

Guest operating systems	32-bit	64-bit	Supported file systems
SLES 9 with SP3	Yes	Yes	reiserfs/LVM
SLES 10 with SP1	No	No	n/a
Solaris 10	N.A.	No	n/a

* Supports increasing allocated storage once.

Note that file systems on raw device maps do not support increasing allocated storage.

Supported applications

VCS provides agents to monitor the following applications that run in virtual machines.

Table 2-3 Supported guest applications

Platform	Applications	Versions
Linux	Apache Web server	1.3, 2.0, and 2.2
" "	IBM HTTP Server	1.3 and 2.0
" "	Oracle	10g
" "	SAP NetWeaver	SAP R/3-4.6C with a 4.6D Kernel, 4.6D, 4.7 Enterprise Version SAP Web AS-6.20, 6.40, 7.00 SAP NetWeaver-2004, 2004s
" "	WebLogic Server	9.0, 9.1, 9.2, and 10.0
Windows	Exchange	Exchange Server 2003 Exchange Server 2007 (SP1)
Windows	SharePoint Portal Server	SharePoint Portal Server 2007
" "	IIS	5.0 and 6.0
" "	SQL	Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (both require SP4) Microsoft SQL Server 2005, 32-bit (SP1 required)

Table 2-3 Supported guest applications

Platform	Applications	Versions
" "	WebSphere Application Server	6.0 WebSphere Application Server is configured as a GenericService agent on virtual machines running Windows.

VCS additionally provides the following agents to monitor other applications:

- Application agent on virtual machines running Linux
- GenericService agent on virtual machines running Windows

Support for detecting intentional offline for specific applications

Certain agents can identify when an application has been intentionally shut down as opposed to when an application has crashed. When VCS detects an intentional offline, VCS does not trigger a failover. This feature allows administrators to manage the applications (start/stop) that run inside the virtual machines without causing additional failovers.

Table 2-4 Agents that support detection of intentional offline of the configured application

Guest operating systems	Applications
Linux	<ul style="list-style-type: none"> ■ Apache ■ Oracle ■ NetIsnr ■ SAP NetWeaver ■ WebLogic Server
Windows	<ul style="list-style-type: none"> ■ Exchange Server ■ Internet Information Services (IIS) ■ SQL Server ■ GenericService ■ WebSphere Application Server

Supported hardware

For the latest information on supported hardware, see the hardware compatibility list published by VMware.

See the documentation published by your array vendor for information about:

- Hardware compatibility with VMware ESX
- Supported microcode or firmware versions
- Supported versions of client software for the array
- Supported versions of the replication and mirroring software
- Recommended array settings

Veritas Cluster Server hardware requirements

Make sure that your hardware meets the following requirements.

Table 2-5 Hardware requirements for a cluster

Item	Description
VCS systems	From one to sixteen ESX Servers that run the supported ESX Server operating system version.
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.
Disks	Typical VCS configurations require that shared disks support applications that migrate between systems in the cluster.
Disk space	See " Veritas Cluster Server required disk space " on page 35.
Network Interface Cards	In addition to the built-in public Network Interface Card (NIC), VCS requires at least one more NIC per system. Symantec recommends two additional NICs.
Fibre Channel or SCSI host bus adapters	Typical VCS configuration requires at least one SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.
RAM	Each VCS system requires at least 256 megabytes in addition to other system and application requirements.

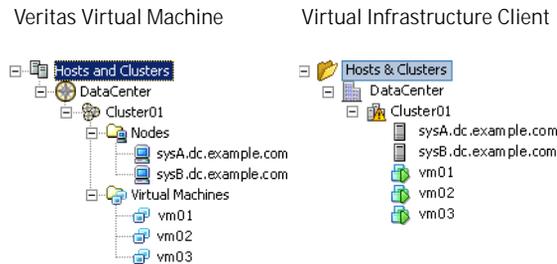
Veritas Cluster Server required disk space

Refer to the *Veritas Cluster Server Release Notes* for updated disk space requirements.

VMware ESX components and configuration requirements

- VMware Tools installed in the guest operating system of each virtual machine. VCS requires VMware Tools for application monitoring.
- VMware VirtualCenter Web Service properly configured to enable SSL communication for the Virtual Machine Deployment wizard.
- VCS for VMware ESX supports both VMotion and DRS. Both of these VMware features require exact parity among the nodes in the VCS and VMware clusters. Both of these VMware features have VCS equivalents, and certain requirements for their proper use.

Figure 2-1 Maintaining an exact correlation between products



Planning to install VCS on an ESX Server

This chapter contains the following topics:

- [About installing VCS on an ESX Server](#)
- [About optional VCS features](#)
- [Performing pre-installation tasks](#)
- [Preparing your VCS installation and configuration information](#)

About installing VCS on an ESX Server

Before you perform the VCS installation, read the following Veritas Technical Support TechNote for the latest information on updates, patches, and software issues:

<http://entsupport.symantec.com/docs/289940>

About optional VCS features

Make sure to install all RPMs when the installation program prompts you to get the optional features. Review the description of each optional feature and decide which features that you want to configure with VCS:

- [Veritas Cluster Server Management Console](#)
- [SMTP email notification for VCS events](#)
- [SNMP trap notification for VCS events](#)
- [Global clusters using VCS](#)

Note: This release does not support configuring clusters in the secure mode. Do not configure the Symantec Product Authentication server while installing or configuring VCS.

Veritas Cluster Server Management Console

The Veritas Cluster Server Management Console is a management interface that enables you to monitor and administer clusters from a web console.

A newer version of Veritas Cluster Server (VCS) Management Console is available manage VCS clusters. VCS Management Console was earlier known as Cluster Management Console.

Refer to the *Veritas Cluster Server Release Notes* for more information.

Cluster Manager (Java Console)

The Cluster Manager (Java Console) is a light-weight Java-based interface. You install the Java Console on a Windows client, and you use it to manage your clusters and the service groups that comprise your cluster. The Java Console is automatically installed when you install the Veritas Virtualization Manager.

For more information, refer to the *Veritas Cluster Server User's Guide*.

See "[Installing the Cluster Manager \(Java Console\)](#)" on page 68.

Notifications

VCS for VMware ESX offers two server-side notification services. You can get notification from SMTP email or using SNMP traps.

For more information about these notification services, refer to the *Veritas Cluster Server User's Guide*.

SMTP email notification for VCS events

You have the option to configure SMTP email notification of VCS events from the VCS Notifier component. If you choose SMTP notification, have the appropriate information ready.

SNMP trap notification for VCS events

You have the option to configure SNMP trap notification of VCS events from the VCS Notifier component. If you choose SNMP notification, have the appropriate information ready.

Global clusters using VCS

Global clusters provide the ability to fail over applications between geographically distributed clusters. You require a separate license to configure global clusters. You must add this license during the installation.

If you choose to configure global clusters, the `installvcs` program enables you to choose whether or not to use the same NIC, virtual IP address, and netmask as are configured for the ClusterService group, which are the defaults. If you choose not to use the same networking information, you must specify appropriate values for the NIC, virtual IP address, and netmask when prompted.

Performing pre-installation tasks

[Table 3-1](#) lists the tasks you must perform before you install VCS.

Table 3-1 Pre-installation tasks

Task	Reference
Set the PATH and MANPATH variables.	“Setting environment variables” on page 41
Set up the private network.	“Setting up the private network” on page 41

Table 3-1 Pre-installation tasks

Task	Reference
Configure the private network.	“Selecting virtual interfaces for use by the private network” on page 42
Enable communication between systems.	“Enabling communication between systems” on page 43
Review basic instructions to optimize LLT media speeds.	“Optimizing LLT media speed settings on private NICs” on page 43
Review guidelines to help you set the LLT interconnects.	“Setting the media speed of the LLT interconnects” on page 43
Set up SSH on cluster systems.	“Enabling password free SSH-communication between systems” on page 44
Obtain license keys.	“Obtaining VCS license keys” on page 46
Mount the product disc	

Setting environment variables

Setting the PATH variable

Installation commands as well as other commands reside in the `/sbin`, `/usr/sbin`, `/opt/VRTS/bin`, and `/opt/VRTSvcs/bin` directories. You need to add these directories to your PATH environment variable to access the commands.

To set the PATH variable

- ◆ Do one of the following:
 - For the Bourne Shell (sh or ksh), type:


```
$ PATH=/usr/sbin:/sbin:/opt/VRTS/bin:/opt/VRTSvcs/bin:\
$PATH; export PATH
```
 - For the C Shell (csh or tcsh), type:


```
% setenv PATH /usr/sbin:/sbin:/opt/VRTS/bin:\
/opt/VRTSvcs/bin:$PATH
```

Setting the MANPATH variable

To set the MANPATH variable

- ◆ Do one of the following:
 - For the Bourne Shell (sh or ksh), type:


```
$ MANPATH=/usr/share/man:/opt/VRTS/man; export MANPATH
```
 - For the C Shell (csh or tcsh), type:


```
% setenv MANPATH /usr/share/man:/opt/VRTS/man
```

If you use the `man` command to access manual pages, set `LC_ALL` to "C" in your shell for correct page display.

```
# export LC_ALL=C
```

Setting up the private network

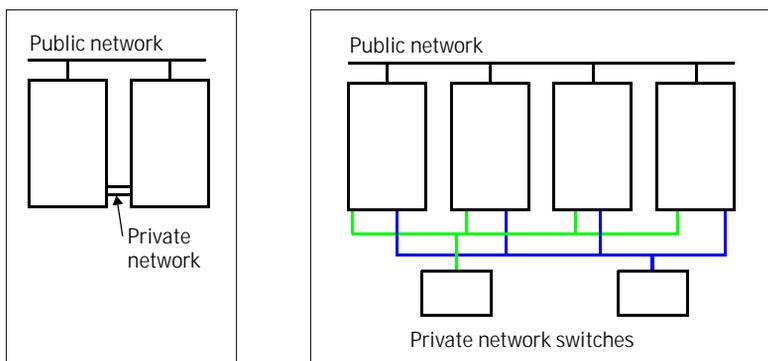
VCS requires you to set up a private network between the systems that form part of a cluster.

To set up the private network hardware

- 1 Install the required network interface cards (NICs).
- 2 Connect the private NICs on each system.
- 3 Use cross-over Ethernet cables (supported only on two systems), or switches for each VCS communication network. (Hubs work too.) Ensure that power to the switches comes from separate sources.

- 4 On each system, use two independent network cards to provide redundancy. When you set up heartbeat connections, if a failure removes all communications between nodes, a chance for shared storage data corruption exists. For this reason, heartbeat network architecture must be as robust as possible, reducing or eliminating potential single points of failure (SPOFs).

Figure 3-1 Private network setups: two-node and four-node clusters



- 5 Test network connections by temporarily assigning network addresses and use `telnet` or `ping` to verify communications. LLT uses its own protocol, and does not use TCP/IP. To ensure the private network connections are used only for LLT communication and not for TCP/IP traffic, unconfigure the temporary addresses after testing. The `installvcs` program configures the private network in the cluster during installation. See “[About installing and configuring VCS](#)” on page 52.

Selecting virtual interfaces for use by the private network

VCS uses LLT private links to monitor network communication. LLT requires virtual interfaces that it can use for private links.

During installation you can either specify physical or virtual interface information for the private links.

- If you choose to specify physical interface information, the `installvcs` program creates a virtual interface (`vswif`) that is mapped to the physical interface that you chose.

- If you choose to specify an existing virtual interface, verify that it is mapped to the correct physical interface. Make sure that all the systems in the cluster have virtual interfaces configured.

On each node, Symantec recommends that you map at least two virtual interfaces to two separate physical interfaces to provide redundancy.

Enabling communication between systems

When you install VCS using the `installvcs` program, to install and configure the entire cluster at one time, make sure that communication between systems exists. By default the `installvcs` program uses `ssh`. You must grant permissions for the system where you run `installvcs` program to issue `ssh` commands as root on all systems in the cluster. If `ssh` is used to communicate between systems, it must be configured in a way such that it operates without requests for passwords or passphrases.

If system communication is not possible between systems using `ssh`, you have recourse.

See [“Performing VCS installation in a secure environment”](#) on page 70.

Optimizing LLT media speed settings on private NICs

For optimal LLT communication among the cluster nodes, the interface cards on each node must use the same media speed settings. Also, the settings for switches or hubs used for the interconnects must match that of the interface cards. Incorrect settings can cause poor network performance or even network failure.

Setting the media speed of the LLT interconnects

If you have switches or hubs for LLT interconnects, Symantec recommends using the `Auto_Negotiation` media speed setting on each Ethernet card on each node. If you do not use `Auto_Negotiation`, you have to set it to the same speed on all nodes for all NICs used by LLT.

If you have switches or hubs for LLT interconnects and you do not use the `Auto_Negotiation` media speed setting, set the hub or switch port to the same setting as that used on the cards on each node.

If you use directly connected Ethernet links (using crossover cables), set the media speed to the highest value common to both cards, typically `100_Full_Duplex`.

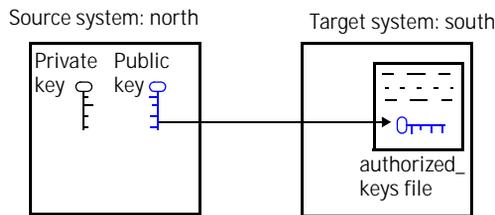
Symantec does not recommend using dissimilar network cards for private links.

Details for setting the media speeds for specific devices are outside of the scope of this manual. Consult the device's documentation for more information.

Enabling password free SSH-communication between systems

VCS requires password free SSH-based communication for installation. The following is an example of how to set up SSH-based password free communication between a source system (north) and a target system (south). In this procedure, you first create a DSA key pair. You then check to see if the .ssh directory is in place on the target system (south). If doesn't exist, create it and set permission to 700. From the key pair, you append the public key from the source system (north) to the authorized_keys file on the target systems.

Figure 3-2 Creating the DSA key pair and appending it to target systems



Visit the OpenSSH website located at: <http://openssh.org> to access online manuals and other resources.

To create the DSA key pair

- 1 On the source system (north), log in as root to generate a DSA key pair. Enter the following command:

```
# ssh-keygen -t dsa
```

Output that resembles the following appears:

```
Generating public/private dsa key pair.
```

```
Enter file in which to save the key (/root/.ssh/id_dsa):
```

- 2 Press the Enter key to accept the default location of `/root/.ssh/id_dsa`.

- 3 When the program asks you to enter the passphrase, press the Enter key twice.

```
Enter passphrase (empty for no passphrase):
```

```
Press the Enter key.
```

```
Enter same passphrase again:
```

```
Press the Enter key again. Output that resembles the following lines appears:
```

```
Your identification has been saved in /root/.ssh/id_dsa.
```

```
Your public key has been saved in /root/.ssh/id_dsa.pub.
```

```
The key fingerprint is:
```

```
1f:00:e0:c2:9b:4e:29:b4:0b:6e:08:f8:50:de:48:d2 root@north
```

To check the .ssh directory on the target system

- 1 Log in to the target system (south).
- 2 Change to the /root/.ssh directory.
cd /root/.ssh
- 3 If the directory doesn't exist, add it.

To add and set permissions for the .ssh directory on the target system

- 1 Log in to the target system (south).
- 2 Change to the root directory.
cd /root
- 3 Create the .ssh directory.
mkdir .ssh
- 4 Change the directory's permission to 700.
chmod 700 .ssh

To append the public key from the source system to the authorized_keys file on the target system

- 1 On the source system (north), run the secure copy command to move the public key to a temporary file on the target system (south):
scp ~/.ssh/id_dsa.pub south:~/.ssh/north_id_dsa.pub
Where north_id_dsa.pub is the name that you have given the temporary public key file.
- 2 On north, enter the root password for south:
south's password: **foobar**
Output that resembles the following line appears:
id_dsa.pub 100% 603 0.6KB/s 00:00
- 3 On north, run the following command to append north's public key, which is now a temporary file on south, to south's authorized keys:
ssh -l root south "cat ~/.ssh/north_id_dsa.pub >> ~/.ssh/authorized_keys"
- 4 On north, enter the root password for south.
south's password: **foobar**

- 5 On south, enter the following command:

```
# cat ~/.ssh/authorized_keys
```

Output that resembles the following appears:

```
ssh-dss  
JJCJB3NzaC1kc3MJJCBLNGnJQQfk91gxKazYarXpjUNyy85sCa5rfaCIVII87  
laGss6NT6pc7N/NeL1cSckc6U0XD5xIGkXpdPW7omH1TJkJKMfIJNTzsY/QrUGz  
.  
.  
.  
oTumbtLjennnd4jnm4oE0MOFJ+ST7wZgsVn1seHPdW3seXr+bUhKI+3bMqvmZs7M  
+Lp36z/YZc0J= root@north
```
- 6 On south, remove the temporary file. Enter the following command:

```
# rm ~/.ssh/north_id_dsa.pub
```
- 7 Repeat this procedure on each target system where you plan to install VCS.

To verify that you can connect to a target system

- ◆ On the source system (north), type the following command:

```
# ssh -l root south uname -a
```

Where south is the name of the target system. The command should execute from the source system (north) to the target system (south) without the system requesting a passphrase or password.

Obtaining VCS license keys

This product includes a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased. A single key lets you install the product on the number and type of systems for which you purchased the license. A key may enable the operation of more products than are specified on the certificate; however, you are legally limited to the number of product licenses purchased.

The product installation procedure describes how to activate the key. If you encounter problems while licensing this product, visit the Symantec licensing support website at:

<https://licensing.symantec.com>

The VRTSvlic package enables product licensing. After the VRTSvlic is installed, the following commands and their manual pages are available on the system:

- vxlicinst Installs a license key for a Symantec product
- vxlicrep Displays currently installed licenses
- vxlictest Retrieves features and their descriptions encoded in a license key

Accessing documentation

You can access all the documentation for VCS on the product disc. Insert the product disc into the system's drive and change directory to the docs directory. Copy the contents of this directory to each system where you need to access the documentation.

Preparing your VCS installation and configuration information

When you perform the installation, ready the following information:

- To install the VCS RPMs, prepare the system names and license keys.
 - The system names
The names of the systems where you plan to install VCS.
 - The license keys
Keys can include: a valid site license, a demo license, a VCS global cluster license key.
See "[Obtaining VCS license keys](#)" on page 46.
- To configure VCS, prepare the cluster's name, the cluster's unique ID, and the names for the private network's NICs.
 - The cluster's name
The cluster name must begin with a letter of the alphabet. The cluster name can contain only the characters "a" through "z", "A" through "Z", the numbers "0" through "9", the hyphen "-", and the underscore "_".
For example: **vcs_cluster27**
 - The cluster's unique ID number
A number in the range of 0-65535. Within the site that contains the cluster, each cluster must have a unique ID.
 - The private network's device names for the NICs
The device names of the NICs that the private networks use among systems.
Do not use the public network's name of the network interface card, which is typically vswif0.
Example: **vswif1, vswif2**
- To add users to VCS, prepare the VCS user's name, password, and privileges.
 - The VCS user's privileges
Users have three levels of privilege: A=Administrator, O=Operator, or G=Guest.

- For the Veritas Cluster Server Management Console to locally manage this cluster (optional), prepare the names of the public NICs for each node in the cluster, and the virtual IP address of the NIC for the console.
 - The name of the public NIC for each node in the cluster
The device name for the NIC that provides public network access.
Example: vswif0
 - A virtual IP address of the NIC for the Veritas Cluster Server Management Console
This virtual IP address becomes a resource for use by the ClusterService group that includes the Veritas Cluster Server Management Console. The cluster virtual IP address can fail over to another cluster system, making the Web console highly available.
- For the configuration of the cluster connector (optional) for the Veritas Cluster Server Management Console, prepare the management server's network address for the console, the console's service account password, and the root hash of the management server.
 - The management server network address for Veritas Cluster Server Management Console
The Veritas Cluster Server Management Console cluster connector requires the management server network address.
For example: mgmtserver1.symantecexample.com
See "[Veritas Cluster Server Management Console](#)" on page 38.
 - A Veritas Cluster Server Management Console service account password
You need to set this account password while you install the management server.
 - The root hash of the management server
You can use `vssat showbrokerhash` command and copy the root hash of the management server.
- To configure SMTP email notification (optional), prepare the domain-based address of the SMTP server, the email addresses recipients, and select the event's severity.
 - The domain-based address of the SMTP server
The SMTP server sends notification emails about the events within the cluster.
Example: smtp.symantecexample.com
 - The email address of each SMTP recipient to be notified
Example: john@symantecexample.com

- To decide the minimum severity of events for SMTP email notification
 Events have four levels of severity: I=Information, W=Warning, E=Error, and S=SevereError.
 Example: E
- To configure SNMP trap notification (optional), prepare the SNMP trap daemon's port, the SNMP console's system name, and select the event's severity.
 - The port for the SNMP trap daemon
 The default port is 162.
 - The system name for each SNMP console
 Example: **saturn**
 - To decide the minimum severity of events for SNMP trap notification
 Events have four levels of severity: I=Information, W=Warning, E=Error, and S=SevereError.
 Example: E
- To configure global clusters (optional), prepare the name of the public NIC, the NIC's virtual IP address, and the netmask for the NIC's virtual IP address.
 - The name of the public NIC
 You can use the same NIC that you configured for the ClusterService group. Otherwise, specify appropriate values for the NIC.
 Example: **vswif0**
 - The virtual IP address of the NIC
 You can use the same virtual IP address that you configured for the ClusterService group. Otherwise, specify appropriate values for the virtual IP address.
 Example: **10.10.12.1**
 - The netmask for the virtual IP address
 You can use the same netmask as configured for the ClusterService group. Otherwise, specify appropriate values for the netmask.
 Example: **255.255.240.0**

Optional VCS RPMs

The optional VCS RPMs include:

- VRTScmccc
Veritas Cluster Server Management Console Cluster Connector
- VRTScmcs
Veritas Cluster Server Management Console
- VRTScssim
VCS Simulator
- VRTSvcsmn
Manual pages for VCS commands

Service console firewall ports that installvcs opens during installation

Note that during installation, the following service console firewall ports are automatically opened. These ports must remain open during normal VCS operations.

21/TCP/FTP	22/TCP/SSH	80/TCP/HTTP
443/TCP/HTTPS	8443/TCP/VRTSweb	8181/TCP/VRTSweb
14300/TCP/VRTSweb	14301/TCP/VRTSweb	8181/TCP/GCM Web Server port
14145/TCP/GCM default port	14151/TCP/GCM default DNS	14152/TCP/GCM default messenger
14147/TCP/GCM slave port	14141/TCP/HAD	14142/TCP/engine test
14143/TCP/GAB simulator	14144/TCP/notifier	14149/TCP/VCSTD
14153/TCP/Simulator	14150/TCP/Command server	14154/TCP/SimServer
14155/TCP/WAC	512/TCP/exec	513/TCP/remote login
514/UDP/syslogd	514/TCP/RSH	

Installing and configuring VCS on ESX Servers

This chapter contains the following topics:

- [About installing and configuring VCS](#)
- [Installing and configuring VCS 5.1 MP2](#)
- [Installing VCS 5.1 MP2](#)
- [Configuring VCS 5.1 MP2](#)
- [Optional installvcs program actions](#)
- [Checking licensing information on the system](#)
- [Updating product licenses using vxlicinst](#)
- [About installvcs command options](#)
- [About the uninstallvcs program](#)
- [Uninstalling VCS](#)

About installing and configuring VCS

You install VCS on ESX Servers. You can install VCS on clusters of up to 16 systems. The `installvcs` program uses `ssh` to install by default.

- For a fresh installation on systems running ESX Server 3.5
See [“For ESX Server 3.5, applying the pre-patch and starting VCS installation”](#) on page 55.
- For a fresh installation on systems running ESX Server 3.0.1 or 3.0.2
See [“For ESX Server 3.0.1 or 3.0.2, starting VCS installation”](#) on page 55.

You can access the `installvcs` program from the command line or through the `installvcs` program. The `installvcs` program is interactive and enables you to install, configure, license, and start VCS and its options on multiple nodes.

Note: This release does not support configuring clusters in the secure mode. Do not configure the Symantec Product Authentication server while installing or configuring VCS.

Installing and configuring VCS 5.1 MP2

You install and configure VCS 5.1 MP2 from different media sources. Make sure that you have either (or both) of the following options available for installation:

- The discs for VCS 5.1 and VCS 5.1 MP2
- The disc for VCS 5.1 and the gzipped tarfile for VCS 5.1 MP2

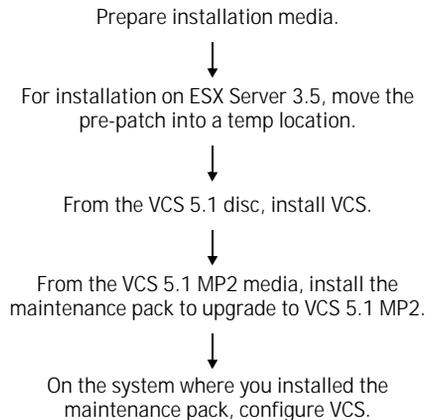
Pre-patching servers running ESX Server 3.5

Systems that run ESX Server 3.5 require a VCS patch before installation. You can find the `esx35patch.pl` patch file in the `cluster_server/scripts` directory of either the gzipped tarfile or on disc.

VCS installation flow overview

Figure 4-2 illustrates VCS installation flow and media use.

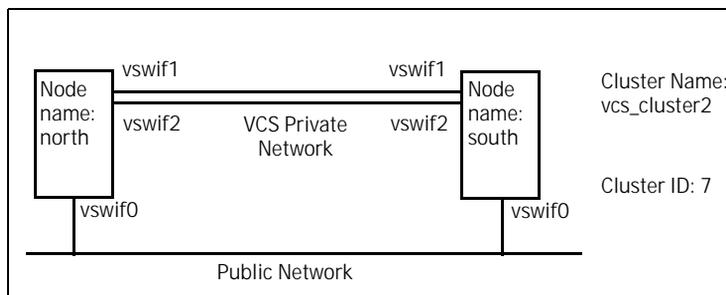
Figure 4-1 Flow for installing VCS



VCS installation example

Figure 4-2 illustrates an example VCS installation that is used in the rest of this chapter.

Figure 4-2 An example of a VCS installation on a two-node cluster



In this installation setup, you install VCS on two systems, install all VCS RPMs, and configure all optional features. For this setup, the DNS resolvable system names where you install VCS are north and south. The cluster's name is vcs_cluster2 and the cluster's ID is 7.

For the purposes of this example, virtual interface vswif1 is mapped to physical interface vmnic1. Similarly vswif2 is mapped to vmnic2.

See ["Selecting virtual interfaces for use by the private network"](#) on page 42.

Checking the systems for installation

Before beginning the installation of VCS software, you can check the readiness of the systems where you plan to install VCS. The command to start the pre-installation check is:

```
installvcs -precheck sysA sysB ...
```

To check the systems

- 1 Navigate to the folder containing the installvcs program.

```
# cd /cdrom/cluster_server
```
- 2 Start the pre-installation check:

```
# ./installvcs -precheck north south
```

The program proceeds in a non-interactive mode, examining the systems for licenses, RPMs, disk space, and system-to-system communications.
- 3 Review the output as the program displays the results of the check and saves the results to a log file.

See ["About installvcs command options"](#) on page 80.

Starting the software installation

Start the installation for ESX Server 3.5 or for ESX Server 3.0.1 or 3.0.2.

For ESX Server 3.5, applying the pre-patch and starting VCS installation

You must perform these steps if you are installing VCS onto a system that runs ESX Server 3.5.

The command to install the patch and to start VCS installation is:

```
./installvcs -require esx35patch.pl -installonly sysA sysB
```

To install the pre-patch and start VCS installation on ESX Server 3.5

- 1 Confirm that you are logged in as the superuser.
- 2 From the 5.1 MP2 disc, locate the `esx35patch.pl` file. It is in the `cluster_server/scripts` directory.

```
# cd cluster_server/scripts
```
- 3 Copy the `esx35patch.pl` file to a temporary location on the system where you intend to install VCS for ESX.
- 4 Mount the 5.1 disc.
- 5 Navigate to the folder that contains the `installvcs` program.

```
# cd /cluster_server
```
- 6 Start the patch and software installation. Use the `-installonly` option for the command to apply the software without configuration. You only perform the configuration steps after you apply the VCS 5.1 MP2 patch.

```
# ./installvcs -require /tmp/esx35patch.pl -installonly north south
```

For ESX Server 3.0.1 or 3.0.2, starting VCS installation

You must perform these steps if you plan to install VCS onto a system that runs ESX Server 3.0.1 or 3.0.2.

The command to start VCS installation is:

```
./installvcs -installonly sysA sysB
```

To install VCS 5.1

- 1 Confirm that you are logged in as the superuser and that you have mounted the 5.1 product disc.
- 2 Navigate to the folder that contains the `installvcs` program.

```
# cd /cluster_server
```

- 3 Use the `-installonly` option for the command to apply the software without configuration. You only perform the configuration steps after you apply the VCS 5.1 MP2 patch.

```
# ./installvcs -installonly north south
```

Specifying systems for installation

The `installvcs` program prompts you for the system names where you want to install VCS. It then performs an initial system check.

To specify system names for installation

- 1 Enter the DNS resolvable names of the systems where you want to install VCS. Note that these names are case sensitive.
Enter the system names separated by spaces on which to install VCS: **north south**
For a single node installation, enter one name for the system.
See [“Starting the installvcs program for the single node cluster”](#) on page 119.
- 2 Review the output as the `installvcs` program verifies the systems, the `installvcs` program performs the following actions:
 - Checks that the local node that runs the `installvcs` program can communicate with remote nodes.
If the `installvcs` program finds `ssh` binaries, it confirms that `ssh` can operate without requests for passwords or passphrases.
See [“Enabling password free SSH-communication between systems”](#) on page 44.
 - Makes sure the systems use the proper operating system.
 - Checks whether a previous version of VCS is installed.

Licensing VCS

The `installvcs` program checks whether VCS license keys are currently in place on each system. If the license keys are absent, the `installvcs` program prompts you for them.

See [“Checking licensing information on the system”](#) on page 78.

To license VCS

- 1 Review the output as the utility checks system licensing and installs the licensing RPM.
- 2 Enter the license key for Veritas Cluster Server as the `installvcs` program prompts for each node.

```
Enter a VCS license key for north: [?] XXXX-XXXX-XXXX-XXXX-XXX  
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on north  
VCS license registered on north
```

3 Enter keys for additional product features.

```
Do you want to enter another license key for north? [y,n,q,?]  
(n) y
```

```
Enter a VCS license key for north: [?] XXXX-XXXX-XXXX-XXXX-XXX  
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on north
```

```
Do you want to enter another license key for north? [y,n,q,?]  
(n)
```

4 Review the output as the installvcs program registers the license key on the other nodes. Enter keys for additional product features on the other nodes when the installvcs program prompts you.

```
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on south  
VCS license registered on south
```

```
Do you want to enter another license key for south? [y,n,q,?]  
(n)
```

Choosing VCS RPMs

The installvcs program checks for any previously installed RPMs. Based on your choice, it then installs only the required RPMs or all the RPMs.

To install VCS RPMs

- 1 Review the output as the installvcs program checks previously installed RPMs.
- 2 Select option 1 or 2 to install either all or required RPMs.
- 3 Review the list of RPMs that the installvcs program installs on each node. If the current version of a RPM is on a system, the installvcs program removes it from the RPM installation list for the system.

Installing VCS 5.1 MP2

Perform the following steps to install VCS 5.1 MP2. At the end of the installation, do not perform the suggested system shut down and reboot, instead follow the instructions provided.

To install VCS 5.1 MP2

- 1 Log in as superuser on one of the systems.
- 2 Insert the disc containing the VCS 5.1 MP2 software into the disc drive of one of the cluster nodes.
- 3 Mount the disc on a suitable mount point and change directory to the location of the disc mount.
- 4 Start the VCS 5.1 MP2 installation.

```
# ./installmp
```

Running the installmp script stops VCS.
- 5 The installer prompts you for the system names where you want to install the maintenance pack.

```
Enter the system names separated by spaces on which to install  
MP1: north south
```
- 6 After an initial system check, the setup program is ready to install VCS 5.1 MP2 and seeks confirmation.

```
Are you sure you want to install MP2? [y,n,q] (y)
```

Enter **y** to begin installation.
Do not reboot after VCS 5.1 MP2 is installed.

Configuring VCS 5.1 MP2

Run the `installvcs` program with the `-configure` option from one of the ESX Server systems where you installed VCS.

To configure VCS 5.1 MP2

- 1 On one of the ESX Server systems where you installed VCS, run the `installvcs` program with the `-configure` option and the node names.

```
# /opt/VRTS/install/installvcs -configure north south
```
- 2 Review the configuration requirements that the `installvcs` program lists. Note the location of the log file.
- 3 Answer the licensing prompts.
- 4 Enter a unique cluster name and cluster ID.
Enter the unique cluster name: [?] **vcs_cluster2**
Enter the unique Cluster ID number between 0-65535: [b,?] **7**
Review the interfaces that are available on the first system as the `installvcs` program discovers and reports them.
- 5 Choose from the virtual or physical interfaces for the private heartbeat links. Note that before you select a virtual interface, make sure that it is mapped to the correct physical interface.
You must not enter the interface that is used for the public network (typically `vswif0`.)
Enter the NIC for the first private heartbeat NIC on north:
[b,?] **vmnic1**
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat NIC on north:
[b,?] **vmnic2**
Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n)
Do you want to configure an additional low priority heartbeat link? [y,n,q,b,?] (n)
- 6 Choose whether to use the same interface on all nodes.
Are you using the same NICs for private heartbeat links on all systems? [y,n,q,b,?] (y)
 - If you want to use the same interfaces for private heartbeat links on all nodes, make sure the same interfaces are available on each system and enter **y**.
 - If the interfaces are not the same on all nodes, enter **n**. You must configure interface information for each node.

- 7 Verify and confirm the information that the `installvcs` program summarizes. For the specified physical interfaces, the `installvcs` program creates virtual interfaces that are mapped to the physical interfaces.

Adding VCS users

On systems operating under an English locale, now add VCS users.

To add VCS users

- 1 Review the required information to add VCS users.
- 2 Reset the password for the Admin user, if necessary.

```
Do you want to set the password for the Admin user
(default password='password')? [y,n,q] (n) y
```



```
Enter New Password:*****
Enter Again:*****
```
- 3 To add a user, enter **y** at the prompt.

```
Do you want to add another user to the cluster? [y,n,q] (y)
```
- 4 Enter the user's name, password, and level of privileges.

```
Enter the user name: [?] smith
Enter New Password:*****
```



```
Enter Again:*****
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [?] a
```
- 5 Enter **n** at the prompt if you have finished adding users.

```
Would you like to add another user? [y,n,q] (n)
```
- 6 Review the summary of the newly added users and confirm the information to proceed.

Configuring cluster connector

If you configured the Cluster Management Console management server to centrally manage this cluster, you can now configure cluster connector. If a firewall exists between the management server and this cluster, then you must configure cluster connector. Make sure you meet the prerequisites to configure cluster connector.

To configure cluster connector

- 1 Review the information to configure Veritas Cluster Server Management Console.
- 2 Choose whether to configure cluster connector or not. Do one of the following:

- To configure cluster connector on the systems, press the **Enter** key to accept **y** (yes) as your default answer.
Do you want this cluster to be managed by a management server? Enter 'y' if you have set up a management server.
[y,n,q] (y)
 - To skip configuring cluster connector and advance to configuring Cluster Management Console for local cluster management, enter **n**.
- 3 If you choose to configure the cluster connector, review the required information to configure cluster connector.
 - 4 Enter the management server network address for the Veritas Cluster Server Management Console.
Enter the network address used by the management server [?]
(north) **mgmtserver1.symantecexample.com**
 - 5 Verify and confirm the management server information.

Configuring Veritas Cluster Server Management Console

Configure Veritas Cluster Server Management Console to use it.

To configure the Cluster Management Console

- 1 Choose whether to configure the Cluster Management Console or not. Do one of the following:
 - To configure the Veritas Cluster Server Management Console on the systems, press **Enter**.
Do you want this cluster to be managed by a management server? Enter 'y' if you have set up a management server.
[y,n,q] (y)
 - To skip configuring the Veritas Cluster Server Management Console and advance to configuring SMTP, enter **n**.
See “[Configuring SMTP email notification](#)” on page 62.
- 2 Confirm whether you want to use the discovered public NIC on the first system. Do one of the following:
 - If the discovered NIC is the one to use, press **Enter**.
 - If you want to use a different NIC, type the name of a NIC to use and press **Enter**.
Active NIC devices discovered on north: vswif0
Enter the NIC for Cluster Management Console to use on north:
[b,?] (vswif0)

- 3 Confirm whether you want to use the same public NIC on all nodes. Do one of the following:

- If all nodes use the same public NIC, enter **y**.
- If unique NICs are used, enter **n** and enter a NIC for each node.

Is vswif0 to be the public NIC used by all systems [y,n,q,b,?] (y)

- 4 Enter the virtual IP address for the Veritas Cluster Server Management Console.

Enter the Virtual IP address for Cluster Management Console: [b,?] **10.10.12.1**

- 5 Confirm the default netmask or enter another one:

Enter the netmask for IP 10.10.12.1: [b,?] (255.255.240.0)

- 6 Verify and confirm the Veritas Cluster Server Management Console information.

Cluster Management Console verification:

```
NIC: vswif0
IP: 10.10.12.1
Netmask: 255.255.240.0
```

Is this information correct? [y,n,q] (y)

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP e-mail services. You need to provide the SMTP server name and email addresses of people to be notified. Note that it is also possible to configure notification after installation. Refer to the *Veritas Cluster Server User's Guide* for more information.

To configure SMTP email notification

- 1 Review the required information to configure the SMTP email notification.
- 2 Specify whether you want to configure the SMTP notification. Do one of the following:

- To configure SMTP notification, press **Enter**.

Do you want to configure SMTP notification? [y,n,q] (y) **y**

- To skip configuring SMTP notification and advance to configuring SNMP notification, enter **n**.

See See ["Configuring SNMP trap notification"](#) on page 63..

3 Provide information to configure SMTP notification.

- Enter the SMTP server's host name.

Enter the domain-based hostname of the SMTP server
 (example: smtp.yourcompany.com): [b,?] **smtp.example.com**

- Enter the email address of each recipient.

Enter the full email address of the SMTP recipient
 (example: user@yourcompany.com): [b,?] **smith@example.com**

- Enter the minimum security level of messages to be sent to each recipient.

Enter the minimum severity of events for which mail should be sent to ozzie@example.com [I=Information, W=Warning, E=Error, S=SevereError]: [b,?] **w**

4 Add more SMTP recipients, if necessary.

- If you want to add another SMTP recipient, enter **y** and provide the required information at the prompt.

Would you like to add another SMTP recipient? [y,n,q,b] (n) **y**

Enter the full email address of the SMTP recipient
 (example: user@yourcompany.com): [b,?] **jones@example.com**

Enter the minimum severity of events for which mail should be sent to harriet@example.com [I=Information, W=Warning, E=Error, S=SevereError]: [b,?] **E**

- If you do not want to add, answer **n**.

Would you like to add another SMTP recipient? [y,n,q,b] (n)

5 Verify and confirm the SMTP notification information.

SMTP Address: smtp.example.com
 Recipient: ozzie@example.com receives email for Warning or higher events
 Recipient: harriet@example.com receives email for Error or higher events

Is this information correct? [y,n,q] (y)

Configuring SNMP trap notification

You can choose to configure VCS to send event notifications to SNMP management consoles. You need to provide the SNMP management console name to be notified and message severity levels. Note that it is also possible to configure notification after installation. Refer to the *Veritas Cluster Server User's Guide* for more information.

To configure the SNMP trap notification

- 1 Review the required information to configure the SNMP notification feature of VCS.
- 2 Specify whether you want to configure the SNMP notification. Do one of the following:
 - To configure SNMP notification, press **Enter**.
Do you want to configure SNMP notification? [y,n,q] (y)
 - To skip configuring SNMP notification and advance to configuring global clustering option, enter **n**.
See “[Configuring global clusters](#)” on page 65..
- 3 Provide information to configure SNMP trap notification.
 - Enter the SNMP trap daemon port.
Enter the SNMP trap daemon port: [b,?] (162)
 - Enter the SNMP console system name.
Enter the SNMP console system name: [b,?] **west**
 - Enter the minimum security level of messages to be sent to each console.
Enter the minimum severity of events for which SNMP traps should be sent to saturn [I=Information, W=Warning, E=Error, S=SevereError]: [b,?] **E**
- 4 Add more SNMP consoles, if necessary.
 - If you want to add another SNMP console, enter **y** and provide the required information at the prompt.
Would you like to add another SNMP console? [y,n,q,b] (n) **y**
Enter the SNMP console system name: [b,?] **jupiter**
Enter the minimum severity of events for which SNMP traps should be sent to jupiter [I=Information, W=Warning, E=Error, S=SevereError]: [b,?] **S**
 - If you do not want to add, answer **n**.
Would you like to add another SNMP console? [y,n,q,b] (n)
- 5 Verify and confirm the SNMP notification information.
SNMP Port: 162
Console: saturn receives SNMP traps for Error or higher events
Console: jupiter receives SNMP traps for SevereError or higher events

Is this information correct? [y,n,q] (y)

Configuring global clusters

You can configure global clusters to connect clusters at separate locations and enable wide-area failover and disaster recovery. Note that you must have entered a valid license key for VCS global clusters.

To configure the global cluster option

- 1 Review the required information to configure the global cluster option. See [“Preparing your VCS installation and configuration information”](#) on page 47.
- 2 Specify whether you want to configure the global cluster option. Do one of the following:
 - To configure global cluster option, press **Enter**.
Do you want to configure the Global Cluster Option? [y,n,q]
(y)
 - To skip configuring global cluster option enter **n**.
- 3 Provide information to configure the Global Cluster option.
If you configured the Veritas Cluster Server Management Console to manage this cluster locally, the `installvcs` program discovers and displays the virtual IP address and netmask used by the Veritas Cluster Server Management Console. You can use the same virtual IP address and netmask. Do one of the following:
 - If you want to use the default values, press **Enter**.
 - If you do not want to use the default value, enter another IP address. The `installvcs` program prompts you for a NIC and value for the netmask.
Enter the Virtual IP address for Global Cluster Option:
[b,?] (10.10.12.1)
- 4 Verify and confirm the configuration of the global cluster.
Global Cluster Option configuration verification:

NIC: vswif0
IP: 10.10.12.1
Netmask: 255.255.240.0

Matching Cluster Management Console Virtual IP configuration

Is this information correct? [y,n,q] (y)

ESX Server network configuration

Configure the ESX Server network.

To configure the ESX network configuration

- 1 When asked if you want to continue with firewall installation, answer **y** or accept the default.
- 2 Enter a name for a new ESX user.
- 3 Enter a password for VCS.
- 4 The program asks if want to use the default SSL certificate that is used to generate the keystore. Answer **y** or accept the default.
Enter pathname of a valid ssl cert on north
(etc/vmware/ssl/rui.crt):
- 5 Enter a password for the keystore for the first system.
Enter keystore password, not less than 6 characters:
Re-type the password if required:
Enter the keystore password again:
- 6 Enter the password for the root user on the first system.
Enter the password for user root on host north:
Re-type the password:
Enter the password for user root again:
- 7 If you have the same root password on all the nodes in the cluster, answer **y**.
The installvcs program generates the security credentials across all the systems in the cluster.
Do you want to use the same root and keystore password on the other hosts? [y,n,q] (y) **y**
If you use different root passwords on any of the nodes in the cluster, installvcs prompts you for those.
Installvcs now generates the security credentials for VCS agents.
Generating security credentials for VCS agents Done

Starting VCS

Start VCS and its components on each system.

To start VCS

- ◆ Confirm to start VCS and its components on each node.
Do you want to start Veritas Cluster Server for ESX VMware processes now? [y,n,q] (y) **y**

Completing the installation

After the installation completes successfully, the `installvcs` program creates summary, log, and response files. The files provide useful information that can assist you with the installation and can also assist future installations.

[Table 4-1](#) specifies the files created at the end of the installation. Review the location of the installation log files, summary file, and response file that the `installvcs` program displays.

Table 4-1 File description

File	Description
summary file	<ul style="list-style-type: none"> ■ Lists RPMs installed on each system. ■ Describes the cluster and its configured resources. ■ Provides information for managing the cluster.
log file	Details the entire installation.
response file	<p>Contains configuration information that can be used to perform secure or unattended installations on other systems.</p> <p>See “Example response file” on page 72.</p>

Enforcing compatibility between VCS and VMware features (DRS and ESX Server maintenance mode)

During VCS installation, the VCS installation program deploys the ESXHost agent on each node in the cluster. The program also creates a service group called the ESXHostServiceGroup. This ESXHostServiceGroup is an internal-use-only service group.

The ESXHost agent enforces a two-way compatibility between VCS and VMware’s DRS and maintenance mode functions. If you bring a ESX Server into VMware maintenance mode to perform system maintenance tasks, the ESXHost agent evacuates all VCS service groups on that ESX Server to other ESX Server nodes in the VCS cluster. VCS then performs a system freeze for the cluster. The agent also provides reverse compatibility—if it detects that a VCS cluster is frozen, and no active virtual machines are on the host, it puts the ESX Server node into VMware maintenance mode.

For the ESXHost agent to work properly, ensure that there is an exact node list in the definitions between the VCS and VMware clusters. Make sure that one cluster does not have more nodes or fewer nodes than the other cluster.

See [“VMware ESX components and configuration requirements”](#) on page 35.

Copying the Implementation Guide to each node

After you install VCS, copy the PDF version of this guide from the installation disc (docs/vcs_implementation.pdf) to the directory /opt/VRTS on each node to make it available for reference.

Installing the Cluster Manager (Java Console)

Note: While you can install the Cluster Manager (Java Console) separately, Symantec recommends that you install Veritas Virtualization Manager (VVM) instead. VVM comes pre-loaded with the Java Console.

You can administer VCS using the Cluster Manager (Java Console), which is a Java-based graphical user interface. After VCS has been installed, install the Java Console on a Windows NT or Windows 2000 Professional system. The system where you run the Java Console can be in the cluster or a client. You can administer each system in the cluster remotely if you install it as a client.

When installing the Java Console, make sure a printer is configured to that system. On a system without a configured printer, printing from the online JavaHelp could cause the Java Console to hang.

For information about using the Cluster Manager and the Configuration Editor components of the Java Console, see the applicable chapter in the *Veritas Cluster Server User's Guide*.

Hardware requirements for the Cluster Manager (Java Console)

Symantec recommends a Pentium III at 400 MHz, with 256 MB RAM.

Minimum hardware requirements follow:

- Pentium II 300 megahertz CPU
- 256 megabytes RAM
- A display capable of at least 800 x 600 resolution
- 8-bit color display depth
- Graphics card capable of 2D images
- Approximately 40 MB of free hard drive space

The version of the Java™ 2 Runtime Environment (JRE) requires 32 megabytes of RAM and is supported on Intel Pentium platforms running the Linux kernel v 2.2.12 and glibc v2.1.2-11 (or later). Symantec recommends using 48 megabytes of RAM, 16-bit color mode, and KDE and KWM window managers used in conjunction with displays set to local hosts.

Installing the Java Console on a Windows workstation

You can install the Cluster Manager (Java Console) on a Windows NT workstation or a Windows 2000 Professional Workstation to administer the cluster.

To install the Java Console on a Windows system

- 1 Insert the software disc with the VCS software into a drive on your Windows system.
- 2 Using Windows Explorer, select the disc drive.
- 3 Go to \windows\ClusterManager.
- 4 Open the language folder of your choice, for example EN.
- 5 Double-click setup.exe.
- 6 The Veritas Cluster Manager Install Wizard guides you through the installation process.

Verifying the cluster after installation

You must verify that your cluster is operating properly after the installation. See [“Verifying VCS on ESX Servers”](#) on page 87.

Optional installvcs program actions

[Table 4-2](#) specifies the optional actions that the installvcs program can perform.

Table 4-2 Optional installvcs program features

Optional action	Reference
Check the systems to verify that they meet the requirements to install VCS.	See “Checking the systems for installation” on page 54.
Install VCS RPMs without configuring VCS.	See “Installing VCS using installonly option” on page 70.
Configure or reconfigure VCS when VCS RPMs are already installed.	See “Configuring VCS using configure option” on page 70.
Perform secure installations using values stored in a configuration file.	See “Performing VCS installation in a secure environment” on page 70.
Perform automated installations using values stored in a configuration file.	See “Performing automated installations” on page 72.

Installing VCS using `installonly` option

In certain situations, you may choose to install the VCS RPMs on a system before it is ready for cluster configuration. During such situations, use the `installvcs -installonly` option. The installation program licenses and installs VCS RPMs on the systems entered without creating any VCS configuration files.

Configuring VCS using `configure` option

If you installed VCS and did not choose to configure VCS immediately, use the `installvcs -configure` option to configure VCS when you are ready for cluster configuration. The `installvcs` program prompts for cluster information, and creates VCS configuration files without performing installation.

The `-configure` option can be used to reconfigure a VCS cluster. VCS must not be running on systems when you reconfigure it.

Performing VCS installation in a secure environment

In secure enterprise environments, ssh or rsh communication is not allowed between systems. In such cases, the `installvcs` program can install and configure VCS only on systems with which it can communicate—most often the local system only. When installation is complete, a “response” file is created.

See “[Example response file](#)” on page 72.

Note that a response file generated by the `installvcs` program contains descriptions and explanations of the variables and their values. By copying this file to the other systems in the cluster and editing it to reflect the current local system, you can use the installation program with the `-responsefile` option to install and configure VCS identically on each system without being prompted.

To use `installvcs` in a secure environment

- 1 On one node in the cluster, start VCS installation using the `installvcs` program.
See “[Starting the software installation](#)” on page 55.
- 2 Review the output as the `installvcs` program performs the initial system checks.
The `installvcs` program detects the inability to communicate between systems.
- 3 Press the **Enter** key to install VCS on one system and create a response file with which you can install on other systems.

```
Would you like to install Cluster Server on systems north only  
and create a responsefile for systems south? [y,n,q] (y)
```

- 4 Enter all cluster information. Proceed with the installation and configuration tasks.
See “[Installing and configuring VCS 5.1 MP2](#)” on page 53.
The installvcs program installs and configures VCS on systems where communication is possible.
- 5 After the installation is complete, review the installvcs program report.
The installvcs program stores the response file within the file `/opt/VRTS/install/logs/installvcs-universaluniqueidentifier/installvcs-universaluniqueidentifier.response`.
- 6 If you start VCS before VCS is installed and started on all nodes in the cluster, you see the output similar to:

```
VCS:11306:Did not receive cluster membership, manual intervention may be needed for seeding
```
- 7 Using a method of your choice (for example, by using NFS, ftp, or a floppy disk), place a copy of the response file in a directory such as `/tmp` on the next system to install VCS.
- 8 On the next system in your cluster, edit the response file.
For the variables described in the example, change the name of the system to reflect the current local system:

```
.  
$CFG{SYSTEMS} = ["east "];  
.  
.  
$CFG{KEYS}{east} = ["XXXX-XXXX-XXXX-XXXX-XXXX-XXX "];  
.
```


For demo or site licenses, the license key need not be changed.
- 9 On the next system:
 - Mount the product disc.
 - Start the software installation using the `installvcs -responsefile` option.

```
# ./installvcs -responsefile /tmp/installvcs-uui.response
```


Where *uui* is the Universal Unique Identifier that the installvcs program automatically assigned to the response file.
See “[Starting the software installation](#)” on page 55.
- 10 Repeat [step 7](#) through [step 9](#) until VCS has been installed on all nodes in the cluster.

Performing automated installations

Using `installvcs` program with the `-responsefile` option is useful not only for installing and configuring VCS within a secure environment, but for conducting unattended installations to other clusters as well. Typically, you can use the response file generated during the installation of VCS on one cluster to install VCS on other clusters. You can copy the file to a system in another cluster and manually edit the file to contain appropriate values.

Assuming the systems are set up and meet the requirements for installation, you can perform unattended installation from one of the cluster systems where you have copied the response file.

To perform unattended installation

- 1 Navigate to the folder containing the `installvcs` program.

```
# cd /mnt/cdrom/cluster_server
```
- 2 Start the installation from one of the cluster systems where you have copied the response file.

```
# ./installvcs -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Syntax used in response file

The syntax of Perl statements included in the response file varies, depending on whether "Scalar" or "List" values are required by the variables.

For example,

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value", "value", "value"];
```

Example response file

The example response file resembles the file created by `installvcs` after the example VCS installation. It is a modified version of the response file generated on `vcs_cluster2` that you can use to install VCS on `vcs_cluster3`. Review the variables required for installation.

See [Table 4-3, "Response file variables."](#)

```
#  
# installvcs configuration values:  
#  
$CPI::CFG{AT_ROOTDOMAIN}="root\@east.symantecexample.com";  
$CPI::CFG{CMC_CC_CONFIGURED}=1;  
$CPI::CFG{CMC_CLUSTERID}{east}=1146235600;
```

```
$CPI::CFG{CMC_MSADDR}{east}="mgmtserver1";
$CPI::CFG{CMC_MSADDR}{west}="mgmtserver1";
$CPI::CFG{CMC_MS_ROOT_HASH}="758a33dbd6fae751630058ace3dedb54e5
62fe98";
$CPI::CFG{CMC_SERVICE_PASSWORD}="U2FsdGVkX18vE5tn0hTSWwodThACc+
rX";
$CPI::CFG{ENCRYPTED}="U2FsdGVkX1+k2DHKvcnW7b6vrVghdh+zW4G0WFj5I
JA=";
$CPI::CFG{KEYS}{east}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{KEYS}{west}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{OBC_IGNOREWARNINGS}=0;
$CPI::CFG{OBC_MODE}="STANDALONE";
$CPI::CFG{OPT}{INSTALL}=1;
$CPI::CFG{OPT}{NOEXTRAPKGS}=1;
$CPI::CFG{OPT}{RSH}=1;
$CPI::CFG{SYSTEMS}=[ qw(east west) ];
$CPI::CFG{UPI}="VCS";
$CPI::CFG{VCS_ALLOWCOMMS}="Y";
$CPI::CFG{VCS_CLUSTERID}=13221;
$CPI::CFG{VCS_CLUSTERNAME}="vcs_cluster3";
$CPI::CFG{VCS_CSGNETMASK}="255.255.240.0";
$CPI::CFG{VCS_CSGNIC}{ALL}="vswif0";
$CPI::CFG{VCS_CSGVIP}="10.10.12.1";
$CPI::CFG{VCS_LLTLINK1}{east}="vswif1";
$CPI::CFG{VCS_LLTLINK1}{west}="vswif1";
$CPI::CFG{VCS_LLTLINK2}{east}="vswif2";
$CPI::CFG{VCS_LLTLINK2}{west}="vswif2";
$CPI::CFG{VCS_SMTPRECP}=[ qw(earnie@symantecexample.com) ];
$CPI::CFG{VCS_SMTPRSEV}=[ qw(SevereError) ];
$CPI::CFG{VCS_SMTPSERVER}="smtp.symantecexample.com";
$CPI::CFG{VCS_SNMPCONS}=[ qw(neptune) ];
$CPI::CFG{VCS_SNMPCSEV}=[ qw(SevereError) ];
$CPI::CFG{VCS_SNMPPORT}=162;
```

Response file variable definitions

Table 4-3 lists the variables used in the response file and their definitions. Note that while some variables are labeled as required and others as optional, some of the optional variables, if used, make it necessary to define other optional variables. For example, all variables related to the cluster service group (CSGNIC, CSGVIP, and CSGNETMASK) must be defined if any are defined. The same is true for the SMTP notification (SMTPSERVER, SMTPRECP, and SMTPRSEV), SNMP trap notification (SNMPPORT, SNMPCONS, and SNMPCSEV), and the Global Cluster Option (CGONIC, GCOVIP, and GCONETMASK).

Table 4-3 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
\$CPI::CFG{OPT}{INSTALL}	Scalar	Req'd	List of systems where VCS must be installed and configured.
\$CPI::CFG{OPT}{INSTALLONLY}	Scalar	Opt'l	List of systems where VCS RPMs must be installed. Configuration can be performed at a later time using the <code>-configure</code> option.
\$CPI::CFG{SYSTEMS}	List	Req'd	List of systems on which the product is to be installed, uninstalled, or configured.
\$CPI::CFG{SYSTEMSCFG}	List	Opt'l	List of systems to be recognized in configuration if secure environment prevents all systems from being installed at once.
\$CPI::CFG{UPI}	Scalar	Req'd	Defines the product to be installed, uninstalled, or configured.
\$CPI::CFG{OPT}{KEYFILE}	Scalar	Opt'l	Defines the location of an ssh keyfile that is used to communicate with all remote systems.
\$CPI::CFG{OPT}{LICENSE}	Scalar	Opt'l	Licenses VCS only.

Table 4-3 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
\$CPI::CFG{OPT}{NOLIC}	Scalar	Opt'l	installs the product without any license.
\$CPI::CFG{AT_ROOTDOMAIN}	List	Opt'l	Defines the name of the system where the root broker is installed.
\$CPI::CFG{OPT}{PKGPATH}	Scalar	Opt'l	Defines a location, typically an NFS mount, from which all remote systems can install product depots. The location must be accessible from all target systems.
\$CPI::CFG{OPT}{TMPPATH}	Scalar	Opt'l	Defines the location where a working directory is created to store temporary files and depots needed during the install. The default location is /var/tmp.
\$CPI::CFG{OPT}{RSH}	Scalar	Opt'l	Defines that rsh must be used instead of ssh as the communication method between systems.
\$CPI::CFG{DONOTINSTALL} {RPM}	List	Opt'l	Instructs the installation to not install the optional RPMs designated in the list.
\$CPI::CFG{DONOTREMOVE} {RPM}	List	Opt'l	Instructs the uninstallation to not remove the optional RPMs designated in the list.
\$CPI::CFG{VCS_CLUSTERNAME}	Scalar	Req'd	Defines the name of the cluster.
\$CPI::CFG{VCS_CLUSTERID}	Scalar	Req'd	An integer between 0 and 65535 that uniquely identifies the cluster.
\$CPI::CFG{KEYS}{SYSTEM}	Scalar	Opt'l	List of keys to be registered on the system.

Table 4-3 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
\$CPI::CFG{OPT_LOGPATH}	Scalar	Opt'l	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.
\$CPI::CFG{CONFIGURE}	Scalar	Opt'l	Performs configuration if the RPMs are already installed using the -installonly option.
\$CPI::CFG{VCS_LLLINK#} {SYSTEM}	Scalar	Req'd	Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (LLTLINK1 and LLTLINK2). Up to four LLT links can be configured.
\$CPI::CFG{VCS_LLLINKLOWPRI} {SYSTEM}	Scalar	Opt'l	Defines a low priority heartbeat link. Typically, LLTLINKLOWPRI is used on a public network link to provide an additional layer of communication.
\$CPI::CFG{VCS_CSGNIC}	Scalar	Opt'l	Defines the NIC for Veritas Cluster Server Management Console to use on a system. 'ALL' can be entered as a system value if the same NIC is used on all systems.
\$CPI::CFG{CSGVIP}	Scalar	Opt'l	Defines the virtual IP address to be used by the Veritas Cluster Server Management Console.
\$CPI::CFG{VCS_CSGNETMASK}	Scalar	Opt'l	Defines the Netmask of the virtual IP address to be used by the Veritas Cluster Server Management Console.

Table 4-3 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
\$CPI::CFG{VCS_SMTPSERVER}	Scalar	Opt'l	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for web notification.
\$CPI::CFG{VCS_SMTPRECP}	List	Opt'l	List of full email addresses (example: user@symantecexample.com) of SMTP recipients.
\$CPI::CFG{VCS_SMTPRSEV}	List	Opt'l	Defines minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients.
\$CPI::CFG{VCS_SNMPPORT}	Scalar	Opt'l	Defines the SNMP trap daemon port (default=162).
\$CPI::CFG{VCS_SNMPCONS}	List	Opt'l	List of SNMP console system names
\$CPI::CFG{VCS_SNMPSEV}	List	Opt'l	Defines minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names.
\$CPI::CFG{VCS_GCONIC} {SYSTEM}	Scalar	Opt'l	Defines the NIC for the Virtual IP used for the Global Cluster Option. 'ALL' can be entered as a system value if the same NIC is used on all systems.
\$CPI::CFG{VCS_GCOVIP}	Scalar	Opt'l	Defines the virtual IP address to be used by the Global Cluster Option.

Table 4-3 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
\$CPI::CFG{VCS_GCONETMASK}	Scalar	Opt'l	Defines the Netmask of the virtual IP address to be used by the Global Cluster Option.
\$CPI::CFG{VCS_USERENPW}	List	Opt'l	List of encoded passwords for users
\$CPI::CFG{VCS_USERNAME}	List	Opt'l	List of names of users
\$CPI::CFG{VCS_USERPRIV}	List	Opt'l	List of privileges for users
\$CPI::CFG{OPT}{UNINSTALL}	Scalar	Opt'l	List of systems where VCS must be uninstalled.

Checking licensing information on the system

You can use the `vxlicrep` program to display information about the licenses on a system.

To check licensing information

- 1 Navigate to the folder containing the `vxlicrep` program and enter:


```
# cd /opt/VRTS/bin
# ./vxlicrep
```
- 2 Review the output to determine:
 - The license key
 - The type of license
 - The product for which it applies
 - Its expiration date, if any. Demo keys have expiration dates, permanent keys and site keys do not.

```
License Key           = xxx-xxx-xxx-xxx-xxx
Product Name         = Veritas Cluster Server
Serial Number        = 1249
License Type         = PERMANENT
OEM ID               = 478
```

```
Features :=
  Platform           = VMware ESX
  Version            = 5.1
  Tier                = 0
  Reserved           = 0

Mode              = VCS
```

Updating product licenses using vxlicinst

You can use the `vxlicinst` command to add the VCS license key on each node. If you have VCS already installed and configured and you are using a demo license, you can replace the demo license.

See [“Replacing a VCS demo license with a permanent license”](#) on page 79.

To update product licenses

- ◆ On each node, enter the license key using the command:

```
# cd /opt/VRTS/bin
# ./vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```

Replacing a VCS demo license with a permanent license

When a VCS demonstration key license expires, you can replace it with a permanent license using the `vxlicinst(1)` program.

To replace a demo key

- 1 Make sure you have permissions to log in as root on each of the nodes in the cluster.
- 2 Shut down VCS on all nodes in the cluster:

```
# hstop -all -force
```

This does not shut down any running applications.
- 3 Enter the permanent license key using the following command on *each* node:

```
# cd /opt/VRTS/bin
# ./vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```
- 4 Make sure demo licenses are replaced on all cluster nodes before starting VCS.
- 5 Start VCS on each node:

```
# hstart
```

About installvcs command options

[Table 4-4](#) lists the `installvcs` command options. In addition to the `-precheck`, `-responsefile`, `-installonly`, and `-configure` options, the `installvcs` program has other useful options.

The `installvcs` command usage takes the following form:

```
installvcs [ sysA sysB... ] [ options ]
```

Table 4-4 installvcs options

Option and Syntax	Description
<code>-configure</code>	Configure VCS after using <code>-installonly</code> option to install VCS. See "Configuring VCS using configure option" on page 70.
<code>-enckeyfile</code> <i>encryption_key_file</i>	See the <code>-responsefile</code> and the <code>-encrypt</code> options.
<code>-encrypt password</code>	Encrypt <i>password</i> using the encryption key provided with the <code>-enckeyfile</code> option so that the encrypted password can be stored in response files.
<code>-installonly</code>	Install product RPMs on systems without configuring VCS. See "Installing VCS using installonly option" on page 70.
<code>-installpkgs</code>	Display VCS packages in correct installation order. Output can be used to create scripts for command line installs, or for installations over a network. See the <code>requiredpkgs</code> option.
<code>-keyfile</code> <i>ssh_key_file</i>	Specifies a key file for SSH. The option passes <code>-i ssh_key_file</code> with each SSH invocation.
<code>-license</code>	Register or update product licenses on the specified systems. Useful for replacing demo license.
<code>-logpath log_path</code>	Specifies that <i>log_path</i> , not <code>/opt/VRTS/install/logs</code> , is the location where <code>installvcs</code> log files, summary file, and response file are saved.
<code>-noextrapkgs</code>	Specifies that additional product RPMs such as VxVM and VxFS need not be installed. Note: VCS product upgrades in the future can be simplified if you do not install additional product RPMs.

Table 4-4 installvcs options

Option and Syntax	Description
-nolic	Install product RPMs on systems without licensing or configuration. License-based features or variants are not installed when using this option.
-nooptionalpkgs	Specifies that the optional product RPMs such as man pages and documentation need not be installed.
-nostart	Bypass starting VCS after completing installation and configuration.
-pkgpath <i>pkg_path</i>	Specifies that <i>pkg_path</i> contains all RPMs to be installed by installvcs program on all systems; <i>pkg_path</i> is the complete path of a directory, usually NFS mounted.
-precheck	Verify that systems meet the installation requirements before proceeding with VCS installation. Symantec recommends doing a precheck before installing VCS. See “Checking the systems for installation” on page 54.
-require	Specifies the VCS RPMs to install on an ESX Server.
-requiredpkgs	Displays all required VCS packages in correct installation order. Optional packages are not listed. Output can be used to create scripts for command line installs, or for installations over a network. See <code>installpkgs</code> option.
-responsefile <i>response_file</i> [-enckeyfile <i>encryption_key_file</i>]	Perform automated VCS installation using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. If not specified, the response file is automatically generated as <code>installer.number.response</code> where <i>number</i> is random. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file. The <code>-enckeyfile</code> option and <i>encryption_key_file</i> name are required with the <code>-responsefile</code> option when the response file contains encrypted passwords. See “Performing VCS installation in a secure environment” on page 70. See “Performing automated installations” on page 72.

Table 4-4 `installvcs` options

Option and Syntax	Description
<code>-rsh</code>	Specifies that <code>rsh</code> and <code>rcp</code> are to be used for communication between systems instead of <code>ssh</code> and <code>scp</code> . This option requires that systems be pre-configured such that <code>rsh</code> commands between systems execute without prompting for passwords or confirmations. Note: By default ESX Server 3.x does not contain the <code>rsh</code> RPMs. Add the RPMs for <code>RSH</code> if you want to install VCS with <code>rsh</code> .
<code>-tmp_path tmp_path</code>	Specifies that <code>tmp_path</code> , not <code>/var/tmp</code> , is the working directory for <code>installvcs</code> program. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.

About the `uninstallvcs` program

You can uninstall VCS from all nodes in the cluster or from specific nodes in the cluster using the `uninstallvcs` program. The `uninstallvcs` program does not automatically uninstall VCS high availability agents, but offers uninstallation if proper RPM dependencies on `VRTSvcs` are found.

If `uninstallvcs` program does not remove an high availability agent, see the documentation for the specific high availability agent for instructions on removing it.

Prerequisites

- Before removing VCS from any node in the cluster, you must shut down applications such as the Java Console or any VCS high availability agents that depend on VCS.
- Before removing VCS from fewer than all nodes in a cluster, make sure that no service groups are running on the nodes from which VCS is uninstalled. You must also reconfigure VCS on the remaining nodes. See [“Adding and removing cluster nodes”](#) on page 103.

Uninstalling VCS

The example demonstrates how to uninstall VCS on two nodes: north and south.

Removing VCS RPMs

The program stops the VCS processes that are currently running during the uninstallation process.

To uninstall VCS

- 1 Do one of the following to begin uninstalling:

- If you can execute commands as superuser on the remote nodes in the cluster using `ssh` without supplying a password, run `uninstallvcs` program on one node to uninstall VCS on all nodes in the cluster.
- If you cannot execute commands as superuser on remote nodes in the cluster using `ssh`, you must run `uninstallvcs` program on each node in the cluster.

- 2 Start `uninstallvcs` program.

```
# cd /opt/VRTS/install  
# ./uninstallvcs
```

The program specifies the directory where the logs are created and displays a copyright notice followed by a description of the cluster:

```
VCS configuration files exist on this system with the following  
information:
```

```
Cluster Name: VCS_cluster2  
Cluster ID Number: 7  
Systems: north south  
Service Groups: ClusterService groupA groupB
```

- 3 Answer the prompt to proceed with uninstalling the software.

- To uninstall VCS on all nodes, press **Enter**.
- To uninstall VCS only on specific nodes, enter **n**.

Note that if you enter **n** or if no VCS configuration files are found on the local node, the `uninstallvcs` program prompts you to enter a list of nodes from which you want to uninstall VCS.

```
Do you want to uninstall VCS from these systems? [y,n,q] (y)
```

- 4 Review the output as the `uninstallvcs` program continues to verify communication between systems and check the installations on each system to determine the RPMs to be uninstalled.
- 5 If RPMs, such as high availability agents, are found to be dependent on a VCS RPM, the uninstaller prompts you on whether you want them removed. Enter **y** to remove the designated RPMs.

- 6 Review the uninstaller report after the verification.
- 7 Press Enter to uninstall the VCS RPMs.
`Are you sure you want to uninstall VCS rpms? [y,n,q] (y)`
- 8 Review the output as the uninstaller stops processes, unloads kernel modules, and removes the RPMs.
- 9 Note the location of summary and log files that the uninstaller creates after removing all the RPMs.

Running `uninstallvcs` from the disc

If you need to uninstall VCS after an incomplete installation, or if the `uninstallvcs` program is not available in `/opt/VRTS/install`, you may need to use the `uninstallvcs` program on the VCS 5.1 disc.

Uninstalling the Cluster Management Console cluster connector

Perform the following procedure to remove the cluster connector from UNIX or Windows systems.

Uninstalling cluster connector from UNIX systems

Use this procedure to remove the Cluster Management Console cluster connector from each cluster.

On UNIX systems, the default `installvcs` program option is `-ssh`. If you are performing a remote uninstallation and `ssh` is not enabled, run the `installvcs` program with the `-rsh` option. Otherwise, the `installvcs` program generates an error during the uninstallation.

To uninstall cluster connector from UNIX systems

- 1 Insert the product disc into the drive on the local system. At the command prompt, type:
`./installer [-rsh]`
- 2 Enter **u** to specify uninstallation.
`Enter a Task: [I,C,L,P,U,D,Q,?] u`
The `installvcs` program displays another menu that lists products that are available for uninstallation.
- 3 Enter the menu number that corresponds to Veritas Cluster Server Management Console.
`Select a product to uninstall:nn`
The `installvcs` program presents a description of the product.

- 4 Enter **2** if you are prompted to select a product component. Otherwise, proceed to [step 6](#).

```
Enter '1' to install the Management Server, '2' to install the  
Cluster Connector: [1-2,q] (1) 2
```

The `installvcs` program presents a message stating that it will uninstall cluster connector.

- 5 The uninstall program prompts you for the name of at least one node in the cluster.

```
Enter one system name from each cluster separated by spaces from  
which to uninstall CMC: north
```

Based on this, it determines the nodes from which to uninstall and perform the necessary checks.

Note: If you get an error message similar to this:

```
Checking ssh communication with sysA Enter passphrase for key  
'/.ssh/id_dsa'
```

You must return and set up ssh.

- 6 Enter **y** to verify that the information up to this point is correct.

```
Is this information correct? [y,n,q] (y)
```

The `installvcs` program performs an initial system check of the cluster nodes and checks for installed packages on the cluster nodes. If these checks are satisfactory, the `installvcs` program lists the packages to be uninstalled.

- 7 Enter **y** to verify that you want to uninstall cluster connector.

```
Are you sure you want to uninstall CMC? [y,n,q] (y)
```

The `installvcs` program lists package dependencies and uninstallation progress percentages. If the uninstallation is successful, the program displays this message followed by the location of the uninstallation logs:
Uninstall completed successfully.

Verifying VCS on ESX Servers

This chapter contains the following topics:

- [About verifying the VCS installation](#)
- [Verifying LLT and GAB configuration files](#)
- [Verifying the main.cf file](#)
- [Verifying LLT, GAB, and cluster operation](#)

About verifying the VCS installation

After successful installation, you can inspect the contents of the key configuration files that you have installed and modified during the process. These files reflect the configuration based on the information you supplied.

Verifying LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires `/etc/llthosts` and `/etc/llttab` files. GAB requires `/etc/gabtab` file.

`/etc/llthosts`

The file `llthosts` is a database, containing one entry per system, that links the LLT system ID (in the first column) with the LLT host name. This file is identical on each node in the cluster.

For example, the file `/etc/llthosts` contains entries that resemble:

```
0 north
1 south
```

`/etc/llttab`

The file `llttab` contains information that is derived during installation and used by the utility `lltconfig`. After installation, this file lists the private network links that correspond to the specific system.

For example, the file `/etc/llttab` contains entries that resemble:

```
set-node north
set-cluster 2
link link1 vswif1 vswif1 - ether - -
link link2 vswif2 vswif2 - ether - -
```

If you use MAC address for the network interface, the file `/etc/llttab` contains entries that resemble:

```
set-node north
set-cluster 2
link vswif1 eth-00:04:23:AC:12:C4 - ether - -
link vswif2 eth-00:04:23:AC:12:C5 - ether - -
```

The first line identifies the system. The second line identifies the cluster (the cluster ID you entered during installation). The next two lines, beginning with the `link` command, identify the two network cards used by the LLT protocol.

Refer to the `llttab(4)` manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the `llttab` file.

/etc/gabtab

After you install VCS, the file `/etc/gabtab` contains a `gabconfig` command that configures the GAB driver for use.

The file `/etc/gabtab` contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

where the `-c` option configures the driver for use and `-nN` specifies that the cluster is not formed until at least `N` nodes are ready to form the cluster. By default, `N` is the number of nodes in the cluster.

Note: The use of the `-c -x` option for `/sbin/gabconfig` is not recommended. Gigabit Ethernet controllers do not support the use of `-c -x`.

Verifying the main.cf file

The VCS configuration file `/etc/VRTSvcs/conf/config/main.cf` is created during the installation process.

- See [“Example main.cf, for clusters without the GCO option”](#) on page 90.
- See [“Example main.cf, for clusters with the GCO option”](#) on page 97.

The `main.cf` file contains the minimum information that defines the cluster and its nodes. In addition, the file `types.cf`, which is listed in the include statement, defines the VCS bundled types for VCS resources. The file `types.cf` is in the directory `/etc/VRTSvcs/conf/config` after installation.

Note the following information about the VCS configuration file after installing and configuring VCS:

- The cluster definition includes the cluster information that you provided during the configuration. This includes the cluster name, cluster address, and the names of users and administrators of the cluster.
Notice that the cluster has an attribute `UserNames`. The `installvcs` program creates a user “admin” whose password is encrypted; the word “password” is the default password.
- The `installvcs` program creates the `ClusterService` service group and includes the following:
 - The `ClusterService` service group includes the IP, NIC, and `VRTSWebApp` resources.
 - If you configured Veritas Cluster Server Management Console to manage this cluster locally, the `main.cf` includes the `VRTSWebApp` resource that includes `AppName = cmc` attribute.

- If you configured Cluster Connector so that Veritas Cluster Server Management Console can centrally manage this cluster, the main.cf includes the CMC service group.

The CMC service group includes the ClusterConnectorConfig and Process resources.

- The service group also includes the notifier resource configuration, which is based on your input to installvcs program prompts about notification.
- The installvcs program also creates a resource dependency tree.
- If you installed VCS with the Global Cluster Option, the ClusterService service group contains an Application resource, wac (wide-area connector), whose attributes contain definitions for controlling the cluster in a Global Cluster environment.

For information about managing VCS global clusters:

See the *Veritas Cluster Server User's Guide*.

Refer to the *Veritas Cluster Server User's Guide* and review the chapter on configuration concepts for descriptions and examples of main.cf and types.cf files for ESX systems.

Example main.cf, for clusters without the GCO option

The following sample main.cf is for a cluster. This example is a stub for nine virtual machine service groups, all that is needed is to add applications.

```
Main.cf:

include "types.cf"

cluster vcs (
    UserNames = { admin = IpqIpkPmqLqqOyqKpn }
    Administrators = { admin }
)

system sysA (
)

system sysB (
)

group vm2 (
    SystemList = { sysA = 0, sysB = 0 }
    AutoStartList = { sysA }
)

ESXVirtualMachine vm2_ESX (
```

```
        CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm2/vm2.v
mx"

        VCserver = "648G771.example.com"
        username = v023968
        password = 123123123
        sslcert = "/etc/my.keystore"
        esxhostdomain = "veritas.com"
    )

    VMFSVolume vm2_vmfs (
        Volume = {

"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
    )

    VSwitch vm2_switch (
        VirtualSwitch = vSwitch0
    )

    vm2_ESX requires vm2_vmfs
    vm2_vmfs requires vm2_switch

// resource dependency tree
//
//     group vm2
//     {
//     ESXVirtualMachine vm2_ESX
//     {
//         VMFSVolume vm2_vmfs
//         {
//             VSwitch vm2_switch
//         }
//     }
//     }group vm3 (
SystemList = { sysA = 0, sysB = 1 }
AutoStartList = { sysA }
    )

    ESXVirtualMachine vm3_ESX (
        CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm3/vm3.v
mx"

        VCserver = "648G771.example.com"
        username = v023968
        password = 123123123
        sslcert = "/etc/my.keystore"
        esxhostdomain = "veritas.com"
    )

    VMFSVolume vm3_vmfs (
```

```
        Volume = {
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
    )

    VSwitch vm3_switch (
        VirtualSwitch = vSwitch0
    )

    vm3_ESX requires vm3_vmfs
    vm3_vmfs requires vm3_switch

// resource dependency tree
//
//     group vm3
//     {
//         ESXVirtualMachine vm3_ESX
//         {
//             VMFSVolume vm3_vmfs
//             {
//                 VSwitch vm3_switch
//             }
//         }
//     }
// }

group vm4 (
    SystemList = { sysA = 0 }
    AutoStartList = { sysA }
)

    ESXVirtualMachine vm4_ESX (
        CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm4/vm4.v
mx"
    )

    VMFSVolume vm4_vmfs (
        Volume = {

"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
    )

    VSwitch vm4_switch (
        VirtualSwitch = vSwitch0
    )

    vm4_ESX requires vm4_vmfs
    vm4_vmfs requires vm4_switch

// resource dependency tree
```

```
//
//     group vm4
//     {
//     ESXVirtualMachine vm4_ESX
//         {
//             VMFSVolume vm4_vmfs
//                 {
//                     VSwitch vm4_switch
//                 }
//             }
//         }
//     }

group vm5 (
    SystemList = { sysA = 0 }
    AutoStartList = { sysA }
)

    ESXVirtualMachine vm5_ESX (
        CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm5/vm5.v
mx"
    )

        VMFSVolume vm5_vmfs (
            Volume = {

"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
        )

            VSwitch vm5_switch (
                VirtualSwitch = vSwitch0
            )

vm5_ESX requires vm5_vmfs
vm5_vmfs requires vm5_switch

// resource dependency tree
//
//     group vm5
//     {
//     ESXVirtualMachine vm5_ESX
//         {
//             VMFSVolume vm5_vmfs
//                 {
//                     VSwitch vm5_switch
//                 }
//             }
//         }
//     }
```

```
group vm6 (
    SystemList = { sysA = 0 }
    AutoStartList = { sysA }
)

    ESXVirtualMachine vm6_ESX (
        CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm6/vm6.v
mx"
    )

    VMFSVolume vm6_vmfs (
        Volume = {

"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
    )

    VSwitch vm6_switch (
        VirtualSwitch = vSwitch0
    )

    vm6_ESX requires vm6_vmfs
    vm6_vmfs requires vm6_switch

    // resource dependency tree
    //
    //     group vm6
    //     {
    //         ESXVirtualMachine vm6_ESX
    //         {
    //             VMFSVolume vm6_vmfs
    //             {
    //                 VSwitch vm6_switch
    //             }
    //         }
    //     }
    //

group vm7 (
    SystemList = { sysA = 0 }
    AutoStartList = { sysA }
)

    ESXVirtualMachine vm7_ESX (
        CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm7/vm7.v
mx"
    )

    VMFSVolume vm7_vmfs (
        Volume = {
```

```
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
)

VSwitch vm7_switch (
    VirtualSwitch = vSwitch0
)

vm7_ESX requires vm7_vmfs
vm7_vmfs requires vm7_switch

// resource dependency tree
//
//     group vm7
//     {
//     ESXVirtualMachine vm7_ESX
//         {
//             VMFSVolume vm7_vmfs
//                 {
//                     VSwitch vm7_switch
//                 }
//             }
//         }
//     }

group vm8 (
    SystemList = { sysA = 0 }
    AutoStartList = { sysA }
)

ESXVirtualMachine vm8_ESX (
    CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm8/vm8.v
mx"
)

VMFSVolume vm8_vmfs (
    Volume = {

"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
)

VSwitch vm8_switch (
    VirtualSwitch = vSwitch0
)

vm8_ESX requires vm8_vmfs
vm8_vmfs requires vm8_switch

// resource dependency tree
```

```
//
//     group vm8
//     {
//     ESXVirtualMachine vm8_ESX
//     {
//         VMFSVolume vm8_vmfs
//         {
//             VSwitch vm8_switch
//         }
//     }
//     }

group vm9 (
    SystemList = { sysA = 0 }
    AutoStartList = { sysA }
)

    ESXVirtualMachine vm9_ESX (
        CfgFile =
"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b/vm9/vm9.v
mx"
    )

    VMFSVolume vm9_vmfs (
        Volume = {

"/vmfs/volumes/44bcfdac-f2362ab4-a390-00093d12792b" }
    )

    VSwitch vm9_switch (
        VirtualSwitch = vSwitch0
    )

vm9_ESX requires vm9_vmfs
vm9_vmfs requires vm9_switch

// resource dependency tree
//
//     group vm9
//     {
//     ESXVirtualMachine vm9_ESX
//     {
//         VMFSVolume vm9_vmfs
//         {
//             VSwitch vm9_switch
//         }
//     }
//     }
//     }
```

Example main.cf, for clusters with the GCO option

If you installed VCS with the Global Cluster option, note that the ClusterService group also contains the Application resource, wac, required to control the cluster in a Global Cluster environment.

```

:
:
group ClusterService (
    SystemList = { north = 0, south = 1 }
    UserStrGlobal = "LocalCluster@https://10.182.2.78:8443;"
    AutoStartList = { north, south }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
)
:
:

```

Verifying LLT, GAB, and cluster operation

Before attempting to verify the operation of LLT, GAB, or the cluster, you must:

- Log in to any node in the cluster as superuser.
- Place the VCS command directory in your PATH variable:
export PATH=\$PATH:/sbin:/usr/sbin:/opt/VRTS/bin

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. This command returns information about the links for LLT for the node on which you typed the command. Refer to the `lltstat(1M)` manual page for more information.

Using lltstat -n

In the following example, `lltstat -n` is typed on each node in the cluster:

Node 1

```
# lltstat -n
```

Output resembles:

```

LLT node information:
Node                State                Links

```

```
*0 north      OPEN      2
 1 south      OPEN      2
```

Node 2

```
# lltstat -n
```

Output resembles:

```
LLT node information:
Node          State      Links
 0 north      OPEN       2
*1 south      OPEN       2
```

Note that each node has two links and that each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

Using lltstat -nvv

With LLT configured correctly, the output of `lltstat -n` shows all the nodes in the cluster and two links for each node. If the output shows otherwise, you can use the verbose option of `lltstat`.

For example, type `lltstat -nvv | more` on a node to view additional information about LLT. In the following example, `lltstat -nvv | more` is typed on node north in a two-node cluster:

```
# lltstat -nvv | more
```

Output resembles:

```
Node          State      Link      Status      Address
*0 north      OPEN
                link1 UP        08:00:20:93:0E:34
                link2 UP        08:00:20:93:0E:34
 1 south      OPEN
                link1 UP        08:00:20:8F:D1:F2
                link2 DOWN
 2            CONNWAIT
                link1 DOWN
                link2 DOWN
 3            CONNWAIT
                link1 DOWN
                link2 DOWN
.
.
.
15           CONNWAIT
                link1 DOWN
                link2 DOWN
```

Note that the output lists 16 nodes. It reports on the two nodes in the cluster, north and south, plus non-existent nodes. For each correctly configured node, the information should show a state of OPEN, a status for each link of UP, and an

address for each link. However, the output in the example shows that for the node south the private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

To obtain information about the ports open for LLT, type `lltstat -p` on any node. In the following example, `lltstat -p` is typed on one node in a two-node cluster:

Node 1

```
# lltstat -p
```

Output resembles:

```
LLT port information:
Port      Usage      Cookie
0         gab        0x0
  opens:   0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
connects: 0 1
7         gab        0x7
  opens:   0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
connects: 0 1
31        gab        0x1F
  opens:   0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
connects: 0 1
```

Verifying GAB

To verify that GAB is operating, type the following command on each node:

```
# /sbin/gabconfig -a
```

If GAB is operating, the following GAB port membership information is returned:

```
GAB Port Memberships
=====
Port a gen a36e0003 membership 01
Port h gen fd570002 membership 01
```

Port a indicates that GAB is communicating, gen a36e0003 is a randomly generated number, and membership 01 indicates that nodes 0 and 1 are connected.

Port h indicates that VCS is started, gen fd570002 is a randomly generated number, and membership 01 indicates that nodes 0 and 1 are both running VCS.

If GAB is not operating, the command does not return any GAB port membership information:

```
GAB Port Memberships
=====
```

If only one network is connected, the command returns the following GAB port membership information:

```
GAB Port Memberships
=====
Port a gen a36e0003 membership 01
Port a gen a36e0003 jeopardy 1
Port h gen fd570002 membership 01
Port h gen fd570002 jeopardy 1
```

For more information on GAB, refer to the *Veritas Cluster Server User's Guide*.

Verifying the cluster

To verify that the cluster is operating, type the following command:

```
# hastatus -summary
```

The output resembles:

```
-- SYSTEM STATE
-- System                State                Frozen

A north                  RUNNING              0
A south                  RUNNING              0

-- GROUP STATE
-- Group                 System              Probed  AutoDisabled  State

B ClusterService north    Y        N                ONLINE
B ClusterService south   Y        N                OFFLINE
```

Note the system state. If the value is `RUNNING`, VCS is successfully installed and running. The group state lists the `ClusterService` group, which is `ONLINE` on north and `OFFLINE` on south. Refer to the `hastatus(1M)` manual page. In the *Veritas Cluster Server User's Guide*, look for a description of system states and the transitions between them.

hasys -display

On one of the nodes, use the `hasys` command:

```
# /opt/VRTSvcs/bin/hasys -display
```

On each node, the output should be similar. For more information on the `hasys -display` command, refer to the `hasys(1M)` manual page. Also refer to the *Veritas Cluster Server User's Guide* for information about administering VCS from the command-line.

The example shows the output when the `hasys -display` command is run on the node north; the list continues with similar information for south (not shown) and any other nodes in the cluster:

```
#System Attribute      Value
north  AgentsStopped      0
north  AvailableCapacity  100
north  CPUUsage            0
north  CPUUsageMonitoring Enabled 0 ActionThreshold 0
        ActionTimeLimit 0 Action NONE
        NotifyThreshold 0 NotifyTimeLimit 0
north  Capacity            100
north  ConfigBlockCount    142
north  ConfigCheckSum      4085
north  ConfigDiskState     CURRENT
north  ConfigFile          /etc/VRTSvcs/conf/config
north  ConfigInfoCnt       0
north  ConfigModDate       Fri May 26 17:22:48 2006
north  ConnectorState      Down
north  CurrentLimits
north  DiskHbStatus
north  DynamicLoad         0
north  EngineRestarted     0
north  EngineVersion       5.0.20.0
north  Frozen              0
north  GUIIPAddr
north  LLTNodeId           0
north  LicenseType         DEMO
north  Limits
north  LinkHbStatus        vswif1 UP vswif2 UP
```

```
north LoadTimeCounter 0
north LoadTimeThreshold 600
north LoadWarningLevel 80
north NoAutoDisable 0
north NodeId 0
north OnGrpCnt 1
north ShutdownTimeout 120
north SourceFile ./main.cf
north SysInfo Linux:north,#1 Fri Apr 22 18:13:58 EDT
2005,2.6.9-34-default,i686
north SysName north
north SysState RUNNING
north SystemLocation
north SystemOwner
north TFrozen 0
north TRSE 0
north UpDownState Up
north UserInt 0
north UserStr
north VCSFeatures DR
north VCSMode VCS
```

Adding and removing cluster nodes

This chapter contains the following topics:

- [About adding and removing nodes](#)
- [Adding a node to a cluster](#)
- [Removing a node from a cluster](#)

About adding and removing nodes

After installing VCS and creating a cluster, you can add and remove nodes from the cluster. You can create a clusters of up to 16 nodes.

Adding nodes

You must add a node to VCS cluster before you add the same node to the VMware cluster. This prevents DRS from automatically moving a virtual machine to a host that does not have VCS.

Removing nodes

You must delete a node from the VMware cluster before you remove the same node from VCS cluster. This prevents DRS from automatically moving a virtual machine to a node that does not have VCS.

In order to remove a node from a DRS cluster, you need to put it into maintenance mode. You then remove it from the VMware DRS cluster. The ESXHost agent detects this state change and internally performs a VCS system freeze command with the evacuate option.

Adding a node to a cluster

The system you add to the cluster must meet the hardware and software requirements.

See [“Planning to install VCS on an ESX Server”](#) on page 37.

[Table 6-1](#) specifies the tasks involved in adding a cluster. The example demonstrates how to add a node east to already existing nodes, north and south.

Table 6-1 Tasks involved in adding a node to a cluster

Task	Reference
Set up the hardware.	“Setting up the hardware” on page 105
Prepare the files—use files from both the 5.1 and 5.1 MP2 discs.	“Installing VCS RPMs for a manual installation” on page 106
Install the software manually.	“Installing VCS RPMs for a manual installation” on page 106
Add a license key.	“Adding a license key” on page 107

Table 6-1 Tasks involved in adding a node to a cluster

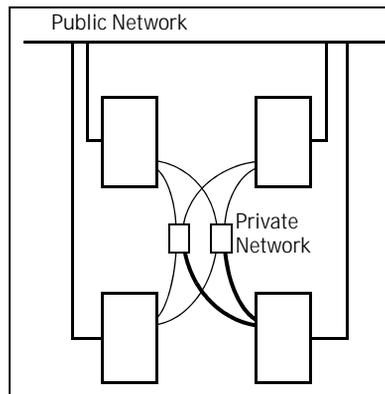
Task	Reference
Configure LLT and GAB.	“Configuring LLT and GAB” on page 107
Add the node to the existing cluster.	“Adding the node to the existing cluster” on page 109
Start VCS and verify the cluster.	“Starting VCS and verifying the cluster” on page 110
Add the node to the VMware cluster.	“Adding the node to the VMware cluster” on page 110

Setting up the hardware

Before configuring a new system to an existing cluster, you must physically add the system to the cluster.

[Figure 6-1](#) illustrates a new node being added to an existing three-node cluster using two independent switches or hubs.

Figure 6-1 Adding a node to a three-node cluster using two independent switches



To set up the hardware

- 1 Connect the VCS private Ethernet controllers.
 - If you are expanding from a two-node cluster, you need to use independent switches or hubs for the private network connections, replacing crossover cables if they are used.
 - If you already use independent switches or hubs, connect the two Ethernet controllers on the new node to the independent switches or hubs.
- 2 Connect the system to the shared storage, if required.

Preparing for a manual installation

Before you install, log in as the superuser. You then mount the disc and put the files in a temporary folder for installation.

Installing VCS RPMs for a manual installation

Install the VCS RPMs.

To install VCS RPMs on a node

- 1 Use the `rpm -i` command to install the required VCS RPMs in the order shown. Do not install any RPMs already installed on the system.

Table 6-2 Perform the `rpm -i` command to install these RPMs in the following order using specific versions

Package	Release to use
VRTSIlt-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSgab-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcs-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcsmg-5.1.20.00-GA_GENERIC.*	5.1
VRTSvcsag-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcsmn-5.1.20.00-GA_GENERIC.*	5.1
VRTSjre15-1.5.1.3-3.*	5.1
VRTScsim-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSweb-5.0.1-GA4_GENERIC.*	5.1
VRTScmcs-5.0.00.00-GA_ESX.*	5.1

Table 6-2 Perform the rpm -i command to install these RPMs in the following order using specific versions

Package	Release to use
VRTScmccc-5.0.00.00-GA_ESX.*	5.1
VRTSvcsvmip-5.1.30.00-5.1_ESX30.*	5.1 MP2
VRTSvcsdns-5.1.20.00-GA_ESX30.*	5.1
VRTSvcsvisdsk-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcсна-5.1.40.0-5.1MP2_GENERIC.*	5.1 MP2
VRTSvcsm-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcstc-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcsex-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcsci-5.1.30.00-5.1_ESX30.*	5.1 MP2
VRTSvcse-5.1.40.0-5.1MP2_GENERIC.*	5.1 MP2

Adding a license key

After you have installed all RPMs on each cluster node, use the `vxlicinst` command to add the VCS license key on each system:

```
# cd /opt/VRTS/bin
# ./vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```

Checking licensing information on the system

Use the `vxlicrep` utility to display information about all Veritas licenses on a system. For example, enter:

```
# cd /opt/VRTS/bin
# ./vxlicrep
```

From the output, you can determine the license key, the type of license, the product for which it applies, and its expiration date, if any. Demo keys have expiration dates, while permanent keys and site keys do not.

Configuring LLT and GAB

Create the LLT and GAB configuration files on the new node and update the files on the existing nodes.

To configure LLT

- 1 Create the file `/etc/llthosts` on the new node. You must also update it on each of the current nodes in the cluster.

For example, suppose you are adding east to a cluster consisting of north and south:

- If the file on one of the existing nodes resembles:

```
0 north
1 south
```

- Update the file for all nodes, including the new one, resembling:

```
0 north
1 south
2 east
```

- 2 Create the file `/etc/llttab` on the new node, making sure that line beginning “`set-node`” specifies the new node.

The file `/etc/llttab` on an existing node can serve as a guide.

See “[/etc/llttab](#)” on page 88.

The following example describes a system where node east is the new node on cluster number 2:

```
set-node east
set-cluster 2
link vswif1 vswif1 - ether - -
link vswif2 vswif2 - ether - -
```

- 3 On the new system, run the command:

```
# /sbin/lltconfig -c
```

To configure GAB

- 1 Create the file `/etc/gabtab` on the new system.

- If the `/etc/gabtab` file on the existing nodes resembles:

```
/sbin/gabconfig -c
```

then the file on the new node should be the same, although it is recommended to use the `-c -nN` option, where *N* is the number of cluster nodes.

- If the `/etc/gabtab` file on the existing nodes resembles:

```
/sbin/gabconfig -c -n2
```

then, the file on all nodes, including the new node, should change to reflect the change in the number of cluster nodes. For example, the new file on each node should resemble:

```
/sbin/gabconfig -c -n3
```

See “[/etc/gabtab](#)” on page 89.

The `-n` flag indicates to VCS the number of nodes required to be ready to form a cluster before VCS starts.

- 2 On the new node, run the command, to configure GAB:

```
# /sbin/gabconfig -c
```

To verify GAB

- 1 On the new node, run the command:

```
# /sbin/gabconfig -a
```

The output should indicate that *Port a* membership shows all nodes including the new node. The output should resemble:

```
GAB Port Memberships
=====
Port a gen a3640003 membership 012
See "Verifying GAB" on page 99.
```

- 2 Run the same command on the other nodes (north and south) to verify that the *Port a* membership includes the new node:

```
# /sbin/gabconfig -a
```

```
GAB Port Memberships
=====
Port a gen a3640003 membership 012
Port h gen fd570002 membership 01
Port h gen fd570002 visible ; 2
```

Adding the node to the existing cluster

Perform the tasks on one of the existing nodes in the cluster.

To add the new node to the existing cluster

- 1 Enter the command:

```
# haconf -makerw
```

- 2 Add the new system to the cluster:

```
# hasys -add east
```

- 3 Stop VCS on the new node:

```
# hastop -sys east
```

- 4 Copy the main.cf file from an existing node to your new node:

```
# scp /etc/VRTSvcs/conf/config/main.cf
east:/etc/VRTSvcs/conf/config/
```

- 5 Start VCS on the new node:

```
# hastart
```

- 6 If necessary, modify any new system attributes.

- 7 Enter the command:

```
# haconf -dump -makero
```

Starting VCS and verifying the cluster

Start VCS after adding the new node to the cluster and verify the cluster.

To start VCS and verify the cluster

- 1 From the new system, start VCS with the new system added to the cluster:
- 2 Run the GAB configuration command on each node to verify that *Port a* and *Port h* include the new node in the membership:

```
# hastart

# /sbin/gabconfig -a
  GAB Port Memberships
  =====
  Port a gen a3640003 membership 012
  Port h gen fd570002 membership 012
```

Adding the node to the VMware cluster

From the VMware Infrastructure client add the exact node that you added to the VCS cluster.

For more information on adding a node to a VMware cluster, refer to the VMware documentation.

Removing a node from a cluster

[Table 6-3](#) specifies the tasks involved in removing a node from a cluster. In the example procedure, the cluster consists of nodes A, B, and C; node C is to leave the cluster.

Table 6-3 Tasks involved in removing a node

Task	Reference
<ul style="list-style-type: none"> ■ Back up the configuration file. ■ Check the status of the nodes and the service groups. 	“Verify the status of nodes and service groups” on page 111
Remove the node from the VMware cluster.	
<ul style="list-style-type: none"> ■ Switch or remove any VCS service groups on the node departing the cluster. ■ Delete the node from VCS configuration. 	“Deleting the departing node from VCS configuration” on page 112
Modify the <i>llhosts</i> and <i>gabtab</i> files to reflect the change.	“Modifying configuration files on each remaining node” on page 114

Table 6-3 Tasks involved in removing a node

Task	Reference
<p>On the node departing the cluster:</p> <ul style="list-style-type: none"> ■ Modify startup scripts for LLT, GAB, and VCS to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. ■ Remove the VCS RPMs. 	<p>“Unloading LLT and GAB and removing VCS on the departing node” on page 114</p>

Verify the status of nodes and service groups

Start by issuing the following commands from one of the nodes to remain, node A or node B.

To verify the status of the nodes and the service groups

- 1 Make a backup copy of the current configuration file, main.cf.

```
# cp -p /etc/VRTSvcs/conf/config/main.cf \
  /etc/VRTSvcs/conf/config/main.cf.goodcopy
```

- 2 Check the status of the systems and the service groups.

```
# hastatus -summary
```

```
-- SYSTEM STATE
-- System      State      Frozen
A A            RUNNING   0
A B            RUNNING   0
A C            RUNNING   0

-- GROUP STATE
-- Group      System    Probed   AutoDisabled   State
B grp1       A         Y        N               ONLINE
B grp1       B         Y        N               OFFLINE
B grp2       A         Y        N               ONLINE
B grp3       B         Y        N               OFFLINE
B grp3       C         Y        N               ONLINE
B grp4       C         Y        N               ONLINE
```

The example output from the `hastatus` command shows that nodes A, B, and C are the nodes in the cluster. Also, service group `grp3` is configured to run on node B and node C, the departing node. Service group `grp4` runs only on node C. Service groups `grp1` and `grp2` do not run on node C.

Deleting the node from the VMware cluster

From the VMware Infrastructure client delete the exact node that you deleted from the VCS cluster.

For more information on deleting a node to a VMware cluster, refer to the VMware documentation.

Deleting the departing node from VCS configuration

Before removing a node from the cluster, you must remove or switch from the departing node the service groups on which other service groups depend.

To remove or switch service groups from the departing node

- 1 Switch failover service groups from the departing node. You can switch grp3 from node C to node B.

```
# hagrps -switch grp3 -to B
```

- 2 Check for any dependencies involving any service groups that run on the departing node; for example, grp4 runs only on the departing node.

```
# hagrps -dep
```

- 3 If the service group on the departing node requires other service groups, that is, if it is a parent to service groups on other nodes, then unlink the service groups.

```
# haconf -makerw
```

```
# hagrps -unlink grp4 grp1
```

These commands enable you to edit the configuration and to remove the requirement grp4 has for grp1.

- 4 Stop VCS on the departing node:

```
# hastop -sys C
```

- 5 Check the status again. The state of the departing node should be EXITED. Also, any service groups set up for failover should be online on other nodes:

```
# hastatus -summary
```

```
-- SYSTEM STATE
-- System      State      Frozen
A  A           RUNNING   0
A  B           RUNNING   0
A  C           EXITED    0
```

```

-- GROUP STATE
-- Group      System      Probed  AutoDisabled  State
B grp1       A          Y       N              ONLINE
B grp1       B          Y       N              OFFLINE
B grp2       A          Y       N              ONLINE
B grp3       B          Y       N              ONLINE
B grp3       C          Y       Y              OFFLINE
B grp4       C          Y       N              OFFLINE
    
```

- 6 Delete the departing node from the SystemList of service groups grp3 and grp4.

```

# hagrps -modify grp3 SystemList -delete C
# hagrps -modify grp4 SystemList -delete C
    
```

- 7 For service groups that run only on the departing node, delete the resources from the group before deleting the group.

```

# hagrps -resources grp4
  processx_grp4
  processy_grp4
# hares -delete processx_grp4
# hares -delete processy_grp4
    
```

- 8 Delete the service group configured to run on the departing node.

```

# hagrps -delete grp4
    
```

- 9 Check the status.

```

# hastatus -summary
-- SYSTEM STATE
-- System      State      Frozen
A A            RUNNING   0
A B            RUNNING   0
A C            EXITED    0

-- GROUP STATE
-- Group      System      Probed  AutoDisabled  State
B grp1       A          Y       N              ONLINE
B grp1       B          Y       N              OFFLINE
B grp2       A          Y       N              ONLINE
B grp3       B          Y       N              ONLINE
    
```

- 10 Delete the node from the cluster.

```

# hasys -delete C
    
```

- 11 Save the configuration, making it read only.

```

# haconf -dump -makero
    
```

Modifying configuration files on each remaining node

Perform the following tasks on each of the remaining nodes of the cluster.

To modify the configuration files on a remaining node

- 1 If necessary, modify the `/etc/gabtab` file.

No change is required to this file if the `/sbin/gabconfig` command has only the argument `-c`, although Symantec recommends using the `-nN` option, where *N* is the number of cluster systems.

If the command has the form `/sbin/gabconfig -c -nN`, where *N* is the number of cluster systems, then make sure that *N* is not greater than the actual number of nodes in the cluster, or GAB does not automatically seed.

Note: Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. The Gigabit Ethernet controller does not support the use of `-c -x`.

- 2 Modify `/etc/llthosts` file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 A
1 B
2 C
```

to:

```
0 A
1 B
```

Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node departing the cluster.

To stop LLT and GAB and remove VCS

- 1 Stop GAB and LLT:

```
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```
- 2 To determine the RPMs to remove, enter:

```
# rpm -qa | grep VRTS
```

- 3 To permanently remove the VCS RPMs from the system, use the `rpm -e` command. Remove the following RPMs in the order shown:

```
# rpm -e VRTSvcs-5.1.40.0-5.1MP2_GENERIC.*
# rpm -e VRTSvcsi-5.1.30.00-5.1_ESX30.*
# rpm -e VRTSvcsesx-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSvcsstc-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSvcsn-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSvcsna-5.1.40.0-5.1MP2_GENERIC.*
# rpm -e VRTSvcsvisdk-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSvcsdns-5.1.20.00-GA_ESX30.*
# rpm -e VRTSvcsvip-5.1.30.00-5.1_ESX30.*
# rpm -e VRTScmccc-5.0.00.00-GA_ESX.*
# rpm -e VRTScmcs-5.0.00.00-GA_ESX.*
# rpm -e VRTSweb-5.0.1-GA4_GENERIC.*
# rpm -e VRTScssim-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSjre15-1.5.1.3-3.*
# rpm -e VRTSvcsmn-5.1.20.00-GA_GENERIC.*
# rpm -e VRTSvcsag-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSvcsmg-5.1.20.00-GA_GENERIC.*
# rpm -e VRTSvcs-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSgab-5.1.40.00-5.1MP2_ESX30.*
# rpm -e VRTSllt-5.1.40.00-5.1MP2_ESX30.*
```

- 4 Remove the LLT and GAB configuration files.

```
# rm /etc/llttab
# rm /etc/gabtab
# rm /etc/llthosts
```


Installing VCS on a single node

This chapter contains the following topics:

- [About installing VCS on a single node](#)
- [Creating a single-node cluster using the installvcs program](#)
- [Creating a single-node cluster manually](#)
- [Adding a node to a single-node cluster](#)

About installing VCS on a single node

You can install VCS 5.1 MP2 on a single node. You can subsequently add another node to the single-node cluster to form a multiple-node cluster. You can also prepare a single node cluster for addition into a multi-node cluster. Single node clusters can be used for testing as well.

You can install VCS onto a single node using the `installvcs` program or you can add it manually.

See [“Creating a single-node cluster using the `installvcs` program”](#) on page 118.

See [“Creating a single-node cluster manually”](#) on page 119.

Creating a single-node cluster using the `installvcs` program

[Table 7-1](#) specifies the tasks involved in installing VCS on a single node using the `installvcs` program.

Table 7-1 Tasks to create a single-node cluster using the `installvcs` program

Task	Reference
Prepare for installation.	“Preparing for a single node installation” on page 118
Install the VCS software on the system using the <code>installvcs</code> program.	“Starting the <code>installvcs</code> program for the single node cluster” on page 119

Preparing for a single node installation

You can use the `installvcs` program to install a cluster on a single system for two purposes:

- To prepare the single node cluster to join a larger cluster
- To prepare the single node cluster to be a standalone single node cluster

When you prepare it to join a larger cluster, install it with LLT and GAB. For a standalone cluster, you do not need to enable LLT and GAB.

For more information about LLT and GAB:

See [“LLT and GAB”](#) on page 21.

Starting the installvcs program for the single node cluster

When you install VCS on a single system, follow the instructions in this guide for installing VCS using the installvcs program.

See [“Installing and configuring VCS 5.1 MP2”](#) on page 53.

During the installation, you need to answer two questions specifically for single node installations. When the installvcs program asks:

```
Enter the system names separated by spaces on which to install
VCS:
```

Enter a single system name. The installvcs program now asks if you want to enable LLT and GAB:

```
If you plan to run VCS on a single node without any need for
adding cluster node online, you have an option to proceed
without starting GAB and LLT.
```

```
Starting GAB and LLT is recommended.
```

```
Do you want to start GAB and LLT? [y,n,q,?] (n)
```

Answer **n** if you want to use the single node cluster as a standalone cluster.

Answer **y** if you plan to incorporate the single node cluster into a multi-node cluster in the future.

Continue with the installation.

See [“Licensing VCS”](#) on page 56.

Creating a single-node cluster manually

[Table 7-2](#) specifies the tasks involved in installing VCS on a single node.

Table 7-2 Tasks to create a single-node cluster manually

Task	Reference
Set the PATH variable	“Setting the PATH variable” on page 120
Install the VCS software manually and add a license key	“Installing VCS RPMs for a manual installation” on page 121 “Adding a license key” on page 122
Remove any LLT or GAB configuration files and rename LLT and GAB startup files.	“Renaming the LLT and GAB startup files” on page 122
A single-node cluster does not require the node-to-node communication service, LLT, or the membership communication service, GAB.	

Table 7-2 Tasks to create a single-node cluster manually

Task	Reference
Modify the VCS startup file for single-node operation.	“Modifying the startup files” on page 122
Create and modify the VCS configuration files.	“Configuring VCS” on page 123
Start VCS and verify single-node operation.	“Verifying single-node operation” on page 124

Setting the PATH variable

Installation commands as well as other commands reside in the `/sbin`, `/usr/sbin`, `/opt/VRTS/bin`, and `/opt/VRTSvcs/bin` directories. Add these directories to your PATH environment variable.

To set the PATH variable

- ◆ Do one of the following:
 - For the Bourne Shell (sh or ksh), type:


```
$ PATH=/usr/sbin:/sbin:/opt/VRTS/bin:/opt/VRTSvcs/bin:\
$PATH; export PATH
```
 - For the C Shell (csh or tcsh), type:


```
% setenv PATH /usr/sbin:/sbin:/opt/VRTS/bin:\
/opt/VRTSvcs/bin:$PATH
```

Installing VCS RPMs for a manual installation

Install the VCS RPMs.

To install VCS RPMs on a node

- 1 Use the `rpm -i` command to install the required VCS RPMs in the order shown. You need to install the files from the indicated media or location.

Table 7-3 Perform the `rpm -i` command to install these RPMs in the following order using specific versions

Package	Release to use
VRTSIlt-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSgab-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcS-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcsmg-5.1.20.00-GA_GENERIC.*	5.1
VRTSvcSag-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcsmn-5.1.20.00-GA_GENERIC.*	5.1
VRTSjre15-1.5.1.3-3.*	5.1
VRTScsim-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSweb-5.0.1-GA4_GENERIC.*	5.1
VRTScmcs-5.0.00.00-GA_ESX.*	5.1
VRTScmccc-5.0.00.00-GA_ESX.*	5.1
VRTSvcsvmip-5.1.30.00-5.1_ESX30.*	5.1 MP2
VRTSvcSdns-5.1.20.00-GA_ESX30.*	5.1
VRTSvcSvisdk-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcSna-5.1.40.0-5.1MP2_GENERIC.*	5.1 MP2
VRTSvcsm-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcstc-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcsex-5.1.40.00-5.1MP2_ESX30.*	5.1 MP2
VRTSvcSi-5.1.30.00-5.1_ESX30.*	5.1 MP2
VRTSvcse-5.1.40.0-5.1MP2_GENERIC.*	5.1 MP2

Adding a license key

After you have installed all RPMs on each cluster node, use the `vxlicinst` command to add the VCS license key on each system:

```
# cd /opt/VRTS/bin
# ./vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```

Checking licensing information on the system

Use the `vxlicrep` utility to display information about all Veritas licenses on a system. For example, enter:

```
# cd /opt/VRTS/bin
# ./vxlicrep
```

From the output, you can determine the license key, the type of license, the product for which it applies, and its expiration date, if any. Demo keys have expiration dates, while permanent keys and site keys do not.

Renaming the LLT and GAB startup files

You may need the LLT and GAB startup files if you need to upgrade the single-node cluster to a multiple-node cluster at a later time.

To rename the LLT and GAB startup files

- ◆ Rename the LLT and GAB startup files.

```
# mv /etc/init.d/llt /etc/init.d/llt.old
# mv /etc/init.d/gab /etc/init.d/gab.old
```

Modifying the startup files

Modify the VCS startup file `/etc/sysconfig/vcs` to include the `-onenode` option as follows:

Change the line:

```
ONENODE=no
```

To:

```
ONENODE=yes
```

Configuring VCS

VCS configuration requires the `types.cf` and `main.cf` files on each system in the cluster. Both of the files are in the `/etc/VRTSvcs/conf/config` directory.

main.cf file

The `main.cf` configuration file requires the following minimum essential elements:

- An “include” statement that specifies the file, `types.cf`, which defines the VCS bundled agent resources.
- The name of the cluster.
- The name of the systems that make up the cluster.

Editing the main.cf file

When you manually install VCS, the file `/etc/VRTSvcs/conf/config/main.cf` contains only the line:

```
include "types.cf"
```

To edit the main.cf file

- 1 Log in as superuser, and move to the directory containing the configuration file:

```
# cd /etc/VRTSvcs/conf/config
```
- 2 Using `vi`, or another text editor, edit the `main.cf` file, defining your cluster name and system names. Refer to the following example.
- 3 Save and close the file.

Refer to the *Veritas Cluster Server User's Guide* for a full description of the `main.cf` file, how to edit it and verify it.

Example, main.cf

An example `main.cf` for a two-node cluster:

```
include "types.cf"
cluster VCSCluster2 ( )
system north ( )
system south ( )
```

An example `main.cf` for a single-node cluster:

```
include "types.cf"
cluster VCSCluster1 ( )
system sn1
```

types.cf file

Note that the “include” statement in main.cf refers to a file named types.cf. This text file describes the VCS bundled agent resources. During new installations, the types.cf file is automatically copied in to the /etc/VRTSvcs/conf/config directory.

Verifying single-node operation

After successfully creating a single-node cluster, start VCS and verify the cluster.

To verify single-node cluster

- 1 Bring up VCS manually as a single-node cluster using `hastart` with the `-onenode` option:

```
# hastart -onenode
```

- 2 Verify that the `had` and `hashadow` daemons are running in single-node mode:

```
# ps -ef | grep ha  
root 285 1 0 14:49:31 ? 0:02 /opt/VRTSvcs/bin/had -onenode  
root 288 1 0 14:49:33 ? 0:00 /opt/VRTSvcs/bin/hashadow
```

Adding a node to a single-node cluster

Table 7-4 specifies the activities involved in adding nodes to a single-node cluster. All nodes in the new cluster must run the same version of VCS. The example procedure refers to the existing single-node VCS node as Node A and the node that is to join Node A to form a multiple-node cluster as Node B.

Table 7-4 Tasks to add a node to a single-node cluster

Task	Reference
Set up Node B to be compatible with Node A.	“Setting up a node to join the single-node cluster” on page 126
<ul style="list-style-type: none"> ■ Add Ethernet cards for private heartbeat network for Node B. ■ If necessary, add Ethernet cards for private heartbeat network for Node A. ■ Make the Ethernet cable connections between the two nodes. 	“Installing and configuring Ethernet cards for private network” on page 126
Connect both nodes to shared storage.	“Configuring the shared storage” on page 127
<ul style="list-style-type: none"> ■ Bring up VCS on Node A. ■ Edit the configuration file. ■ Edit the startup scripts. 	“Bringing up the existing node” on page 127
If necessary, install VCS on Node B and add a license key.	“Installing VCS RPMs for a manual installation” on page 121
Make sure Node B is running the same version of VCS as the version on Node A.	“Adding a license key” on page 122
Edit the configuration files on Node B.	“Configuring LLT and GAB” on page 128
Start LLT and GAB on Node B.	“Starting LLT and GAB” on page 130
<ul style="list-style-type: none"> ■ Start LLT and GAB on Node A. ■ Restart VCS on Node A. ■ Modify service groups for two nodes. 	“Reconfiguring VCS on the existing node” on page 130
<ul style="list-style-type: none"> ■ Start VCS on Node B. ■ Verify the two-node cluster. 	“Verifying configuration on both nodes” on page 131

Setting up a node to join the single-node cluster

The new node to join the existing single node running VCS must run the same version of operating system and patch level.

To set up a node to join the single-node cluster

- ◆ Do one of the following:
 - If the node you plan to add as Node B is currently part of an existing cluster, remove the node from the cluster. After removing the node from the cluster, remove the VCS RPMs and configuration files. See [“Removing a node from a cluster”](#) on page 110.
 - If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS. See [“Uninstalling VCS”](#) on page 83.
 - If you renamed the LLT and GAB startup files, remove them. See [“Renaming the LLT and GAB startup files”](#) on page 122.

Installing and configuring Ethernet cards for private network

Both nodes require ethernet cards (NICs) that enable the private network. If both Node A and Node B have ethernet cards installed, you can ignore this step.

For high availability, two separate NICs on each node should be used, such that the failure of one NIC does not restrict heartbeating between the nodes.

See [“Setting up the private network”](#) on page 41.

To install and configure ethernet cards for private network

- 1 Shut down VCS on Node A.
`# hastop -local`
- 2 Shut down the node to get to the OK prompt:
`# sync;sync;init 0`
- 3 Install the ethernet card on Node A.
- 4 Install the ethernet card on Node B.
- 5 Configure the ethernet card on both nodes.
- 6 Make the two ethernet cable connections from Node A to Node B for the private networks.
- 7 Restart the nodes.

Configuring the shared storage

Make the connection to shared storage from Node B. Reboot the node when you are prompted.

Bringing up the existing node

- 1 Restart Node A.
- 2 Log in as superuser.
- 3 Make the VCS configuration writable.

```
# haconf -makerw
```
- 4 Display the service groups currently configured.

```
# hagrps -list
```
- 5 Freeze the service groups.

```
# hagrps -freeze group -persistent
```

Repeat this command for each service group listed in [step 4](#).
- 6 Make the configuration read-only.

```
# haconf -dump -makero
```
- 7 Stop VCS on Node A.

```
# hastop -local -force
```
- 8 Edit the VCS system configuration file `/etc/sysconfig/vcs`, and remove the “-onenode” option.
Change the line:

```
ONENODE=yes
```

To:

```
ONENODE=no
```
- 9 Rename the GAB and LLT startup files so they can be used.

```
# mv /etc/init.d/gab.old /etc/init.d/gab  
# mv /etc/init.d/llt.old /etc/init.d/llt
```

Installing the VCS RPMs and license key

Install the VCS 5.1 RPMs manually and install the license key.

See [“Installing VCS RPMs for a manual installation”](#) on page 121.

See [“Adding a license key”](#) on page 122.

Configuring LLT and GAB

VCS uses LLT and GAB to replace the functions of TCP/IP for VCS private network communications. LLT and GAB provide the performance and reliability required by VCS for these and other functions.

LLT and GAB must be configured as described in the following sections.

Configuring low latency transport (LLT)

To configure LLT, set up two files: `/etc/llthosts` and `/etc/llttab` on each node in the cluster.

Setting up `/etc/llthosts`

The file `llthosts` is a database, containing one entry per system that links the LLT system ID (in the first column) with the LLT host name. You must create an identical file on each node in the cluster.

Use `vi`, or another editor, to create the file `/etc/llthosts` that contains entries that resemble:

```
0 north
1 south
```

Setting Up `/etc/llttab`

The `/etc/llttab` file must specify the system's ID number (or, its node name), and the network links that correspond to the system. In addition, the file can contain other directives. Refer also to the sample `llttab` file in `/opt/VRTSllt`.

See [“LLT directives”](#) on page 129.

Using `vi` or another editor, create the file `/etc/llttab` that contains entries that resemble:

```
set-node north
set-cluster 2
link vswif1 eth-00:50:56:41:07:c1 - ether - -
link vswif2 eth-00:50:56:46:58:54 - ether - -
```

The first line must identify the system on which the file exists. In the example above, the value for `set-node` could be `north`, `0`, or the file name `/etc/nodename`, provided the file contains the name of the system (`north` in this example). The next two lines, beginning with the `link` command, identify the two private

network cards that the LLT protocol uses. The order of directives must be the same as in the sample file `/opt/VRTSllt/llttab`.

LLT directives

For more information about LLT directives, refer to the `llttab(4)` manual page.

Table 7-5 LLT directives

Directive	Description
<code>set-node</code>	<p>Assigns the system ID or symbolic name. The system ID number must be unique for each system in the cluster, and must be in the range 0-31. The symbolic name corresponds to the system ID listed in <code>/etc/llthosts</code> file.</p> <p>Note that LLT fails to operate if any systems share the same ID.</p>
<code>link</code>	<p>Attaches LLT to a network interface. At least one link is required, and up to eight are supported. The first argument to <code>link</code> is a user-defined tag shown in the <code>lltstat</code> output to identify the link. It may also be used in <code>llttab</code> to set optional static MAC addresses.</p> <p>The second argument to <code>link</code> is the device name of the network interface. Its format is <code>device_name:device_instance_number</code>. The remaining four arguments to <code>link</code> are defaults; these arguments should be modified only in advanced configurations. There should be one <code>link</code> directive for each network interface. LLT uses an unregistered Ethernet SAP of 0xCAFÉ. If the SAP is unacceptable, refer to the <code>llttab(4)</code> manual page for information on how to customize SAP. Note that IP addresses do not need to be assigned to the network device; LLT does not use IP addresses.</p>
<code>set-cluster</code>	<p>Assigns a unique cluster number. Use this directive when more than one cluster is configured on the same physical network connection. LLT uses a default cluster number of zero.</p>
<code>link-lowpri</code>	<p>Use this directive in place of <code>link</code> for public network interfaces. This directive prevents VCS communication on the public network until the network is the last link, and reduces the rate of heartbeat broadcasts. Note that LLT distributes network traffic evenly across all available network connections and, in addition to enabling VCS communication, broadcasts heartbeats to monitor each network connection.</p>

For more information about LLT directives, refer to the `llttab(4)` manual page.

Additional considerations for LLT

You must attach each network interface that is configured for LLT to a separate and distinct physical network.

Configuring group membership and atomic broadcast (GAB)

To configure GAB, use vi or another editor to set up an `/etc/gabtab` configuration file on each node in the cluster. The following example shows a simple `/etc/gabtab` file:

```
/sbin/gabconfig -c -nN
```

Where the `-c` option configures the driver for use and `-nN` specifies that the cluster is not formed until at least N systems are ready to form the cluster. By default, N is the number of systems in the cluster.

Note: Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. Using `-c -x` dramatically increases configuration time for the Gigabit Ethernet controller and can lead to a split-brain condition.

Starting LLT and GAB

On the new node, start LLT and GAB.

To start LLT and GAB

- 1 Start LLT on Node B.
`# /etc/init.d/llt start`
- 2 Start GAB on Node B.
`# /etc/init.d/gab start`

Reconfiguring VCS on the existing node

- 1 On Node A, create the files `/etc/llttab`, `/etc/llthosts`, and `/etc/gabtab`. Use the files created on Node B as a guide, customizing the `/etc/llttab` for Node A.
- 2 Start LLT on Node A.
`# /etc/init.d/llt start`
- 3 Start GAB on Node A.
`# /etc/init.d/gab start`
- 4 Check the membership of the cluster.
`# gabconfig -a`
- 5 Start VCS on Node A.
`# hstart`

- 6 Make the VCS configuration writable.
haconf -makerw
- 7 Add Node B to the cluster.
hasys -add sysB
- 8 Add Node B to the system list of each service group.
 - List the service groups.
hagrp -list
 - For each service group listed, add the node.
hagrp -modify group SystemList -add sysB 1

Verifying configuration on both nodes

- 1 On Node B, check the cluster membership.
gabconfig -a
- 2 Start the VCS on Node B.
hastart
- 3 Verify that VCS is up on both nodes.
hastatus
- 4 List the service groups.
hagrp -list
- 5 Unfreeze the service groups.
hagrp -unfreeze group -persistent
- 6 Implement the new two-node configuration.
haconf -dump -makero

Upgrading to VCS 5.1 MP2

This chapter includes the following topics:

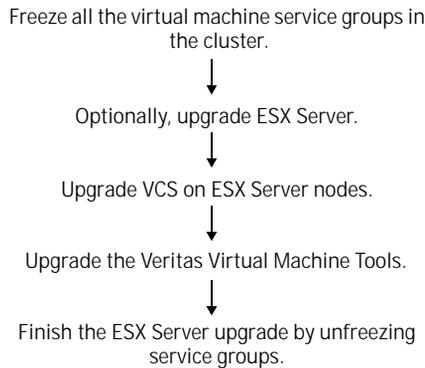
- [About upgrading to VCS 5.1 MP2](#)
- [Upgrading VCS on ESX Server nodes](#)
- [Upgrading Veritas Virtual Machine Tools in virtual machines running Windows](#)
- [Upgrading Veritas Virtual Machine Tools in virtual machines running Linux](#)
- [Upgrading Veritas Virtualization Manager \(VVM\) on clients](#)

About upgrading to VCS 5.1 MP2

Upgrade VCS on ESX Server nodes, the Veritas Virtualization Manager on clients, and the Veritas Virtual Machine Tools in virtual machines. You can upgrade from VCS 5.1 or VCS 5.1 MP1. You can also upgrade the operating system from ESX Server 3.0.1 or 3.0.2 to 3.5. You must perform an operating system upgrade before you perform

[Figure 8-1](#) on page 134 illustrates a VCS for VMware ESX upgrade.

Figure 8-1 Upgrade flow overview



[Table 8-1](#) on page 134 describes where to look for the pertinent upgrade information.

Table 8-1 Upgrade flow

Upgrade task description	Section
Freeze the service groups.	See “Freeze all the virtual machine service groups in the cluster” on page 135.
Upgrade ESX Server (optional).	See “Upgrading ESX Server” on page 135.
On ESX Server nodes, upgrade VCS.	See “Upgrading VCS on ESX Server nodes” on page 135.
On clients running VVM, upgrade them.	See “Upgrading Veritas Virtualization Manager (VVM) on clients” on page 140.

Table 8-1 Upgrade flow

Upgrade task description	Section
On each virtual machine, upgrade Veritas Virtual Machine Tools.	<ul style="list-style-type: none"> ■ See “Upgrading Veritas Virtual Machine Tools in virtual machines running Windows” on page 138. ■ See “Upgrading Veritas Virtual Machine Tools in virtual machines running Linux” on page 139.
Unfreeze frozen service groups	See “Unfreezing the service groups” on page 139.

Symantec recommends that you install all required patches prior to upgrade.

Freeze all the virtual machine service groups in the cluster

Perform the following tasks for each service group in the cluster.

To freeze all the virtual machine service groups in the cluster

- ◆ On one of the nodes in the cluster, freeze the service group. At a prompt, type:


```
# hagrp -freeze service_group_name
```

 Where *service_group_name* is the name of the service group that the resource resides in.

Upgrading ESX Server

You can optionally perform the operating system upgrade from ESX Server 3.0.1 or 3.0.2 to 3.5. Refer to the relevant VMware documentation for more information.

Upgrading VCS on ESX Server nodes

To upgrade from VCS 5.1 or VCS 5.1 MP1 to VCS 5.1 MP2 perform the following procedures.

- See [“Upgrading to VCS 5.1 MP2”](#) on page 136.
- See [“Upgrading agents’ configuration”](#) on page 137.

Upgrading to VCS 5.1 MP2

To upgrade to VCS 5.1 MP2, perform the following instructions. At the end of the upgrade, do not perform the suggested system shut down and reboot, instead follow the instructions provided.

For operating system upgrades from ESX Server 3.0.1 and 3.0.2 to ESX Server 3.5, perform the operating system upgrade before you upgrade VCS.

Do not install VCS 5.1 MP1 on ESX Server 3.5.

To upgrade to VCS 5.1 MP2

- 1 Log in as superuser on one of the systems for installation.
- 2 Insert the disc containing the VCS 5.1 MP2 software into the disc drive of one of the cluster nodes.
- 3 On the node, freeze the service group. At a prompt, type:

```
# hagr -freeze service_group_name
```

Where *service_group_name* is the name of the service group that the resource resides in.
- 4 Mount the disc on a suitable mount point and change directory to the location of the disc mount's `cluster_server` directory.

```
# cd cluster_server
```
- 5 Install VCS 5.1 MP2 using the `installmp` script:

```
# ./installmp
```

Running the `installmp` script stops VCS.
- 6 The installer prompts for the system names on which you want to install the maintenance pack.
Enter the system names separated by spaces on which to install MP1: **sysA sysB**
Running the `installmp` script stops VCS but leaves the applications under VCS control in their current state.
- 7 After an initial system check, the setup program is ready to install VCS 5.1 MP2 and seeks confirmation.
Are you sure you want to install MP1? [y,n,q] (y)
Enter **y** to begin installation.
- 8 Review the output as VCS 5.1 MP2 installs.

- 9 Once the installation is complete verify if VCS restarts and it is at the VCS 5.1 MP2 level. Run the `rpm -qa | grep VRTSvcs` and the `hasys -state` commands. Example output follows:

```
# rpm -qa | grep VRTSvcs
VRTSvcs-5.1.40.00-MP2_ESX30
# hasys -state
#System      Attribute      Value
sysA         SysState      RUNNING
```

- 10 Run the following command on each node of the cluster to start VCS.

```
$ service llt start
$ service gab start
$ /etc/init.d/vcs start
$ hasys -state vcs_system_name
```

Make sure the system is in the running state before moving on to start VCS on the next system.

Upgrading agents' configuration

You must now upgrade agents' configuration. This section provides information on how to upgrade your configuration from VCS 5.1 to VCS 5.1 MP2. Use the `upgradevcsconf.sh` script (on the disc in `cluster_server/tools/`) to automate the upgrade.

The `upgradevcsconf.sh` script performs the following tasks, it:

- Adds the MetroMirror types file to the configuration.
- Adds the NICConf attribute to the configuration.
- Removes all existing VSwitch resources to follow the network infrastructure scheme.
- Detects any duplicate VSwitch resources in the configuration and Creates a Proxy resource that points to a common VSwitch resource in the Network Infrastructure service group.

To upgrade your configuration

- 1 Run the `upgradevcsconf.sh` script to upgrade the VCS configuration to the 5.1 MP2 level.


```
# ./upgradevcsconf.sh
```
- 2 The script prompts you to start configuration.


```
Are you ready to configure the VCS 5.1MP1 configuration (Y/N) [Y] ?
Enter y to begin the configuration.
```
- 3 As the script prompts you, provide answers, and make sure that the script completes with no errors.

Upgrading Veritas Virtual Machine Tools in virtual machines running Windows

You must upgrade each virtual machine individually.

The utility is on the disc inside the installvcsvm_tools directory. It is in the ISO format:

- For x86 (32-bit) architectures, use:
win-x86-vcsvm-tools.iso
- For x64 architectures, use:
win-x64-vcsvm-tools.iso

To add the Tools .iso file

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Right-click the virtual machine where you want to add the .iso file. Select **Add ISO Image**. VVM automatically selects the proper .iso file to match the operating system.
The Add CD-ROM ISO window appears.
- 3 Click the **OK** button to add the .iso file.
The ISO file is now available for your use.

To upgrade the Veritas Virtual Machine Tools

- 1 Once you have mounted the appropriate ISO file to your virtual machine, run the vcsvm-tools.exe file to upgrade the Tools.
- 2 Review the Welcome screen and click **Next**.
- 3 Click **Finish** to close the installer.
- 4 Verify the installation. Check to see if the VCSAgMD service is present in the Services panel. (**Start > Programs > Administrative Tools > Services**)
- 5 To confirm that the upgrade is successful, you can go to **Start > Settings > Control Panel > Add Remove Programs** and click **Support Information** for vcsvm-tools-dotnet or vcsvmm-tools-winx64 to verify that the version is 5.11.2000.

Upgrading Veritas Virtual Machine Tools in virtual machines running Linux

You must upgrade each virtual machine individually.

To upgrade VCS Virtual Machine Tools from VCS 5.1 to 5.1 MP2 on virtual machines running Linux

- 1 Mount the product disc.
- 2 Navigate to the `installvcsvm-tools` location.

```
# cd /media/cdrom/
```
- 3 On the virtual machine, enter the `installvcsvm-tools -n` command.

```
# ./installvcsvm-tools -n
```
- 4 To confirm that the upgrade is successful, run the following command inside the Linux virtual machine:

```
$ /opt/VRTSvcs/bin/vcsvm-tools -v
```

Verify that the version is 5.1.40.00.
- 5 If the result of [step 4](#) is unsuccessful, perform the following procedure: [“Configuring Veritas Virtual Machine Tools”](#) on page 231.

Unfreezing the service groups

To complete the upgrade, unfreeze the frozen service groups.

To unfreeze the service group

- ◆ On one of the nodes in the cluster, unfreeze the service group. At a prompt, type:

```
# hagrps -unfreeze service_group_name
```

Where *service_group_name* is the name of the service group that the resource resides in.

Upgrading Veritas Virtualization Manager (VVM) on clients

If you do not have the Veritas Virtualization Manager installed on your system, you can install or upgrade it from the maintenance pack disc.

To install or upgrade Veritas Virtualization Manager (VVM)

- 1 Insert the product disc into a drive on the client system.
- 2 Change the directory to windows/vvm.
- 3 Run the vcsvm.msi file.
- 4 Review the Welcome screen and click **Next**.
- 5 Click **Install** to begin the installation of the Veritas Virtualization Manager. If the installer cannot remove the previous version, go to **Start > Settings > Control Panel > Add Remove Programs** and remove the program. After removing the program, restart this procedure.
- 6 After the InstallShield wizard completes the installation, click **Finish** to exit the wizard.
- 7 To confirm that the upgrade is successful, you can go to **Start > Settings > Control Panel > Add Remove Programs** and click on Support Information for Veritas Virtualization Manager to verify that the version is 5.1.2000. If the installer cannot remove the previous version, go to **Start > Settings > Control Panel > Add Remove Programs** and remove the program. After removing the program, restart this procedure.



Configuring VCS for virtual machines

This section contains the following chapters:

- [Chapter 9, “Installing the Veritas Virtualization Manager \(VVM\)”](#) on page 143
- [Chapter 10, “Configuring virtual machines for high availability”](#) on page 153
- [Chapter 11, “Configuring virtual machines for disaster recovery”](#) on page 159

Installing the Veritas Virtualization Manager (VVM)

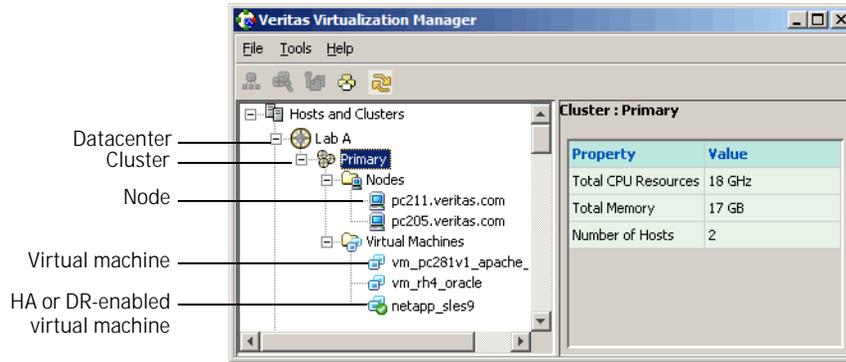
This chapter contains the following topics:

- [About Veritas Virtualization Manager \(VVM\)](#)
- [Installing the Veritas Virtualization Manager \(VVM\)](#)
- [Preparing keystores](#)
- [Starting Veritas Virtualization Manager \(VVM\)](#)
- [Removing Veritas Virtualization Manager \(VVM\)](#)

About Veritas Virtualization Manager (VVM)

The Veritas Virtualization Manager (VVM) is a user interface that enables you to configure virtual machines for high availability, disaster recovery, and to increase allocated storage for a virtual machine.

Figure 9-1 The Veritas Virtualization Manager interface



After you start the application, VVM presents a hierarchy of objects in the left side of the pane, and the selected object's properties and values in the right. The hierarchy matches the datacenters, clusters, nodes, and virtual machines that are inside the cluster that you created in VMware. VVM also provides information about the selected object.

You can use VVM to perform the following tasks:

- Configure virtual machines for high availability.
See ["Configuring virtual machines for high availability using the Veritas Virtualization Manager"](#) on page 154.
- Configure a virtual machine for disaster recovery.
See ["Configuring virtual machines for disaster recovery using the Veritas Virtualization Manager"](#) on page 169.
- Manage service groups using the Cluster Server (Java Console).
Right-click a cluster and select Launch VCS Console. This full-featured Java Console is part of VVM. You cannot access this version of the Java Console through the Start menu.
- Grow storage size for your virtual machines.
See ["Increasing allocated storage"](#) on page 308.

- Add .iso image files to your virtual machines. This gives you easy access to the Veritas Virtual Machine Tools.
 - See “[Supported operating systems for increasing allocated storage](#)” on page 31.
 - See “[Mounting Veritas Virtual Machine Tools](#)” on page 230.
 - See “[To add the tools .iso file](#)” on page 245.
- Disable VMware HA
To disable VMware HA, right-click a cluster and select **Disable VMware HA**.

Installing the Veritas Virtualization Manager (VVM)

Install the Veritas Virtualization Manager on a Windows computer. Install it on a standalone system, which is outside of the cluster.

Veritas Virtualization Manager hardware requirements

Symantec recommends a Pentium III at 400 MHz, with 256 MB RAM.

Minimum hardware requirements follow:

- Pentium II 300 megahertz CPU
- 256 megabytes RAM
- 800 x 600 display resolution
- 8-bit color display depth
- Graphics card capable of 2D images
- Approximately 40 MB of free hard drive space

Installing the Veritas Virtualization Manager

Install the Veritas Virtualization Manager on a standalone Windows system that is outside of the cluster.

To install the Veritas Virtualization Manager

- 1 Insert the product disc into a drive on the client system.
- 2 Change directory to windows/vvm.
- 3 Run the vcsvm.msi file.
- 4 Review the Welcome screen and click **Next**.
- 5 Read the license agreement. If you choose to accept it, click the **I accept the terms in the license agreement** radio button. Click **Next**.

- 6 Enter your user name, and your organization. Select the radio button that reflects your usage on the client. Click **Next**.
- 7 Accept the complete setup default and click **Next**.
- 8 Click the **Install** button to install VVM.
- 9 Click the **Finish** button.

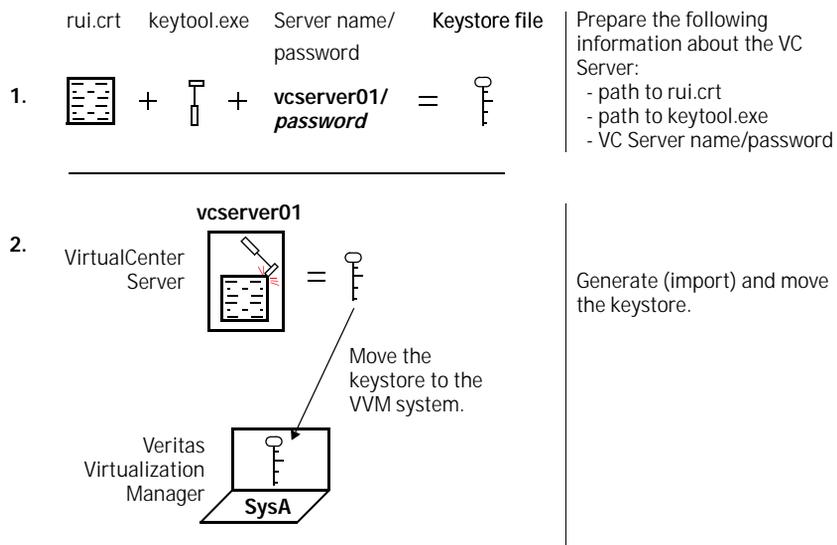
When the installation program finishes, prepare the SSL certificate for the VirtualCenter Servers that you want to log in to. Note that VVM now comes with an option that allows you to log in without a keystore.

Preparing keystores

The Veritas Virtualization Manager (VVM) requires a VMware keystore (generated from a SSL certificate) to connect and communicate with the VirtualCenter Server. The keystore secures communication between the client where VVM runs and the VirtualCenter Server.

The following is an example of how to set up the required keystore. For each VirtualCenter Server that you have deployed, you need a keystore file for that server on each client where you plan to run VVM. You must also copy the keystore to all of the ESX Server nodes in the VCS cluster.

Figure 9-2 Preparing the keystore file (SSL Certificate)



Keystore generation and population quickstart

Use this procedure if you familiar with creating a keystore file.

To generate and populate the keystore file

- 1 Create a directory called certificate-store in the following locations:
 - On each client where you plan to install VVM
 - On each ESX Server node in the VCS cluster
- 2 On each VirtualCenter Server where you plan to use VVM, create a directory called temp-certificate-store.
- 3 For each VirtualCenter Server prepare the following information:
 - The path to the rui.crt file
See [step 5](#) for the recommended path.
 - The path to the keytool.exe file
See [step 5](#) for the recommended path.
 - The VirtualCenter Server's name and password
- 4 On one VirtualCenter Server, open a Command Prompt window. Navigate into the temp-certificate-store directory.

```
C: \>cd temp-certificate-store
```

You don't need to supply a path for the keystore file when you create it from the same directory where you run the following command.
- 5 Run the following command to generate the keystore file. Make the appropriate changes to: paths, server names, and server name passwords.

```
C: \>"C:\Program Files\Common Files\VERITAS Shared\VRTSjre\jre1.5\bin\keytool.exe" -import -file "C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt" -alias vcserver01 -keystore vcserver01.keystore
```
- 6 Repeat [step 4](#) and [step 5](#) for each VirtualCenter Server that you want to manage with VVM.
- 7 Copy each keystore file from each VirtualCenter Server's temp-certificate-store to the following locations:
 - The certificate-store directory on each client where VVM runs
 - The certificate-store directory on each ESX Server node in the VCS cluster
- 8 You can now start VVM. Use the path to the certificate store on the client.

Finding the path information on the VirtualCenter Server

You need some basic information for the `keytool.exe` and the `ruicert` files on the VirtualCenter Server.

To find and note the path to the `keytool.exe` executable

- 1 On the VirtualCenter Server (`vcserver01`), locate the `keytool.exe` executable. The `keytool.exe` file is commonly in: `C:\Program Files\Common Files\VERITAS Shared\VRTSjre\jre1.5\bin`.
Some systems may have this file in a different directory. In this situation, use the Windows search feature to locate the file. If multiple `jre` directories exist, find the file that is in the version 1.5 directory.
- 2 On the VirtualCenter Server (`vcserver01`), copy the path information for the file into a temporary text file. In this example, the path information resembles:

```
"C:\Program Files\Common Files\VERITAS Shared\  
VRTSjre\jre1.5\bin\keytool.exe"
```

Remember to add the file name—`keytool.exe`—and quotes at the beginning and end of the path.

To find and note the path to the `ruicert` certificate file

- 1 On the VirtualCenter Server (`vcserver01`), locate the `ruicert` certificate file. Symantec recommends that you use the `ruicert` file in: `C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL\ruicert`.
Note that this file location can move, and that multiple `ruicert` files can exist on a VirtualCenter Server. If you cannot locate the `ruicert` file, you can use the Windows search feature to locate the file.
- 2 On the VirtualCenter Server (`vcserver01`), copy the path information for the file to a temporary text file. In this example, the path information resembles:

```
"C:\Documents and Settings\All Users\Application  
Data\VMware\VMware VirtualCenter\SSL\ruicert"
```

Remember to add the file name—`ruicert`—and quotes at the beginning and end of the path. If for any reason, the `ruicert` file does not exist, refer to VMware SDK documentation.

To create the keystore

- 1 On the VirtualCenter Server (`vcserver01`), open a command line prompt.
- 2 Click **Start > Run**.
- 3 In the Run dialog, enter `cmd` in the Open field. Click the **OK** button.

4 From the command line, make sure that you are at root. At the prompt, type:
C:\>**cd **

5 To make it easy to find the keystore, create a temporary directory and open it.
C:\>**mkdir \temp-certificate-store**
C:\>**cd temp-certificate-store**

6 At the prompt type the following command.

```
"path_to_keytool\keytool.exe" -import -file  
"path_to_rui.crt\rui.crt" -alias server-name -keystore  
"path_to_keystorefile\vmware.keystore"
```

A full command example is:

```
"C:\Program Files\Common Files\VERITAS Shared\VRTSjre\jre1.5\  
bin\keytool.exe" -import -file "C:\Documents and Settings\All  
Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt"  
-alias vcserver01 -keystore vcserver01.keystore
```

■ **"path_to_keytool\keytool.exe"**

The pathname and the executable for the keytool (keytool.exe) is the path from root to the keytool.exe file that you found in the previous procedure. "C:\Program Files\Common Files\VERITAS Shared\VRTSjre\jre1.5\bin\keytool.exe" in this example. Start and end the path and file name with quotes.

■ **"path_to_rui.crt\rui.crt"**

The pathname and the certificate file leads to the rui.crt certificate file ("C:\Documents and Settings\All Users\Application Data\VMware\VMware VirtualCenter\SSL\rui.crt" in this example). Start and end this path and file name with quotes.

■ **server-name**

The server-name is the name of the VirtualCenter Server (vcserver01 in this example).

■ **vmware.keystore**

The certificate store that you are creating (vcserver01.keystore in this example). If you do not specify the absolute path to this file, it is created in the current directory.

7 When you enter the password for the keystore use the same password as the one that you use to log on to the VirtualCenter Server. Press the **Enter** key.

8 When asked if you want to trust the certificate, answer **yes**.

Creating the certificate store for the Veritas Virtualization Manager

The certificate store is a directory that you create to hold keystore files. For convenience's sake, consider using identical certificate store names on each system where you plan to install and use VVM.

To create the certificate store

- 1 On the client where you have installed VVM, create the new certificate store directory.
- 2 Open a command prompt on the VirtualCenter Server. Click **Start > Run**.
- 3 In the Run dialog, enter **cmd** in the Open field. Click the **OK** button.
- 4 From the command line, make sure that you are at root. At the prompt, type:

```
C:\> cd \
```
- 5 Create a directory to hold the keystore file. For example, at the prompt type:

```
C:\> mkdir certificate-store
```

Where `certificate-store` is the name of the new directory and the certificate store.

Copying the keystore file from the VirtualCenter Server to the Veritas Virtualization Manager (VVM)

You now need to copy the keystore file onto the client where you plan to use VVM. When you log into VVM, you need the path to the keystore.

To copy the file from the VirtualCenter Server to the VVM client

- ◆ Copy the keystore file that you have created on the VirtualCenter Server (`vcserver01.keystore` in this example) to the new directory that you created on the system where you plan to run VVM.

Copying the keystore file from the VirtualCenter Server to each of the ESX Server nodes in the VCS cluster

For VMotion to work in a VCS environment, the VCS nodes need access to the keystore file. Copy the keystore file to the same location on each of the ESX Server nodes that are in the cluster. For example, place the keystore file in the `/etc/VRTSvcs/conf/` directory. When you configure the ESXVirtualMachine agent, you must configure the `sslcert` attribute to point to this keystore file.

Starting Veritas Virtualization Manager (VVM)

Figure 9-3 shows the VirtualCenter Server name, the username and password for the VirtualCenter Server, and the SSL certificate path to the keystore file. You can also use VVM without a keystore.

Warning: Using VVM without a keystore is a security risk, so use this option judiciously and only behind secure firewalls. Do not use VVM without a keystore in a production environment.

Figure 9-3 Entering the information required to log in to VVM



Starting VVM with or with an SSL certificate

You can start VVM with an SSL certificate or without. If you disable the SSL certificate, make sure that your cluster is behind a firewall and is not a production environment. For instructions on starting VVM with or with an SSL certificate, see the following:

- See [“To start Veritas Virtualization Manager”](#) on page 151.
- See [“To start Veritas Virtualization Manager without a keystore”](#) on page 152.

To start Veritas Virtualization Manager

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Enter the name of the VirtualCenter Server, your user name and password.

- 3 Enter the path information to the keystore. For example, enter:
C:\certificate-store\vcserver01.keystore.
See [“Finding the path information on the VirtualCenter Server”](#) on page 148.

To start Veritas Virtualization Manager without a keystore

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Enter the name of the VirtualCenter Server, your user name and password.
- 3 Clear the **Use Keystore** box and click **OK**.
- 4 Review the cautionary message, and if you accept it click **Yes**.

For more information about the tasks that VVM can perform:

See [“About Veritas Virtualization Manager \(VVM\)”](#) on page 144.

Removing Veritas Virtualization Manager (VVM)

Use standard Windows Add or Remove program functionality to remove the Veritas Virtualization Manager.

To remove the Veritas Virtualization Manager

- 1 From the Windows system, open the Control Panel.
- 2 Double-click the **Add or Remove Programs** icon.
- 3 Scroll down, and click **Veritas Virtualization Manager** to select it.
- 4 Click the **Remove** button.

Configuring virtual machines for high availability

This chapter contains the following topics:

- [About configuring virtual machines](#)
- [Configuring virtual machines for high availability using the Veritas Virtualization Manager](#)

About configuring virtual machines

After you install and configure VCS on the ESX Servers, you can configure the virtual machines for high availability. You create the virtual machines using VMware tools and then use the Veritas Virtualization Manager (VVM) to configure them.

A tight link exists between VMware clusters and VCS clusters. You must have an existing VMware cluster before you can use VVM to deploy virtual machines. VVM checks for the existence of the VMware cluster, and uses that foundation to configure highly available virtual machines.

Configuring virtual machines for high availability using the Veritas Virtualization Manager

You can take existing virtual machines that are part of a VMware cluster, and configure them into virtual machines that run under VCS. Use one of the following procedures:

- See [“To configure one virtual machine for high availability”](#) on page 155.
- See [“To configure multiple virtual machines for high availability”](#) on page 155.

Prerequisites for configuring virtual machines for high availability

Before you start, ensure that:

- The hardware for the cluster ready to use, which includes shared storage and networks that are visible from all nodes.
- Enough computing power is available for the virtual machines where you plan to add to the nodes.

You need to prepare the VMware configuration, in the following list ensure that:

- The VirtualCenter Server is configured and running.
- You have administrative access to the VirtualCenter Server.
- You have disabled VMware HA on the target clusters.
- The nodes are part of a VMware cluster.
- The nodes that are defined in the VCS cluster and the nodes that are defined in the VMware cluster must be identical.

You must also ensure that:

- VCS for VMware ESX is configured and running on your nodes.
- If you use the SSL certificate, it must be available for the Veritas Virtualization Manager.
See [“Preparing keystores”](#) on page 146.

To start the virtual machine deployment

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Enter the name of the VirtualCenter Server, your user name and password. If you use the SSL certificate the full path to it.
See [“To start Veritas Virtualization Manager”](#) on page 151.

To configure one virtual machine for high availability

- 1 Right-click the virtual machine that you want to make highly available. From the pull-down menu, select **Configure HA**.
- 2 Review the next window, and click the **Next** button.
- 3 In the next window, enter the user name and password for the VCS cluster. You only need to connect to a cluster once per VVM session.
- 4 Choose to configure a new administrator. This user is the person who you want to administer the application and the service group that VVM creates when you complete this wizard. Note the name and password for future reference.
Click the **Next** button to proceed.
- 5 Review the configuration summary.
 - Click on the summary item in the right-column to edit it. Note that some items cannot be edited.
 - Click the **Back** button to change any values.
 - Click the **Finish** button to finalize the configuration for the virtual machine.You have made the virtual machine highly available.

To configure multiple virtual machines for high availability

- 1 Right-click the cluster that contains the virtual machines that you want to make highly available. From the pull-down menu, select **Configure HA**.
- 2 Review the next window, and click the **Next** button.

- 3 Check the boxes next to the virtual machines that you want to configure. Click the **Next** button.
- 4 Review the names.
 - Click on an item to edit it. Note that some items cannot be edited.
 - Click the **Back** button to go back and change any values.
 - Click the **Next** button to finalize the configuration for the virtual machines.

You have now made the selected virtual machines highly available.

Reviewing the generated service groups

After you have used the Veritas Virtualization Manager to configure a virtual machine for high availability, VVM creates a service group. The high availability service group for a virtual machine has the following resources:

- an ESXVirtualMachine resource, which monitors the virtual machine and its operating system.
- a VMFSVolume resource, which monitors the virtual machine storage
- a VSwitch which monitors the virtual machine network
(or a Proxy resource if you already have the same VSwitch resource that is monitored through a different service group)

The ESXVirtualMachine resource depends on its storage and network resources. If any of the resources in this group fails for any reason, VCS moves the service group to next available node in the cluster. For example, in situations where the operating system freezes or crashes, the ESXVirtualMachine resource moves the entire virtual machine to another node.

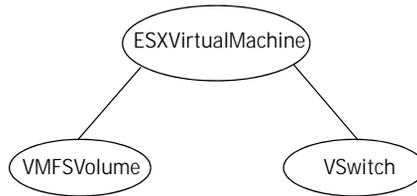
VVM also creates a service group administrator account for this service group, if you provided the administrator's information while using VVM.

Note that fail over of the virtual machine does not include applications that you want to monitor within the virtual machines. For application monitoring, see:

- [“Configuring applications and resources in Windows virtual machines”](#) on page 239
- [“Configuring applications and resources in Linux virtual machines”](#) on page 187

In [Figure 10-1](#), you can see a basic example of a resource dependency graph for a service group.

Figure 10-1 A service group with the ESXVirtualMachine, VSwitch, and VMFSVolume resources



Accessing the service groups

Although the Veritas Virtualization Manager creates a service group for you, use command line tools, Veritas Cluster Server Management Console, Cluster Server (Java Console) to manage resources. Note that you can only manage resources that are outside of virtual machines with these tools.

For information about using service groups, either through a CLI or GUI, refer to the *Veritas Cluster Server User's Guide*.

Verifying virtual machine failover

Verify the configuration in different situations.

Using a switch command

Switch the virtual machine to another node in the cluster to make sure the service group fails over. If all the resources are properly configured, the service group shuts down on the first node and comes up on the second node.

Other verification scenarios

In all of these verification scenarios, you are stopping or moving a virtual machine, or stopping a resource for that virtual machine. VCS should detect the failure, or the movement, and either fail over the effected virtual machine or take no action. The following list presents some quick testing scenarios:

- From outside of VCS control, stop the virtual machine. VCS should fail the virtual machine over to the other node.
- Boot the virtual machine through VCS by entering a `hagrp -online` command. Move the virtual machine to another node by shutting it down through VCS on the node where the virtual machine is running. Boot the virtual machine outside of VCS control on the other node—the service group comes online on that node.

- Trigger a VMotion for a virtual machine. When you trigger a VMotion for a virtual machine, VCS marks the service group, which contains the virtual machine, as offline on the first node. It then marks the service group as online on the target node.

Configuring virtual machines for disaster recovery

This chapter contains the following topics:

- [About VCS global clusters](#)
- [Setting up a global cluster manually](#)
- [Configuring virtual machines for disaster recovery using the Veritas Virtualization Manager](#)
- [Verifying virtual machine failover](#)
- [Disaster recovery best practices](#)

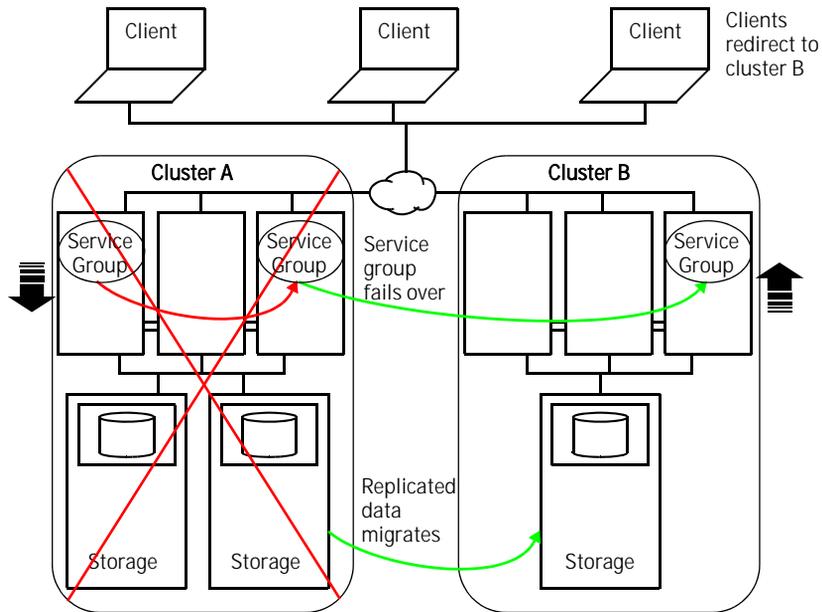
About VCS global clusters

Local clustering provides local failover for each site or building. Local clusters do not, however, provide protection against large-scale disasters for an entire city or region.

VCS global clusters ensure application availability by migrating service groups to remote clusters located considerable distances apart.

Take the example of an Oracle database configured in a VCS global cluster. Oracle is installed and configured in a virtual machine. The virtual machine is configured in a VCS service group. The service group is online on a system in cluster A and is configured to fail over globally, on clusters A and B.

Figure 11-1 Sample global cluster setup



VCS continuously monitors and communicates events between clusters. Inter-cluster communication ensures that the global cluster is aware of the state of global service group at all times.

In the event of a system or application failure, VCS fails over the virtual machine service group to another system in the same cluster. If the entire cluster fails, VCS fails over the service group to the remote cluster, which is part of the global cluster. VCS also redirects clients once the application is online on the new location.

Note: You cannot use raw device mapping (RDM) in a disaster recovery-enabled environment.

VCS global clusters: The building blocks

VCS extends clustering concepts to wide-area high availability and disaster recovery with the following building blocks.

Global service groups

A global service group is a regular VCS group with additional properties to enable wide-area failover. The global service group attribute `ClusterList` defines the list of clusters to which the group can fail over. The service group must be configured on all participating clusters and must have the same name on each cluster.

VCS agents manage replication during cross-cluster failover. You must configure a resource of type DNS to perform a canonical name update if cross-cluster failover spans subnets.

Global cluster management

VCS enables you to perform operations (online, offline, switch) on global service groups from any system in any cluster. You must log on with adequate privileges for cluster operations.

You can bring service groups online or switch them to any system in any cluster. If you do not specify a target system, VCS uses the `FailOverPolicy` to determine the system.

Management of remote cluster objects is aided by inter-cluster communication enabled by the wide-area connector (wac) process.

Resiliency and right-of-way

VCS global clusters maintain resiliency using the wide-area connector process and the ClusterService group. The wide-area connector process runs as long as there is at least one surviving node in a cluster.

The wide-area connector and notifier are components of the ClusterService group.

VCS framework

VCS agents manage external objects that are part of wide-area failover. These objects include replication, DNS updates, and so on. These agents provide a robust framework for specifying attributes and restarts, and can be brought online upon fail over.

Wide-area heartbeats

VCS requires at least one wide-area heartbeat going from each cluster to every other cluster. VCS starts communicating with a cluster only after the heartbeat reports an alive state. VCS uses the ICMP ping by default, the infrastructure for which is bundled with the product.

DNS agent

The DNS agent updates the canonical name-mapping in the domain name server after a wide-area failover. For more agent information, refer to the *Veritas Cluster Server Bundled Agents Reference Guide*.

VCS agents to manage replication

VCS supports several replication technologies. VCS agents manage the replication status between primary and secondary sites. Contact your Symantec sales representative for a list of replication technologies supported with this release of VCS.

Prerequisites for global clusters

This section describes the prerequisites for configuring global clusters.

Cluster setup

You must have at least two clusters to set up a global cluster. Every cluster must have the required licenses. A cluster can be part of one global cluster. VCS supports a maximum of four clusters participating in a global cluster.

Clusters must be running on the same platform; the operating system versions can be different. Clusters must be using the same VCS version.

Cluster names must be unique within each global cluster; system and resource names need not be unique across clusters. Service group names need not be unique across clusters; however, global service groups must have identical names.

Every cluster must have a valid virtual IP address, which is tied to the cluster. Define this IP address in the cluster's ClusterAddress attribute. This address is normally configured as part of the initial VCS installation. The IP address must have a DNS entry.

For remote cluster operations, you must configure a VCS user with the same name and privileges in each cluster.

ClusterService group

The ClusterService group must be configured with the wac, VSwitch, and IP resources. It is configured automatically when VCS is installed or upgraded, or by the GCO configuration wizard. The service group may contain additional resources for Veritas Cluster Server Management Console and notification, if these components are configured.

If you entered a license that includes VCS global cluster support during the VCS install or upgrade, the ClusterService group, including the wide-area connector process, is automatically configured.

Replication setup

You must set up real-time data replication between clusters. Contact your sales representative for a list of replication technologies supported with this release of VCS.

Setting up a global cluster manually

This section describes the steps for planning, configuring, and testing a global cluster. It describes an example of converting a virtual machine for local high availability in a VCS cluster to a highly available, disaster-protected infrastructure using a second cluster.

- [Configuring the ClusterService group](#)
- [Configuring replication](#)
- [Configuring the second cluster](#)
- [Creating the global service group](#)

Note: The procedure assumes your local cluster is set up and that you are replicating data between the local and remote clusters.

Configuring the ClusterService group

Configure the ClusterService group as described in this section.

To configure the ClusterService group

- 1 Create a service group called ClusterService.
- 2 Add a resource of type VSwitch to the service group. Name the resource csgnic. Set the value of the Device attribute to the name of the VSwitch. Configure other attributes, if desired.
- 3 Add a resource of type IP to the service group. Name the resource gcoip. Configure the following attributes for the resource:
 - Address—A virtual IP address assigned to the cluster.
 - Device—The name of the switch on the system. The device is defined as a local attribute for each system in the cluster.
 - NetMask—The subnet to which the virtual IP address belongs.
- 4 Link the VSwitch and IP resources such that the IP resource depends on the VSwitch resource.
- 5 Add a resource of type Application to the service group. Name the resource wac. Configure the following attributes for the resource:
 - StartProgram—"/opt/VRTSvcs/bin/wacstart"
 - StopProgram—"/opt/VRTSvcs/bin/wacstop"
 - MonitorProcesses— {" /opt/VRTSvcs/bin/wac" }
 - RestartLimit—3
- 6 Link the Application and IP resources, making Application the parent resource.
- 7 Enable both resources.
- 8 Bring the ClusterService service group online.

Sample configuration

```
group ClusterService (
  SystemList = { thorpc132 = 1, thorpc136 = 2 }
  PrintTree = 0
  AutoStartList = { thorpc132 }
  OnlineRetryLimit = 3
  OnlineRetryInterval = 120
)

Application wac (
  StartProgram = "/opt/VRTSvcs/bin/wacstart"
  StopProgram = "/opt/VRTSvcs/bin/wacstop"
  MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
  RestartLimit = 3
)

IP webip (
  Device @thorpc132 = vswif0
  Device @thorpc136 = vswif0
  Address = "10.182.146.154"
  NetMask = "255.255.252.0"
)

VSwitch csgnic (
  VirtualSwitch @thorpc132 = vSwitch0
  VirtualSwitch @thorpc136 = vSwitch0
)

wac requires webip
webip requires csgnic
```

Configuring replication

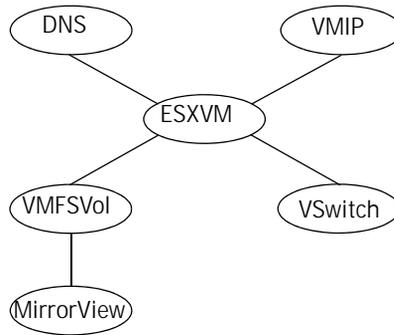
VCS supports several replication solutions for global clustering. Contact your Symantec sales representative for the solutions supported by VCS. This example describes how to set up replication using VCS agent for MirrorView.

Adding the resources for replication

- 1 Add a resource of type MirrorView to the service group.
- 2 Configure the attributes of the MirrorView resource.
 - NaviCliHome—NaviCLI installation directory
 - LocalArraySPNames—The list of storage processors within the array to which the local hosts are connected. Can be names or IP addresses.
 - RemoteArraySPNames—The list of storage processors within the array to which the remote hosts are connected. Can be names or IP addresses.
 - Mode—The replication mode, which is either: sync or async.

- GrpName—The name of the consistency group to which the mirrors belong. This function applies only if the mode is async.
 - MirNames—This function lists individual mirrors that are a part of the replication relationship and managed by VCS. This attribute is ignored if you specify the GrpName attribute.
 - SplitTakeover—This integer indicates whether VCS should forcefully promote a secondary to a primary.
- 3 Add a resource of type DNS to the service group and configure its attributes:
- Domain—Domain name. For example, veritas.com.
 - Alias—Alias to the canonical name. For example, www.
 - Hostname—Canonical name of a system or its IP address. For example, mtv.veritas.com.
 - TTL—Time To Live (in seconds) for the DNS entries in the zone being updated. Default value: 86400.
 - StealthMasters—List of primary master name servers in the domain. This attribute is optional if the primary master name server is listed in the zone's NS record. If the primary master name server is a stealth server, the attribute must be defined.
Note that a stealth server is a name server that is authoritative for a zone but is not listed in the zone's NS records.
Optionally, configure the TsigKeyFile attribute for secure DNS updates.
See the *Veritas Cluster Server Bundled Agents Reference Guide*.
- 4 Add a resource of type VMIP and configure its attributes.
- VMwareResName—The name of the VCS resource that manages the virtual machine.
 - IPAddress—The IP address that is assigned to the virtual machine interface.
 - MACAddress—The MAC address of the virtual NIC
 - NetMask—The subnet mask that is associated with the IP address. You must specify this value in decimal (base 10).
 - Gateway—The default gateway for the virtual machine.
 - DNS—List of DNS servers in the required search order.
- 5 Create the following resource dependencies:
- VMFSVol resource depends on the MirrorView resource.
 - DNS resource depends on the ESXVirtualMachine resource.
 - VMIP resource depends on the ESXVirtualMachine resource.

Figure 11-2 Sample dependency graph where ESXVM stands for the ESXVirtualMachine resource and VMFSVol stands for the VMFSVolume resource



Configuring the second cluster

- 1 Configure the ClusterService group in the second cluster.
 See “[Configuring the ClusterService group](#)” on page 164.
- 2 Create a configuration that is similar to the one in the first cluster. You can do this by copying the configuration of the service group from the main.cf file in the primary cluster to the secondary cluster.
- 3 Make appropriate changes to the configuration. For example, you must modify the SystemList attribute to reflect the systems in the secondary cluster. Also, you must modify the VMFSVolume and replication resources to point to the local storage at the remote cluster.
 Make sure that the name of the service group is identical in both clusters.
- 4 To assign remote administration privileges to users, configure users with the same name and privileges on both clusters.

Linking clusters

Once the VCS and the replication infrastructure has been set up at both sites, you must link the two clusters.

Before linking clusters, verify the virtual IP address for the ClusterAddress attribute for each cluster is set. Use the same IP address as the one assigned to the IP resource in the ClusterService group.

To add a remote cluster to a global environment

- 1 In the **Cluster:Summary** view, in the **Configuration** task panel, click **Add/Delete Remote Cluster**.
- 2 In the **Remote Cluster Configuration** wizard, read the introductory information and then click **Next**.
- 3 In the **Configuration Options** dialog box, click **Add Cluster** and then click **Next**.
- 4 Provide cluster connection information:
In the **Connection Details** dialog box, specify the following details for the connection to the remote cluster and then click **Next**:
 - A name or address
Enter the IP address of the cluster, the IP address of a cluster system, or the name of a cluster system.
 - The port
Verify the port number. The default is 14141.
 - An administrator user name and password.
Enter an administrator-level user name and password that is valid on the remote cluster.
- 5 Click **Finish**.
The cluster icon changes to indicate that the cluster is a global cluster.

Creating the global service group

Use the Veritas Cluster Server Management Console to configure the global service group. The Global Group Configuration wizard configures a service group in a local cluster as a global service group.

To convert a service group on a local cluster to a global service group

- 1 Start Veritas Cluster Server Management Console and log on to the cluster.
- 2 In the **Cluster:Summary** view, in the **Groups Listing** table, click the linked name of the service group that you want to convert.
This service group should already have been commonly configured on at least one local and one remote cluster.
- 3 In the **Group:Summary** view, in the **Configuration** task panel, click **Configure Global Group**.
- 4 In the **Global Group Configuration** wizard, read the introductory information and click **Next**.

- 5 In the **Cluster List Configuration** dialog box, under **Available Clusters**, select the clusters on which the global service group can come online. To select a cluster, click the right-arrow button to move the cluster name under **Selected Clusters**.
- 6 Select the policy for service group failover and then click **Next**:
 - **Manual** prevents a service group from automatically failing over to another cluster.
 - **Auto** enables a service group to automatically fail over to another cluster if it is unable to fail over within the cluster, or if the entire cluster faults.
 - **Connected** enables a service group to automatically fail over to another cluster if it is unable to fail over within the cluster.
- 7 In this step, you update the cluster list of remaining instances of the selected global service group. To perform the update, you must first verify or supply the authentication credentials for each remaining global cluster in the list. The Cluster Management Console can then connect to those clusters and update the lists.

In the **Remote Cluster Configuration** dialog box, verify the required information for the remaining remote clusters and then click **Next**. To change authentication information, click a cluster name under **Existing Clusters** and then enter the authentication information in the fields to the right. The requisite information in this dialog box varies depending upon whether or not the cluster is secure (uses an authentication broker).
- 8 Click **No** if you want the operation to be completed only if the wizard can connect to all selected clusters.
- 9 Click **Next**.
- 10 Click **Finish**.

Configuring virtual machines for disaster recovery using the Veritas Virtualization Manager

You can enable existing virtual machines for disaster recovery using the Veritas Virtualization Manager. Before you start you need to make sure that replication exists between the primary site and the secondary site.

See [“Prerequisites for global clusters”](#) on page 162.

Overview of tasks

Table 11-1 Configuration tasks

Task	Reference
Review prerequisites.	See “Prerequisites for configuring virtual machines for disaster recovery” on page 171.
Set up secure DNS update.	See “Setting up secure DNS update” on page 172.
Use VVM to configure virtual machines for disaster recovery.	See “Using Veritas Virtualization Manager to configure virtual machines for disaster recovery” on page 173.
On the primary site, deploy the Tools and applications.	See “Deploying VCS components on the virtual machines in the primary site” on page 175.
Confirm the availability of the service group.	See “Confirming service group availability” on page 175.
Reverse the direction of replication.	See “Reversing the direction of replication” on page 176.
On the secondary site, use VVM to configure the virtual machines.	See “Using VVM to configure virtual machines for disaster recovery on the secondary site” on page 176.
Deploy VCS components.	See “Deploying VCS components on virtual machines in the secondary site” on page 177.
Verify the service group and use the Global Wizard.	See “Verifying the service group on the secondary site and using the Global Wizard” on page 177.
Perform post-failover actions.	See “Post-failover actions” on page 177.
Review the service group that VVM creates.	See “Reviewing the generated service groups” on page 178.
Work with the service groups.	See “Accessing the service groups” on page 179.
Verify machine fail over functionality.	See “Verifying virtual machine failover” on page 179.

Prerequisites for configuring virtual machines for disaster recovery

Before you start ensure that:

- The hardware for the cluster is set up and ready to use, which includes shared storage and networks that are visible from all nodes.
- The replication is set up and ready to use. Note that a replication solution is required for disaster recovery. Contact your Symantec representative for a list of supported replication arrays.

You need to prepare the VMware configuration, in the following list ensure that:

- The VirtualCenter Server is configured and running.
- You have administrative access to the VirtualCenter Server.
- You have disabled VMware HA on the VMware clusters.
- The nodes are part of a VMware cluster.
- Ensure that all the ESX nodes that form a VCS ESX cluster are present and configured in the VirtualCenter Server as a corresponding VC Server cluster.
- See [“Disaster recovery best practices”](#) on page 180.

You must also ensure that:

- VCS for VMware ESX is configured and running on your nodes.
- The SSL certificate for the Veritas Virtualization Manager is available. See [“Preparing keystores”](#) on page 146.
- You have the array-specific information ready for the following arrays:
 - EMC MirrorView
 - Local SP Address 1—The first local storage processor’s name or IP address
 - Local SP Address 2—The second local storage processor’s name or IP address
 - Remote SP Address 1—The first remote storage processor’s name or IP address
 - Remote SP Address 2—The second remote storage processor’s name or IP address
 - EMC SRDF
 - Symmetrix CLI Home—The path to the bin directory that contains the Symmetrix command line interface.
 - Hitachi TrueCopy does not need any input.

- IBM MetroMirror
 - Hostname for HMC1—The IP address or host name of the primary management console
 - IBM CLI location—The path to the command line interface
 - Password file—The password file that contains your password.
 - Admin User—The user name to issue commands from the command line.
 - Remote Storage Image ID—The image ID of the remote storage

Setting up secure DNS update

VVM requires a secure key to ensure security and thwart spoofing. You need to create a TSIG key (Transaction Signature) as specified in RFC 2845. TSIG is a shared key message authentication mechanism available in DNS. A TSIG key provides a means to authenticate and verify the validity of DNS data exchanged, using a shared secret key between a resolver and either one or two servers.

Setting up secure updates using TSIG keys

In the following example, the domain is veritas.com.

To use secure updates using TSIG keys

- 1 Run the `dnssec-keygen` command with the HMAC-MD5 option to generate a pair of files that contain the TSIG key:

```
# dnssec-keygen -a HMAC-MD5 -b 512 -n HOST veritas.com.  
Kveritas.com.+157+00000
```

- 2 Open the `Kveritas.com.+157+00000.key` file. After running the `cat` command, the contents of the file resembles:

```
# cat Kveritas.com.+157+00000.key  
veritas.com. IN KEY 512 3 157 +Cdjlkef9ZTSeixERZ433Q==
```

- 3 Copy the shared secret (the TSIG key), which looks like:

```
+Cdjlkef9ZTSeixERZ433Q==
```

- 4 Configure the DNS server to only allow TSIG updates using the generated key. Open the `named.conf` file and add these lines.

```
key veritas.com. {  
    algorithm hmac-md5;  
    secret "+Cdjlkef9ZTSeixERZ433Q==";  
};
```

Where `+Cdjlkef9ZTSeixERZ433Q==` is the key.

- 5 In the `named.conf` file, edit the appropriate zone section and add the `allow-update` sub-statement to reference the key:

```
allow-update { key veritas.com. ; } ;
```

- 6 Save and restart the named process.
- 7 Place the files containing the keys on each of the nodes that is listed in your group's SystemList. The DNS agent uses this key to update the name server. Copy both the private and public key files on to the node. A good location is in the /var/tsig/ directory.
- 8 Set the TSIGKeyFile attribute for the DNS resource to specify the file containing the private key.

```
DNS www (
Domain = "veritas.com"
Alias = www
Hostname = north
TSIGKeyFile a= "/var/tsig/Kveritas.com.+157+00000.private"
)
```

Using Veritas Virtualization Manager to configure virtual machines for disaster recovery

Use the Veritas Virtualization Manager (VVM) to configure the virtual machines for disaster recovery.

To start the virtual machine deployment

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Enter the name of the VirtualCenter Server, your user name and password, and the full path to the SSL certificate. See ["To start Veritas Virtualization Manager"](#) on page 151.

To configure a virtual machine for disaster recovery

- 1 Right-click the virtual machine. From the pull-down menu, select **Configure DR**.
- 2 In the next window, enter the user name and password for the VCS cluster. You only need to connect to a cluster once per VVM session. Click the **OK** button.
- 3 Review the information page. Click **Next**.

- 4 Select the array from the pull-down menu. For the other arrays, enter the following information:
 - EMC MirrorView (Clariion)
 - Local SP Address 1—The first local storage processor's name or IP address
 - Local SP Address 2—The second local storage processor's name or IP address
 - Remote SP Address 1—The first remote storage processor's name or IP address
 - Remote SP Address 2—The second remote storage processor's name or IP address
 - EMC SRDF
 - Symmetrix CLI Home—The path to the bin directory that contains the Symmetrix command line interface (for example `/usr/symcli`).
 - Hitachi True Copy does not need any input.
 - IBM MetroMirror
 - Hostname for HMC1—The IP address or host name of the primary management console
 - IBM CLI location—The path to the command line interface
 - Password file—The password file that contains your password.
 - Admin User—The user name to issue commands from the command line.
 - Remote Storage Image ID—The image ID of the remote storage
- 5 Enter the requested information for the DNS agent.
 - Domain—The string representing the domain name.
 - Alias—The string representing the alias to the canonical name.
 - Hostname—A string representing canonical name of a system.
 - IPAddress—Specifies the IP address that is assigned to the hostname.
 - TSIGKeyFile—Specifies the absolute path to the file containing the private TSIG (Transaction Signature) key.

- 6 Enter the required information for VMIP agent.
 - Gateway—The default gateway for the virtual machine.
 - Primary DNS—The primary DNS for the virtual machine.
 - Secondary DNS—The secondary DNS for the virtual machine.
 - Tertiary DNS—Other DNS for the virtual machine are optional.
 - Additional MACs—if the virtual machine has multiple virtual NICs, enter the IP address and subnet masks for the additional MAC addresses.
- 7 Choose to configure a new VCS user. This user administers the service group that VVM creates when you complete this wizard. Note the name and password for future reference.
- 8 Review the deployment summary.
 - Click on a summary item to edit it. Note that some items cannot be edited.
 - Click the **Back** button to change any previous values.
 - Click the **Finish** button to finalize the VCS configuration for the virtual machine.

Repeat this process for each virtual machine that you want in the cluster.

Deploying VCS components on the virtual machines in the primary site

You must install the Veritas Virtual Machine Tools and configure the applications on the virtual machine. Perform the procedures in the following chapter depending on operating system:

- [Chapter 12, “Configuring applications and resources in Linux virtual machines”](#) on page 187.
- [Chapter 13, “Configuring applications and resources in Windows virtual machines”](#) on page 239.

Confirming service group availability

Check to see if the service group, which you have just created with VVM, can come online. Once you have established that it comes online, bring the service group offline.

See [“Accessing the service groups”](#) on page 179.

Reversing the direction of replication

You now need to reverse the direction of replication between your primary and secondary arrays. Make sure that the service group for this virtual machine is offline on the primary site. You must reverse replication direction for all the datastores used by the virtual machines that you have configured.

Contact your Symantec sales representative for a list of replication technologies supported with this release of VCS. Consult your replication solution's manual for information on reversing replication direction.

Using VVM to configure virtual machines for disaster recovery on the secondary site

Before you configure the virtual machines on the secondary site, perform the following tasks using the VMware commands:

- Rescan all the storage adapters.
- Verify that all datastores are visible on all nodes on the secondary site.
- Register a virtual machine on a node in the secondary site.

Use VVM to log into the cluster on the secondary site and to configure the virtual machines for disaster recovery.

Refer back to the previous procedure, as follows:

- ["Using Veritas Virtualization Manager to configure virtual machines for disaster recovery"](#) on page 173

VVM creates a new service group on the secondary site.

Note: You must give this new service group the same name that you used for your primary site.

Deploying VCS components on virtual machines in the secondary site

You must now configure Veritas Virtual Machine Tools on the secondary site.

- For Linux:
 - See [“Configuring Veritas Virtual Machine Tools”](#) on page 231.
 - See [“Validating the configuration of Veritas Virtual Machine Tools”](#) on page 232.
- For Windows:
 - See [“Configuring Veritas Virtual Machine Tools”](#) on page 247.
 - See [“Validating the configuration of Veritas Virtual Machine Tools”](#) on page 250.

Verifying the service group on the secondary site and using the Global Wizard

Verify that the service group can come online on the secondary site. You are now ready to use the Global Wizard, the wizard configures the service group as a global group.

See [“Creating the global service group”](#) on page 168.

Post-failover actions

After a disaster recovery, or after switching the global service group between the clusters, you must reconfigure Veritas Virtual Machine Tools on the site that the cluster has failed over to.

- For Linux:
 - See [“Configuring Veritas Virtual Machine Tools”](#) on page 231.
 - See [“Validating the configuration of Veritas Virtual Machine Tools”](#) on page 232.
 - Apply the changes to the configuration by running the `vcsag_config.pl` program with the `-apply` option, as follows:

```
# /opt/VRTSvcs/bin/vcsag_config.pl -apply
```
- For Windows:
 - See [“Configuring Veritas Virtual Machine Tools”](#) on page 247.
 - See [“Validating the configuration of Veritas Virtual Machine Tools”](#) on page 250.

- Apply the changes to the configuration by running the `vcsag_config.pl` program with the `-apply` option, as follows:

```
C:\> "%VRTS_PERL_BIN%\perl" "%VCS_HOME%\bin\vcsag_config.pl" -apply
The default for VRTS_PERL_BIN is C:\Program
Files\Veritas\VRTSPer\bin. The default for the VCS_HOME is
C:\Program Files\Veritas\Cluster Server.
```

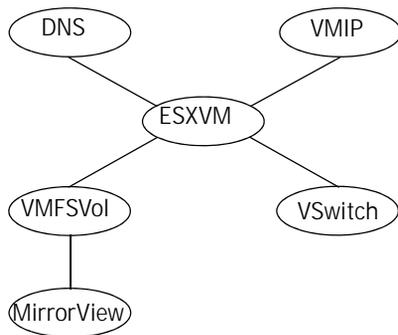
Reviewing the generated service groups

After you have used the Veritas Virtualization Manager to configure a virtual machine for disaster recovery, VVM creates a service group. VVM also creates a service group administrator if you provided the administrator's information while using VVM.

The service group that VVM creates has an `ESXVirtualMachine` resource that depends on its storage (`VMFSVolume`) and network (`VSwitch`) resources. Further, the `VMFSVolume` resource depends on the replicated array. The `VMIP` and `DNS` resources are critical for moving the service group across networks and subnets. These two resources depend on the `ESXVirtualMachine` resource itself. If any of the resources in this group fails for any reason, VCS moves the entire service group to another available node. Note that fail over of the virtual machine does not include applications that you want to monitor within the virtual machines.

In [Figure 11-3](#), you can see a basic example of a resource dependency graph for a service group.

Figure 11-3 Sample dependency graph where `ESXVM` stands for the `ESXVirtualMachine` resource and `VMFSVol` stands for the `VMFSVolume` resource



Accessing the service groups

Although the Veritas Virtualization Manager creates a service group for you, use command line tools, Veritas Cluster Server Management Console, Cluster Server (Java Console) to manage resources. Note that you can only manage resources that are outside of virtual machines with these tools.

For information about using service groups, either through a CLI or GUI, see the *Veritas Cluster Server User's Guide*.

Verifying virtual machine failover

Verify the configuration in different situations.

Using the switch command

Switch the virtual machine to another node in the cluster to make sure the service group fails over. If all the resources are properly configured, the service group shuts down on the first node and comes up on the second node.

Other verification scenarios

In all of these verification scenarios, you are stopping or moving a virtual machine, or stopping a resource for that virtual machine. VCS should detect the failure, or the movement, and either fail over the effected virtual machine or take no action. The following list presents some quick testing scenarios:

- From outside of VCS control, stop the virtual machine. VCS should fail the virtual machine over to the other node.
- Enter a `hagrps -online` command to boot the virtual machine through VCS. Move the virtual machine to another node by shutting it down through VCS on the node where the virtual machine is running. Boot the virtual machine outside of VCS control on the other node—the service group comes online on that node.
- Trigger a VMotion for a virtual machine. When you trigger a VMotion for a virtual machine, VCS marks the service group that contains the virtual machine, as offline on the first node. It then marks the service group as online on the target node.

Disaster recovery best practices

The following are suggestions for best practices in VCS for ESX disaster recovery environment.

General best practices

The following are general best practices:

- All the LUNs that a datastore spans must belong to a single replication unit. Examples of replication units are the device groups used in SRDF or Hitachi TrueCopy.
- A replication unit can contain one or more datastores.
- All VCS global service groups must have ClusterFailoverPolicy attribute set to Manual.

Supported configurations

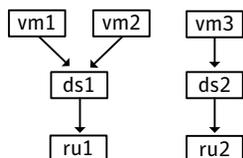
Symantec recommends that you use Veritas Virtualization Manager (VVM) to generate the VCS configuration. When VVM generates the configuration, it creates a service group for each virtual machine.

In the VMware configurations that follow, virtual machines share the datastores. For virtual machines that share a datastore, the service groups have identical VMFSVolume resources. The same holds true for the replication unit (RU). For identical resources, keep attribute values identical across service groups.

Scenario one

In this configuration, virtual machines vm1 and vm2 are on datastore ds1 and a single virtual machine vm3 is on ds2. The replication unit ru1 has the same set of LUNs that the datastore ds1 spans. The same relationship applies to ru2 and ds2.

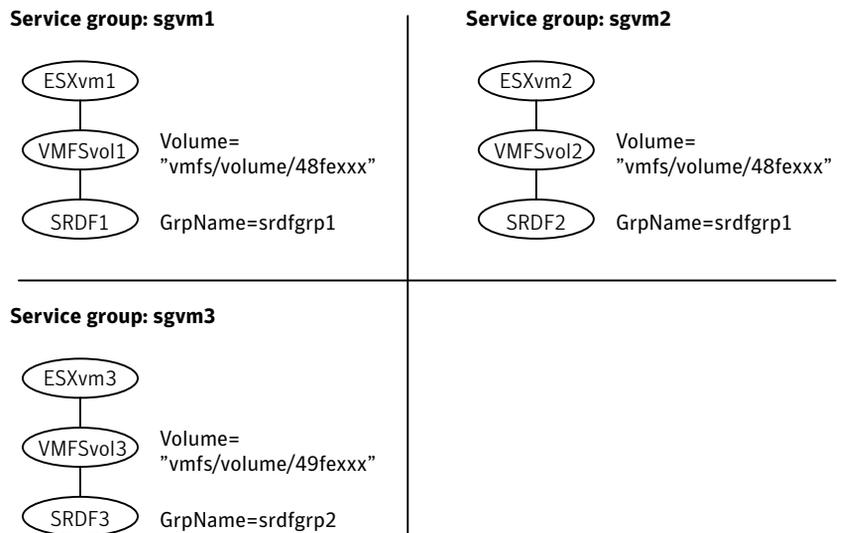
Figure 11-4 One or more virtual machines to a single 1-1 replication unit mapped datastore



The service group `sgvm1` is the VCS configuration that corresponds to this VMware configuration. Service groups `sgvm1` and `sgvm2` have ESXVirtualMachine resources `ESXvm1` and `ESXvm2` corresponding to the virtual machines `vm1` and `vm2`. Since they share the same datastore and the same replication unit, the VMFSVolume and the replication resources (SRDF in this case) have identical attribute values.

The virtual machine `vm3` is on a different datastore and replication unit and has different attribute values for the VMFSVolume and replication resources.

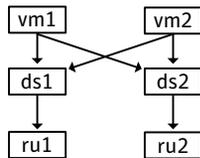
Figure 11-5 Service groups for this configuration



Scenario two

In this configuration, virtual machines vm1 and vm2 share datastores ds1 and ds2. The replication unit ru1 has the same set of LUNs that the datastore ds1 spans. The same relationship applies to ru2 and ds2.

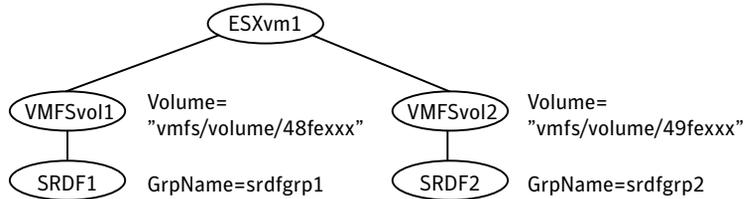
Figure 11-6 One or more virtual machines to multiple 1-1 replication unit mapped datastores



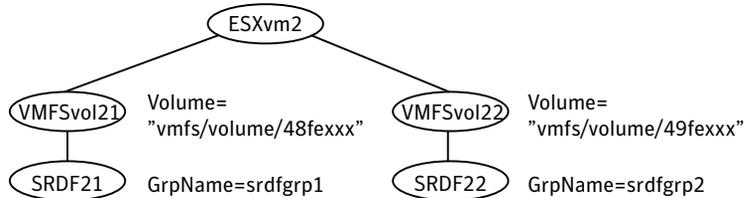
The VCS configuration corresponding to this VMware configuration is shown in service group: sgvm1. Service groups sgvm1 and sgvm2 have ESXVirtualMachine resources ESXvm1 and ESXvm2 corresponding to the virtual machines vm1 and vm2. Since they share the same datastores and the same replication units, the VMFSVolume and the replication resources (SRDF in this case) have identical attribute values.

Figure 11-7 Service groups for this configuration

Service group: sgvm1



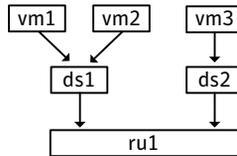
Service group: sgvm2



Scenario three

In this configuration, virtual machines vm1 and vm2 share datastores ds1 and virtual machine vm3 is on datastore ds2. All the datastores share the same replication unit ru1.

Figure 11-8 One or more virtual machines to a single datastore; multiple datastores to a replication unit

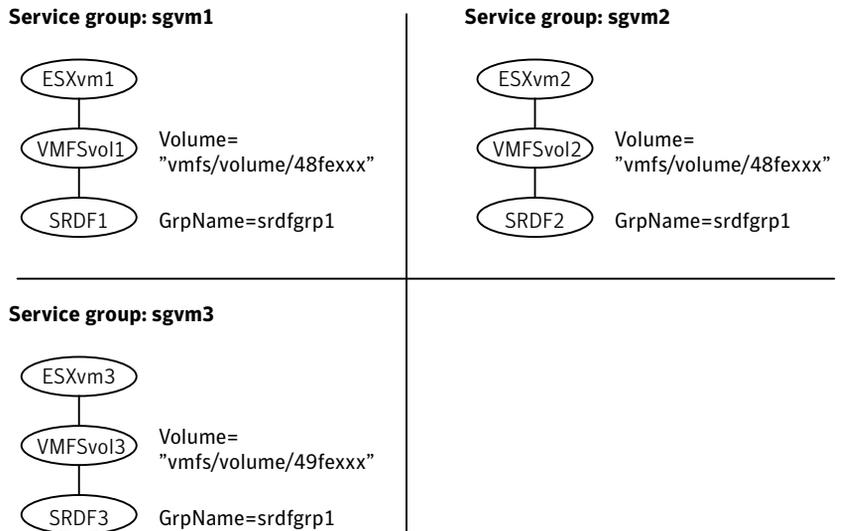


The VCS configuration corresponding to this VMware configuration is shown in service group: sgvml. Service groups sgvml and sgvml2 have ESXVirtualMachine resources ESXvm1 and ESXvm2 corresponding to the virtual machines vm1 and vm2. Since they share the same datastore, the VMFSVolume resources have identical attribute values.

The virtual machine vm3 is on a different datastore and has different attribute values for the VMFSVolume resource.

All the virtual machines share the same replication unit and have the same attribute values for the replication resources.

Figure 11-9 Service groups for this configuration



Planned failover across sites

A planned failover of a service group across sites requires the following steps:

- At the primary site where the service group is currently online, take all the other service groups offline that share a datastore or a replication unit with this service group.
- At the secondary site (the destination), bring the service group online. Wait for the service group to be completely online before you bring any of the other service that groups that share a datastore or replication unit online.

You must perform this sequence of steps to ensure that multiple replication resources do not simultaneously try to reverse the direction of the replication during the failover.

If you want the virtual machines to failover and failback between sites independently, put them on separate replication units.

Failover after a disaster

After a disaster at the primary site, global service groups that were online at the primary site are not automatically brought online at the secondary due to the attribute "ClusterFailoverPolicy = Manual" setting. Refer to the "Planned failover across sites" section above for steps on how to bring the service groups online at the secondary site.



Configuring applications in virtual machines

This section contains the following chapters:

- [Chapter 12, “Configuring applications and resources in Linux virtual machines”](#) on page 187
- [Chapter 13, “Configuring applications and resources in Windows virtual machines”](#) on page 239

Configuring applications and resources in Linux virtual machines

This chapter contains the following topics:

- [About VCS components for virtual machines running Linux](#)
- [How VCS monitors applications and resources on virtual machines](#)
- [Installing the applications](#)
- [Installing Veritas Virtual Machine Tools](#)
- [Configuring application and resource monitoring inside of virtual machines](#)
- [Applying the configuration and creating the corresponding GuestOSApp resource](#)
- [Removing the Veritas Virtual Machine Tools](#)

About VCS components for virtual machines running Linux

VCS for VMware ESX provides agents to monitor applications that run inside virtual machines. When the agent detects an application or resource fault, the agent takes actions to communicate the state of the resource to VCS running on the ESX Server node.

Certain VCS agents also support the ability to detect administrative intervention. When an administrator gracefully shuts down an application, VCS correctly does not initiate failover.

VCS provides agents to monitor the following applications that run on virtual machines running Linux:

Application	Agent information
Oracle	<ul style="list-style-type: none">■ Monitors Oracle instances.■ Detects a graceful administrative shutdown of Oracle.■ Supports detailed monitoring.■ See “About the VCS agent for Oracle” on page 190.
Apache Web Server	<ul style="list-style-type: none">■ Monitors Apache processes.■ Detects a graceful administrative shutdown of Apache processes.■ Supports detailed monitoring.■ See “About the VCS agent for the Apache Web server” on page 199.
SAP NetWeaver	<ul style="list-style-type: none">■ Monitors SAP NetWeaver instances.■ Detects a graceful administrative shutdown of SAP NetWeaver processes.■ Supports detailed monitoring.■ See “About the VCS agent for SAP NetWeaver” on page 205.
WebLogic Server	<ul style="list-style-type: none">■ Monitors WebLogic Server instances.■ Detects a graceful administrative shutdown of WebLogic Server processes.■ Supports detailed monitoring.■ See “About the VCS agent for WebLogic Server” on page 212.

VCS provides the following agents to monitor mount points and applications:

Agent	Agent information
Mount	<ul style="list-style-type: none"> ■ Monitors mount points. ■ Supports detailed monitoring. ■ See “About the Mount agent” on page 224.
Application	<ul style="list-style-type: none"> ■ Monitors applications. ■ See “About the Application agent” on page 220.

If a monitored resource or application fails, the corresponding agent communicates this state change to VCS on the ESX Server node. VCS can then fail over the virtual machine that runs the application onto another node.

About monitoring levels

VCS agents provide two levels of monitoring:

- **Basic monitoring**
Checks for running processes.
- **Detailed monitoring**
Performs application-specific tasks to check the application’s health. For example, the VCS agent for Oracle performs a transaction on the database and checks to see if the transaction succeeds.

All VCS agents provide basic monitoring capabilities. Some agents provide detailed monitoring capabilities.

Supported software

VCS 5.1 MP2 for VMware ESX supports the following software for virtual machines:

- | | |
|-------------------------|--|
| Guest operating systems | <ul style="list-style-type: none"> ■ Red Hat Enterprise Linux 4 (Update 5) ■ SUSE Linux Enterprise Server 9 with SP4 ■ SUSE Linux Enterprise Server 10 with SP1 |
|-------------------------|--|

- Applications
- Oracle 10g
 - Apache Web server:
 - Apache HTTP server 1.3, 2.0, and 2.2
 - IBM HTTP Server 1.3 and 2.0
 - SAP R/3-4.6C with a 4.6D Kernel, 4.6D, 4.7 Enterprise Version
 - SAP Web AS-6.20, 6.40, 7.00
 - SAP NetWeaver-2004, 2004s
 - WebLogic Server 9.0, 9.1, 9.2, 10.0

About the VCS agent for Oracle

The VCS agent for Oracle contains two agents that work together to make Oracle highly available:

- The Oracle agent that monitors the Oracle database processes.
- The Netlsnr agent that monitors the listener process.

VCS agent for Oracle detects graceful shutdown for Oracle 10g only. When an administrator brings down Oracle 10g gracefully, the agent does not trigger a resource fault even though Oracle is down. The agent provides this intentional offline functionality only when the Health check monitoring is enabled.

Agent functions

The VCS agent for Oracle supports the Monitor agent function.

- Oracle agent
 - The Monitor agent function verifies the status of the Oracle processes. The Oracle agent provides two levels of monitoring: basic and detailed. By default, the agent performs basic monitoring. Set the DetailMonitor attribute to 1 to enable detailed monitoring for Oracle.
 - The basic monitoring mode has two options: Process and Health check. The value of MonitorOption attribute is set to 1 by default for health check monitoring. If you want to use process monitoring, you must change the MonitorOption attribute value to 0.

MonitorOption value	Description
---------------------	-------------

- | | |
|---|---------------|
| 0 | Process check |
|---|---------------|
- The agent scans the process table for the ora_dbw, ora_smon, ora_pmon, and ora_lgwr processes to verify that Oracle is running.

MonitorOption value	Description
---------------------	-------------

1	Health check
(Default)	Retain the default value of the MonitorOption attribute (1) to use intentional offline functionality. The agent uses the Health Check APIs from Oracle to monitor the SGA and retrieve the information about the instance.

In the detailed monitoring mode, the agent performs a transaction on a test table in the database to ensure that Oracle functions properly.

■ **Netlsnr agent**

The Monitor agent function verifies the status of the listener process. The Netlsnr agent provides two levels of monitoring: basic and detailed.

- In the basic monitoring mode, which is the default behavior, the agent scans the process table for the tnslsnr process to verify the listener process is running.
- In the detailed monitoring mode, the agent uses the lsnrctl status \$LISTENER command to verify the status of the Listener process.

How the agent handles Oracle error codes during detail monitoring

The VCS agent for Oracle handles Oracle errors during detail monitoring. The agent classifies Oracle errors according to their severity and associates predefined actions with each error code.

The agent includes a reference file called oraerror.dat, which lists Oracle errors and the action to be taken when the error is encountered.

The file stores information in the following format:

```
Oracle_error_string:action_to_be_taken
```

For example:

```
01035:WARN
01034:FAILOVER
```

Table 12-1 lists the predefined actions that the agent takes when an Oracle error is encountered.

Table 12-1 Predefined agent actions for Oracle errors

Action	Description
IGNORE	<p>Ignores the error.</p> <p>When the VCS Agent for Oracle encounters an error that does not have a matching error code in the oraerror.dat file, then the agent ignores the error.</p>
UNKNOWN	<p>Marks the resource state as UNKNOWN and sends a notification if the Notifier resource is configured. See the <i>Veritas Cluster Server User's Guide</i> for more information about VCS notification.</p> <p>This action is typically associated with configuration errors or program interface errors.</p>
WARN	<p>Marks the resource state as ONLINE and sends a notification if the Notifier resource is configured.</p> <p>This action is typically associated with errors due to exceeded quota limits, session limits/restricted sessions so on.</p>
FAILOVER (Default)	<p>Marks the resource state as OFFLINE. This faults the resource. The GuestOSApp agent fails over the virtual machine.</p> <p>This is agent's default behavior. If the file oraerror.dat is not available, the agent assumes this default behavior for every Oracle error encountered.</p>
NOFAILOVER	<p>Freezes the service group temporarily and marks the resource state as OFFLINE. The agent also sends a notification if the Notifier resource is configured.</p> <p>This action is typically associated with errors that are not system-specific. For example, if a database does not open from a node due to corrupt Oracle files, failing it over to another node does not help.</p>

State definitions

- ONLINE
Indicates that Oracle is running.
- OFFLINE
Indicates that Oracle is not running.
Can also indicate that the administrator stopped Oracle gracefully. The Oracle agent can detect that the administrator has gracefully stopped it when the MonitorOption attribute is set to 1 (when health check monitoring is enabled).

- **FAULTED**
 Indicates that the application crashed or unexpectedly went offline.
- **UNKNOWN**
 Indicates that a problem exists with the configuration.

Oracle agent attributes

[Table 12-2](#) lists the required attributes for Oracle agent.

Table 12-2 Oracle agent required attributes

Required attributes	Definition
Sid	The variable \$ORACLE_SID that represents the Oracle instance. The Sid is considered case-sensitive by the Oracle agent and by the Oracle database server. Type and dimension: string-scalar
Owner	The Oracle user, as the defined owner of executables and database files in /etc/passwd. Type and dimension: string-scalar
Home	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Note: Do not append a slash (/) at the end of the path. Type and dimension: string-scalar

[Table 12-3](#) lists the optional attributes for Oracle agent.

Table 12-3 Oracle agent optional attributes

Optional Attributes	Definition
StartUpOpt	This attribute is disabled. Accept the default setting.
ShutDownOpt	This attribute is disabled. Accept the default setting.

Table 12-3 Oracle agent optional attributes

Optional Attributes	Definition
EnvFile	<p>The full path name of the file that is sourced by the agent function scripts. This file contains the environment variables set by the user for the Oracle database server environment such as LD_LIBRARY_PATH, NLS_DATE_FORMAT, and so on.</p> <p>The syntax for the contents of the file depends on the login shell of Owner. File must be readable by Owner. The file must not contain any prompts for user input.</p> <p>Type and dimension: string-scalar</p>
Pfile	<p>The name of the initialization parameter file with the complete path of the startup profile.</p> <p>You can also use the server parameter file. Create a one-line text initialization parameter file that contains only the SPFILE parameter. See the Oracle documentation for more information.</p> <p>Type and dimension: string-scalar</p>
AutoEndBkup	<p>This attribute is disabled. Accept the default setting.</p>
MonitorOption	<p>Monitor options for the Oracle instance. This attribute can take values 0 or 1.</p> <ul style="list-style-type: none"> ■ 0 - Process check monitoring (recommended) ■ 1 - Health check monitoring <p>Default: 0</p> <p>Note: You must set the value of this attribute to 1 to enable the intentional offline functionality of the agent.</p> <p>See “Agent functions” on page 190.</p> <p>Type and dimension: integer-scalar</p>

Table 12-3 Oracle agent optional attributes

Optional Attributes	Definition
DetailMonitor	<p>Setting this attribute to a non-zero value enables detailed monitoring for Oracle. The value indicates the number of monitor cycles after which the agent monitors Oracle in detail. For example, the value 5 indicates that the agent monitors Oracle in detail every five monitor intervals. The monitor interval is 60 seconds by default.</p> <p>Default: 0</p> <p>Type and dimension: integer-scalar</p>
MonScript	<p>Pathname to the script provided for detailed monitoring. The default, basic monitoring, monitors the database PIDs only.</p> <p>Note: Detailed monitoring is disabled if the value of the attribute MonScript is invalid or is set to an empty string.</p> <p>The pathname to the supplied detail monitor script is /opt/VRTSagents/ha/bin/Oracle/SqlTest.pl.</p> <p>MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with "./", as in the path ./bin/Oracle/SqlTest.pl.</p> <p>Type and dimension: string-scalar</p>
User	<p>Internal database user. Connects to the database for detail monitoring.</p> <p>Type and dimension: string-scalar</p>
Pword	<p>Encrypted password for internal database-user authentication.</p> <p>You must encrypt passwords using the VCS Encrypt utility before you configure this attribute.</p> <p>See “Prerequisites for configuring Oracle and NetIsnr resources” on page 233.</p> <p>Note: The VCS Encrypt utility is installed as part of the Veritas Virtual Machine tools.</p> <p>Type and dimension: string-scalar</p>

Table 12-3 Oracle agent optional attributes

Optional Attributes	Definition
Table	Table for update by <code>User/Pword</code> . Type and dimension: string-scalar
Encoding	This attribute is disabled. Accept the default setting.
AgentDebug	Additional debug messages are logged when this flag is set. Default: 0 Type and dimension: boolean-scalar

[Table 12-4](#) lists the internal attribute for Oracle agent. This attribute is for internal use only. Symantec recommends that you do not modify the value of this attribute.

Table 12-4 Oracle agent internal attributes

Optional Attributes	Definition
AgentDirectory	Specifies the location of binaries, scripts, and other files related to the Oracle agent. Default: <code>/opt/VRTSagents/ha/bin/Oracle</code> Type and dimension: static-string

Netlsnr agent attributes

[Table 12-5](#) lists the required attributes for Netlsnr agent.

Table 12-5 Netlsnr agent required attributes

Required attributes	Definition
Owner	Oracle user, as the defined owner of executables and database files in /etc/passwd. Type and dimension: string-scalar
Home	The \$ORACLE_HOME path to Oracle binaries and configuration files. For example, you could specify the path as /opt/ora_home. Do not append a slash (/) at the end of the path. Type and dimension: string-scalar

[Table 12-6](#) lists the optional attributes for Netlsnr agent.

Table 12-6 Netlsnr agent optional attributes

Optional attributes	Definition
TnsAdmin	The \$TNS_ADMIN path to directory in which the Listener configuration file resides (listener.ora). Default: /var/opt/oracle Type and dimension: string-scalar
Listener	Name of Listener. The name for Listener is considered case-insensitive by the Netlsnr agent and the Oracle database server. Default: LISTENER Type and dimension: string-scalar

Table 12-6 Netlsnr agent optional attributes

Optional attributes	Definition
LsnrPwd	<p>The VCS encrypted password used to stop and monitor the listener. This password is set in the Listener configuration file.</p> <p>You must encrypt passwords using the VCS Encrypt utility before you configure this attribute.</p> <p>See “Prerequisites for configuring Oracle and Netlsnr resources” on page 233.</p> <p>Note: The VCS Encrypt utility is installed as part of the Veritas Virtual Machine tools.</p> <p>Type and dimension: string-scalar</p>
EnvFile	<p>Specifies the full path name of the file that is sourced by the agent function scripts. This file contains the environment variables set by the user for the Oracle listener environment such as LD_LIBRARY_PATH and so on.</p> <p>The syntax for the contents of the file depends on the login shell of Owner. This file must be readable by Owner. The file must not contain any prompts for user input.</p> <p>Type and dimension: string-scalar</p>
MonScript	<p>Pathname to the script provided for detail monitoring. The default (basic monitoring) is to monitor the listener process only.</p> <p>Note: Detail monitoring is disabled if the value of the attribute MonScript is invalid or is set to an empty string.</p> <p>The pathname to the supplied detail monitoring script is /opt/VRTSagents/ha/bin/Netlsnr/LsnrTest.pl.</p> <p>MonScript also accepts a pathname relative to /opt/VRTSagents/ha. A relative pathname should start with “./”, as in the path ./bin/Netlsnr/LsnrTest.pl.</p> <p>Type and dimension: string-scalar</p>

Table 12-6 Netlsnr agent optional attributes

Optional attributes	Definition
Encoding	Specifies operating system encoding that corresponds to Oracle encoding for the displayed Oracle output. Default: "" Type and dimension: string-scalar
AgentDebug	Additional debug messages are logged when this flag is set. Default: 0 Type and dimension: boolean

[Table 12-7](#) lists the internal attribute for Netlsnr agent. This attribute is for internal use only. Symantec recommends that you do not modify the value of this attribute.

Table 12-7 Netlsnr agent internal attributes

Optional Attributes	Definition
AgentDirectory	Specifies the location of binaries, scripts, and other files related to the Netlsnr agent. Default: /opt/VRTSagents/ha/bin/Netlsnr Type and dimension: static-string

About the VCS agent for the Apache Web server

The Apache Web server agent monitors the Apache server processes.

The agent can detect when Apache is brought down gracefully by an administrator. When Apache is brought down gracefully, the agent does not trigger a resource fault even though Apache is down.

This agent supports detailed monitoring, refer to the SecondLevelMonitor optional attribute for more information on how to activate detailed monitoring.

Agent functions

- **Monitor**
Monitors the state of the Apache server. First it checks for the processes, next it can perform an optional state check.
For detailed monitoring, the agent can perform a deeper, more thorough state check of the Apache HTTP server performed by issuing an HTTP GET request on the web server's root directory.

State definitions

- **ONLINE**
Indicates that the Apache server is running.
- **OFFLINE**
Indicates that the Apache server is not running.
Can also indicate that the administrator stopped the Apache server gracefully. Note that the agent uses the PidFile attribute for the intentional offline detection.
- **FAULTED**
Indicates that the Apache server unexpectedly went offline.
- **UNKNOWN**
Indicates that a problem exists with the configuration.

Apache Web server agent attributes

Table 12-8 lists the required attributes for the Apache agent.

Table 12-8 Apache Web server agent required attributes

Required attribute	Description
ConfigFile	Full path and file name of the main configuration file for the Apache server. See " About bringing an Apache Web server online outside of VCS control " on page 204. Type and dimension: string-scalar Example: "/apache/server1/conf/httpd.conf"

Table 12-8 Apache Web server agent required attributes

Required attribute	Description
httpdDir	<p>Full path of the directory to the httpd binary file.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/apache/server1/bin"</p>
SecondLevelMonitor	<p>Valid attribute values are true (1) and false (0).</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p> <p>Example: "1"</p>
ResLogLevel	<p>Controls the agent's logging detail for a specific instance of a resource. Values are:</p> <ul style="list-style-type: none"> ■ ERROR: Logs error messages. ■ WARN: Logs error and warning messages. ■ INFO: Logs error, warning, and informational messages. ■ TRACE: Logs error, warning, informational, and trace messages. Trace logging is verbose. Use for initial configuration or troubleshooting. <p>Type and dimension: string-scalar</p> <p>Default: INFO</p> <p>Example: "TRACE"</p>
PidFile	<p>This attribute is required when you want to enable the detection of a graceful shutdown outside of VCS control.</p> <p>See "PidFile" on page 203.</p>

Table 12-9 lists the optional attributes for the Apache agent.

Table 12-9 Apache Web server agent optional attributes

Optional attribute	Description
DirectiveAfter	A list of directives that httpd processes after reading the configuration file. Type and dimension: string-association Example: DirectiveAfter{} = { KeepAlive=On }
DirectiveBefore	A list of directives that httpd processes before it reads the configuration file. Type and dimension: string-association Example: DirectiveBefore{} = { User=nobody, Group=nobody }
User	Account name the agent uses to execute the httpd program. If you do not specify this value, the agent executes httpd as the root user. Type and dimension: string-scalar Example: "apache1"
EnableSSL	Set to 1 (true) to have the Online agent function add support for SSL by including the option <code>-DSSL</code> in the start command. For example: <code>/usr/sbin/httpd -k start -DSSL</code> Set to 0 (false) to exclude the <code>-DSSL</code> option from the command. Type and dimension: boolean-scalar Default: 0 Example: "1"

Table 12-9 Apache Web server agent optional attributes

Optional attribute	Description
HostName	<p>The virtual host name that is assigned to the Apache server instance. The host name is used in second-level monitoring to establish a socket connection with the Apache HTTP server.</p> <p>Note: The HostName attribute is only required when the value of SecondLevelMonitor is 1 (true).</p> <p>Type and dimension: string-scalar</p> <p>Example: "web1.veritas.com"</p>
Port	<p>Port number where the Apache HTTP server instance listens. The port number is used in second-level monitoring to establish a socket connection with the server. Specify this attribute only if SecondLevelMonitor is set to 1 (true).</p> <p>Type and dimension: integer-scalar</p> <p>Default: 80</p> <p>Example: "80"</p>
EnvFile	<p>Full path and file name of the file that is sourced prior to executing httpdDir/httpd. With Apache 2.0, the file <i>ServerRoot/bin/envvars</i>, which is supplied in most Apache 2.0 distributions, is commonly used to set the environment prior to executing httpd. If EnvFile is specified, the login shell for user root must be Bourne, Korn, or C shell.</p> <p>Type and dimension: string-scalar</p> <p>Example: "/apache/server1/bin/envvars"</p>
PidFile	<p>The PidFile attribute sets the file to which the server records the process ID of the daemon. The value of PidFile attribute must be the absolute path where the Apache instance records the pid.</p> <p>This attribute is required when you want the agent to detect the graceful shutdown of the Web server. For the agent to detect the graceful shutdown of the Web server, the value of the IntentionalOffline resource type attribute must be 1 (true).</p> <p>Type and dimension: string-scalar</p> <p>Example: /var/run/httpd.pid</p>

Table 12-9 Apache Web server agent optional attributes

Optional attribute	Description
SharedObjDir	<p>Full path of the directory in which the Apache HTTP shared object files are located. It is used when the HTTP Server is compiled using the SHARED_CORE rule. If specified, the directory is passed to the <code>-R</code> option when executing the <code>httpd</code> program. Refer to the <code>httpd</code> man pages for more information about the <code>-R</code> option.</p> <p>Type and dimension: boolean-scalar</p> <p>Example: <code>"/apache/server1/libexec"</code></p>
SecondLevelTimeout	<p>Number of seconds the Monitor agent function waits for the second-level monitor to complete. If the second-level monitor program does not respond within the <code>SecondLevelTimeout</code> value, the Monitor agent function stops blocking on the program sub-process and reports that the resource is offline. The value for the attribute must be less than 60 seconds.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 30</p>

About bringing an Apache Web server online outside of VCS control

When you bring an Apache Web server online outside of VCS control, first source its environment file. Start the server with the `-f` option so the server knows which instance to start. You can then specify additional options (such as `EnableSSL` or `SharedObjDir`) that you want the server to use at start.

To start an Apache Web server outside of VCS control

- 1 Source the environment file if required.
- 2 Start the Apache Web server. You must use the `-f` option so that the agent can distinguish different instances of the server.

```
httpdDir/httpd -f ConfigFile -k start
```

Where *httpdDir* is `/apache/v2.2/bin` *ConfigFile* is `/apache/v2.2/conf/httpd.conf`. When fully formed, the start example looks like:

```
/apache/v2.2/bin/httpd -f /apache/v2.2/conf/httpd.conf -k start
```

- 3 Specify additional options such as EnableSSL or SharedObjDir that you want to use when you start server. When you add EnableSSL to the command, it resembles:

```
httpdDir/httpd -f ConfigFile -k start -DSSL
```

About the VCS agent for SAP NetWeaver

The Veritas agent for SAP NetWeaver provides high availability for SAP R/3 and SAP NetWeaver in a cluster. The SAP NetWeaver agent supports a wide range of SAP environments, including the traditional Basis architecture and the SAP J2EE Web Application Server architecture (NetWeaver). The agent also supports Standalone Enqueue Servers in a distributed SAP installation.

The SAP components are:

- Central instance
- Dialog instance
- Standalone Enqueue Server
 Standalone Enqueue Server is also known as SAP Central Services (SCS).

The agent supports the following SAP Web Application Server architectures:

- ABAP
- Java
- Java Add-In (ABAP + Java)

The sample resource configurations for SAP NetWeaver agent are shown in the following figures. In these configurations, three SAP components are configured in three different virtual machines.

Figure 12-1 shows the resource configuration for SCS instance.

Figure 12-1 Resource configuration for SCS instance

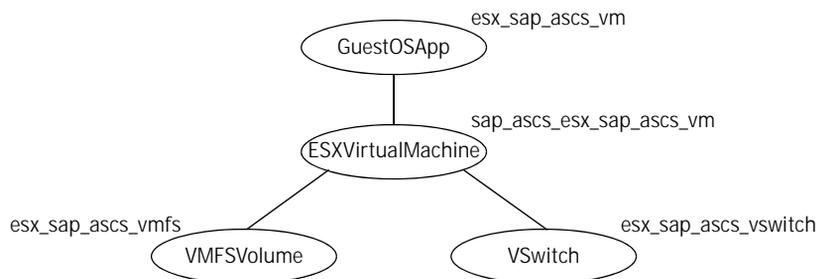


Figure 12-2 shows the resource configuration for Central instance.

Figure 12-2 Resource configuration for Central instance

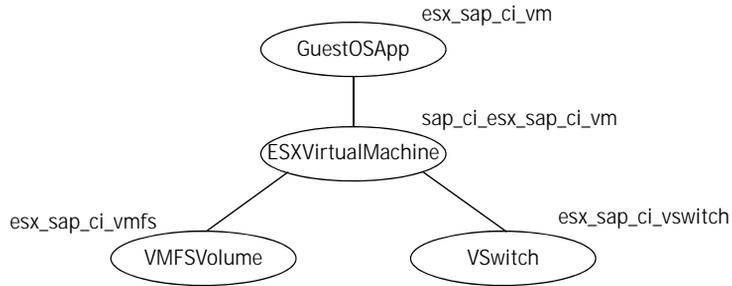
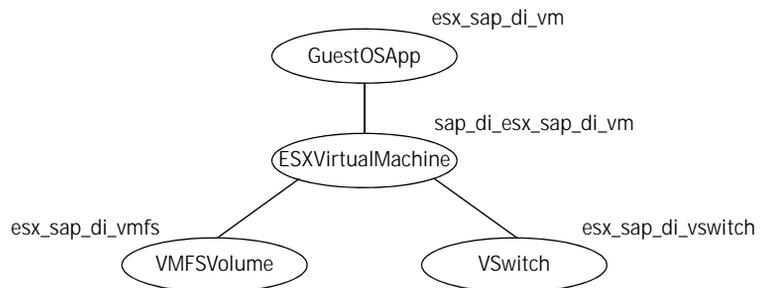


Figure 12-3 shows the resource configuration for Dialog instance.

Figure 12-3 Resource configuration for Dialog instance



VCS supports detection of graceful shutdown for SAP instances. When SAP NetWeaver is brought down gracefully, the agent does not trigger a resource fault even though SAP NetWeaver is down.

Agent functions

The SAP NetWeaver agent supports the Monitor function, which is described as follows:

- **Monitor**

The monitor operation monitors the state of the SAP instance on all nodes in the cluster. The operation performs the following tasks:

- Depending upon the search criteria that the ProcMon attribute specifies, the monitor operation scans the process table to verify that the SAP instance processes are running. Review the information for setting the ProcMon attribute.

See [“SAP NetWeaver agent attributes”](#) on page 208.

- If the SecondLevelMonitor attribute is greater than 0, the monitor operation performs a thorough check of the application instance.
 - For Central or Dialog instances, the operation uses the following utilities to perform this check:

Server architecture	SAP utility used
SAP Web Application Server as ABAP	sapinfo
SAP Web Application Server as Java	jcmon
SAP Web Application Server as Java Add-In	sapinfo and jcmon

- For Enqueue Server instances, the operation uses the ensmon utility for all the architectures (ABAP, Java, and Java Add-In).
- The monitor operation can also execute a custom monitor utility that the MonitorProgram attribute specifies.

State definitions

- **ONLINE**
Indicates that the SAP server is running.
- **OFFLINE**
Indicates that the SAP server is not running.
Can also indicate that the administrator has intervened to stop the SAP server.
- **FAULTED**
Indicates that the SAP server unexpectedly went offline.
- **UNKNOWN**
Indicates that a problem exists with the configuration.

SAP NetWeaver agent attributes

Table 12-10 shows the required attributes for the SAP NetWeaver agent.

Table 12-10 SAP NetWeaver agent required attributes

Required attributes	Description
EnvFile	<p>The absolute path to the file that must be sourced with the UNIX shell. You must source this file to set the environment before executing the SAP scripts for monitor operation.</p> <p>Supported shell environments are ksh, sh, and csh.</p> <p>Note: Ensure that the syntax of this file is in accordance with the user shell that the SAPAdmin attribute specifies.</p> <p>Symantec recommends that you store this file on shared disk so that the file is always available to an online system.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: /usr/sap/EC1/DVEBMGS00/envfile</p>
InstName	<p>Identifies a SAP server instance.</p> <p>Type and dimension: string-scalar</p> <p>Default: DVEBMGS00</p> <p>Example: DVBGS01</p>
InstType	<p>An identifier that classifies and describes the SAP server instance type. Valid values are:</p> <ul style="list-style-type: none"> ■ CENTRAL: SAP Central Services instance ■ DIALOG: SAP Dialog instance ■ ENQUEUE: Standalone Enqueue instance ■ AENQUEUE: ABAP SAP Standalone Enqueue instance ■ JENQUEUE: Java SAP Standalone Enqueue instance <p>Note: The value of this attribute is not case sensitive.</p> <p>Type and dimension: string-scalar</p> <p>Default: CENTRAL</p> <p>Example: DIALOG</p>

Table 12-10 SAP NetWeaver agent required attributes

Required attributes	Description
ProcMon	<p>The list of SAP processes to monitor. Use a space to separate the entities in this list. The entities can appear in any order.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: dw se jc</p>
ResLogLevel	<p>Controls the agent's logging detail for a specific instance of a resource. Values are:</p> <ul style="list-style-type: none"> ■ ERROR: Logs error messages. ■ WARN: Logs error and warning messages. ■ INFO: Logs error, warning, and informational messages. ■ TRACE: Logs error, warning, informational, and trace messages. Trace logging is verbose. Use for initial configuration or troubleshooting operations. <p>Type and dimension: string-scalar</p> <p>Default: INFO</p> <p>Example: TRACE</p>
SAPAdmin	<p>UNIX user, as the defined administrator of executables of application directories and executables. This user name is usually a concatenation of the SAPSID attribute and the adm string.</p> <p>One or more system naming services store this user name, for example, NIS, NIS+, and LDAP servers. The agent functions use this user name to execute their respective core subroutines.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: ec1adm</p>

Table 12-10 SAP NetWeaver agent required attributes

Required attributes	Description
SAPMonHome	<p>The location of the directory that contains the binary used for second level monitoring process.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: /usr/sap/\${SAPSID}/SYS/exe/runU</p>
SAPSID	<p>SAP system name.</p> <p>The value of this attribute is three characters in length, and must begin with an alphabetical character. The value of this attribute is defined during the SAP installation.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: EC1</p>
StartProfile	<p>The full path and file name of the StartProfile instance.</p> <p>The StartProfile instance is found in /usr/sap/SAPSID/SYS/profile directory. The value of the instance is START_InstName_hostname. The hostname must resolve into a valid IP address that is used to cluster the SAP instance.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: /usr/sap/EC1/SYS/profile/START_DVEBMGS01_sunabap</p>

Table 12-11 lists the optional attributes for the SAP NetWeaver agent.

Table 12-11 SAP NetWeaver agent optional attributes

Optional attribute	Description
MonitorProgram	<p>Absolute path name of an external, user-supplied monitor executable.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example 1: /usr/sap/EC1/DVEBMGS00/work/myMonitor.sh</p> <p>Example 2: /usr/sap/EC1/DVEBMGS00/work/myMonitor.sh arg1 arg2</p>
SecondLevelMonitor	<p>Used to enable second-level monitoring. Second-level monitoring is a deeper, more thorough state check of the SAP instance.</p> <p>By default, the agent performs basic monitoring. If the value of the SecondLevelMonitor attribute is greater than 0, then the agent performs detailed monitoring. The numeric value specifies how often the monitoring routines must run. For example, if you set the value of this attribute as 2, then the agent monitors the SAP instance in detail every second monitor interval. The monitor interval is 60 seconds by default.</p> <p>Note: Exercise caution while setting SecondLevelMonitor to large numbers. For example, if you set the value of SecondLevelMonitor to 100, then sapinfo is executed every 100 minutes, which may not be as often as intended. For maximum flexibility, no upper limit is defined for SecondLevelMonitor.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p> <p>Example: 1</p>

About the VCS agent for WebLogic Server

The agent for WebLogic Server is named WebLogic9. It consists of a resource type declaration and the agent executables. The agent is responsible for monitoring and detecting failures of WebLogic Server (WLS) components.

WebLogic Servers fall into two categories: Administrative and Managed. The Administrative Server provides a central point from which you can manage the domain. It provides access to WebLogic Server administration tools. All other servers are considered as Managed Servers.

Note that this release supports only non-node manager based configurations. Node manager and node manager based configurations are not supported.

Agent functions

The Monitor agent function performs the following tasks:

- Conducts a first level check on the WebLogic Server component to ensure that the WebLogic Server component's process is running. The agent identifies the process for the WebLogic Server component by applying the pattern matching on command lines of processes running in the system.
- Depending on the settings that you make, the monitor operation can conduct a second level check on the WebLogic Server component. The second level check uses the `wlst.sh` scripting utility to attempt to connect to the WebLogic Server component. [Table 12-12](#) lists the `wlst` commands used to connect to the WebLogic Server component.

Table 12-12 Commands to connect to the WebLogic Server component

Resource Configuration	Mechanism used for second level monitoring
Administrative server	Uses the <code>wlst</code> command <code>connect</code> .
Managed server	Uses the <code>wlst</code> command <code>connect</code> .

- Depending upon the value of the `MonitorProgram` attribute, the monitor operation can perform a customized check using a user-supplied monitoring utility.

State Definitions

- **ONLINE**
Indicates that the WebLogic server is running.
- **OFFLINE**
Indicates that the WebLogic server is not running.
Can also indicate that the administrator stopped the WebLogic server gracefully.
- **FAULTED**
Indicates that the WebLogic server unexpectedly went offline.
- **UNKNOWN**
Indicates that a problem exists with the configuration.

WebLogic Server agent attributes

Refer to the following required and optional attributes while configuring the agent for WebLogic Server.

[Table 12-13](#) lists the required attributes for the agent for WebLogic Server.

Table 12-13 Required attributes

Required Attribute	Description
BEA_HOME	The absolute path to BEA home directory of WebLogic Server installation. BEA_HOME is used to uniquely identify the ServerRole processes. Type and dimension: string-scalar Default: "" Example: /bea/wls90/admin
DomainDir	The domain directory of the WebLogic Server domain to which the instance belongs. Specify this attribute for Administrative and Managed Servers. Type and dimension: string-scalar Default: "" Example: /bea/wls90/admin/user_projects/domains/WLS90Domain

Table 12-13 Required attributes

Required Attribute	Description
DomainName	<p>The name of the WebLogic Server domain to which the instance belongs.</p> <p>Specify this attribute for Administrative and Managed Servers.</p> <p>Review the section “Uniquely identifying WebLogic Server instances” in the <i>Veritas™ High Availability Agent for WebLogic Server Installation and Configuration Guide</i> for tips on defining a unique name for WebLogic Server instances.</p> <p>Type and dimension: string-scalar</p> <p>Default: “”</p> <p>Example: WLS90Domain</p>
ListenAddressPort	<p>The Listen Address and port of the WebLogic instance. The format is ListenAddress:port. Ensure that the ListenAddress string resolves to the proper IP Address, using the network name service that you used on the host. The WebLogic Server connects to the ListenAddress on the specified port through the wlst.sh API.</p> <p>Specify this attribute for Administrative and Managed Servers only.</p> <p>Type and dimension: string-scalar</p> <p>Default: “”</p> <p>Examples: wls90adminsolv.veritas.com:7001 or wls90adminsolv.veritas.com:5556</p>
ResLogLevel	<p>The logging detail performed by the agent for WebLogic Server for the resource. Valid values include:</p> <p>ERROR: Only logs error messages.</p> <p>WARN: Logs above plus warning messages.</p> <p>INFO: Logs above plus informational messages.</p> <p>TRACE: Logs above plus trace messages. TRACE is very verbose and should only be used during initial configuration or for troubleshooting and diagnostic operations.</p> <p>Type and dimension: string-scalar</p> <p>Default: INFO</p> <p>Example: TRACE</p>

Table 12-13 Required attributes

Required Attribute	Description
ServerName	<p>The name of the WebLogic Server. You must specify this attribute for Administrative and Managed Servers only.</p> <p>Review the section "Uniquely identifying WebLogic Server instances" in the <i>Veritas™ High Availability Agent for WebLogic Server Installation and Configuration Guide</i> for tips on defining a unique name for WebLogic Server instances.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: AdminServer</p>
WL_HOME	<p>The absolute path to the product installation directory of the WebLogic Server. The agent for WebLogic Server uses this attribute to locate the wlst.sh utility and the Node Manager home directory.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: /bea/wls90/admin/weblogic90</p>
WLSUser	<p>The user name of the user that is connecting the wlst.sh utility to the server running the WebLogic Server instance, along with WLSPassword.</p> <p>Specify this attribute for Administrative and Managed Servers only.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: weblogic</p>
ServerRole	<p>Type of WebLogic Server. Valid values include:</p> <ul style="list-style-type: none"> ■ Administrative: Online operation executes wlst.sh script with nmConnect() and nmStart() API. Example: nmStart ('AdminServer1') ■ Managed: Online operation executes wlst.sh script with nmConnect() and nmStart() API. Example: nmStart ('ManagedServer1') <p>Type and dimension: string-scalar</p> <p>Default: ""</p>

Table 12-13 Required attributes

Required Attribute	Description
User	<p>The UNIX user name used to start and stop the WebLogic Server instance. If MonitorProgram is specified, the agent for WebLogic Server uses this user's credentials to run the defined program.</p> <p>You must synchronize the user name across the systems within the cluster. This user name must resolve to the same UID and have the same default shell on each system in the cluster. The agent operations use the getpwnam(3C) function system call to obtain UNIX user attributes. Hence you can define the user name locally or in a common repository, such as NIS, NIS+, or LDAP).</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: wlsadmin</p>
WLSPassword	<p>The password of user connecting WLST to ServerRole Application Server, along with WLSUser.</p> <ul style="list-style-type: none"> ■ For VCS, encrypt the value of this attribute using the \$VCS_HOME/bin/vcscrypt utility that VCS provides. <p>While encrypting the password, use the -agent option.</p> <p>Specify this attribute for Administrative and Managed Servers only.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: weblogic</p>

Table 12-13 Required attributes

Required Attribute	Description
ServerStartProgram	<p>The complete command line of the script used to start WebLogic Server.</p> <p>If command line arguments are passed to ServerStartProgram, the agent uses the command and arguments as it is.</p> <p>Example: ServerStartProgram = "/wls/my_domain/startManagedWebLogic.sh Managed1"</p> <p>If no arguments are passed (for example, ServerStartProgram ="/wls/my_domain/startManagedWebLogic.sh"), the agent forms the command line as follows:</p> <ul style="list-style-type: none"> ■ For Managed Server: \$ServerStartProgram \$ServerName \$AdminURL ■ For Administrative Server: \$ServerStartProgramType and dimension: string-scalar <p>Type and dimension: string-scalar</p> <p>Example: /bea/user_projects/domains/WLS90Domain/bin/startManagedWebLogic.sh</p> <p>Default: No default value</p>
ServerStopProgram	<p>The complete command line of the script used to stop WebLogic Server.</p> <p>If command line arguments are passed to ServerStopProgram, the agent uses the command and arguments as it is.</p> <p>Example: ServerStopProgram = "wls/my_domain/stopManagedWebLogic.sh Managed1 t3://adminurl:7001 weblogic passwd"</p> <p>If no arguments are passed (for example, ServerStopProgram ="/wls/my_domain/stopManagedWebLogic.sh", the agent forms the command line as follows:</p> <pre>\$ServerStopProgram \$ServerName \$AdminURL \$WLSUser \$WLSPassword</pre> <p>Type and dimension: string-scalar</p> <p>Example: /bea/user_projects/domains/WLS90Domain/bin/stopManagedWebLogic.sh</p> <p>Default: No default value</p>

Review the *Veritas High Availability Agent for WebLogic Server Installation and Configuration Guide* for more information on the attributes required in non-Node Manager based configurations.

[Table 12-14](#) lists the optional attributes.

Table 12-14 Optional attributes

Optional attribute	Description
AdminUrl	<p>The URL of the Managed Server's Administrative Server. Set this attribute only for resources whose ServerRole attribute is Managed.</p> <p>Ensure that the value of this attribute is the same as management.server that appears in the long listing of processes for the Managed Server.</p> <p>If the RequireAdminServer attribute is set to one, AdminUrl is used to connect to the Administrative Server for the domain to determine if the server is fully online. Managed Servers also use this URL to connect to the Administrative Server and download its web applications and services (JMS, JDBC Connection Pool, etc.) configuration.</p> <p>Type and dimension: string-scalar</p> <p>Default: ""</p> <p>Example: http://wlsadmin:7001</p>
AdminServerMaxWait	<p>The maximum number of seconds that a Managed Server waits for an Administrative Server to respond to a test probe.</p> <p>For more information on how this attribute is used to delay the Managed Server startup process, see the section "Delaying Managed Server startup process" in the <i>Veritas™ High Availability Agent for WebLogic Server Installation and Configuration Guide</i>.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 60</p> <p>Example: 90</p>

Table 12-14 Optional attributes

Optional attribute	Description
MonitorProgram	<p>The full path name and command-line arguments for an externally provided monitor program.</p> <p>For information on executing the customized monitoring program see the section “Executing a customized monitoring program” in the <i>Veritas™ High Availability Agent for WebLogic Server Installation and Configuration Guide</i>.</p> <p>Type and dimension: string-scalar</p> <p>Default: “”</p> <p>Example 1: /bea/wls90/admin/mymonitor.sh</p> <p>Example 2: /usr/local/bin/MyMonitor.sh myWLS.foo.com 8080</p>
SecondLevelMonitor	<p>Used to enable second-level monitoring. Second-level monitoring is a deeper, more thorough state check of the configured ServerRole attribute.</p> <p>By default, the agent performs basic monitoring. If the value of the SecondLevelMonitor attribute is greater than zero, then the agent performs detailed monitoring. The numeric value specifies how often the monitoring routines must run. For example, if you set the value of this attribute as two, the agent monitors the WebLogic Server instance in detail every second monitor interval. The monitor interval is 60 seconds by default.</p> <p>The numeric value specifies how often the monitoring routines must run.</p> <p>For WebLogic Server versions 9.0, 9.1, 9.2, and 10.0 the second-level monitoring is performed as follows: The agent for WebLogic Server uses the BEA supplied WebLogic Server scripting tool <code>wlst.sh</code> to perform second-level monitoring. Depending upon the ServerRole, <code>wlst.sh</code> uses <code>api</code> commands <code>connect()</code>, <code>nmConnect()</code> and <code>nmServerStatus()</code> to perform monitoring routines.</p> <p>Note: Exercise caution while setting SecondLevelMonitor to large numbers. For example, if the MonitorInterval is set to 60 seconds and the SecondLevelMonitor is set to 100, then <code>wlst.sh</code> is executed every 100 minutes, which may not be as often as intended. For maximum flexibility, no upper limit is defined for SecondLevelMonitor.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 0</p> <p>Example: 1</p>

Table 12-14 Optional attributes

Optional attribute	Description
RequireAdminServer	<p>The flag that is used to control the startup behavior of a WebLogic Server instance.</p> <p>When the RequireAdminServer attribute is set to 1 (true), the Managed Server resource is not allowed to complete an initiated online operation until the Administrative Server is ready to accept connections.</p> <p>If the RequireAdminServer attribute is set to zero and the AdminServerMaxWait is set to a value that is greater than five, the online function first probes the Administrative Server instance to see if it is ready to accept connections. If the server is not ready, the operation waits for five seconds and then probes the server again to determine its state. This cycle of probe and wait repeats until either the Administrative Server is ready or the AdminServerMaxWait time expires.</p> <p>Specify this attribute for Managed Server only.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0 (false)</p> <p>Example: 1 (true)</p>

About the Application agent

The Application agent monitors the status of applications. An application runs in the default context of root. Configure the User attribute to run the application in that user's context (for example jsmith and not root).

You can monitor the application in the following ways:

- Use the monitor program
- Specify a list of processes
- Specify a list of process ID files
- Any combination of the above

Agent functions

- **Monitor**

If you specify the `MonitorProgram`, the agent executes the user-defined `MonitorProgram` in the user-specified context. If you specify `PidFiles`, the routine verifies that the process ID found in each listed file is running. If you specify `MonitorProcesses`, the routine verifies that each listed process is running in the context you specify.

Use any one, two, or three of these attributes to monitor the application. If any one process specified in either `PidFiles` or `MonitorProcesses` is determined not to be running, the monitor returns `OFFLINE`. If the process terminates ungracefully, the monitor returns `OFFLINE` and failover occurs.
- **Clean**

Terminates processes specified in `PidFiles` or `MonitorProcesses`. Ensures that only those processes (specified in `MonitorProcesses`) running with the user ID specified in the `User` attribute are killed. If the `CleanProgram` is defined, the agent executes the `CleanProgram`.

State definitions

- **ONLINE**

Indicates that all processes specified in `PidFiles` and `MonitorProcesses` are running and that the `MonitorProgram` returns `ONLINE`.
- **OFFLINE**

Indicates that at least one process specified in `PidFiles` or `MonitorProcesses` is not running, or that the `MonitorProgram` returns `OFFLINE`.
- **FAULTED**

Indicates this resource state if an application process terminates.
- **UNKNOWN**

Indicates an indeterminable application state or invalid configuration.

Application agent attributes

Table 12-15 lists the required attributes for the Application agent.

Table 12-15 Application agent required attributes

Required attribute	Description
At least one of the following attributes: <ul style="list-style-type: none"> ■ MonitorProcesses ■ MonitorProgram ■ PidFiles 	See " Application agent optional attributes " on page 222.
StartProgram	This attribute is disabled. Accept the default setting.
StopProgram	This attribute is disabled. Accept the default setting.

Table 12-16 lists the optional attributes for the VCS agent for the Application agent.

Table 12-16 Application agent optional attributes

Optional attribute	Description
CleanProgram	The executable, created locally on each node, which forcibly stops the application. Specify the complete path of the executable. Applicable command line arguments follow the name of the executable and are separated by spaces. Type and dimension: string-scalar
MonitorProcesses	A list of processes that you want monitored and cleaned. Each process name is the name of an executable. If the executable requires the complete path to start, specify the complete path. The process name must be the name displayed by the <code>ps -ef</code> command for the process. Type and dimension: string-vector Example: "nmbd"

Table 12-16 Application agent optional attributes

Optional attribute	Description
MonitorProgram	<p>The executable, created locally on each node, which monitors the application. Specify the complete path of the executable. Applicable command line arguments follow the name of the executable and are separated by spaces.</p> <p>MonitorProgram can return the following VCSAgResState values: OFFLINE value is 100; ONLINE values range from 101 to 110 (depending on the confidence level); 110 equals confidence level of 100%. Any other value = UNKNOWN.</p> <p>Type and dimension: string-scalar</p>
PidFiles	<p>A list of PID files that contain the process ID (PID) of the processes that you want monitored and cleaned. These are application generated files. Each PID file contains one monitored PID. Specify the complete path of each PID file in the list.</p> <p>The process ID can change when the process restarts. If the application takes time to update the PID file, the agent's monitor script may return an incorrect result.</p> <p>Type and dimension: string-vector</p>
User	<p>The user ID for running StartProgram, StopProgram, MonitorProgram, and CleanProgram. The processes specified in the MonitorProcesses list must run in the context of the specified user. Monitor checks the processes to make sure they run in this context.</p> <p>Type and dimension: string-scalar</p> <p>Default: root</p>

About the Mount agent

Use the Mount agent to monitor an NFS mount point. If the mount point fails, VCS detects the fault, and fails over the service group to another node.

This agent supports detailed monitoring, refer to the `SecondLevelMonitor` optional attribute for more information.

Agent functions

- **Monitor**
Determines if the file system is mounted.
Supports detailed monitoring of the mount point.

State definitions

- **ONLINE**
Indicates that the file system is properly mounted on the given mount point.
- **OFFLINE**
Indicates that the file system is not mounted properly on the mount point.
- **FAULTED**
Indicates that the file system unexpectedly unmounted.
- **UNKNOWN**
Indicates that a problem exists either with the configuration or the ability to determine the status of the resource.

Mount agent attributes

Table 12-17 lists the required attributes for the Mount agent.

Table 12-17 Mount agent required attributes

Required attribute	Description
BlockDevice	<p>The block device for the mount point.</p> <p>Type and dimension: string-scalar</p> <p>Examples:</p> <p><code>/dev/sdc1</code></p> <p>If the device is an LVM2 volume, specify the BlockDevice as:</p> <p><code>/dev/mapper/volume-group-logical-volume</code></p> <p>For example:</p> <p><code>/dev/mapper/voldg-lvol0</code></p> <p>If the file system type is NFS, then specify the BlockDevice as:</p> <p><code>server:/path/to/share</code></p> <p>NFS device example:</p> <p><code>vcslnx1.veritas.com:/usr/share1</code></p>
FsckOpt	Specify -n for this value.
FSType	<p>Type of file system.</p> <p>Supports ext2, ext3, nfs, or reiserfs.</p> <p>Type and dimension: string-scalar</p>
MountPoint	<p>Directory for mount point.</p> <p>Type and dimension: string-scalar</p> <p>Example: <code>"/mnt/apache1"</code></p>
VxFSMountLock	This attribute is not supported. Accept the default setting.

Table 12-18 lists the optional attributes for the Mount agent.

Table 12-18 Mount agent optional attributes

Optional attribute	Description
CkptUmount	This attribute is disabled. Accept the default setting.
MountOpt	This attribute is disabled. Accept the default setting.
SecondLevelMonitor	<p>This attribute is only applicable to NFS client mounts.</p> <p>If set to 1, this attribute enables detailed monitoring of a NFS mounted file system.</p> <p>When you set up fstab (for mounting NFS while the guest operating system boots) mount nfs with the <code>-o intr</code> option.</p> <p>Type and dimension: boolean-scalar</p> <p>Default: 0</p>
SecondLevelTimeout	<p>This attribute is only applicable for a NFS client mount.</p> <p>This is the timeout (in seconds) for the SecondLevelMonitor/Detail monitoring of NFS Mounts.</p> <p>Number of seconds the Monitor agent function waits for the second-level monitor to complete. If the second-level monitor program does not respond within the SecondLevelTimeout value, the Monitor agent function stops blocking on the program sub-process and reports that the resource is offline. The value for the attribute must be less than sixty.</p> <p>Type and dimension: integer-scalar</p> <p>Default: 30</p>
SnapUmount	This attribute is disabled. Accept the default setting.

Sample configuration

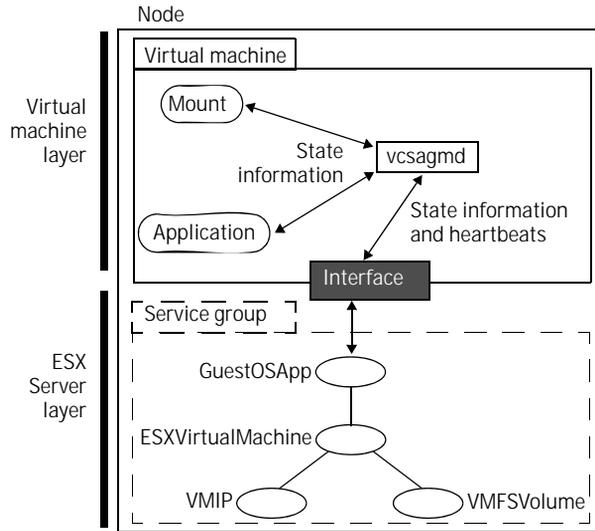
You can use the Mount agent to monitor an application that uses a network file system. For example, Oracle uses a network file system that is mounted at `/mnt/oracle`. The Mount agent monitors `/mnt/oracle`. If `/mnt/oracle` goes offline, then the Mount resource would fault, and the virtual machine would fail over to another cluster node. Like all other applications that VCS monitors in virtual machines, the Mount agent is represented as a `GuestOSApp` resource at the VMware ESX Server level.

How VCS monitors applications and resources on virtual machines

The VCS agent management daemon (`vcsagmd`) is a very light-weight daemon that resides in the virtual machine. It manages the agents that monitor virtual machine resources. While the agents monitor the resources, the daemon monitors the agents. The daemon sends heartbeats to VCS on the ESX Server layer. The heartbeats inform VCS at the ESX Server layer if VCS at the virtual machine layer is functioning properly or not. In cases where the guest operating system does not boot, heartbeats are interrupted. An interrupted heartbeat can result in an `UNKNOWN` state and can also result in a failover.

When a resource (like Oracle or a mount point) fails, the resource in the virtual machine reports an `OFFLINE` state. The `OFFLINE` state is reported to VCS at the ESX Server layer. VCS at the ESX Server layer then decides where the application's virtual machine needs to move to re-continue services.

Figure 12-4 State information passing between the virtual machine and the VCS node where the machine resides



Communication between GuestOSApp and application agents on virtual machines

The GuestOSApp agent that runs in the ESX Server layer must communicate with the application agents that run in virtual machines. The agent monitors the application that runs in the virtual machine and indicates its state to the VCS Agent Management Daemon (vcsagmd) on the virtual machine.

In case of a change in the state of the application, vcsagmd reports it to VMware's guestinfo interface. The GuestOSApp resource monitors the VMware guestinfo interface for any change in the state of the application and reports this to VCS that runs on the ESX Server layer. VCS on the ESX Server layer takes the appropriate action depending on the state change reported.

If state changes are not reflected, make sure the resource name used for the application inside the virtual machine does not have a 0 followed by another number.

You can test VMware's guestinfo interface. From the virtual machine, perform the following command to set the value of the mytestvariable1 guestinfo variable:

```
[/usr/sbin/] vmware-guestd --cmd "info-set guestinfo.mytestvariable1 myvalue1"
```

On the ESX Server where the virtual machine is online, perform the following command to see if the ESX Server can read the `mytestvariable1` variable:

```
vmware-cmd path_to_vmx_config_file getguestinfo mytestvariable1
```

To obtain the state of the application, perform the following command on the ESX Server where the virtual machine is online :

```
/opt/VRTSvcs/bin/hares -action GuestOSApp-resource getappstate \  
-sys sysA
```

This provides the following information in the given order:

- The state of the resource as reported by the agent running inside the virtual machine
- The contents of the `vcsagmd` heartbeat file, which tells you whether the `vcsagmd` heartbeats are functioning properly or not
- The current `vcsagmd` heartbeat value

If `vcsagmd` heartbeats are not updated, check the virtual machine to see if `vcsagmd` was stopped for maintenance and never restarted. Also, check if the Veritas Virtual Machine Tools package was upgraded to the current version.

Check the `vcsagmd` log file (`/var/VRTSvcs/log/vcsagmd_A.log`) to verify if the agents inside the virtual machine report the state to `vcsagmd` .

Installing the applications

Install the application on the virtual machine. This is the application that you want VCS to monitor in the virtual machine. Refer to the application's instructions for installation information. Note that you can install VCS for VMware ESX on virtual machines that already contain applications.

You must set up the application to automatically start when the virtual machine starts. VCS only monitors the application—it does not start or stop the applications. In case of a failover, VCS moves the entire virtual machine. When the virtual machine starts on the other node, the application must start automatically when the virtual machine boots.

Installing Veritas Virtual Machine Tools

Install Veritas Virtual Machine Tools on a virtual machine where you want to monitor resources for high availability. These tools contain different programs that enable high availability and monitoring. The tools installation is in a `.iso` file that you can mount as a disc on the virtual machine using the Veritas Virtualization Manager (VVM).

See [“Mounting Veritas Virtual Machine Tools”](#) on page 230.

Mounting, installing, and configuring Veritas Virtual Machine Tools on the virtual machine

Before you install Veritas Virtual Machine Tools, you need to mount the .iso file for the virtual machine. Before you mount or install the tools, prepare the following information:

- The virtual IP address for the cluster.
- The username and password required to administer the service group.
- The name of the ESXVirtualMachine resource that is associated with the virtual machine.
- In disaster recovery environment only, you need the device path for the location of the pagefile datastore on another storage device.

Mounting Veritas Virtual Machine Tools

VVM can make the Veritas Virtual Machine Tools installation program available to you for easy access.

To mount the tools .iso file

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Right-click the virtual machine where you want to mount the .iso file. Select **Add VCSVM-Tools ISO**.
If you receive this message, "This virtual machine has the ISO image mounted already," perform the following troubleshooting task:
["To resolve the "virtual machine has the ISO image mounted already" error"](#) on page 230
- 3 Click the **OK** button to add the .iso file.
The Add CD-ROM ISO window appears. The .iso file is now available for your use.

To resolve the "virtual machine has the ISO image mounted already" error

- 1 Open the VMware Infrastructure Client.
- 2 Click the virtual machine, and go to the Summary tab.
- 3 On the Summary tab, click **Edit Settings**.
- 4 Check **Client Device** in the Device Type section, and click **OK** to save.
- 5 In VVM, right-click the virtual machine and select **Add VCSVM-Tools ISO**.

Installing Veritas Virtual Machine Tools

Install the Veritas Virtual Machine Tools.

To install the Veritas Virtual Machine Tools

- 1 Navigate to the `installvcsvm-tools` program location.

```
# cd /media/cdrom
```
- 2 On the virtual machine, enter the `installvcsvm-tools -i` command.

```
# ./installvcsvm-tools -i
```
- 3 When you are prompted to install the tools, press the **y** key to proceed.
- 4 The `installvcsvm-tools` program prints some information and asks if you want to configure the tools. Press the **y** key to configure the tools.

Configuring Veritas Virtual Machine Tools

Configure Veritas Virtual Machine Tools.

To configure the Veritas Virtual Machine Tools

- 1 On the virtual machine, enter the `installvcsvm-tools -c` command.

```
./installvcsvm-tools -c
```

On a virtual machine that already has the tools mounted, use the full path for the command. At the prompt, enter:

```
/opt/VRTSvcs/bin/installvcsvm-tools -c
```
- 2 When asked if you are ready to configure the tools, answer **y**.
- 3 Enter the virtual IP address of the VCS cluster that the virtual machine belongs to.
- 4 Enter the username and password for the cluster that the virtual machine belongs to. This is the same administrator that you created when you configured the virtual machine for high availability or disaster recovery.
 - See [“To configure one virtual machine for high availability”](#) on page 155.
 - See [“To configure a virtual machine for disaster recovery”](#) on page 173.
- 5 Enter the name of the ESXVirtualMachine resource that is associated with the virtual machine.
- 6 In a disaster recovery environment, enter the device path for the location of the pagefile datastore on the secondary storage device. Note that this step is not required for a high availability environment.
Veritas Virtual Machine Tools configuration is now complete.

Validating the configuration of Veritas Virtual Machine Tools

You can verify that the tools are properly configured.

To validate the tools' configuration

- 1 Get the virtual machine resource's name, which is in the file `/etc/VRTSvcs/.vcsvmresname`. To get the name, type:

```
# cat /etc/VRTSvcs/.vcsvmresname
```

Warning: Ensure that the `.vcsvmresname` file does not get deleted. This file is critical to convey application faults to the ESX layer.

- 2 Use the virtual machine resource's name to run the following command and make sure that the command completes. At the prompt, type:

```
# /opt/VRTSvcs/bin/hares -value vmres_name Type
```

Where `vmres_name` is the virtual machine resource's name.

- 3 Again use the virtual machine resource's name and run the following command:

```
# /opt/VRTSvcs/bin/hares -state vmres_name
```

Where `vmres_name` is the virtual machine resource's name.

Configuring application and resource monitoring inside of virtual machines

After you have installed the Veritas Virtual Machine Tools, configure resources within a virtual machine. You must use the resource configuration (`vcsag_config.pl`) program to configure the resources inside virtual machines.

Before you run the resource configuration program, make sure you review the attributes for the agent that you want to configure.

For the Oracle and NetIsnr agents, see the following attributes:

- See [“Oracle agent attributes”](#) on page 193.
- See [“NetIsnr agent attributes”](#) on page 197.
- See [“Apache Web server agent attributes”](#) on page 200.
- See [“SAP NetWeaver agent attributes”](#) on page 208.
- See [“WebLogic Server agent attributes”](#) on page 213.
- See [“Application agent attributes”](#) on page 222.
- See [“Mount agent attributes”](#) on page 225.

Prerequisites

These prerequisites are to make sure the high availability setup in the virtual machine is correct. Review the following:

- Before you configure your applications on the virtual machine, ensure that the applications are running.
- Each virtual machine has VMware Tools installed on it.
- The guestinfo interfaces of the VMware Virtual Machine tools are enabled. Note that these interfaces are enabled by default.
- The virtual machine or application administrator has VCS privileges to configure the application resources.
- The ESX Server administrator has provided you with the username and password credentials for the service group that contains the virtual machine resource.
- Set up the applications to automatically start when the virtual machine starts. VCS only monitors the applications. VCS does not start or stop the applications. In case of a failover, VCS moves the entire virtual machine to another node in the cluster.
- Symantec recommends that you configure one virtual machine in one service group. You can configure multiple virtual machines in one service group if you have a multi-tier application that required the multiple virtual machines to fail over together.

Prerequisites for configuring Oracle and Netlsnr resources

You must encrypt the Pword attribute in the Oracle agent and the LsnrPwd attribute in the Netlsnr agent before configuring these attributes.

Encrypting passwords for VCS agent for Oracle

VCS provides a utility to encrypt the Oracle database user passwords and listener passwords.

Oracle provides the option of storing the listener password in the listener.ora file, in both clear text and encrypted formats. Irrespective of the format in which the password is stored in Oracle, you must encrypt the password using the `vcscrypt` utility before configuring the `LsnrPwd` attribute. If you encrypted the listener password using the Oracle `lsnrctl` utility, make sure that you pass the encrypted password to the `vcscrypt` utility. You can find the Oracle `lsnrctl` encrypted password from the following line in the listener.ora file:

```
PASSWORDS_ENCRYPTED = XXXXXXXX
```

When the agent decrypts this password, the decrypted password for the listener must be of the same format as stored in the listener.ora file.

To encrypt passwords

- 1 From the path `$VCS_HOME/bin/`, run the `vcseencrypt` utility.
 - Type the following command:

```
# vcseencrypt -agent
```
 - Enter the password and confirm it by entering it again. Press Enter.

```
# Enter New Password:  
# Enter Again:
```
- 2 Review as the utility encrypts the password and displays the encrypted password.
- 3 Enter this encrypted password as the value for the attribute.
- 4 Copy the encrypted password for future reference.

Configuring resources inside virtual machines

You need to configure the resources that you want to monitor inside the virtual machine.

To configure the resources that agents monitor inside of virtual machines

- 1 In the virtual machine, change directory to `/opt/VRTSvcs/bin`.
- 2 To start the configuration program, enter the command:

```
./vcsag_config.pl
```
- 3 Enter the name of the resource or application that you want to monitor. Your choices are:
 - **Apache**
 - **Application**
 - **Mount**
 - **Netlsnr**
 - **Oracle**
 - **SAPNW04**
 - **WebLogic9**The program lists the configured resources. If you have no configured resources, it displays a message.
- 4 To reconfigure or delete an existing resource, enter the name of the resource from the list of the displayed resources.
- 5 To configure a new resource, enter a name for that resource.

- 6 When prompted, enter a value for each attribute.
- 7 You need to enter these values in the formats requested. For definitions for these data types:
 See [“Resource data types”](#) on page 236.
- 8 Enter information for each attribute.
 - See [“Oracle agent attributes”](#) on page 193.
 - See [“Netlsnr agent attributes”](#) on page 197.
 - See [“Apache Web server agent attributes”](#) on page 200.
 - See [“SAP NetWeaver agent attributes”](#) on page 208.
 - See [“WebLogic Server agent attributes”](#) on page 213.
 - See [“Application agent attributes”](#) on page 222.
 - See [“Mount agent attributes”](#) on page 225.
- 9 When you are done, enter **done** to complete the configuration for that resource type.
- 10 You can now choose to configure more resource types, or end the configuration tasks. Enter **done**, which:
 - Saves the final configuration.
 - Sets up the corresponding configuration on the ESX Server node.
 - Restarts the VCS agent management daemon (vcsagmd), which applies the configuration on the virtual machine. The daemon starts the agents on the virtual machine for the resource types that you have configured. It also starts the GrowFS and VMIP agent processes for the virtual machine, which are used for internal purposes.
- 11 If for any reason [step 10](#) does not complete, you can apply the changes to the configuration by running the vcsag_config.pl program with the -apply option, as follows:

```
# /opt/VRTSvcs/bin/vcsag_config.pl -apply
```

Resource data types

Table 12-19 shows the data types for information that you need to enter when running `vcsag_config.pl`.

Table 12-19 Resource data types and acceptable values

Data type	Description
string	String values can include any character. If you use a space character in the string, you must enclose the string within double-quotes. You can use two double-quotes to represent an empty string.
integer	Integer values include the numbers 0 through 9.
boolean	Boolean values include 0 and 1. Zero is false, or off. One is true, or on.
association	Association values are pairs of keys and their values. The keys are always strings, while the data types of the values are provided by the association's data type. All key value pairs are separated by spaces. Two kinds of association data types exist, they are as follows: <ul style="list-style-type: none">■ Integer association—in this kind of association, values must be integers.■ String association—in this kind of association, values are strings. Examples: <ul style="list-style-type: none">■ <code>keyA 1 keyB 3 0 2</code> In this integer association, <code>keyA</code>, <code>keyB</code>, and <code>0</code> are the keys.■ <code>keyA valA keyB valB keyC valC</code> In this string association, <code>keyA</code>, <code>keyB</code>, and <code>keyC</code> are the keys.
vector	A vector is an ordered list of values. Each value is indexed using a positive integer beginning with zero.

Verifying that VCS is running

You can verify if VCS is now running. A simple test is to check for the process.

```
# ps -ef | grep agentname
```

Grep for the agent name, for example: `OracleAgent`, `NetlsnrAgent`, `ApacheAgent`, or `Sapnw04Agent`.

Applying the configuration and creating the corresponding GuestOSApp resource

The `vcsag_config.pl` program automatically takes care of a couple of tasks for you. After configuring the resource in the virtual machine, it creates a

corresponding GuestOSApp resource on the ESX Server. The GuestOSApp resource mirrors the state of your application resource inside of the virtual machine.

After VCS creates the corresponding GuestOSApp resource, it restarts the VCS agent management daemon that applies the new configuration. The GuestOSApp agent on the ESX Server then listens for updates on the GuestOSApp resource. This agent enables virtual machine failover when the resource in the virtual machine faults.

Removing the Veritas Virtual Machine Tools

To remove the Veritas Virtual Machine Tools

- 1 On the virtual machine where you want to remove the tools, enter the command `installvcsvm-tools -u`.

```
# installvcsvm-tools -u
```
- 2 When asked if you are ready to uninstall the tools, answer **y**.
The `installvcsvm-tools` program prints some information and the location of the log files.

Configuring applications and resources in Windows virtual machines

- About VCS components on virtual machines running Windows
- How VCS monitors applications and resources on virtual machines
- Overview of deploying applications in Windows virtual machines
- Installing the applications
- Installing Veritas Virtual Machine Tools
- Prerequisites for configuring agents in Windows virtual machines
- Configuring the agents for SQL Server in a Windows virtual machine
- Configuring the agent for Internet Information Services in a Windows virtual machine
- Configuring the agent for Exchange 2003 in a Windows virtual machine
- Configuring the agent for Exchange 2007 in a Windows virtual machine
- Configuring WebSphere Application Server in a Windows virtual machine
- Configuring SharePoint Portal Server in a Windows virtual machine
- Configuring generic services in a Windows virtual machine
- Removing Veritas Virtual Machine Tools from the virtual machine running Windows

About VCS components on virtual machines running Windows

VCS for VMware ESX provides agents to monitor applications that run inside virtual machines. When the agent detects an application or resource fault, the agent takes actions to communicate the state of the resource to VCS running on the ESX Server node.

Certain VCS agents also support the ability to detect administrative intervention. When an administrator gracefully shuts down an application, VCS does not initiate failover.

VCS provides agents to monitor the following applications that run on virtual machines running Windows:

Application	Agent information
SQL Server	<ul style="list-style-type: none">■ Monitors SQL Server instances.■ Detects a graceful shutdown of SQL Server.■ Supports detailed monitoring.■ See “About the VCS agents for SQL Server” on page 252.
Internet Information Services (IIS)	<ul style="list-style-type: none">■ Monitors IIS servers and services■ Detects a graceful shutdown of IIS servers.■ Supports detailed monitoring.■ See “About the VCS agent for Internet Information Services” on page 263.
Exchange Server 2003	<ul style="list-style-type: none">■ Monitors Exchange services and protocol servers.■ Detects a graceful shutdown of Exchange protocol servers.■ Supports detailed monitoring.■ See “About the VCS agent for Exchange Server 2003” on page 269.
Exchange Server 2007	<ul style="list-style-type: none">■ Monitors Exchange services and protocol servers.■ Detects a graceful shutdown of Exchange protocol servers.■ Supports detailed monitoring.■ See “About the VCS agent for Exchange Server 2007” on page 276.
SharePoint Portal Server 2007	<ul style="list-style-type: none">■ Monitors the backend SQL database.■ See “About SharePoint Portal Server 2007” on page 286.
WebSphere Application Server 6.0	<ul style="list-style-type: none">■ Monitors WebSphere Application Servers using the GenericService agent.■ See “About WebSphere Application Server 6.0” on page 284.

VCS provides the following agent to monitor generic services:

Agent	Agent information
GenericService	<ul style="list-style-type: none"> ■ Monitors generic services ■ Detects a graceful shutdown of a service ■ See “About the VCS agent for generic services” on page 287.

If a monitored resource or application fails, the corresponding agent communicates this state change to VCS on the ESX Server node. VCS can then fail over the virtual machine that runs the application onto another node.

About monitoring levels

VCS agents provide two levels of monitoring:

- **Basic monitoring**
Checks for running processes.
- **Detailed monitoring**
Performs application-specific tasks to check the application’s health. For example, the VCS agent for Oracle performs a transaction on the database and checks to see if the transaction succeeds.

All VCS agents provide basic monitoring capabilities. Some agents provide detailed monitoring capabilities.

Supported software

VCS supports the following software for virtual machines:

Guest operating systems	<ul style="list-style-type: none"> ■ Windows 2000 Server or Advanced Server with Service Pack 4 ■ Windows Server 2003: Standard Edition or Enterprise Edition (SP1 required) on either x86 (32-bit) or x86 (64-bit) ■ Windows Server 2008 (32-bit): Standard Edition, Enterprise Edition and Datacenter Edition ■ Windows Server 2008 (x64): Standard Edition, Enterprise Edition and Datacenter Edition
Required software	<ul style="list-style-type: none"> ■ Microsoft .NET Framework version 2.0 ■ VMware Tools ■ Veritas Virtual Machine Tools

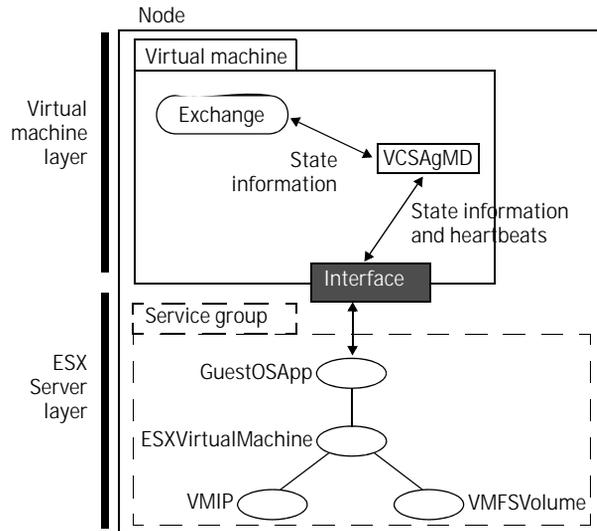
Required software for Veritas Virtual Machine Tools on Windows 2000	Veritas Virtual Machine Tools requires the sc.exe utility. It is a separate download for Windows 2000 Server. To download it, perform the following steps: <ol style="list-style-type: none">1 Download the sc.exe utility: ftp://ftp.microsoft.com/reskit/win2000/sc.zip2 Extract the sc.exe file to the system folder (C:\WINNT is the default).
Microsoft SQL servers and their operating systems	<ul style="list-style-type: none">■ Microsoft SQL Server 2000 Standard Edition or Enterprise Edition (both require SP4) with Windows Server 2003: Standard Edition or Enterprise Edition (SP1 required) on either x86 (32-bit) or x86 (64-bit)■ Windows 2000 Server or Windows 2000 Advanced Server (both require SP4)■ Microsoft SQL Server 2005 (SP1 required) on either x86 (32-bit) or x86 (64-bit) with Windows Server 2003: Enterprise Edition (SP1 required) on either x86 (32-bit) or x86 (64-bit)
Applications	<ul style="list-style-type: none">■ Internet Information Services (IIS) 5.0 and 6.0■ Microsoft Exchange Server 2003■ Microsoft Exchange Server 2007 SP1■ SharePoint Portal Server 2007■ WebSphere Application Server 6.0

How VCS monitors applications and resources on virtual machines

The VCS agent management daemon (vcsagmd) is a very light-weight daemon that resides in the virtual machine. It manages the agents that monitor virtual machine resources. While the agents monitor the resources, the daemon monitors the agents. The daemon sends heartbeats to VCS on the ESX Server layer. The heartbeats inform VCS at the ESX Server layer if VCS at the virtual machine layer is functioning properly or not. In cases where the guest operating system does not boot, heartbeats are interrupted. An interrupted heartbeat can result in an UNKNOWN state and can also result in a failover.

When a resource (like Exchange) fails, the resource in the virtual machine reports an OFFLINE state. The OFFLINE state is reported to VCS at the ESX Server layer. VCS at the ESX Server layer then decides where the application's virtual machine needs to move to re-continue services.

Figure 13-1 State information passing between the virtual machine and the VCS node where the machine resides



Overview of deploying applications in Windows virtual machines

Deploying applications in a Windows virtual machine involves the following steps:

- I Make sure the virtual machine meets the requirements for supported software.
See "[Supported software](#)" on page 241.
- II Install and configure Veritas Virtual Machine Tools
See "[Installing Veritas Virtual Machine Tools](#)" on page 245.
- II Install the application in the virtual machine
See "[Installing the applications](#)" on page 245.

- IV Configure the application resources.
- Review the prerequisites for configuring agents.
- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- Configure the agent for the application that you want to monitor:
- See [“Configuring the agents for SQL Server in a Windows virtual machine”](#) on page 252.
 - [“Configuring the agent for Internet Information Services in a Windows virtual machine”](#) on page 263.
 - See [“Configuring the agent for Exchange 2003 in a Windows virtual machine”](#) on page 269.
 - See [“Configuring the agent for Exchange 2007 in a Windows virtual machine”](#) on page 276.
 - See [“Configuring WebSphere Application Server in a Windows virtual machine”](#) on page 284.
 - See [“Configuring SharePoint Portal Server in a Windows virtual machine”](#) on page 286.
 - See [“Configuring generic services in a Windows virtual machine”](#) on page 287.
- V Verify the configuration.
- See [“Verifying the configuration for application monitoring”](#) on page 293.
- VI Apply the configuration and create corresponding resources on the ESX host.
- See [“Applying the configuration and creating the corresponding GuestOSApp resource”](#) on page 293.

Installing the applications

Install the application on the virtual machine. This is the application that you want VCS to monitor in the virtual machine. Refer to the application's instructions for installation information. Note that you can install VCS for VMware ESX on virtual machines that already contain applications.

You must set up the application to automatically start when the virtual machine starts. VCS only monitors the application—it does not start or stop the applications. In case of a failover, VCS moves the entire virtual machine. When the virtual machine starts on the other node, the application must start automatically when the virtual machine boots.

While installing WebSphere Application Server 6.0, during Profile creation you must specify the Profile name in the correct format. The Profile Name must not consist of "0" followed by another digit.

Installing Veritas Virtual Machine Tools

Install Veritas Virtual Machine Tools on a Windows virtual machine where you want to monitor an application for high availability.

Adding the tools .iso file

The utility is on the disc inside the `vcsvm_tools` directory. It is in the ISO format as `win-x86-vcsvm-tools.iso`.

To add the tools .iso file

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Right-click the virtual machine where you want to add the .iso file. Select **Add VCSVM-Tools ISO**. VVM automatically selects the appropriate .iso file to match the operating system.
- 3 The Add CD-ROM ISO window appears.
- 4 Click the **OK** button to add the .iso file.
The ISO file is now available for your use.

Connecting to the virtual machine using the Remote Desktop Connection program

If you are planning to run the Veritas Virtual Machine Tools installer for Windows virtual machines using the Remote Desktop Connection program,

make sure you use the /console option. This ensures that environment variable-related updates that the installer adds are sent to the operating system and applications. Installing the tools using the Remote Desktop Connection program with the /console option allows you to run the agent configuration wizards without the need to reboot the virtual machine.

To use a remote desktop session to connect to the virtual machine

- 1 Click the **Start** button and select **Run**.
- 2 In the command prompt, run the following command:
`%SystemRoot%\system32\mstsc.exe /console`
- 3 Specify the IP address of the virtual machine.
- 4 Specify administrator privilege credentials to login.

Installing Veritas Virtual Machine Tools

To install Veritas Virtual Machine Tools

- 1 Once you have mounted the appropriate ISO file to your virtual machine, run the vcsvm-tools.exe file to install the tools.
- 2 Review the Welcome screen and click **Next**.
- 3 Review the License Agreement, choose to accept, and then click **Next**.
- 4 On the Destination Folder screen, either accept the default location or click **Browse** to choose another location. Click **Next** when done.
- 5 On the HALogin Configuration screen, enter the cluster login credentials.

Note: You must configure halogin before running the SQL and IIS agent wizards; otherwise, the wizards won't be able to complete the configuration.

- IP address or DNS name of your VCS ESX cluster
 - User name and password for your VCS ESX cluster
 - VCS virtual machine resource associated with this system
- 6 Click **Next**.

- 7 On the Convert Basic Disks to Dynamic screen, check the box if you have basic disks that you need to convert to dynamic disks. Make sure that if you have any volumes mounted on a basic disk, unmount the volumes before converting the disks to dynamic.

Note: This applies to Windows 2003 users only. Windows 2000 users won't see this screen.

Click **Next** when done.

- 8 On the PageFile Drive Selection screen, you are presented with a list of available drives, the size and type of the pagefile if present, and the available space on the drive.
Select a drive and
 - Delete the pagefile, if it exists on a replicated volume. Click **Delete Pagefile**.
 - Create a pagefile. You can choose to either create a Custom pagefile or a System-managed pagefile.
For a custom pagefile, you must enter the initial size (in MB) and the maximum size (in MB), and then click **Create**. Note that the maximum size is constrained by the Windows maximum size limit of 4096 MB.
or
Check the System Managed check box and click **Create**.
- Click **Next** when done.
See ["Pagefile configuration in Windows Server 2008"](#) on page 249.
- 9 On the Ready to Install screen, click **Install**.
 - 10 Click **Finish** to close the installer.
 - 11 Verify the installation. Check to see if the VCSAgMD service is present in the Services panel. (**Start > Programs > Administrative Tools > Services**)

Configuring Veritas Virtual Machine Tools

If you skipped the configuration of the tools, you can return to configure it.

To configure Veritas Virtual Machine Tools

- 1 On the virtual machine, open the **Control Panel** and go to **Add or Remove Programs**.
- 2 Select the **Veritas Virtual Machine Tools** and click the **Change** button.
- 3 Click the **Next** button.
- 4 Select the **Modify** radio button and click the **Next** button.

- 5 On the HALogin Configuration screen, enter the cluster login credentials.

Note: You must configure halogin before running the SQL and IIS agent wizards; otherwise, the wizards won't be able to complete the configuration.

- IP address or DNS name of your VCS ESX cluster
- User name and password for your VCS ESX cluster
- VCS virtual machine resource associated with this system

- 6 Click **Next**.

- 7 In the Convert Basic Disks to Dynamic screen, check the box if you have basic disks that you need to convert to dynamic disks. Make sure that if you have any volumes mounted on a basic disk, unmount the volumes before converting the disks to dynamic.

Note that this applies for Windows 2003 users only, as Windows 2000 users won't see this screen.

Click **Next** when done.

- 8 In the PageFile Drive Selection screen, you are presented with a list of available drives, the size and type of the pagefile if present, and the available space on the drive.

Select a drive and

- Delete the pagefile, if it exists on a replicated volume. Click **Delete Pagefile**.
- Create a pagefile. You can choose to either create a Custom pagefile or a System-managed pagefile.
For a custom pagefile, you must enter the initial size (in MB) and the maximum size (in MB), and then click **Create**. Note that the maximum size is constrained by the Windows maximum size limit of 4096 MB.

or

Check the System Managed check box and click **Create**.

Click **Next** when done.

See "[Pagefile configuration in Windows Server 2008](#)" on page 249.

- 9 In the Ready to Install screen, click **Install**.
- 10 Click **Finish** to close the installer.
- 11 Verify the configuration. Check to see if the VCSAgMD service is present in the Services panel. (**Start > Programs > Administrative Tools > Services**)

Pagefile configuration in Windows Server 2008

Windows Server 2008 manages pagefiles by default. This requires that you perform several extra steps to enable the installer for the Veritas Virtual Machine Tools to configure pagefiles. These procedures enable you to set pagefile configuration to manual mode so that the installer can either delete existing pagefiles or create new pagefiles. Pagefiles can be accessed but not manipulated in System Managed mode.

Perform the following commands to see which mode the pagefiles are in.

To list the pagefiles that are in the System Managed mode

- ◆ List files that are in System Managed mode:

```
$wmic pagefile list /format:list
```

To list the pagefiles that are in Manual mode

- ◆ List the pagefiles that are in manual mode:

```
$wmic pagefileset list /format:list
```

After you have determined the mode that the pagefiles are in, you need to set the operating system managed pagefiles to manual mode. Perform the following procedure to set the pagefile configuration to manual.

To set pagefile configuration mode to manual

- 1 Create a placeholder pagefile for each drive:

```
$wmic computersystem set automaticmanagedpagefile=FALSE
```

- 2 Verify that you have switched to manual mode.

```
$wmic computersystem get automaticmanagedpagefile
```

- 3 Delete the placeholder pagefiles for each drive:

```
$wmic pagefileset where Name='C:\\pagefile.sys' delete
```

- 4 Reboot the server to prepare for the manual creation of pagefiles by the installer for the Veritas Virtual Machine Tools.

Note that you need to reboot the Windows Server 2008 whenever you want to:

- Change the pagefile management mode
- Delete pagefiles in manual mode

The installer uses the following command to create a pagefile.

```
$wmic pagefileset create  
Name="x:\pagefile.sys",InitialSize=xxx,MaximumSize=xxx
```

Validating the configuration of Veritas Virtual Machine Tools

Use the following procedure to verify that the tools are properly configured.

To validate Veritas Virtual Machine Tools configuration

- 1 Get the virtual machine resource name, which is in the file `$VCS_HOME\vcsvmresname`. Typically, you can find the `vcsvmresname` file in the `C:\Program Files\Veritas\cluster server\`. To get the name, type the following:

```
C:\Program Files\Veritas\cluster server\vcsvmresname
```

Warning: Ensure that the `vcsvmresname` file does not get deleted. This file is critical to convey application faults to the ESX layer.

- 2 Use the virtual machine resource name to run the following command and make sure that the command completes. At the prompt, type:

```
C:\Program Files\Veritas\cluster server\hares -value  
vmres_name Type
```

Where `vmres_name` is the virtual machine resource name.

- 3 Again, use the virtual machine resource name and run the following command:

```
C:\Program Files\Veritas\cluster server\hares -state  
vmres_name
```

Where `vmres_name` is the virtual machine resource name.

Prerequisites for configuring agents in Windows virtual machines

- Each guest operating system must have VMware Tools installed on it.
- The guestinfo interfaces of the VMware Virtual Machine tools must be enabled. Note that these interfaces are enabled by default.
- The Veritas Virtual Machine Tools must be installed on a virtual machine running Windows.
- The virtual machine or application administrator must have VCS privileges to configure the application resources.
- The ESX Server administrator should provide the username and password credentials for the service group to the virtual machine or application administrator.
- If you have configured Windows Firewall, add the following to the Firewall Exceptions list:
 - Port 14150 or the VCS Command Server service,
%vcs_home%\bin\CmdServer.exe.
Here, %vcs_home% is the installation directory for VCS, typically
C:\Program Files\Veritas\Cluster Server.
 - Port 14141
- If User Access Control (UAC) is enabled on Windows Server 2008 systems, you must launch the command prompt in the *Run as administrator mode* and then run the VCS commands.

Configuring the agents for SQL Server in a Windows virtual machine

This section describes the agent for SQL Server and provides instructions for configuring the agent.

- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- See [“About the VCS agents for SQL Server”](#) on page 252.
- See [“Configuring the agents for SQL Server”](#) on page 260.

About the VCS agents for SQL Server

Microsoft SQL Server is a Relational Database Management System (RDBMS) used for building, managing, and deploying business applications. The SQL Server infrastructure provides services such as jobs, notification, and in-built replication.

The SQL Server agents monitor Microsoft SQL Server and its services on a VCS cluster to ensure high availability. VCS provides separate agents for SQL Server 2000 and SQL Server 2005.

The agents can detect when the SQL Server is brought down gracefully. When the SQL Server is brought down gracefully, the agents do not trigger a resource fault even though the SQL Server is down.

This agent supports detailed monitoring, refer to the optional DetailMonitor attribute for more information.

VCS agents for SQL Server 2000 and SQL Server 2005

The SQL agents are as follows:

- **Agent for SQL Server 2000 service.** The agent monitors SQL Server 2000 service.
- **Agent for SQL Server 2005 service.** The agent monitors SQL Server 2005 service.
- **Agent for SQL Server 2005 Agent service.** The agent monitors SQL Server 2005 agent service.
- **Agent for SQL Server 2005 Analysis service.** The agent monitors SQL Server 2005 Analysis service.
- **Agent for SQL Server 2005 Search service.** The agent provides high availability for full-text search indices with a clustered SQL instance.

Agent functions

- **Monitor**
Verifies the configured SQL Server instance is running.
Can perform detailed monitoring for SQL Server 2000 and SQL Server 2005.

State definitions

- **ONLINE**
Indicates the configured SQL Server instance is available.
- **OFFLINE**
Indicates the configured SQL Server instance is not available.
Can also indicate that the administrator gracefully stopped the SQL Server instance.
- **UNKNOWN**
Indicates the agent could not determine the status of SQL Server.

SQL Server agent attributes

Review the following information to familiarize yourself with the various agent attributes. This information will assist you during the agent configuration.

Agent for SQL Server 2000

The agent for SQL Server 2000 is represented by the SQLServer2000 resource type. The attributes for this agent are as follows:

Table 13-1 SQL Server 2000 agent required attributes

Required attributes	Description
Instance	Name of instance to monitor. If the attribute is blank, the agent monitors the default instance. Type and dimension: string-scalar
LanmanResName	The Lanman resource name on which the SQL Server 2000 resource depends. Type and dimension: string-scalar
MountResName	The mount resource name on which the SQL Server 2000 resource depends. Type and dimension: string-scalar
SQLOnlineTimeout	Number of seconds that can elapse before the Online agent function aborts. Default is 90. Type and dimension: integer-scalar
SQLOfflineTimeout	Number of seconds that can elapse before Offline agent function aborts. Default is 90. Type and dimension: integer-scalar
IsGuestOS	A flag that indicates whether SQL is deployed in a VCS for VMware environment. This attribute controls the agent behavior. Set this attribute to 1 (True) if you have installed SQL on a VMware virtual machine in a VCS for VMware cluster environment. Default value is 1 (True). Type and dimension: boolean-scalar

Table 13-2 SQL Server 2000 agent optional attributes

Optional Attributes	Description
DetailMonitor	<p>Defines whether the agent performs detail monitoring of SQL Server 2000 database. The value 0 indicates the agent does not monitor the database in detail. The value 1 indicates the agent does.</p> <p>Default = 0</p> <p>Note: If the attribute is set to 1, the attributes Username, Password, Domain, SQLDetailMonitorTimeOut, and SQLFile must be assigned appropriate values.</p> <p>Type and dimension: boolean-scalar</p>
FaultOnDMScriptFailure	<p>Defines whether the agent fails over the service group if the detail monitoring script execution fails.</p> <p>Default = 1</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.</p> <p>Type and dimension: boolean-scalar</p>
SQLDetailMonitorTimeout	<p>Number of seconds that can elapse before the detail monitor routine aborts. Default is 30.</p> <p>Type and dimension: integer-scalar</p>
Username	<p>The Microsoft Windows authentication name when logging in to a database for detail monitoring. This attribute must not be null if the DetailMonitor attribute is set to 1.</p> <p>Type and dimension: string-scalar</p>
Domain	<p>Domain for the user account. This attribute is used to create a trusted connection to the SQL Server 2000 instance if the DetailMonitor attribute is set to 1.</p> <p>Type and dimension: string-scalar</p>
Password	<p>Password for logging in to a database for detail monitoring. This attribute must not be null if the DetailMonitor attribute is set to 1.</p> <p>Type and dimension: string-scalar</p>

Table 13-2 SQL Server 2000 agent optional attributes (continued)

Optional Attributes	Description
SQLFile	The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the the DetailMonitor attribute is set to 1. Type and dimension: string-scalar

Agent for MSSearch service

The agent for MSSearch service is represented by the MSSearch resource type. The attribute for this agent is as follows:

Table 13-3 MSSearch service agent required attribute

Required Attribute	Description
AppName	The name of MSSearch instance to be monitored. Type and dimension: string-scalar

Agent for SQL Server 2005

The agent for SQL Server 2005 is represented by the SQLServer2005 resource type. The attributes for this agent are as follows:

Table 13-4 SQL Server 2005 agent required attributes

Required Attributes	Description
Instance	Name of instance to monitor. If the attribute is blank, the agent monitors the default instance. Type and dimension: string-scalar
LanmanResName	The Lanman resource name on which the SQL Server 2005 resource depends. Type and dimension: string-scalar
SQLOnlineTimeout	Number of seconds that can elapse before Online agent function aborts. Default is 90. Type and dimension: integer-scalar
SQLOfflineTimeout	Number of seconds that can elapse before Offline agent function aborts. Default is 90. Type and dimension: integer-scalar

Table 13-4 SQL Server 2005 agent required attributes (continued)

Required Attributes	Description
IsGuestOS	<p>A flag that indicates whether SQL is deployed in a VCS for VMware environment. This attribute controls the agent behavior. Set this attribute to 1 (True) if you have installed SQL on a VMware virtual machine in a VCS for VMware cluster environment.</p> <p>Default value is 1 (True).</p> <p>Type and dimension: boolean-scalar</p>

Table 13-5 SQL Server 2005 agent optional attributes

Optional Attributes	Description
DetailMonitor	<p>Defines whether the agent performs detail monitoring of SQL Server 2005 database. The value 0 indicates the agent will not monitor the database in detail. The value 1 indicates the agent will.</p> <p>Default = 0.</p> <p>Note: If the attribute is set to 1, the attributes Username, Password, Domain, SQLDetailMonitorTimeOut, and SQLFile must be assigned appropriate values.</p> <p>Type and dimension: boolean-scalar</p>
FaultOnDMScriptFailure	<p>Defines whether the agent fails over the service group if the detail monitoring script execution fails.</p> <p>Default = 1</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.</p> <p>Type and dimension: boolean-scalar</p>
SQLDetailMonitorTimeout	<p>Number of seconds that can elapse before the detail monitor routine aborts. Default is 30.</p> <p>Type and dimension: integer-scalar</p>

Table 13-5 SQL Server 2005 agent optional attributes (continued)

Optional Attributes	Description
Username	<p>The Microsoft Windows authentication name when logging in to a database for detail monitoring. This attribute must not be null if the DetailMonitor attribute is set to 1.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
Domain	<p>Domain for the user account. This attribute is used to create a trusted connection to the SQL Server 2005 instance if the DetailMonitor attribute is set to 1.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
Password	<p>Password for logging in to a database for in-depth monitoring. This attribute must not be null if the DetailMonitor attribute is set to 1.</p> <p>Type and dimension: string-scalar</p>
SQLFile	<p>The location of the SQLFile executed during a monitor cycle. This attribute must not be null if the the DetailMonitor attribute is set to 1.</p> <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>

Agent for SQL Server 2005 Agent service

The agent for SQL Server 2005 Agent service is represented by the SQLAgService2005 resource type. The attributes for this agent are as follows:

Table 13-6 SQL Server 2005 Agent service agent required attributes

Required Attributes	Description
SQLServer2005ResName	<p>The name of the SQLServer2005 resource on which the SQL Server 2005 Agent service resource depends.</p> <p>Type and dimension: string-scalar</p>
LanmanResName	<p>The Lanman resource name on which the SQL Server 2005 resource depends.</p> <p>Type and dimension: string-scalar</p>

Agent for SQL Server 2005 Analysis service

The agent for SQL Server 2005 Analysis service is represented by the SQLLapService2005 resource type. The attributes for this agent are as follows:

Table 13-7 SQL Server 2005 Analysis service agent required attributes

Required Attributes	Description
SQLServer2005ResName	The name of the SQLServer2005 resource on which the SQL Server 2005 Analysis service resource depends. Type and dimension: string-scalar
LanmanResName	The Lanman resource name on which the SQL Server 2005 resource depends. Type and dimension: string-scalar

MSDTC agent

The MSDTC agent is represented by the MSDTC resource type. The attributes for this agent are as follows:

Table 13-8 MSDTC agent required attributes

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

Configuring the agents for SQL Server

The SQL Server agents can be configured using the SQL Server Configuration wizard.

Prerequisites for configuring the SQL Server agents

Make sure the following prerequisites have been met before running the wizard.

- Local administrator privileges are assigned to the user running the wizard.
- SQL Server 2000 or SQL Server 2005 are installed on the computer to be monitored.
- SQL Server and SQL Server Browser services for each instance must be online.

To configure the SQL Server Agent

- 1 Navigate to the configuration wizard by clicking **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > SQL Server Configuration Wizard**.
- 2 Review the Welcome screen and click **Next**.
- 3 In the Instance Selection screen, select the SQL Server instance or instances that you want to monitor. By default, once you have selected a SQL Server 2005 instance, the boxes under the Search, SQLAgent, and Analysis fields are selected. You can deselect any of these fields if you do not want that agent monitoring the corresponding SQL Server service. If you select a SQL Server 2000 instance, no fields are displayed.

The Instance Selection screen in the SQL Server Agent Configuration Wizard corresponds to the Microsoft SQL Server services as follows.

Table 13-9 List of SQL Server services and agents in Instance Selection screen

Microsoft SQL Server Services	SQL Server Agents
SQL Server [instance name] Is either SQL Server 2000 or 2005	SQL Server (2000 or 2005) instance selected for monitoring
SQL Server (2005) Full-Text Search Creates full-text indexes on content and properties of structured and semi-structured data	Search Monitors the SQL Server Full-Text Search function

Table 13-9 List of SQL Server services and agents in Instance Selection screen

Microsoft SQL Server Services	SQL Server Agents
SQL Server (2005) Agent Executes jobs, monitors SQL Server, fires alerts, and allows some automation of administrative functions	SQLAgent Monitors the SQL Server Agent
SQL Server (2005) Analysis Services Provides online analytical processing (OLAP) and data mining functionality	Analysis Monitors the Server Analysis Services

- 4 Select **Configure detailed monitoring for selected instances** to enable detailed monitoring of the selected agents.
- 5 When you have selected all of the SQL Server instances that you wish to monitor, click **Next**.
If you have selected detailed monitoring for the SQL Server instance, the Detailed Monitoring screen appears.
- 6 From the SQL Instance List, select an instance and specify the path name, if different from the default path name given.
- 7 Specify whether you want to logon as the local system account (default) or as a domain account. If you choose to logon as a domain account, you need to enter your username, password, and domain information. Click **Apply**.
- 8 If you have more than one instance in the list, repeat this step for each instance, and click **Apply** after each selection.
Note that if you do not apply a path name to an instance, you get an error message that reminds you to specify the path name for that instance.
- 9 When you have finished, click **Next**. The Failure Actions screen appears. The Failure Action screen lists the Microsoft SQL Server Services that you have chosen to monitor (in step 3).

Table 13-10 List of chosen SQL Server services

Service Name	Description
SQL Server [instance name]	Provides storage, processing and controlled access of data...
SQL Server Full Text Search	Quickly creates full-text indexes on content...
SQL Server Analysis Services	Supplies online analytical processing (OLAP)...

Table 13-10 List of chosen SQL Server services

Service Name	Description
SQL Server Agent	Executes jobs, monitors SQL Server, and fires alerts

- 10 Double-click one of the Services to open the Recovery Action dialog box. Specify the action to take after the first failure, the second failure, and subsequent failures. The options are:
 - Take no action (default for all three)
 - Restart the Service
 - Run a Script
- 11 Specify when to reset the fail count (after 1 day is the default value).
- 12 Specify the script path name. You can accept the default value, or click **Browse** to select a different path name.
- 13 Enter any command line parameters in the text box (optional) and click **OK**. You are returned to the Failure Actions screen.
- 14 Repeat [step 10](#) through [step 13](#) for each of the Services listed. Click **Next** when you are done.
- 15 The Resource Summary page lists the agents, or resources, you have selected: [Instance name]_SQLServer, [Instance name]_SearchService, [Instance name]_SQLAgService, and [Instance name]_SQLAnalysisService. Select a resource and click it to view its attributes.

In this window, you can rename the resources that you have created. Make sure that these resource names match the GuestOSApp resources that will be monitored at the ESX cluster level.
- 16 When you are done reviewing the resources, click **Next**.
- 17 Click **Finish** to close the configuration wizard.

Configuring the agent for Internet Information Services in a Windows virtual machine

This section describes the agent for Internet Information Services (IIS) and provides instructions for configuring the agent.

- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- See [“About the VCS agent for Internet Information Services”](#) on page 263.
- See [“Configuring the agent for Internet Information Services”](#) on page 266.

About the VCS agent for Internet Information Services

The Internet Information Services (IIS) agent monitors the status of sites configured using IIS 5.0 and 6.0.

The agent provides two ways of monitoring application pools associated with IIS Web sites:

- One IIS resource configures a Web site and sets monitoring options for application pools associated with the site
- One IIS resource configures a Web site; other resources configure individual application pools.

The agent can detect when IIS is brought down gracefully. When the IIS is brought down gracefully, the agent does not trigger a resource fault even though IIS is down.

This agent supports detailed monitoring, refer to the optional `DetailMonitor` attribute for more information.

See [“Configuring the agent for Internet Information Services”](#) on page 266.

Agent functions

- **Monitor**
Verifies the configured sites or application pools are running.
Monitors each site in detail by attempting an actual socket connection to the port.

State definitions

- **ONLINE**
Indicates the configured site or application pool is available.

- OFFLINE
Indicates the configured site or application pool is not available.
Can also indicate that the administrator gracefully stopped the application.
- UNKNOWN
Indicates the agent could not determine the status of the resource.

Internet Information Services (IIS) agent attributes

To configure the agent to monitor an application pool, configure the SiteType and SiteName attributes only. The agent ignores other attributes when it is configured to monitor an application pool.

Table 13-11 Internet Information Services (IIS) required attributes

Required Attributes	Description
SiteType	<p>Defines whether the resource is configured to monitor an IIS site or an application pool.</p> <p>If the resource is configured to monitor an application pool, set the attribute to APPPOOL.</p> <p>If the resource is configured to monitor an IIS site, set this attribute to the name of the IIS service associated with the site. The attribute can take any of the following values:</p> <ul style="list-style-type: none"> ■ W3SVC ■ MSFTPSVC ■ SMTPSVC ■ NNTPSVC <p>Type and dimension: string-scalar</p>
SiteName	<p>The name of the IIS site, the virtual server, or the application pool to be monitored by the agent.</p> <p>The value of this attribute depends on that of the SiteType attribute. The SiteName attribute can take the following values:</p> <ul style="list-style-type: none"> ■ The name of a site, if SiteType is W3SVC or MSFTPSVC ■ The name of a virtual server, if SiteType is SMTPSVC or NNTPSVC ■ The name of an application pool, if SiteType is APPPOOL <p>Note: This attribute can take localized values.</p> <p>Type and dimension: string-scalar</p>
IPResName	<p>The name of the IP resource configured for the IP to which the site is bound.</p> <p>Type and dimension: string-scalar</p>
PortNumber	<p>The port to which the site is bound.</p> <p>Type and dimension: string-scalar</p>

Table 13-12 Internet Information Services (IIS) optional attributes

Optional Attributes	Description
AppPoolMon	<p>Defines the monitoring modes for the application pool associated with the Web site being monitored.</p> <p>Configure this attribute only if SiteType is W3SVC and IIS is configured to run in the Worker Process Isolation mode.</p> <p>The attribute can take one of the following values:</p> <ul style="list-style-type: none"> ■ NONE: Indicates that the agent does not monitor the application pool associated with the Web site. ■ DEFAULT: Indicates that the agent monitors the <i>root</i> application pool associated with the Web site. If this attribute is set, the agent starts, stops, and monitors the root application pool associated with the Web site. If the root application pool is stopped externally, the agent fails over the service group. ■ ALL: Indicates the agent starts all application pools associated with the Web site, but monitors and stop the <i>root</i> application pool only. <p>Type and dimension: integer-scalar</p>
DetailMonitor	<p>A flag that defines whether the agent monitors the site in detail. The value 1 indicates the agent monitors each site in detail by attempting an actual socket connection to the port.</p> <p>Type and dimension: boolean-scalar</p>
DetailMonitorInterval	<p>The number of monitor cycles after which the agent attempts detail monitoring. For example, the value 5 indicates that the agent monitors the resource in detail after every 5 monitor cycles.</p> <p>Type and dimension: integer-scalar</p>

Configuring the agent for Internet Information Services

The IIS agent can be configured using the IIS Configuration wizard.

Prerequisites for configuring the IIS agent

Make sure the following prerequisites have been met before running the wizard.

- IP address must have a forward and reverse entry in the DNS.
- The Site Name for the site to be monitored must be unique.

To configure the IIS agent

- 1 Navigate to the configuration wizard by clicking **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > IIS Configuration Wizard**.
- 2 Review the Welcome screen and click **Next**.
- 3 In the Configure IIS Sites screen, select the sites that you want to monitor. By default, all listed sites are selected. Note that when you deselect a site, the corresponding check box for detailed monitoring is also deselected. The Site Name entries correspond to the resources listed under the IIS Manager in Windows Administrative Tools (**Control Panel > Administrative Tools > IIS Manager**). There are four services associated with the IIS Manager:
 - FTP sites
 - Web Sites
 - Default SMTP Virtual Server
 - Default NNTP Virtual ServerEach of the four services has resources, or sites, associated with it. For example, under Web Sites you could have Default Web Site, Administration, and Microsoft SharePoint.
- 4 For each selected site, type the IP address for that site. If you do not enter an IP address, an error message appears that prompts you to enter an IP address.
When you are done, click **Next**.
- 5 In the Application Pool Configuration screen, select the application pools that you want to monitor, if any. Using the drop-down list, select one of the following options and click **Next** when done:
 - DEFAULT to monitor the root application pool.
 - NONE for no application pool monitoring.
 - ALL to monitor all application pools associated with the site.
- 6 The Failure Action screen lists the IIS services that you have chosen to monitor.
- 7 Double-click one of the sites to open the Recovery Action dialog box. Specify the action to take after the first failure, the second failure, and subsequent failures. The options are:
 - Take no action (default for all three)
 - Restart the Service
 - Run a Script
- 8 Specify when to reset the fail count (after 1 day is the default value).

- 9 Specify the script path name. You can accept the default value, or click **Browse** to select a different path name.
- 10 Enter any command line parameters in the text box (optional) and click **OK**. You are returned to the Failure Actions screen.
- 11 Repeat [step 7](#) through [step 10](#) for each of the Sites listed. Click **Next** when you are done.
- 12 The Resource Summary page lists the sites, or resources, you have selected. Select a resource and click it to view its attributes.
In this window, you can rename the resources that you have created. Make sure that these resource names match the GuestOSApp resources that will be monitored at the ESX cluster level.
- 13 When you are done reviewing the resources, click **Next**.
- 14 Click **Finish** to close the configuration wizard.

Configuring the agent for Exchange 2003 in a Windows virtual machine

This section describes the agent for Exchange 2003 and provides instructions for configuring the agent.

- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- See [“About the VCS agent for Exchange Server 2003”](#) on page 269.
- See [“Configuring the agent for Exchange 2003”](#) on page 273.

About the VCS agent for Exchange Server 2003

Microsoft Exchange Server 2003 provides a reliable and scalable messaging and collaboration infrastructure. The VCS for VMware agent for Microsoft Exchange Server 2003 provides high availability for Microsoft Exchange in a VCS for VMware cluster environment.

The VCS for VMware agent for Microsoft Exchange contains two agents:

- **Exchange Service agent**
Monitors core Exchange services.
- **Exchange Protocol agent**
Monitors Exchange protocol servers configured under the Exchange protocol services.

The agents for Microsoft Exchange monitor the configured Exchange services and Exchange protocol servers. Both agents work in conjunction to provide high availability for Microsoft Exchange.

This agent supports detailed monitoring, refer to the optional `DetailMonitor` attribute for more information.

Exchange Service agent

The Exchange Service agent monitors the status of the following Exchange services:

- **Microsoft Exchange Information Store (MSEExchangeIS)**
The Exchange storage used to hold messages in users' mailboxes and in public folders.
- **Microsoft Exchange System Attendant (MSEExchangeSA)**
The Exchange component responsible for maintenance and ensuring that operations run smoothly.

- Microsoft Exchange Message Transfer Agent (MSEExchangeMTA)
The Exchange component responsible for routing messages.
- Microsoft Exchange Routing Engine (RESvc)
The Exchange routing engine service.
- Microsoft Exchange Management Service (MSEExchangeMGMT)
Provides Exchange management information through WMI.

Each Microsoft Exchange service is configured as a VCS resource of type ExchService.

Agent functions

- Monitor
Determines the state of the configured Exchange service by querying the Service Control Manager (SCM).
The agent verifies the state of the enabled databases (databases that are automatically mounted when the service starts up). If an enabled database is dismounted, the agent returns UNKNOWN state.

Note: The VCS for VMware agent for Microsoft Exchange monitors only the enabled databases. To enable databases, run Microsoft Exchange System Manager and clear the **Do not mount this store at start-up** check box in database properties. If the agent detects that an enabled database is not mounted, it returns an UNKNOWN state. So, to disable the database, check the **Do not mount this store at start-up** check box if you want to dismount a database.

State definitions

- ONLINE
Indicates the configured Exchange service is available.
Can perform detailed monitoring.
- OFFLINE
Indicates the configured Exchange service is not available.
Can also indicate that the administrator gracefully stopped the Exchange service.
- UNKNOWN
Indicates the agent could not determine the status of Exchange service.

Exchange Protocol agent

The Exchange Protocol agent monitors the protocol servers configured under the following Exchange protocols:

- **Post Office Protocol (POP3SVC)**
Internet messaging protocol used to access email from a remote location.
- **Simple Mail Transfer Protocol (SMTPSVC)**
TCP/IP protocol used to transfer email over the internet, which is also the native mail transport in Microsoft Exchange.
- **Internet Message Access Protocol (IMAP4SVC)**
Internet messaging protocol used to access email messages stored on a remote server.
- **World Wide Web (W3SVC)**
World Wide Web service.

Each protocol server to be monitored is configured as a VCS resource of type ExchProtocol.

Agent functions

- **Monitor**
Determines the state of the configured Exchange protocol servers.

State definitions

- **ONLINE**
Indicates the configured Exchange protocol service is available.
- **OFFLINE**
Indicates the configured Exchange protocol service is not available.
Can also indicate that the administrator gracefully stopped the Exchange protocol service.
- **UNKNOWN**
Indicates the agent could not determine the status of Exchange protocol service.

Exchange Server 2003 agent attributes

Review the following information to familiarize yourself with the Exchange agent attributes. Use this information during the agent’s configuration.

Exchange service agent attributes

Review the following information to familiarize yourself with the required agent attributes for an ExchService resource type. This information assists you during the agent configuration.

Table 13-13 Exchange Service agent required attributes

Required Attributes	Description
Service	<p>The name of the Exchange service to be monitored. By default this attribute takes the following values:</p> <ul style="list-style-type: none"> ■ MExchangeIS ■ MExchangeSA ■ RESvc <p>Additionally, you can also configure the following services:</p> <ul style="list-style-type: none"> ■ MExchangeMTA ■ MExchangeMGMT <p>Type and dimension: string-scalar</p>
LanmanResName	Do not modify this attribute. For internal use only.

Table 13-14 Exchange Service agent optional attributes

Optional Attribute	Definition
DetailMonitor	<p>A flag that determines whether the agent monitors the MExchangeIS service in detail. The value 1 indicates that the agent monitors the service in detail; the value 0 indicates it does not.</p> <p>Set this attribute only for resources configured to monitor the MExchangeIS service; the attribute is ignored for other services.</p> <p>Type and dimension: integer-scalar</p>

Exchange protocol agent attributes

Review the following information to familiarize yourself with the required agent attributes for an ExchProtocol resource type. This information will assist you during the agent configuration.

Table 13-15 Exchange Protocol agent required attributes

Required Attributes	Description
VirtualServer	The name of the Exchange protocol server to be monitored. This attribute can take localized values. Type and dimension: string-scalar
LanmanResName	Do not modify this attribute. For internal use only.
Protocol	The Exchange protocol for which the Exchange protocol server is configured. This attribute could take any of the following values: <ul style="list-style-type: none"> ■ POP3SVC ■ W3SVC ■ IMAP4SVC ■ SMTPSVC Type and dimension: string-scalar

Table 13-16 Exchange Protocol agent optional attribute

Optional Attribute	Description
DetailMonitor	Do not modify this attribute. For internal use.

Configuring the agent for Exchange 2003

The Exchange agent for Exchange Server 2003 can be configured using the Exchange Server Agent Configuration Wizard.

Prerequisites for configuring the Exchange agents

Make sure the following prerequisites are met before configuring the Exchange agent:

- The Exchange Server services must be running or in the “Started” state.
- The Protocol Virtual Servers must be running or in “Started” state.

To check the status of the Exchange Server services

- 1 Open the Control Panel, select **Administrative Tools**, and click **Services**.
- 2 In the Services window, scroll down the list of services until you reach the Microsoft Exchange services and verify that their status is "Started".
- 3 If the Exchange services are not in the "Started" state, right-click each service and click **Start**.

To check the status of the Exchange Protocol virtual servers

- 1 Open the Microsoft Exchange System Manager (**Start > Microsoft Exchange > System Manager**).
- 2 In the left panel of the window (treeview), select the appropriate Exchange Server that you plan to configure and expand its subdirectory.
- 3 The Protocols folder contains the folders for each of the virtual servers. Double click each folder and then right click the virtual server inside.
- 4 Click **Start** to start the virtual server, if it is not already running.

Once you have verified that the Exchange Server services and the Protocol virtual servers are enabled you can configure the Exchange agents.

To configure the Exchange agents

- 1 Navigate to the configuration wizard by clicking **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Exchange Server Agent Configuration Wizard**.
- 2 Review the Welcome screen and click **Next**. The wizard checks whether the Exchange Server is online and the version of the Exchange Server.
- 3 On the Detailed Monitoring and Configuring Exchange Services panel, if you select the Detailed Monitoring option, the following Compulsory services are automatically selected and monitored:

■ Compulsory Services

MSExchangeIS	Microsoft Exchange Information Store
MSExchange SA	Microsoft Exchange System Attendant
RESvc	Microsoft Exchange Routing Engine

You can select any one or both or none of the Optional services given below for Detailed Monitoring

■ **Optional Services**

MSEExchangeMTA	Microsoft Exchange Transfer Agent
MSEExchangeMGMT	Microsoft Exchange Management

Click **Next** when done.

- 4 On the Protocol Configuration panel, the four protocol virtual servers — POP3, SMTP, IMAP4, and W3 — are selected by default. If you do not want to configure any of the protocol servers, deselect it. If any of the protocol servers are not running, the corresponding selection will be grayed out and not available for selection.
Click **Next** when done.
- 5 The Failure Action screen lists the Exchange services and protocol servers that you have chosen to monitor.
- 6 Double-click one of the services to open the Recovery Action dialog box. Specify the action to take after the first failure, the second failure, and subsequent failures. The options are:
 - Take no action (default for all three)
 - Restart the Service
 - Run a Script
- 7 If you select the Run a Script option, then you must specify the script path name. You can accept the default value, or click **Browse** to select a different path name.
- 8 Enter any command line parameters in the text box (optional) and click **OK**. You are returned to the Failure Actions panel.
- 9 Specify when to reset the fail count. The default value is after 1 day.
- 10 Repeat [step 6](#) through [step 9](#) for each of the services listed. Click **Next** when done.
- 11 The Summary screen lists the resources you have selected. Select a resource to view its attributes.
In this window, you can rename the resources that you have created. Double click the selected resource and type the new name. If you choose to change the name, make sure that the new name is unique.
Click **Next**.
- 12 Click **Finish** to close the wizard.

Configuring the agent for Exchange 2007 in a Windows virtual machine

This section describes the agent for Exchange 2007 and provides instructions for configuring the agent.

- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- See [“About the VCS agent for Exchange Server 2007”](#) on page 276.
- See [“Configuring the Exchange 2007 agent”](#) on page 278.

About the VCS agent for Exchange Server 2007

The VCS application agent for Microsoft Exchange monitors Exchange services in a VCS cluster.

The VCS application agent for Microsoft Exchange contains the following agent:

- Exchange Service agent—Monitors core Exchange services.

The agent provides high availability for Microsoft Exchange Server 2007 in a VCS cluster.

This agent supports detailed monitoring.

See [“Detail monitoring and agent behavior”](#) on page 282.

Exchange Service agent

The Exchange Service (ExchService2007) agent monitors the status of the following Exchange services:

- Microsoft Exchange AD Topology service (MSEExchangeADTopology):
This service provides Active Directory topology information to the Exchange services. If this service is stopped, most Exchange services are unable to start.
- Microsoft Exchange System Attendant (MSEExchangeSA):
The Exchange component responsible for monitoring, maintenance, and Active Directory lookup services, and ensuring that operations run smoothly.
- Microsoft Exchange Information Store (MSEExchangeIS):
The Exchange storage used to hold messages in users’ mailboxes and in public folders.

- Microsoft Exchange Mail Submission (MSEExchangeMailSubmission):
This service submits messages from the Mailbox Server to the Hub Transport Server.

In addition, you can also configure the agent to monitor the following optional services:

- Microsoft Exchange Mailbox Assistants (MSEExchangeMailboxAssistants):
This service performs background processing of mailboxes in the Exchange store.
- Microsoft Exchange Monitoring (MSEExchangeMonitoring):
This service allows applications to call the Exchange diagnostic cmdlets (pronounced "command-lets").
- Microsoft Exchange Replication Service (MSEExchangeRepl):
This service provides replication functionality for Mailbox Server role databases and is used by Local Continuous Replication (LCR) and Cluster Continuous Replication (CCR).
- Microsoft Exchange Search Indexer (MSEExchangeSearch):
This service performs indexing of mailbox content, which improves the performance of content search.
- Microsoft Exchange Service Host (MSEExchangeServiceHost):
This service provides a host for several Microsoft Exchange services.
- Microsoft Exchange Transport Log Search (MSEExchangeTransportLogSearch):
This service provides remote search capability for Microsoft Exchange Transport log files.
- Microsoft Search (msftesql-Exchange):
This service creates full-text indexes on content and properties of structured and semi-structured data to allow fast linguistic searches on the data.

Each Microsoft Exchange service is configured as a VCS resource of type ExchService2007.

Note: The agent does not support the Active Directory Connector and the Site Replication Service. Do not run these services on the Exchange node.

Agent Operations

- **Monitor**—Determines the state of the configured Exchange service by querying the Service Control Manager (SCM).
The agent monitors and verifies the state of all the databases that are selected for detail monitoring. The agent behavior varies depending on how the attributes are configured.
See [“Detail monitoring and agent behavior”](#) on page 282.

State definitions

- **Online**—Indicates that the configured Exchange service has started.
- **Offline**—Indicates that the configured Exchange service has stopped.
- **Unknown**—Indicates that the agent is unable to determine the state of the configured Exchange service.

Configuring the Exchange 2007 agent

The Exchange agents can be configured using the Exchange Server 2007 Configuration wizard.

Prerequisites for configuring the Exchange agents

Before configuring the Exchange agents, make sure the Exchange Server services are running, or in the “Started” state. If they are not running, you will not successfully configure the Exchange services.

To check the status of the Exchange Server services

- 1 Open the Control Panel, select **Administrative Tools**, and click **Services**.
- 2 In the Services window, scroll down the list of services until you reach the Microsoft Exchange services and verify that their status is “Started”.
- 3 If the Exchange services are not in the “Started” state, right-click each non-started service and click **Start**.

To configure the Exchange agents

- 1 Navigate to the configuration wizard by clicking **Start > Programs > Symantec > Veritas Cluster Server > Configuration Wizards > Exchange Server 2007 Agent Configuration Wizard**.
- 2 Review the Welcome screen and click **Next**.

- 3 On the Detailed Monitoring and Configuring Exchange Services panel, if you select the Detailed Monitoring option, the following Compulsory services are automatically selected and monitored:

- **Compulsory Services**

MSEExchangeIS	Microsoft Exchange Information Store
MSEExchangeSA	Microsoft Exchange System Attendant
MSEExchangeADTopology	Microsoft Exchange Active Directory Topology
MSEMailSubmission	Microsoft Exchange Mail Submission

You can select any of the Optional services given below for Detailed Monitoring.

- **Optional Services.**

MSEMailboxAssistants	Microsoft Exchange Mailbox Assistants
MSEMonitoring	Microsoft Exchange Monitoring
MSEReplication Service	Microsoft Exchange Replication Service
MSESearch Indexer	Microsoft Exchange Search Indexer
MSEServiceHost	Microsoft Exchange Service Host
MSETransportLogSearch	Microsoft Exchange Transport Log Search
msftesql-Exchange	Microsoft Search (Exchange)

- 4 The Detail Monitoring panel shows a list of the available Exchange databases. Select the databases that you wish to configure for detailed monitoring.
 You can specify the detailed monitoring interval. The default value is 5. This means that every fifth cycle is monitored in detail.
 You can also specify whether the service group should be faulted if any of the selected databases are not mounted. By default the option is unchecked. Click **Next** when done.

- 5 The Failure Actions panel lists the Exchange services and protocol servers that you have chosen to monitor. Double-click one of the services to open the Recovery Action dialog box. Specify the action to take after the first failure, the second failure, and subsequent failures. The options are:
 - Take no action (default for all three)
 - Restart the Service
 - Run a Script
- 6 If you select the Run a Script option, then you must specify the script path name. You can accept the default value, or click **Browse** to select a different path name.
- 7 Enter any command line parameters in the text box (optional) and click **OK**. You are returned to the Failure Actions panel.
- 8 Specify when to reset the fail count. The default value is after 1 day.
- 9 The Summary screen lists the resources you have selected. Select a resource and click it to view its attributes.

In this window, you can rename the resources that you have created. Double click the selected resource and type the new name. If you choose to change the name, make sure that the new name is unique.

Click **Next**.
- 10 Click **Finish** to close the wizard.

Exchange Server 2007 agent attributes

Review the following information to familiarize yourself with the required agent attributes for an ExchService2007 resource type. This information will assist you during the agent configuration.

Table 13-17 Exchange Service agent required attributes

Required Attributes	Description
Service	<p>The name of the Exchange service to be monitored. This attribute could take any of the following values:</p> <ul style="list-style-type: none"> ■ MExchangeADTopology ■ MExchangeIS ■ MExchangeMailSubmission ■ MExchangeSA ■ MExchangeMailboxAssistants ■ MExchangeServiceHost ■ MExchangeTransportLogSearch ■ MExchangeSearch ■ msftesql-Exchange ■ MExchangeMonitoring ■ MExchangeRepl <p>Type and dimension: string-scalar</p>
LanmanResName	<p>The name of the Lanman resource on which the ExchService2007 resource depends.</p> <p>Type and dimension: string-scalar</p>

Table 13-18 Exchange Service agent optional attributes

Optional Attribute	Description
DetailMonitor	<p>The interval at which the agent performs detail monitoring on the databases specified in the DBList attribute.</p> <p>The default value 5 indicates that the agent performs detail monitoring on every 5th monitor cycle.</p> <p>Setting this value to 0 disables detail monitoring.</p> <p>Type and dimension: integer-scalar</p>

Table 13-18 Exchange Service agent optional attributes

Optional Attribute	Description
FaultOnMountFailure	<p>This flag is used to control the agent behavior in case of detail monitoring. It is applicable to Exchange databases that are selected for detail monitoring.</p> <p>If this flag is set to true and a database that is set to mount automatically on startup is dismounted, the agent will fault the service group.</p> <p>The default value is 0 (false).</p> <p>Type and dimension: boolean-scalar</p>
DBList	<p>List of databases for which the agent will perform detail monitoring.</p> <p>Type and dimension: string-vector</p>

Detail monitoring and agent behavior

You can configure the VCS agent for Exchange Server 2007 to perform detail monitoring on Exchange databases by specifying the desired databases in the DBList attribute. The frequency at which the agent monitors the database is determined by the Detail Monitor attribute.

If you have selected certain databases but do not want the agent to perform detail monitoring on those databases, you do not have to delete the selected databases from the DBList attribute. You can disable detail monitoring by just setting the value of the Detail Monitor attribute to 0. That way, you do not have to select the databases again.

[Table 13-19](#) describes the agent behavior depending on the state of the databases selected for detail monitoring and the FaultonMountFailure attribute settings.

Table 13-19 Detail monitoring and agent behavior

Exchange database set to mount on startup	Exchange database state	FaultonMountFailure attribute value	Agent state
Yes	Mounted	Does not matter	Online
Yes	Dismounted	1 (True)	Offline (Service group will fault)

Table 13-19 Detail monitoring and agent behavior

Exchange database set to mount on startup	Exchange database state	FaultonMountFailure attribute value	Agent state
		0 (False)	Unknown (Administrative action required)
No	Mounted	Does not matter	Online
No	Dismounted	Does not matter	Unknown (Administrative action required)

You may want to dismount the Exchange databases for performing certain administrative operations. In such cases, to avoid the agent from faulting the service group, you can set the FaultonMountFailure attribute value to 0 (false), and then dismount the database and perform the operations.

Once done, you can again mount the databases, set the FaultonMountFailure attribute to 1 (true) and restore the agent behavior to fault the service group if a database is dismounted.

Configuring WebSphere Application Server in a Windows virtual machine

This section describes how you can configure the GenericService agent to monitor WebSphere Application Server.

- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- See [“About WebSphere Application Server 6.0”](#) on page 284.
- See [“Configuring the GenericService agent to monitor WebSphere Application Server”](#) on page 284.

About WebSphere Application Server 6.0

IBM WebSphere Application Server 6.0 ensures automated performance optimization and centralized management and monitoring for business-critical applications. VCS for VMware 5.1 MP2 offers high availability to IBM WebSphere Application Server 6.0. The IBM WebSphere Application Server services are monitored using the GenericService agent. You can configure the GenericService agent to monitor services by defining a resource for each service to be monitored. You can monitor a service in a user-context by specifying the user name, password, and domain.

Configuring the GenericService agent to monitor WebSphere Application Server

You can use the GenericService Agent configuration wizard to configure the GenericService Agent. The GenericService Agent configures all the services listed in the Services panel (**Control Panel > Administrative Tools > Services**).

To configure the GenericService agent follow these steps:

- 1 From the Start menu, select **All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > GenericService Agent**.
- 2 Review the prerequisites listed in the Welcome screen and then click **Next**.
- 3 On the Services Configuration panel select the following services that control the IBM WAS V6 servers:
 - Cell Manager (IBMWebSphere Application Server V6-websphrCellManagerOne)
 - Application Server Node (IBMWebSphere Application Server V6-websphrNodeOne)

- Node Agent (IBMWebSphere Application Server V6-websphrNodeOne_nodeagent)

Click **Next**.

- 4 The Failure Action screen lists the services that you have chosen to monitor. Double-click one of the services to open the Recovery Action dialog box. Specify the action to take after the first failure, the second failure, and subsequent failures. The options are:
 - Take no action (default for all three)
 - Restart the Service
 - Run a Script
- 5 Specify when to reset the fail count (after 1 day is the default value).
- 6 Specify the script path name. You can accept the default value, or click **Browse** to select a different path name.
- 7 Enter any command line parameters in the text box (optional) and click **OK**. You are returned to the Failure Actions screen.
- 8 Repeat [step 4](#) through [step 7](#) for each of the services listed. Click **Next** when you are done.
- 9 The Summary screen lists the resources you have selected. Select a resource and click it to view its attributes.
In this window, you can rename the resources that you have created. Double click the selected resource and type the new name. If you choose to change the name, make sure that the new name is unique.
Click **Next**.
- 10 Click **Finish** to close the wizard.

Configuring SharePoint Portal Server in a Windows virtual machine

This section describes how to configure the GenericService agent to monitor SharePoint Portal Server in a Windows virtual machine.

- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- See [“About SharePoint Portal Server 2007”](#) on page 286.
- See [“Configuring SharePoint Portal Server”](#) on page 286.

About SharePoint Portal Server 2007

VCS for VMware 5.1 MP2 support for SharePoint Portal Server 2007 is limited to the SQL database server. Failure of the SQL host causes the host to fail over to a standby host. The other SPS services like index, search and Web access continue to run normally.

Configuring SharePoint Portal Server

Symantec recommends that you configure the SharePoint Portal Server and the SQL database server on separate virtual machines.

The following steps will help you setup the configuration:

- 1 Install SQL Server 2005 on a virtual machine. Refer to the application's instructions for installation information. Note that you can install VCS for VMware on virtual machines that already contain applications.
- 2 Install Veritas Virtual Machine Tools on the virtual machine where the application to be monitored is installed. Review the instructions for the installing Veritas Virtual Machine Toolkit.
See [“Installing Veritas Virtual Machine Tools”](#) on page 245.
- 3 Install SharePoint Portal Server 2007 on another virtual machine. Refer to the application's instructions for the installation information.
- 4 Configure SharePoint Portal Server 2007 using its configuration wizard. Specify the virtual machine on which you installed SQL Server in [step 1](#), as the database server.
- 5 Using the SQL Agent Configuration Wizard, configure the SQL database for high availability. Review the procedure to be followed for configuration.
See [“Configuring the agents for SQL Server”](#) on page 260.

Configuring generic services in a Windows virtual machine

This section describes the agent for generic services and provides instructions for configuring the agent.

- See [“Prerequisites for configuring agents in Windows virtual machines”](#) on page 251.
- See [“About the VCS agent for generic services”](#) on page 287.
- See [“Configuring the agent for generic services”](#) on page 289.

About the VCS agent for generic services

The GenericService agent monitors the status of a service. Note that a service is an application type supported by Windows that conforms to the interface rules of the Service Control Manager (SCM).

Services are defined as resources of type GenericService. You can configure the GenericService agent to monitor multiple services by defining a resource for each service to be monitored. You can monitor a service in a user-context by specifying the user name, password, and domain.

The agent can detect when a service is brought down gracefully. When the service is brought down gracefully, the agent does not trigger a resource fault even though the service is down.

Note: The service to be configured using the GenericService agent must have the status as Started and the startup type as Automatic.

See [“Configuring the agent for generic services”](#) on page 289.

Agent functions

- **Monitor**
Retrieves the current state of the configured service. It also verifies the user context, if applicable.

State definitions

- **ONLINE**
Indicates the service being monitored is online.
- **OFFLINE**
Indicates the service being monitored is offline.

- UNKNOWN
 Indicates the service operation is in a pending state, or that the agent could not determine the state of the resource.

GenericService agent attributes

Review the following information to familiarize yourself with the agent attributes for the GenericService resource type. Use this information during the agent's configuration.

Table 13-20 GenericService agent required attribute

Required Attribute	Description
ServiceName	Name of the service to be monitored. The service name can be the Service Display Name or the Service Key Name. Note: This attribute can take localized values. Type and dimension: string-scalar

Table 13-21 GenericService agent optional attributes

Optional Attributes	Description
DelayAfterOffline	Number of seconds the offline routine waits for the service to go offline. Default is 10 seconds. Type and dimension: integer-scalar
DelayAfterOnline	Number of seconds the online routine waits for the service to go online. Default is 10 seconds. Type and dimension: integer-scalar
Domain	The domain name to which the user specified in the UserAccount attribute belongs. Note: This attribute can take localized values. Type and dimension: string-scalar
Password	The password of the user, in whose context, the service would be started. If the UserAccount attribute is empty or contains a built-in service account, this attribute is ignored. Type and dimension: string-scalar

Table 13-21 GenericService agent optional attributes (continued)

Optional Attributes	Description
service_arg	An array of arguments passed to the service. Type and dimension: string-vector
UserAccount	A valid user account in whose context the service will be monitored. Username can be of the form username@domain.com or domain.com\username. If you do not specify a value for this attribute, then the user account of the service in the SCM is ignored. To monitor service under built-in accounts, you must provide explicit values. For example: <ul style="list-style-type: none"> ■ On Windows 2000: UserAccount='LocalSystem' ■ On Windows 2003: UserAccount='LocalSystem', 'Local Service', or 'Network Service'. Domain='NT Authority'. The 'NT Authority' domain is not applicable for the 'LocalSystem' account. <p>Note: This attribute can take localized values.</p> Type and dimension: string-scalar
UseVirtualName	Do not modify this attribute. For internal use only.
LanmanResName	Do not modify this attribute. For internal use only.

Configuring the agent for generic services

You can use the GenericService Agent configuration wizard to configure the GenericService Agent. The GenericService Agent can configure all of the services listed in the Services panel (**Control Panel > Administrative Tools > Services**). However, for the SQL Server, Exchange, and IIS services, you should use the individual configuration wizards provided for detailed configuration and monitoring, as the GenericService Agent configuration wizard only provides minimal configuration and monitoring.

In addition, you must run the GenericService Agent configuration wizard locally.

To configure the GenericService agent

- 1 From the Start menu, select **All Programs > Symantec > Veritas Cluster Server > Configuration Wizards > GenericService Agent**.
- 2 Review the prerequisites listed in the Welcome screen and then click **Next**.
- 3 From the list provided, select the services that you want to configure. Click **Next**. You can click the **Check All** button to select all services or the **Uncheck All** button to deselect all the services.
- 4 The Failure Action screen lists the services that you have chosen to monitor. Double-click one of the services to open the Recovery Action dialog box. Specify the action to take after the first failure, the second failure, and subsequent failures. The options are:
 - Take no action (default for all three)
 - Restart the Service
 - Run a Script
- 5 Specify when to reset the fail count (after 1 day is the default value).
- 6 Specify the script path name. You can accept the default value, or click **Browse** to select a different path name.
- 7 Enter any command line parameters in the text box (optional) and click **OK**. You are returned to the Failure Actions screen.
- 8 Repeat [step 4](#) through [step 7](#) for each of the services listed. Click **Next** when you are done.
- 9 The Summary screen lists the resources you have selected. Select a resource and click it to view its attributes.

In this window, you can rename the resources that you have created. Double click the selected resource and type the new name. If you choose to change the name, make sure that the new name is unique.

Click **Next**.
- 10 Click **Finish** to close the wizard.

Troubleshooting communication between the GuestOSApp and application agents on virtual machines

The GuestOSApp agent that runs in the ESX Server layer must communicate with the application agents that run in virtual machines. The agent monitors the application that runs in the virtual machine and indicates its state to the VCS Agent Management Daemon (vcsagmd) on the virtual machine.

In case of a change in the state of the application, vcsagmd reports it to VMware's guestinfo interface. The GuestOSApp resource monitors the VMware guestinfo interface for any change in the state of the application and reports this to VCS that runs on the ESX Server layer. VCS on the ESX Server layer takes the appropriate action depending on the state change reported.

If state changes are not reflected, make sure the resource name used for the application inside the virtual machine does not have a 0 followed by another number.

You can test VMware's guestinfo interface. From the virtual machine, perform the following command to set the value of the mytestvariable1 guestinfo variable:

```
[C:\program files\vmware\vmware tools\] vmwareservice.exe --cmd "info-set guestinfo.mytestvariable1 myvalue1"
```

On the ESX Server where the virtual machine is online, perform the following command to see if the ESX Server can read the mytestvariable1 variable:

```
vmware-cmd path_to_vmx_config_file getguestinfo mytestvariable1
```

To obtain the state of the application, perform the following command on the ESX Server where the virtual machine is online :

```
/opt/VRTSvcs/bin/hares -action GuestOSApp-resource getappstate \
-sys sysA
```

This provides the following information in the given order:

- The state of the resource as reported by the agent running inside the virtual machine
- The contents of the vcsagmd heartbeat file, which tells you whether the vcsagmd heartbeats are functioning properly or not
- The current vcsagmd heartbeat value

If vcsagmd heartbeats are not updated, check the virtual machine to see if vcsagmd was stopped for maintenance and never restarted. Also, check if the Veritas Virtual Machine Tools package was upgraded to the current version.

Check the vcsagmd log file at the following locations to verify if the agents inside the virtual machine report the state to vcsagmd.

- C:\Program Files\VERITAS\cluster server\log\vcsagmd_A.txt
- For Windows 2003 64bit:
C:\Program Files(x86)\VERITAS\cluster server\log\vcsagmd_A.txt

Configuring application monitoring

After you have installed Veritas Virtual Machine Tools, configure resources within a virtual machine. You use the configuration wizards to configure the resources inside virtual machines.

Before you run the configuration wizards, make sure that you review the attributes for the agent that you want to configure.

- See [“SQL Server agent attributes”](#) on page 253.
- See [“Internet Information Services \(IIS\) agent attributes”](#) on page 265.
- See [“Exchange Server 2003 agent attributes”](#) on page 272.
- See [“Exchange Server 2007 agent attributes”](#) on page 281.
- See [“GenericService agent attributes”](#) on page 288.

Once you have reviewed the agent attributes, use the configuration wizard for your application. Use the following wizards:

- See [“Configuring the agents for SQL Server”](#) on page 260.
- See [“Configuring the agent for Internet Information Services”](#) on page 266.
- See [“Configuring the agent for Exchange 2003”](#) on page 273.
- See [“Configuring the Exchange 2007 agent”](#) on page 278.
- See [“Configuring the GenericService agent to monitor WebSphere Application Server”](#) on page 284.
- See [“Configuring the agent for generic services”](#) on page 289.

Note: Every time you run a VCS configuration wizard on a virtual machine, the process creates a new configuration. To preserve your earlier configuration, you must recreate it when running the wizard.

Verifying the configuration for application monitoring

You can verify your VCS configuration for application monitoring by opening the log file and checking the state of the resource (such as SQL Server, Exchange, or IIS) being logged. The states are:

- **ONLINE**
Indicates the configured site or application pool is available.
- **OFFLINE**
Indicates the configured site or application pool is not available.
- **UNKNOWN**-Indicates the agent could not determine the status of the resource.

Example 2006/09/27 16:01:15 VCS INFO V-16-2-50017
Resource(VMIP) is in UNKNOWN state

Example 2006/09/27 16:01:18 VCS INFO V-16-2-13352
Resource(PRSPOOL) is ONLINE

Applying the configuration and creating the corresponding GuestOSApp resource

After configuring the application in the virtual machine, you must update the VCS configuration on the ESX Server node. For each application configured as a VCS resource in the virtual machine, you must add a resource of type GuestOSApp to the configuration on the ESX server. The GuestOSApp agent then listens for updates on the corresponding application resource configured inside the virtual machine. The agent enables virtual machine failover when an application in the virtual machine faults.

If you use the wizard to configure a resource, VCS adds the corresponding resources of type GuestOSApp to the configuration on the ESX Server. VCS also restarts the VCS Agent Management Daemon (vcsagmd) on the virtual machine. This restart applies the newly created configuration.

If you configure a generic service using the Windows Service Control Manager, an ESX administrator must add a resource of type GuestOSApp to the configuration on the ESX server.

To manually configure the GuestOSApp resource on the ESX server

- 1 On the ESX Server, edit the service group that contains the virtual machine configuration.
- 2 For each service configured using Service Control Manager in the virtual machine, add a resource of type GuestOSApp to the service group.
- 3 Make sure that the name of the GuestOSApp resource uses the following naming convention:
 ServiceName_ESXVirtualMachineResourceName
- 4 If you want the virtual machine to fail over when the application faults, set the Critical attribute of the GuestOSApp resource to 1.

Removing Veritas Virtual Machine Tools from the virtual machine running Windows

This section describes steps for uninstalling Veritas Virtual Machine Tools.

To remove Veritas Virtual Machine Tools

- 1 Open the Windows Control Panel and click **Add or Remove Programs**.
- 2 Select Veritas Virtual Machine Tools and click **Remove**.
- 3 When asked if you are sure you want to remove the program, click **Yes**.
- 4 The installer displays the status of uninstallation. When the removal of the Veritas Virtual Machine Tools are complete, you can close the Add or Remove Programs screen.



Administering VCS for VMware ESX

This section contains the following chapter:

- [Chapter 14, “Administration”](#) on page 297.

Administration

This chapter contains the following topics:

- [Administering a VCS cluster](#)
- [How VCS handles the graceful shutdown of applications outside of VCS control](#)
- [Using VMware features and commands in a VCS environment](#)
- [Increasing allocated storage](#)
- [Preserving the last-known good copy of your configuration](#)
- [Using raw devices for the virtual machine's boot image](#)
- [Performing maintenance on virtual machines and applications in virtual machines](#)
- [Atypical VCS configuration with the virtual machine configuration file on local storage](#)

Administering a VCS cluster

Use the following tools to administer a VCS cluster:

- The command line interface
- Cluster Manager (Java Console)
- Veritas Cluster Server Management Console (web-based)

See the *Veritas Cluster Server User's Guide* for more information about administering clusters.

To launch the Java Console from the Veritas Virtualization Manager

- 1 From the Veritas Virtualization Manager (VVM), right-click the cluster that you want to manage.
- 2 Right-click a cluster icon and select **Launch VCS Console**.

How VCS handles the graceful shutdown of applications outside of VCS control

VCS can now recognize when you bring down certain types of resources gracefully and intentionally. It uses a combination of three different attributes to make sure that the service group does not show a fault. These attributes also determine how VCS reacts to unfreezing of a node or a service group. These attributes follow:

- IntentionalOffline
- ExternalStateChange
- OnlineAtUnfreeze

To configure a resource to recognize that it can be brought offline intentionally, first check to see if the agent supports a graceful shutdown (the IntentionalOffline attribute). If the agent supports a graceful shutdown, toggle the value of its IntentionalOffline attribute to 1. Note that you only have to leave the IntentionalOffline attribute to its default value of 0 to disable this functionality.

When the value of this attribute is true, and if the agent detects that it is brought offline intentionally, it reports the return code from the monitor agent function to the agent and the engine. The engine determines that the resource is offline intentionally, and reports the resource's state as OFFLINE instead of FAULTED.

If the value of the IntentionalOffline attribute is zero and the resource is brought offline intentionally it enters a FAULTED state instead of an OFFLINE state.

For different resource types the method that the agent uses to detect intentional offline of the resource varies.

When a resource is brought down intentionally, you can control how you want its corresponding service group to behave using the `ExternalStateChange` attribute. When the resource goes online outside of VCS control, the value of the `ExternalStateChange` attribute controls the behavior of the resource's service group.

The `ExternalStateChange` attribute performs the following tasks, it can:

- Bring the corresponding service group online when resource goes online outside VCS control when its `OnlineGroup` key is set.
- Take the corresponding service group offline when resource is brought down intentionally when its `OfflineGroup` key is set.
- Not bring any parent resources offline if the resource is brought down intentionally when its `OfflineHold` key is set.

Note: The `OfflineGroup` key and the `OfflineHold` key are mutually exclusive.

The `OnlineAtUnfreeze` attribute specifies how an offline service group reacts after it or a node is unfrozen. This behavior requires that the service group have at least one resource that uses the `IntentionalOffline` attribute. The value of the `ExternalStateChange` attribute for the resource must include the `OnlineGroup` key.

Two use cases exist for the `OnlineAtUnfreeze` attribute. In the first case, the service group is frozen; in the second case the node is frozen. When a resource is brought online outside VCS control and happens to be part of a frozen service group or is brought online on a frozen system, the service group cannot come online on the node, even though the `OnlineGroup` key is set.

For a resource that uses the `IntentionalOffline` attribute, and has `OnlineGroup` specified for one of the keys of the `ExternalStateChange` attribute the engine sets the value of `OnlineAtUnFreeze` to true. When the service group or the node is unfrozen, the engine determines if the value of `OnlineAtUnFreeze` is true for the service group on the node. If the value is true, then the engine brings the service group online on the node.

Additional resource type attributes

The following resource type attributes are in addition to those in the *Veritas Cluster Server User's Guide*. These attributes support the ability to take an application offline gracefully outside of VCS control.

Table 14-1 lists the resource type attributes.

Table 14-1 Resource type attributes

Resource type attributes	Definitions
<p>ExternalStateChange (user-defined)</p> <p>Note: This attribute can be overridden.</p>	<p>Defines how VCS handles service group state when resources are intentionally brought online or taken offline outside of VCS control.</p> <p>The attribute can take the following values:</p> <p>OnlineGroup: If the configured application is started outside of VCS control, VCS brings the corresponding service group online.</p> <p>OfflineGroup: If the configured application is stopped outside of VCS control, VCS takes the corresponding service group offline.</p> <p>OfflineHold: If a configured application is stopped outside of VCS control, VCS sets the state of the corresponding VCS resource as offline. VCS does not take any parent resources or the service group offline.</p> <p>OfflineHold and OfflineGroup are mutually exclusive.</p>
<p>IntentionalOffline (user-defined)</p>	<p>Defines how VCS reacts when a configured application is intentionally stopped outside of VCS control.</p> <p>Add this attribute for agents that support detection of an intentional offline outside of VCS control. Note that the intentional offline feature is available for agents registered as V51 or later.</p> <p>The value 0 instructs the agent to register a fault and initiate the failover of a service group when the supported resource is taken offline outside of VCS control.</p> <p>The value 1 instructs VCS to take the resource offline when the corresponding application is stopped outside of VCS control.</p> <ul style="list-style-type: none"> ■ Type and dimension: boolean-scalar ■ Default: 0

Additional service group attribute

The following resource type attribute is in addition to those in the *Veritas Cluster Server User's Guide*. This attribute supports the ability to take an application offline gracefully outside of VCS control.

Table 14-2 lists the service group attributes.

Table 14-2 Service group attribute

Service group attribute	Definition
OnlineAtUnfreeze (system use only)	When a node or a service group is frozen, the OnlineAtUnfreeze attribute specifies how an offline service group reacts after it or the node is unfrozen. Type and dimension: integer scalar Default: not applicable

Using VMware features and commands in a VCS environment

VCS is compatible with VMware's VMotion, DRS, and maintenance mode. The ESXHost agent, which VCS automatically configures during installation, ensures compatibility with these features.

See the following topics for more information:

- ["Using VMotion in a VCS environment"](#) on page 301
- ["Using maintenance mode in VCS environment"](#) on page 306
- ["Using DRS in a VCS environment"](#) on page 304

Using VMotion in a VCS environment

You can use VMotion in a VCS environment. When you trigger VMotion through the Virtual Infrastructure Client, VCS accommodates the state changes that occur to the virtual machine.

You can also trigger VMotion through VCS using the service group migration feature. You can use all of the standard VCS management clients to perform service group migrations.

Prerequisites for accommodating VMotion triggered through the VirtualCenter Infrastructure

For VMotion to work properly in a VCS environment, make sure that you have met the following prerequisites:

- You have installed and configured VCS on all nodes in the VCS cluster.
- You have installed and configured VMware clustering on the exact same nodes that are in the VCS cluster.
- You have enabled the VMware's VMotion feature.
- You have disabled VMware's HA feature.
- You have configured a VCS service group for the virtual machine to ensure the machine's high availability.

Prerequisites for setting up service group migration

You must meet the prerequisites from the previous section and the following prerequisites for service group migration:

- Ensure that you have set the following attributes for the ESXVirtualMachine resource in the service group for a virtual machine:
 - username
 - password
 - esxhostdomain
 - vmname
 - sslcert

For more information on the above ESXVirtualMachine agent attributes, refer to the *Veritas Cluster Server Bundled Agents Reference Guide*.

For more information on using the Cluster Manager (Java Console) to set attributes for the ESXVirtualMachine resource, see the *Veritas Cluster Server User's Guide*.

- You need to set the sslcert attribute above to point to a keystore file. To create and copy the keystore file, refer to:
 - See [“Preparing keystores”](#) on page 146.
 - See [“Copying the keystore file from the VirtualCenter Server to each of the ESX Sever nodes in the VCS cluster”](#) on page 150.

Verifying if a service group can be migrated

Use the testVCCoconnect action to verify the connectivity from the cluster nodes to the VirtualCenter Server.

Run the testVCCoconnect action on each node in the cluster. You can use any of the standard VCS management clients to invoke an Action function a particular VCS resource. The following example uses a command line interface.

Run the testVCCoconnect on each node of the cluster:

```
# hares -action resname token -sys system
```

The following line is an example:

```
# hares -action evm testVCCoconnect -sys esxNode1
```

Where *evm* is the name of the ESXVirtualMachine resource, *testVCCoconnect* is the name of the token, and *esxNode1* is the name of the node where you want to test the connection from.

Migrating service groups

Perform the following procedure to migrate a service group. Note that the service group must have a resource with the Migratable attribute set.

To migrate a service group from the command line

- ◆ Enter the following command on any VCS node:

```
# hagrp -migrate service_group -to system
```

Where the *service_group* variable is the name of the service group that you want to move, and the *system* variable is the node where you want the service group to move to.

To migrate a service group from the Cluster Manager (Java Console)

- 1 From the Veritas Virtualization Manager (VVM), right-click the cluster that you want to migrate.
- 2 Right-click a cluster icon and select **Launch VCS Console**.
- 3 In the Service Groups tree, right-click the service group that you want to migrate.
- 4 Select bring your mouse down to **Migrate to**, and select the server where you want to migrate the service group.

Restrictions for service group migration

A few restrictions exist for service group migration. These restrictions follow:

- You can only migrate a service group that is completely online. VCS does not support cold migration.
- The `-migrate` option is not supported for migrating parallel service groups, for migrating hybrid service groups across system zones or across clusters.
- The `-migrate` option is not supported if the VCS service group contains more than one `ESXVirtualMachine` resource that can be migrated.
- The ability of VMotion to migrate virtual machines (either using the `hagrp -migrate` command or through the VirtualCenter Server) is not supported when you have multiple virtual machines that are configured in a single service group. VMotion is not supported if you have tiered applications configured in a single VCS service group. Tiered applications that are configured in different service groups are acceptable.
- VCS does not support running the `hagrp -migrate` command to trigger VMotion if you have configured ESX hosts in VMware VirtualCenter using IP addresses instead of fully qualified host names.

Using DRS in a VCS environment

You can use DRS in a VCS environment. See the following topics for more information:

- [“Enforcing compatibility between VCS and VMware features \(DRS and ESX Server maintenance mode\)”](#) on page 67
- [“Prerequisites for DRS in a VCS environment”](#) on page 305
- [“Restrictions for DRS in a VCS environment”](#) on page 305
- [“Preventing DRS from moving a virtual machine to a node that has no VSwitch activity”](#) on page 305
- [“About adding and removing nodes”](#) on page 104

Prerequisites for DRS in a VCS environment

For DRS to properly work in a VCS environment, make sure that you have met the following prerequisites:

- You have installed and configured VCS on all nodes in the VCS cluster.
- You have installed and configured VMware clustering on the exact same nodes that are in the VCS cluster.
- You have enabled the VMware's DRS feature.
- You have disabled VMware's HA feature.

Restrictions for DRS in a VCS environment

A few restrictions exist for DRS in a VCS environment. Adding and removing nodes from a cluster require special consideration in a DRS environment.

Review the following restrictions:

- You must add a node to the VCS cluster before you add the same node to the VMware cluster. This prevents DRS from moving a virtual machine to a host that does not have VCS. If you add a node to VMware cluster with DRS enabled, it needs to be added to the VCS SystemList attribute and vice-versa.
- You must delete a node from VMware cluster before you remove the same node from VCS cluster. This prevents DRS from moving a virtual machine to a host that does not have VCS.
- In order to remove a node from DRS cluster, you need to put it into maintenance mode. You then remove it from VMware DRS cluster. The ESXHost agent detects this state change and internally performs a VCS system `freeze` command with the `evacuate` option. No automatic way exists to delete a node that is deleted from DRS cluster and vice-versa.

Preventing DRS from moving a virtual machine to a node that has no VSwitch activity

You can disable the VMware Distributed Resource Scheduler (DRS) from trying to move a virtual machine to a node that has lost virtual network functionality. When VCS detects that the virtual network has gone down, then the VSwitch resource changes state to FAULTED. The VCS engine then fires a command to the `resstatechange` trigger. This trigger changes the name of the HostPortGroup from `host_port_group_name` to `host_port_group_name_faulted`. If DRS then tries to move a virtual machine to the node, it fails validation, and thus cannot do the DRS move.

When functionality is restored to the virtual network, VCS changes the VSwitch resource state back to ONLINE. The VCS engine then fires the `harestatementchange` trigger that reinstates the previous `HostPortGroup` name.

To configure VCS to prevent DRS from trying to use a faulted VSwitch resource

- 1 On a node where you have installed VCS, open a command prompt and navigate to:

```
# cd /opt/VRTSvcs/bin/sample_triggers/esx-vswitch-resstatechange
```
- 2 Copy the contents of the file.
- 3 From another prompt on the same node, open the trigger file in a text editor:

```
# cd /opt/VRTSvcs/bin/triggers
```
- 4 Merge the contents of the file that you copied in [step 2](#) to the end of the `resstatechange` file.
- 5 Save and close the file.
- 6 Run the `hagrp` command with the following options:

```
# hagrp -modify grp_name TriggerResStateChange 1
```

Where `grp_name` can be `Network-Infra` if you have the group there. Otherwise the `grp_name` is the service group that contains the VSwitch Resource.
- 7 Save the VCS configuration.

Using maintenance mode in VCS environment

The VCS equivalent of maintenance mode is to perform a VCS system `freeze` command with the `evacuate` option (`hasys -freeze -evacuate sysname`). `Evacuate` fails over all the service groups that are online or active on a node to other nodes of the cluster. VCS then prevents service groups from coming online or failing over to a system that is already frozen. The sum of this command is to move service groups off the selected node, and to prevent further service groups from failing back to this node.

You can run the VCS system `freeze` command with the `evacuate` option from all of the standard VCS management clients. For more information on performing a system `freeze` command with the `evacuate` option, refer to the *Veritas Cluster Server User's Guide*.

Maintenance mode notes

Review the following notes about using maintenance mode:

- The ESXHost agent, which VCS automatically creates, exists to maintain parity between a VMware DRS cluster and VCS.
- When you put a VMware node into maintenance mode with the Virtual Infrastructure Client, maintenance mode eliminates that virtual machine from being a virtual machine target under DRS. As soon as the ESXHost agent detects this mode, it performs a VCS system freeze with the evacuate option. Once this freeze and evacuate are successful, the node is no longer a target for VCS failover.
- If you use a VCS system `freeze` command with the `evacuate` option, the ESXHost agent puts the node into VMware maintenance mode.
- If you put a node in VCS system freeze without using the evacuate option, the ESXHost agent attempts to put the node in VMware maintenance mode. If the first attempt times out, the ESXHost agent keeps trying to put the node in maintenance mode in the subsequent monitor cycles. The agent tries to put the node into maintenance mode until no virtual machines run on the node.
Symantec strongly recommends that you always run VCS system freeze with the evacuate option. The freeze command with the evacuate option automatically moves all the service groups from the node. The freeze command with the evacuate option also ensures that no running virtual machines exist on the node, so that the node can be safely brought down into the maintenance mode.
- If you put a node in VCS system `freeze` without evacuating it first, the ESXHost agent puts the node in VMware maintenance mode, if no virtual machines are powered on the node. In short, you must do a manual VCS switch of service groups that contain the virtual machines, to a node in the VCS cluster. If you do not do that, the ESXHost agent does not react to the VCS `freeze`.
- Symantec strongly recommends that you always run the VCS system `freeze` command with the `evacuate` option. This ensures VCS cluster compatibility with maintenance mode.

Increasing allocated storage

You can increase the amount of application datastore storage that you have allocated for use with a virtual machine.

Different operating systems can grow storage on different files systems, for more information on supported file systems:

See [“Supported operating systems”](#) on page 30.

Prerequisites

- The virtual machine must be configured for VCS.
- VMware Tools must be installed in the virtual machine.
- Veritas Virtual Machine Tools must be installed on the virtual machine.
- Existing disk space or file system space must be available to increase the storage, with:
 - Non-replicated disk space for virtual machines with high availability
 - Replicated disk space for virtual machines with disaster recovery
- For Linux file systems:
 - The storage must reside on an LVM logical volume
 - The tools to grow file systems (`ext2online` or `resize_reiserfs`) must be installed in the virtual machine
- For Windows systems, the disks you want to grow must be:
 - Dynamic disks
 - On the NTFS file system

Increasing storage

From a Windows client, start the Veritas Virtualization Manager.

Note: While VVM is open, if a virtual machine gets migrated to a different host than the one it was connected to, the Grow filesystem operation from VVM fails. You need to refresh VVM before you perform the operation.

To increase allocated storage

- 1 From a Windows client, click **Start > Programs > Symantec > Veritas Virtualization Manager**.
- 2 Enter the name of the VirtualCenter Server, your user name and password, and the full path to the SSL certificate. See [“To start Veritas Virtualization Manager”](#) on page 151.
- 3 Right-click the virtual machine where you want to increase storage. Select **Grow filesystem**.
- 4 Review this screen and click the **Next** button.
- 5 Select the amount of storage that you want to add, note that you can use small number such as 0.03.
- 6 Enter the mount point for the storage (for example /apache). Click the **Next** button.
- 7 Click a datastore to select it. If the current LUN does not have space for a new datastore, select an available LUN for a new datastore. Click the **Next** button.
- 8 Review the summary. Click the **Back** button to return and change settings.
- 9 Click the **Finish** button to increase the allocated storage.

Preserving the last-known good copy of your configuration

VCS replication agents provide the option of running a fire drill to test whether your applications can fail over to a remote site in case of a disaster. When running VCS for VMware, you can run a fire drill in the local cluster to take a snapshot of your application data. See the documentation for the replication agent used in your configuration.

After you run a fire drill, you can keep the final snapshot that passed your basic testing as the last-known good copy of your application as a backup. This copy is preserved until you perform a manual resynchronization on the array, or until you perform another fire drill.

Using raw devices for the virtual machine's boot image

VMware provides the ability to create virtual machines that uses a shared raw device for its boot image. The VMware infrastructure leverages its raw device mapping (RDM) feature to provide this functionality. VCS supports this use of raw devices and raw device mapping in certain high-availability configurations with some limitations.

Although VMotion is supported in this configuration, the possibility of data corruption can exist. This can occur when two virtual machines try to access the same shared raw device at the same time on two different physical nodes.

For more information about RDM, refer to the VMware Infrastructure Documentation.

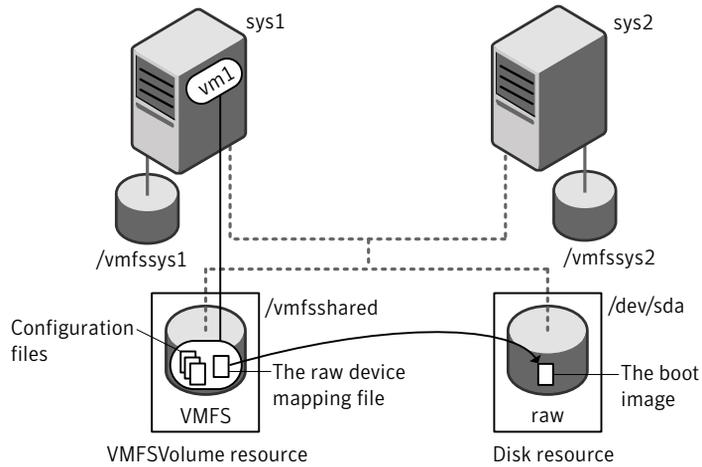
Note: You cannot use raw device mapping (RDM) in a disaster recovery-enabled environment.

Setting up shared raw device storage under VCS

After you have set up a VCS for VMware ESX cluster, you can configure virtual machines to use raw devices. In this configuration, VMotion is allowed. When you set up the virtual machine, you need to configure the shared raw device as a Disk resource in a service group.

[Figure 14-1](#) shows a two-node cluster. Each system has its own local VMFS mounted storage. The first system, `sys1`, has the virtual machine `vm1`. Virtual machine `vm1`'s configuration files, and its raw device mapping (RDM) file are located on `/vmfsshared`. The RDM file points to `/dev/sda` where the virtual machine's boot image resides.

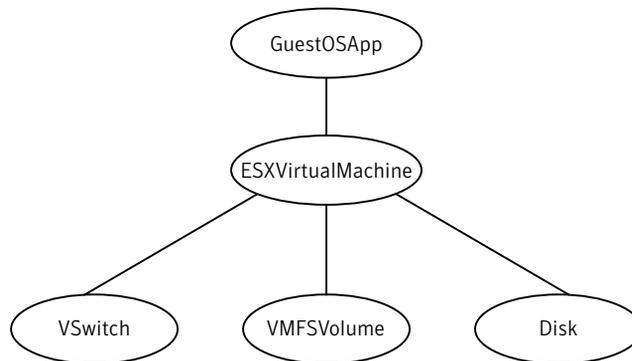
Figure 14-1 A cluster with two systems using RDM for the boot image



The service group for the shared raw device

Figure 14-2 shows the service group that supports the virtual machine. In this service group, the VMFSVolume resource monitors the VMFS volume where you put the virtual machine configuration files and the raw device mapping file. The Disk resource monitors the boot image file that resides on the raw device.

Figure 14-2 A service group configured for RDM



For more information on creating and using service groups refer to the *Veritas Cluster Server User's Guide*.

Performing maintenance on virtual machines and applications in virtual machines

When you need to perform maintenance on a virtual machine, or an application that does not have the ability to perform a graceful shutdown, use following procedures.

Performing maintenance on a virtual machine

You can perform maintenance on a virtual machine by temporarily freezing the service group that contains the ESXVirtualMachine resource.

Perform the following steps from inside the virtual machine that runs these resources.

To perform virtual machine maintenance

- 1 For Linux, open a terminal window, find the home/install directory of the VCS Virtual Machine Tools, and make sure it is in the PATH.

```
# export PATH=$PATH:/opt/VRTSvcs/bin
```
- 2 Find the VCS resource name and the service group name that is associated with this virtual machine.
 - To find the resource name on Windows, type the following command:

```
C:\> type "%VCS_HOME%\vcsvmresname"
```
 - To find the resource name on Linux, type the following command:

```
# cat /etc/VRTSvcs/.vcsvmresname
```
 - To find the service group that the resource resides in, type the following command at the prompt:

```
hares -value vcs_vm_res_name Group
```

 Where `vcs_vm_res_name` is the name of the resource that you found. Record the output of the `hares` command, you need to use this service group name in the following step.
- 3 Freeze the service group. At the prompt, type:

```
# hagr -freeze service_group_name
```

 Where `service_group_name` is the name of the service group that the resource resides in.
- 4 Perform the virtual machine maintenance.
- 5 Unfreeze the service group. At the prompt, type:

```
# hagr -unfreeze service_group_name
```

 Where `service_group_name` is the name of the service group that the resource resides in.

Performing maintenance on applications inside the virtual machine

For VCS agents that have the ability to detect a graceful shutdown, you can directly shut down the application, perform maintenance, and then re-start it. For agents that do not have the ability to detect a graceful shut down (the Mount and Application agents) use the following procedure.

Perform the follow steps from inside the virtual machine that runs the Application or Mount resources.

To bring down the Mount or Application resource for maintenance

- 1 For Linux, open a terminal window, find the home/install directory of the VCS Virtual Machine Tools, and make sure it is in the PATH.

```
# export PATH=$PATH:/opt/VRTSvcs/bin
```
- 2 Find the VCS resource name and the service group name that is associated with this virtual machine.
 - To find the resource name on Windows, type the following command:

```
C:\> type "%VCS_HOME%\vcsvmresname"
```
 - To find the resource name on Linux, type the following command:

```
# cat /etc/VRTSvcs/.vcsvmresname
```
 - To find the service group that the resource resides in, type the following command at the prompt:

```
hares -value vcs_vm_res_name Group
```

Where `vcs_vm_res_name` is the name of the resource that you found.
Record the output of the `hares` command, you need to use this service group name in the following step.
- 3 Freeze the service group. At the prompt, type:

```
# hagr -freeze service_group_name
```

Where `service_group_name` is the name of the service group that the resource resides in.
- 4 Stop the VCS agent management daemon. At the prompt, type:
 - On Windows:

```
C:\> net stop vcsagmd
```
 - On Linux:

```
# /etc/init.d/vcsagmd stop
```
- 5 Perform the application maintenance on the Mount or the Application resource.
- 6 Start the VCS agent management daemon. At the prompt, type:

```
# /etc/init.d/vcsagmd start
```

- 7 Unfreeze the service group. At the prompt, type:

```
# hagr -unfreeze service_group_name
```

Where `service_group_name` is the name of the service group that the resource resides in.

Troubleshooting maintenance

If you find that the commands from the preceding procedures do not work correctly or return errors, reconfigure the Veritas Virtual Machine Tools.

For Linux virtual machines:

See [“Configuring Veritas Virtual Machine Tools”](#) on page 231.

For Windows virtual machines:

See [“Configuring Veritas Virtual Machine Tools”](#) on page 247.

Atypical VCS configuration with the virtual machine configuration file on local storage

VCS supports the following non-standard high availability ESX cluster configuration where:

- An identical virtual machine configuration (*.vmx) file is present on local storage for each node in the cluster.
- The virtual machine swap (*.swp) file is present on local storage for each node in the cluster. You need to add a `sched.swap.dir` entry to the .vmx file. By adding the line, you are specifying a directory on the local storage where the swap file gets created when the virtual machine is powered up. For details refer to the VMware Resource Management Guide.
- All other virtual machine files are present on the shared storage that are accessible to all nodes in the cluster. You need to edit the .vmx file to point the virtual machine's SCSI file name entries to the corresponding full path names of the virtual disks on shared storage.
- You need to make the virtual machine's MAC addresses static. Follow the steps in the VMware Server Configuration Guide.

One caveat is that VMware VMotion does not work in this configuration, which means the VCS migrate interfaces (the `hagr -migrate` command and the GUI equivalent interfaces that trigger VMotion) are not supported

This is an advanced configuration. Contact Symantec support for assistance.

Index

A

- about
 - global clusters 160
 - Veritas Virtualization Manager 144
- accessing service groups 157, 179
- adding
 - users 60
- adding node
 - to a cluster 103
 - to a one-node cluster 125
- administering
 - tools 298
- agent
 - GenericService 287
 - SQL Server 2000 service 252
 - SQL Server 2005 Agent service 252
 - SQL Server 2005 Analysis service 252
 - SQL Server 2005 Search service 252
 - SQL Server 2005 service 252
- agent functions
 - Application agent 221
 - GenericService 287
 - Mount agent 224
 - VCS agent for Apache Web server 200
 - VCS agent for Oracle 190
 - VCS agent for SAP NetWeaver 207
 - VCS agent for WebLogic Server 212
- agent operations
 - ExchProtocol agent 271
 - ExchService agent 270
 - ExchService2007 agent 278
- agent state definition
 - ExchService2007 agent 278
- agents
 - Application 220
 - DNS 162
 - Exchange agents, agent
 - Exchange 269, 276
 - GenericService 287
 - Mount 224
 - VCS agent for Apache Web server 199
 - VCS agent for Oracle 190
 - VCS agent for SAP NetWeaver 205
 - VCS agent for WebLogic Server 212
- allocated storage
 - prerequisites for increasing 308
- Apache agent attributes
 - ConfigFile 200
 - DirectiveAfter 202
 - DirectiveBefore 202
 - EnableSSL 202
 - EnvFile 203
 - HostName 203
 - httpdDir 201
 - Port 203
 - ResLogLevel 201
 - SecondLevelMonitor 201
 - SecondLevelTimeout 204
 - SharedObjDir 204
 - User 202
- Application agent
 - about 220
 - agent functions 221
 - state definitions 221
- application agent. See Exchange Server 2007 agent
- application monitoring 227, 242
- attribute definitions
 - GenericService agent 288
 - IIS agent 265
 - MSDTC agent 259
 - MSSearch agent 256
 - SQL Server 2000 agent 254
 - SQL Server 2005 agent 256
 - SQL Server 2005 Agent service agent 258
 - SQL Server 2005 Analysis service agent 259
- attributes
 - for ExchProtocol agent 273
 - for ExchService agent 272
 - for ExchService2007 agent 281
 - Mount agent 225

B

- basic monitoring
 - health check 191
 - process 190
- bundled agents
 - types.cf file 123

C

- cables
 - cross-over Ethernet 105
- cables, crossover Ethernet 41
- centralized cluster management 60
- cluster
 - creating a single-node cluster, installer 118
 - creating a single-node cluster, manual 119
 - four-node configuration 19
 - removing a node from 110
 - verifying 69
 - verifying operation 100
- cluster connector
 - uninstalling, UNIX 84
- cluster management 61
- Cluster Manager
 - installing Java Console 68
- ClusterService group
 - configuring 164
 - sample configuration 165
- command-line options 80
- commands
 - gabconfig 99, 130
 - hastart 110
 - hastatus 100
 - hasys 100
 - lltconfig 88
 - lltstat 97
 - vxlicinst 79, 107, 122
 - vxlicrep 78, 107, 122
- communication channels 21
- communication disk 21
- configuration files
 - main.cf 89
 - types.cf 89, 124
- configuring
 - disaster recovery 159, 169
 - GAB 130
 - GenericService agent 284
 - hardware 34
 - LLT, manual 128

- private network 41
- configuring application monitoring
 - Linux guest operating system 232
 - Windows guest operating system 292
- configuring replication 165
- configuring VCS
 - adding users 60
 - Cluster Connector 60
 - Cluster Management Console 61
 - event notification 62, 63
 - global clusters 65
 - Veritas Cluster Management Console 60
- configuring virtual machines
 - disaster recovery 173
 - disaster recovery, prerequisites for 171
 - for high availability 154
 - secure DNS update 172
- creating
 - global service group 168
 - service groups, disaster recovery 178
 - service groups, high availability 156
- crossover cables 41

D

- DBList attribute
 - ExchService2007 agent 282
- DetailMonitor attribute
 - ExchProtocol agent 273
 - ExchService agent 272
 - ExchService2007 agent 281
- directives, LLT 129
- disaster recovery 159
 - global clusters 160
 - service groups 178
- disk space
 - directories 34
- disk space, required 34
- DNS agent 162
- DRS 304
 - prerequisites 305

E

- encrypting passwords 233
- entry points. See agent operations
- Ethernet controllers 106
- Exchange agent
 - supported protocols 271
 - supported services 269

- Exchange agents
 - about 269
- Exchange Protocol agent 271
- Exchange Server 2007 agent
 - supported services 276
- Exchange Service agent 269, 276
- ExchProtocol agent
 - attributes 273
 - operations 271
 - state definitions 271
- ExchProtocol agent attributes
 - DetailMonitor 273
 - LanmanResName 273
 - Protocol 273
 - VirtualServer 273
- ExchService agent
 - attributes 272
 - operations 270
 - state definitions 270
- ExchService agent attributes
 - DetailMonitor 272
 - LanmanResName 272
 - Service 272
- ExchService2007 agent
 - attributes 281
 - operations 278
 - state definition 278
- ExchService2007 agent attributes
 - DBList 282
 - DetailMonitor 281
 - FaultOnMountFailure 282
 - LanmanResName 281
 - Service 281
- ExternalStateChange attribute 300

F

- FaultOnMountFailure attribute
 - ExchService2007 agent 282
- fibre channel 34
- file system
 - grow 308
- fire drill
 - last-known copy 309

G

- GAB
 - description 21
 - manual configuration 130

- port membership information 99
- verifying 99
- gabconfig command 99, 130
 - a (verifying GAB) 99
 - in gabtab file 89
- gabtab file
 - creating 130
 - verifying after installation 89
- GenericService agent
 - attribute definitions 288
 - configuring 284
 - description 287
 - state definitions 287
- GenericService agent attributes
 - DelayAfterOffline 288
 - DelayAfterOnline 288
 - Domain 288
 - LanmanResName 289
 - Password 288
 - service_arg 289
 - ServiceName 288
 - UserAccount 289
 - UseVirtualName 289
- Global Cluster option 39
- global clusters 39
 - about 160
 - ClusterService group 163, 164
 - framework 162
 - management 161
 - overview 160
 - prerequisites for 162
 - resiliency 162
 - setting up 163
 - wide-area heartbeats 162
- global clusters service groups
 - service groups
 - global clusters 161
- global clusters, cluster set up 162
- global clusters, configuration 65
- global service group
 - creating 168
- growing file system 308
- GuestOSApp agent 236, 293

H

- hardware
 - configuration 20
 - configuring network and storage 34
- hastart 110

- hastatus -summary command 100
- hasys -display command 100
- health check APIs 191
- health check monitoring 191
- heartbeats
 - wide area 162
- high availability
 - service groups 156
- hubs 41
- hubs, independent 105

I

- IIS agent
 - attribute definitions 265
 - configuring 266
 - prerequisites 266
- IIS agent attributes
 - AppPoolMon 266
 - DetailMonitor 266
 - DetailMonitorInterval 266
 - IPResName 265
 - PortNumber 265
 - SiteName 265
 - SiteType 265
- increasing
 - allocated storage 308
- installing
 - required disk space 34
 - Veritas Virtual Machine Tools 245
 - Veritas Virtualization Manager 143
 - VVM 231
- installing applications
 - Linux guest operating system 229
 - Windows guest operating system 245
- installing VCS
 - checking systems 54
 - choosing packages 57
 - licensing 56
 - required information 47
 - starting 55
 - utilities 52
- installing VCS, example 53
- installing Veritas Virtual Machine Tools
 - Linux guest operating system 229
- installvcs
 - options 69
- intentional offline
 - requirement for Oracle agent 190
- IntentionalOffline attribute 300

J

- Java Console
 - installing 68
 - installing on UNIX 68
 - installing on Windows workstation 69

L

- LanmanResName attribute
 - ExchProtocol agent 273
 - ExchService agent 272
 - ExchService2007 agent 281
- last-known good copy 309
- launch
 - Java Console 298
- license keys
 - adding with vxlicinst 79, 107, 122
 - obtaining 46
 - replacing demo key 79
- licenses, information about 78
- licenses, showing information 107, 122
- licensing commands
 - vxlicinst 46
 - vxlicrep 46
 - vxlictest 46
- licensing VCS 56
- links, private network 41, 88
- Linux applications
 - supported software 189
- Linux guest operating system
 - about 188
 - application monitoring 232
 - installing applications 229
 - installing Veritas Virtual Machine Tools 229
 - removing Veritas Virtual Machine Tools 237
 - supported software 189
- Linux guest operating system application
 - monitoring
 - prerequisites 233
- Linux guest operating system applications
 - verifying 236
- LLT
 - description 21
 - directives 129
 - interconnects 43
 - manual configuration 128
 - verifying 97
- LLT directives
 - link 129

- link-lowpri 129
- set-cluster 129
- set-node 129
- lltconfig command 88
- llthosts file, verifying after installation 88
- lltstat command 97
- llttab file, verifying after installation 88

M

- main.cf file 89
 - example 89
- maintenance
 - applications 312
 - applications in VMs 313
 - virtual machine 312
 - virtual machines 312
- maintenance mode 306
- managing
 - global clusters 161
- managing clusters, centrally 60
- MANPATH variable, setting 41
- media speed 43
 - optimizing 43
- membership information 100
- migrating service groups 303
- monitoring
 - applications 227, 242
- monitoring levels 189
- Mount agent
 - agent functions 224
 - attributes 225
 - description 224
- MSDTC agent
 - attributes 259
- MSDTC agent attributes
 - LanmanResName 259
 - LogPath 259
 - MountResName 259
- MSSearch agent
 - attributes 256
- MSSearch agent attribute
 - AppName 256

N

- NetIsnr agent attributes
 - AgentDebug 199
 - AgentDirectory 199
 - Encoding 199

- EnvFile 198
- Home 197
- Listener 197
- LsnrPwd 198
- MonScript 198
- Owner 197
- TnsAdmin 197
- network partition
 - protecting against 20
- Network partitions
 - protecting against 21

O

- OnlineAtUnfreeze attribute 301
- operations
 - ExchProtocol agent 271
 - ExchService agent 270
 - ExchService2007 agent 278
- optimizing
 - media speed 43
- Oracle
 - error handling 191
- Oracle agent attributes
 - AgentDebug 196
 - AgentDirectory 196
 - AutoEndBkup 194
 - DetailMonitor 195
 - Encoding 196
 - EnvFile 194
 - Home 193
 - MonitorOption 194
 - MonScript 195
 - Owner 193
 - Pfile 194
 - Pword 195
 - ShutDownOpt 193
 - Sid 193
 - StartUpOpt 193
 - Table 196
 - User 195
- overview
 - Veritas Virtualization Manager 144

P

- passwords, encrypting 233
- PATH variable
 - setting 41, 120
- VCS commands 97

- performing
 - maintenance 306
- port a
 - membership 99
- port h
 - membership 99
- port membership information 100
- Preparing 37
- Prerequisites
 - Windows guest OS 252
- prerequisites
 - DRS 305
 - Linux guest operating system application
 - monitoring 233
 - service group migration 302
 - VMotion 302
- private network, configuring 41
- process monitoring 190
- Protocol attribute 273

R

- RAM, installation requirement 34
- raw devices 310
- RDM 310
 - using 310
- removing
 - Veritas Virtual Machine Tools 294
 - Veritas Virtualization Manager 152
- removing a system from a cluster 110
- removing Veritas Virtual Machine Tools
 - Linux guest operating system 237
- replication
 - configuring 165
 - configuring, second cluster 167
 - linking clusters 167
 - reversing 176
- replication agents
 - agents
 - replication 162
- replication setup 163
- requirements
 - Ethernet controllers 34
 - fibre channel 34
 - hardware 34
 - RAM Ethernet controllers 34
 - SCSI host bus adapter 34
 - Veritas Virtualization Manager 145
 - VMware components 35
- resources

- GuestOSApp 236, 293
- restrictions
 - DRS 305
 - service group migration 304
- rpm -e command 115
- rsh 43, 56, 70, 82

S

- sample configuration
 - ClusterService group 165
- SAP NetWeaver agent attributes
 - EnvFile 208
 - InstName 208
 - InstType 208
 - MonitorProgram 211
 - ProcMon 209
 - ResLogLevel 209
 - SAPAdmin 209
 - SAPMonHome 210
 - SAPSID 210
 - SecondLevelMonitor 211
 - StartProfile 210
- SCSI host bus adapter 34
- secure DNS update
 - configuration 172
- See 43, 177
- Service attribute 272, 281
- service group migration
 - prerequisites 302
 - restrictions 304
 - verifying 303
- service groups
 - disaster recovery 178
 - high availability 156
 - migration, restrictions for 304
- setting
 - MANPATH variable 41
 - PATH variable 41, 120
- setting up
 - replication 163
- single-node cluster
 - adding a node to 125
- single-system cluster
 - creating 118, 119
 - modifying startup files 122
- SMTP email notification 62
- SMTP notifications 39
- SNMP notifications 39
- SNMP trap notification 63

- SQL Server 2000 agent attributes
 - DetailMonitor 255
 - Domain 255
 - FaultOnDMScryptFailure 255
 - Instance 254
 - IsGuestOS 254
 - LanmanResName 254
 - MountResName 254
 - Password 255
 - SQLDetailMonitorTimeout 255
 - SQLFile 256
 - SQLOfflineTimeout 254
 - SQLOnlineTimeout 254
 - Username 255
 - SQL Server 2005 agent
 - attributes 256
 - SQL Server 2005 agent attributes
 - DetailMonitor 257
 - Domain 258
 - FaultOnDMScryptFailure 257
 - Instance 256
 - IsGuestOS 257
 - LanmanResName 256
 - Password 258
 - SQLDetailMonitorTimeout 257
 - SQLFile 258
 - SQLOfflineTimeout 256
 - SQLOnlineTimeout 256
 - Username 258
 - SQL Server 2005 Agent service agent
 - attributes 258
 - SQL Server 2005 Agent service agent attributes
 - LanmanResName 258
 - SQLServer2005ResName 258
 - SQL Server 2005 Analysis service agent
 - attributes 259
 - SQL Server 2005 Analysis service agent attributes
 - LanmanResName 259
 - SQLServer2005ResName 259
 - SQL Server agent
 - attributes 254
 - SQL Server agents
 - about 252
 - agent for SQL Server 2000 service 252
 - agent for SQL Server 2005 Agent service 252
 - agent for SQL Server 2005 Analysis
 - service 252
 - agent for SQL Server 2005 Search service 252
 - agent for SQL Server 2005 service 252
 - ssh 43, 56, 70
 - start
 - Java Console 298
 - starting installation
 - installvcs program 55
 - starting VCS 66
 - state definition
 - ExchService2007 agent 278
 - state definitions
 - Application agent 221
 - ExchProtocol agent 271
 - ExchService agent 270
 - GenericService agent 287
 - Mount agent 224
 - VCS agent for Apache Web server 200
 - VCS agent for Oracle 192
 - VCS agent for SAP NetWeaver 207
 - VCS agent for WebLogic Server 213
 - storage
 - fully shared vs. distributed 20
 - shared 20
 - storage, increase 308
 - supported protocols 271
 - supported services 269, 276
 - supported software
 - Linux applications 189
 - Linux guest operating system 189
 - system communication using rsh, ssh 43
 - system state attribute value 100
- ## T
- types.cf 123
 - bundled agents 123
 - types.cf file 124
 - included in main.cf 89
- ## U
- uninstalling
 - cluster connector, UNIX 84
 - uninstalling, VCS 82
 - uninstallvcs 82
 - using RDM 310
- ## V
- validating
 - Veritas Virtual Machine Tools 250
 - variables

- MANPATH 41
 - PATH 41, 120
 - VCS
 - command directory path variable 97
 - configuration files
 - main.cf 89
 - types.cf 89
 - example installation 53
 - global clusters 39
 - installation example 53
 - installing 53
 - replicated states on each system 20
 - VCS agent for Apache Web server
 - about 199
 - agent functions 200
 - state definitions 200
 - VCS agent for Oracle
 - about 190
 - agent functions 190
 - detecting intentional offline 190
 - state definitions 192
 - VCS agent for SAP NetWeaver
 - about 205
 - agent functions 207
 - state definitions 207
 - VCS agent for WebLogic Server
 - about 212
 - agent functions 212
 - state definitions 213
 - VCS Agent Management Daemon 293
 - VCS Agent Management Deaemon 236
 - vcsagmd 236, 293
 - vcsencrypt utility 233
 - verifying
 - cluster 69
 - Linux guest operating system applications 236
 - service group migration 303
 - virtual machine failover 179
 - verifying the configuration
 - Windows guest OS 293
 - Veritas Virtual Machine Tools
 - configuring 245
 - installing 245
 - installing and configuring 230
 - removing 237
 - removing from Windows guest operating system 294
 - validating 250
 - Veritas Virtual Machine Tools ISO file
 - mounting 230
 - Veritas Virtualization Manager
 - disaster recovery configuration 169
 - installation 143
 - removing 152
 - requirements 145
 - virtual machine fail over
 - verification 179
 - virtual machines
 - high availability, prerequisites for 154
 - maintaining 312
 - running Linux 188
 - virtual machines, creation 173
 - virtual machines, high availability 154
 - VirtualServer attribute 273
 - VMotion 301
 - prerequisites 302
 - VMware requirements 35
 - vswif1 54
 - VVM
 - disaster recovery 173
 - installing 231
 - vxlicinst 46
 - vxlicinst command 79, 107, 122
 - vxlicrep 46
 - vxlicrep command 78, 107, 122
 - vxlictest 46
- ## W
- WebLogic Server agent attributes
 - AdminServerMaxWait 218
 - AdminUrl 218
 - BEA_HOME 213
 - DomainDir 213
 - DomainName 214
 - ListenAddressPort 214
 - MonitorProgram 219
 - RequireAdminServer 220
 - ResLogLevel 214
 - SecondLevelMonitor 219
 - ServerName 215
 - ServerRole 215
 - ServerStartProgram 217
 - ServerStopProgram 217
 - User 216
 - WL_HOME 215
 - WLSPassword 216
 - WLSUser 215
 - Windows guest operating system

- application monitoring 292
- installing applications 245
- installing Veritas Virtual Machine Tools 245
- prerequisites 252
- removing Veritas Virtual Machine Tools
from 294
- verifying the configuration 293

