

Veritas™ Cluster Server Installation Guide

Linux

5.0 Maintenance Pack 3



Veritas Cluster Server Installation Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.0 MP3

Document version: 5.0MP3.0

Legal Notice

Copyright © 2008 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, Veritas and Veritas Storage Foundation 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.

THE DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID. 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 as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

<http://www.symantec.com>

Technical Support

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

Symantec's maintenance offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers automatic software upgrade protection
- Global support that is available 24 hours a day, 7 days a week
- Advanced features, including Account Management Services

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

www.symantec.com/techsupp/

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

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

Licensing and registration

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

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

Customer Service is available to assist with the following types of issues:

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

Documentation feedback

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

Maintenance agreement resources

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

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

Additional enterprise services

Symantec offers a comprehensive set of services that allow you to maximize your investment in Symantec products and to develop your knowledge, expertise, and global insight, which enable you to manage your business risks proactively.

Enterprise services that are available include the following:

Symantec Early Warning Solutions	These solutions provide early warning of cyber attacks, comprehensive threat analysis, and countermeasures to prevent attacks before they occur.
Managed Security Services	These services remove the burden of managing and monitoring security devices and events, ensuring rapid response to real threats.
Consulting Services	Symantec Consulting Services provide on-site technical expertise from Symantec and its trusted partners. Symantec Consulting Services offer a variety of prepackaged and customizable options that include assessment, design, implementation, monitoring, and management capabilities. Each is focused on establishing and maintaining the integrity and availability of your IT resources.
Educational Services	Educational Services provide a full array of technical training, security education, security certification, and awareness communication programs.

To access more information about Enterprise services, please visit our Web site at the following URL:

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4
Chapter 1 Introducing Veritas Cluster Server	13
About Veritas Cluster Server	13
About VCS basics	13
About multiple nodes	14
About shared storage	14
About LLT and GAB	15
About network channels for heartbeating	15
About preexisting network partitions	16
About VCS seeding	16
About VCS optional features	17
Veritas Installation Assessment Service	17
About VCS notifications	17
About global clusters	17
About I/O fencing	18
About VCS optional components	18
About Symantec Product Authentication Service (AT)	19
About Veritas Cluster Server Management Console	20
About Cluster Manager (Java Console)	20
Chapter 2 Planning to install VCS	21
About planning to install VCS	21
Hardware requirements	21
Required disk space	22
Supported operating systems	23
Supported software	25
Chapter 3 Preparing to install VCS	27
About preparing to install VCS	27
Preparing to configure the clusters in secure mode	27
Installing the root broker for the security infrastructure	31
Creating authentication broker accounts on root broker system	32

Creating encrypted files for the security infrastructure	33
Preparing the installation system for the security infrastructure	35
Performing preinstallation tasks	36
Obtaining VCS license keys	37
Setting up the private network	38
Configuring SuSE network interfaces	39
Setting up inter-system communication	41
Setting up shared storage	43
Setting the PATH variable	46
Setting the MANPATH variable	47
Optimizing LLT media speed settings on private NICs	47
Guidelines for setting the media speed of the LLT interconnects	47
Mounting the product disc	48
Performing automated pre-installation check	48

Chapter 4	Installing and configuring VCS	51
	About installing and configuring VCS	51
	Getting your VCS installation and configuration information ready	52
	Optional VCS RPMs	55
	About the VCS installation program	55
	Optional features of the installvcs program	55
	Interacting with the installvcs program	56
	About installvcs program command options	57
	Installing and configuring VCS 5.0MP3	61
	Overview of tasks	62
	Starting the software installation	63
	Specifying systems for installation	64
	Licensing VCS	65
	Choosing VCS RPMs for installation	65
	Choosing to install VCS RPMs or configure VCS	66
	Starting the software configuration	67
	Specifying systems for configuration	68
	Configuring the basic cluster	68
	Configuring the cluster in secure mode	70
	Adding VCS users	72
	Configuring SMTP email notification	72
	Configuring SNMP trap notification	74
	Configuring global clusters	75
	Installing VCS RPMs	76

	Creating VCS configuration files	77
	Starting VCS	78
	Completing the installation	79
	Enabling LDAP authentication for clusters that run in secure mode	79
	Installing the Java Console	80
	Verifying the cluster after installation	82
	Verifying and updating licenses on the system	82
	Checking licensing information on the system	82
	Updating product licenses using vxlicinst	83
	Accessing the VCS documentation	84
Chapter 5	Configuring VCS clusters for data integrity	85
	About configuring VCS clusters for data integrity	85
	About I/O fencing components	86
	About data disks	86
	About coordination points	86
	About setting up I/O fencing	87
	Preparing to configure I/O fencing	90
	Initializing disks as VxVM disks	90
	Identifying disks to use as coordinator disks	92
	Checking shared disks for I/O fencing	92
	Setting up I/O fencing	95
	Setting up coordinator disk groups	96
	Configuring I/O fencing	96
	Modifying VCS configuration to use I/O fencing	98
	Verifying I/O fencing configuration	99
	Removing permissions for communication	100
Chapter 6	Verifying the VCS installation	101
	About verifying the VCS installation	101
	About the LLT and GAB configuration files	101
	About the VCS configuration file main.cf	103
	Sample main.cf file for VCS clusters	104
	Sample main.cf file for global clusters	106
	Verifying the LLT, GAB, and VCS configuration files	107
	Verifying LLT, GAB, and cluster operation	107
	Verifying LLT	108
	Verifying GAB	110
	Verifying the cluster	111
	Verifying the cluster nodes	112

Chapter 7	Upgrading VCS	115
	About VCS 5.0 MP3 upgrade	115
	VCS supported upgrade paths	115
	Upgrading VCS in secure enterprise environments	119
	About minimal downtime upgrade	119
	Prerequisites for a minimal downtime upgrade	119
	Planning for the minimal downtime upgrade	119
	Minimal downtime upgrade limitations	120
	Minimal downtime upgrade example	120
	About changes to VCS bundled agents	121
	Deprecated agents	121
	New agents	123
	New and modified attributes for VCS 5.0 MP3 agents	123
	Upgrading to VCS 5.0 MP3	129
	Upgrading VCS to version 5.0 MP3	129
	Upgrading the VCS agents	136
	Upgrading the Cluster Manager (Java Console)	136
	Upgrading the VCS Simulator	136
Chapter 8	Adding and removing cluster nodes	137
	About adding and removing nodes	137
	Adding a node to a cluster	137
	Setting up the hardware	138
	Preparing for a manual installation when adding a node	140
	Installing VCS RPMs for a manual installation	140
	Adding a license key	142
	Verifying the existing security setup on the node	143
	Configuring LLT and GAB	145
	Adding the node to the existing cluster	146
	Starting VCS and verifying the cluster	147
	Removing a node from a cluster	148
	Verifying the status of nodes and service groups	148
	Deleting the departing node from VCS configuration	149
	Modifying configuration files on each remaining node	152
	Removing security credentials from the leaving node	152
	Unloading LLT and GAB and removing VCS on the departing node	153
Chapter 9	Installing VCS on a single node	155
	About installing VCS on a single node	155
	Creating a single-node cluster using the installer program	155

	Preparing for a single node installation	156
	Starting the installer for the single node cluster	156
	Creating a single-node cluster manually	157
	Setting the path variable for a manual single node installation	157
	Installing the VCS software manually on a single node	157
	Renaming the LLT and GAB startup files	158
	Modifying the startup files	158
	Verifying single-node operation	158
	Adding a node to a single-node cluster	159
	Setting up a node to join the single-node cluster	160
	Installing and configuring Ethernet cards for private network	160
	Configuring the shared storage	161
	Bringing up the existing node	161
	Installing the VCS software manually when adding a node to a single node cluster	162
	Configuring LLT	163
	Configuring GAB when adding a node to a single node cluster	165
	Starting LLT and GAB	165
	Reconfiguring VCS on the existing node	165
	Verifying configuration on both nodes	166
Chapter 10	Uninstalling VCS	169
	About the uninstallvcs program	169
	Prerequisites for using the uninstallvcs program	169
	Uninstalling VCS 5.0MP3	170
	Removing VCS 5.0MP3 RPMs	170
	Running uninstallvcs from the VCS 5.0MP3 disc	172
Appendix A	Advanced VCS installation topics	173
	Using the UDP layer for LLT	173
	When to use LLT over UDP	173
	Configuring LLT over UDP	173
	Performing automated VCS installations	180
	Syntax in the response file	181
	Example response file	181
	Response file variable definitions	182
	Installing VCS with a response file where ssh or rsh are disabled	187
Index		191

Introducing Veritas Cluster Server

This chapter includes the following topics:

- [About Veritas Cluster Server](#)
- [About VCS basics](#)
- [About VCS optional features](#)
- [About VCS optional components](#)

About Veritas Cluster Server

Veritas™ Cluster Server by Symantec is a high-availability solution for cluster configurations. Veritas Cluster Server (VCS) monitors systems and application services, and restarts services when hardware or software fails.


About VCS basics

A single VCS cluster consists of multiple systems that are connected in various combinations to shared storage devices. When a system is part of a VCS cluster, it is a node. VCS monitors and controls applications running in the cluster on nodes, and restarts applications in response to a variety of hardware or software faults.

Applications can continue to operate with little or no downtime. In some cases, such as NFS, this continuation is transparent to high-level applications and users. In other cases, a user might have to retry an operation, such as a Web server reloading a page.

Figure 1-1 illustrates a typical VCS configuration of four nodes that are connected to shared storage.

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



Client workstations receive service over the public network from applications running on VCS nodes. VCS monitors the nodes and their services. VCS nodes in the cluster communicate over a private network.

About multiple nodes


VCS runs in a replicated state on each node in the cluster. A private network enables the nodes to share identical state information about all resources. The private network also recognizes active nodes, the nodes that join or leaving the cluster, and failed nodes. The private network requires two communication channels to guard against network partitions.

About shared storage

A VCS hardware configuration typically consists of multiple nodes that are connected to shared storage through I/O channels. Shared storage provides multiple systems with an access path to the same data. It also enables VCS to restart applications on alternate nodes when a node fails, which ensures high availability.

VCS nodes can only access physically-attached storage.

Figure 1-2 illustrates the flexibility of VCS shared storage configurations.

Figure 1-2 Two examples of shared storage configurations

About LLT and GAB

VCS uses two components, LLT and GAB, to share data over private networks among systems. These components provide the performance and reliability that VCS requires.

LLT (Low Latency Transport) provides fast, kernel-to-kernel communications, and monitors network connections.

LLT configuration files are as follows:

- `/etc/llthosts`—lists all the nodes in the cluster
- `/etc/llttab` file—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 that is required to maintain a synchronized state among the nodes. It monitors disk communications such as the VCS heartbeat utility. The `/etc/gabtab` file is the GAB configuration file.

See [“About the LLT and GAB configuration files”](#) on page 101.


About network channels for heartbeating

For the VCS private network, two network channels must be available to carry heartbeat information. These network connections also transmit other VCS-related information.

Each Linux cluster configuration requires at least two network channels between the systems. The requirement for two channels protects your cluster against network partitioning. For more information on network partitioning, refer to the *Veritas Cluster Server User's Guide*.

Figure 1-3 illustrates a two-node VCS cluster where the nodes galaxy and nebula have two private network connections.

Figure 1-3 Two Ethernet connections connecting two nodes



About preexisting network partitions

A preexisting network partition refers to a failure in the communication channels that occurs while the systems are down and VCS cannot respond. When the systems start, VCS is vulnerable to network partitioning, regardless of the cause of the failure.

About VCS seeding

To protect your cluster from a preexisting network partition, VCS uses a seed. A seed is a function of GAB that determines whether or not all nodes have joined a cluster. For this determination, GAB requires that you declare the number of nodes in the cluster. Note that only seeded nodes can run VCS.

GAB automatically seeds nodes under the following conditions:

- An unseeded node communicates with a seeded node
- All nodes in the cluster are unseeded but can communicate with each other

When the last system starts and joins the cluster, the cluster seeds and starts VCS on all nodes. You can then bring down and restart nodes in any combination. Seeding remains in effect as long as at least one instance of VCS is running somewhere in the cluster.

Perform a manual seed to run VCS from a cold start when one or more systems of the cluster are unavailable. VCS does not start service groups on a system until it has a seed.

About VCS optional features

You can use the Veritas Installation Assessment Service to assess your setup for VCS installation.

See “[Veritas Installation Assessment Service](#)” on page 17.

To configure the optional features of the VCS components, make sure to install all RPMs when the installation program prompts you. Review the description of the optional features and decide the features that you want to configure with VCS:

VCS notifications See “[About VCS notifications](#)” on page 17.

VCS global clusters See “[About global clusters](#)” on page 17.

I/O fencing See “[About I/O fencing](#)” on page 18.

Veritas Installation Assessment Service

The Veritas Installation Assessment Service (IAS) utility assists you in getting ready for a Veritas Storage Foundation and High Availability Solutions installation or upgrade. The IAS utility allows the preinstallation evaluation of a configuration, to validate it prior to starting an installation or upgrade.

<https://vias.symantec.com/>

About VCS notifications

You can configure both SNMP and SMTP notifications for VCS. Symantec recommends you to configure one of these notifications. You have the following options:

- Configure SNMP trap notification of VCS events using the VCS Notifier component
- Configure SMTP email notification of VCS events using the VCS Notifier component.

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

About global clusters

Global clusters provide the ability to fail over applications between geographically distributed clusters when disaster occurs. You require a separate license to configure global clusters. You must add this license during the installation. The installer only asks about configuring global clusters if you have used the global cluster license.

See Veritas Cluster Server User's Guide.

About I/O fencing

I/O fencing protects the data on shared disks when nodes in a cluster detect a change in the cluster membership that indicates a split brain condition.

See *Veritas Cluster Server User's Guide*.

The fencing operation determines the following:

- The nodes that must retain access to the shared storage
- The nodes that must be ejected from the cluster

This decision prevents possible data corruption. The `installvcs` program installs the VCS I/O fencing driver, `VRTSvxfen`. If you want to protect data on shared disks, you must configure I/O fencing after you install and configure VCS.

I/O fencing technology uses coordination points for arbitration in the event of a network partition.


Note: Symantec recommends that you use I/O fencing to protect your cluster against split-brain situations.

About VCS optional components

You can add the following optional components to VCS:

Symantec Product Authentication Service	See “ About Symantec Product Authentication Service (AT) ” on page 19.
Veritas Cluster Server Management Console	See “ About Veritas Cluster Server Management Console ” on page 20.
Cluster Manager (Java console)	See “ About Cluster Manager (Java Console) ” on page 20.

[Figure 1-4](#) illustrates a sample VCS deployment with the optional components configured.

Figure 1-4 Typical VCS setup with optional components

About Symantec Product Authentication Service (AT)

VCS uses Symantec Product Authentication Service (AT) to provide secure communication between cluster nodes and clients. It uses digital certificates for authentication and SSL to encrypt communication over the public network to secure communications.

AT uses the following brokers to establish trust relationship between the cluster components:

- **Root broker**

A root broker serves as the main registration and certification authority; it has a self-signed certificate and can authenticate other brokers. The root broker is only used during initial creation of an authentication broker.

A root broker can serve multiple clusters. Symantec recommends that you install a single root broker on a utility system. The utility system, such as an email server or domain controller, can be highly available.

- **Authentication brokers**

Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have root-signed certificates. Each node in VCS serves as an authentication broker.

See Symantec Product Authentication Service documentation for more information.

See [“Preparing to configure the clusters in secure mode”](#) on page 27.

About Veritas Cluster Server Management Console

Veritas Cluster Server Management Console is a high availability management solution that enables monitoring and administering clusters from a single Web console.

You can configure Veritas Cluster Server Management Console to manage a single cluster, multiple clusters, or both.

See *Veritas Cluster Server Management Console Implementation Guide*.

About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers complete administration capabilities for your cluster. Use the different views in the Java Console to monitor clusters and VCS objects, including service groups, systems, resources, and resource types. You can perform many administrative operations using the Java Console. You can also perform these operations using the command line interface or using the Veritas Cluster Server Management Console.

See [“Installing the Java Console”](#) on page 80.

See *Veritas Cluster Server User's Guide*.

Planning to install VCS

This chapter includes the following topics:

- [About planning to install VCS](#)
- [Hardware requirements](#)
- [Supported operating systems](#)
- [Supported software](#)

About planning to install VCS

Every node where you want to install VCS must meet the hardware and software requirements.

For the latest information on updates, patches, and software issues, read the following Veritas Technical Support TechNote:

<http://entsupport.symantec.com/docs/281993>

To find information on supported hardware, see the hardware compatibility list (HCL) in the following TechNote:

<http://support.veritas.com/docs/283282>

Hardware requirements

[Table 2-1](#) lists the hardware requirements for a VCS cluster.

Table 2-1 Hardware requirements for a VCS cluster

Item	Description
VCS nodes	From 1 to 32 Linux systems that run the supported Linux 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 the applications that migrate between systems in the cluster. The VCS I/O fencing feature requires that all data and coordinator disks support SCSI-3 Persistent Reservations (PR). See “About setting up I/O fencing” on page 87.
Disk space	See “Required disk space” on page 22. Note: VCS may require more temporary disk space during installation than the specified disk space.
Network Interface Cards (NICs)	In addition to the built-in public NIC, VCS requires at least one more NIC per system. Symantec recommends two additional NICs. You can also configure aggregated interfaces.
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 node requires at least 256 megabytes.

Required disk space

Confirm that your system has enough free disk space to install VCS.

[Table 2-2](#) shows the approximate disk space usage by directory for the Veritas Cluster Server RPMs.

Table 2-2 Disk space requirements and totals

Packages	/	/opt	/usr	/var	Totals
Required	3 MB	271 MB	8 MB	1 MB	283 MB
Optional	1 MB	52 MB	0 MB	7 MB	60 MB
Required and optional total	4 MB	323 MB	8 MB	8 MB	343 MB

Note: If you do not have enough free space in /var, then use the `installvcs` command with `tmppath` option. Make sure that the specified `tmppath` file system has the required free space.

Supported operating systems

VCS operates on the Linux operating systems and kernels distributed by Oracle, Red Hat, and SUSE.

[Table 2-3](#) lists the supported operating system versions for Oracle Enterprise Linux (OEL), Red Hat Enterprise Linux (RHEL), and SUSE Linux Enterprise Server (SLES). The table also lists the supported kernel versions and the architecture.

Table 2-3 Supported Linux operating system and kernel versions

Operating System	Kernel	Architecture
OEL based on RHEL 4 Update 4	2.6.9-42.EL	x86 (32-bit) Intel Xeon (32-bit, 64-bit) AMD Opteron (32-bit,64-bit)
OEL based on RHEL 4 Update 5	2.6.9-55.EL	x86 (32-bit) Intel Xeon (32-bit, 64-bit) AMD Opteron (32-bit,64-bit)
OEL based on RHEL 4 Update 6	2.6.9-67.0.0.0.1.ELhugemem	x86 (32-bit) Intel Xeon (32-bit, 64-bit) AMD Opteron (32-bit,64-bit)
OEL based on RHEL 5 Update 1	2.6.18-8.el5	x86 (32-bit) Intel Xeon (32-bit, 64-bit) AMD Opteron (32-bit,64-bit)
RHEL 4 Update 3	2.6.9-34.ELsmp 2.6.9.34.EL 2.6.9.34.ELlargesmp	x86 (32-bit) Intel Xeon (32-bit, 64-bit) AMD Opteron (32-bit,64-bit)
RHEL 4 Update 4	2.6.9-42.ELsmp 2.6.9.42.EL 2.6.9.42.ELlargesmp	x86 (32-bit) Intel Xeon (32-bit, 64-bit) AMD Opteron (32-bit,64-bit)

Table 2-3 Supported Linux operating system and kernel versions (*continued*)

Operating System	Kernel	Architecture
RHEL 4 Update 5	2.6.9-55.ELsmp	x86 (32-bit)
	2.6.9-55.EL	Intel Xeon (32-bit, 64-bit)
	2.6.9-55.ELlargesmp	AMD Opteron (32-bit,64-bit)
RHEL 4 Update 6	2.6.9-67.ELsmp	x86 (32-bit)
	2.6.9-67.EL	Intel Xeon (32-bit, 64-bit)
	2.6.9-67.ELlargesmp	AMD Opteron (32-bit,64-bit)
RHEL 5 Update 1	2.6.18-8.el5	x86 (32-bit)
		Intel Xeon (32-bit, 64-bit)
		AMD Opteron (32-bit,64-bit)
RHEL 5 Update 2	2.6.18-92.el5	x86 (32-bit)
		Intel Xeon (32-bit, 64-bit)
		AMD Opteron (32-bit,64-bit)
SLES 9 with SP3	2.6.5-7.244 EL	x86 (32-bit)
		Intel Xeon (32-bit, 64-bit)
		AMD Opteron (32-bit,64-bit)
SLES 9 with SP4	2.6.5-7.308-default	x86 (32-bit)
	2.6.5-7.308-smp	Intel Xeon (32-bit, 64-bit)
		AMD Opteron (32-bit,64-bit)
SLES 10 with SP1	2.6.16.46-0.12-default	x86 (32-bit)
	2.6.16.46-0.12-smp	Intel Xeon (32-bit, 64-bit)
		AMD Opteron (32-bit,64-bit)
SLES 10 with SP2	2.6.16.60-0.21-default	x86 (32-bit)
	2.6.16.60-0.21-smp	Intel Xeon (32-bit, 64-bit)
		AMD Opteron (32-bit,64-bit)

Note: If your system runs an older version of either Red Hat Enterprise Linux or SUSE Linux Enterprise Server, you must upgrade the operating system before you attempt to install the VCS software. Refer to the Oracle, Red Hat, or SUSE documentation for more information on upgrading your system.

Symantec supports only Oracle, Red Hat, and SUSE distributed kernel binaries.

Symantec products operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported Red Hat erratas and updates and SUSE service packs is available in the following TechNote. Read this TechNote before you install Symantec products.

<http://entsupport.symantec.com/docs/281993>

Supported software

VCS supports the following volume managers and files systems:

- ext2, ext3, reiserfs, NFS, and bind on LVM2, Veritas Volume Manager (VxVM) 4.1 and 5.0, and raw disks.
- Veritas Volume Manager (VxVM) with Veritas File System (VxFS)
 - VxVM 4.1 with VxFS 4.1
 - VxVM 5.0 with VxFS 5.0
(On RHEL and SLES only)
 - VxVM 5.0 MP1 with VxFS 5.0 MP1
(On RHEL and SLES only)
 - VxVM 5.0 MP2 with VxFS 5.0 MP2
 - VxVM 5.0 MP3 with VxFS 5.0 MP3

Note: Veritas Storage Foundation 5.0 supports only 64-bit architecture on Linux. See *Veritas Storage Foundation Release Notes* for more details.

Preparing to install VCS

This chapter includes the following topics:

- [About preparing to install VCS](#)
- [Preparing to configure the clusters in secure mode](#)
- [Performing preinstallation tasks](#)

About preparing to install VCS

Before you perform the preinstallation tasks, make sure you reviewed the installation requirements, set up the basic hardware, and planned your VCS setup.

See [“About planning to install VCS”](#) on page 21.

Preparing to configure the clusters in secure mode

You can set up Symantec Product Authentication Service (AT) for the cluster during the VCS installation or after the installation.

Refer to the *Veritas Cluster Server User's Guide* for instructions to configure AT in a cluster that does not run in secure mode.

The prerequisites to configure a cluster in secure mode are as follows:

- A system in your enterprise is configured as root broker.
If a root broker system does not exist, install and configure root broker on a system.
See [“Installing the root broker for the security infrastructure”](#) on page 31.
- An authentication broker (AB) account for each node in the cluster is set up on the root broker system.
See [“Creating authentication broker accounts on root broker system”](#) on page 32.

- The system clocks of the rook broker and authentication brokers must be in sync.

The `installvcs` program provides the following configuration modes:

Automatic mode	The root broker system must allow rsh or ssh passwordless login to use this mode.
Semi-automatic mode	This mode requires encrypted files (BLOB files) from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See “Setting up inter-system communication” on page 41.
Manual mode	This mode requires <code>root_hash</code> file and the root broker information from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See “Setting up inter-system communication” on page 41.

[Figure 3-1](#) depicts the flow of configuring VCS cluster in secure mode.

Figure 3-1 Workflow to configure VCS cluster in secure mode




Table 3-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform.

Table 3-1 Preparatory tasks to configure a cluster in secure mode

Tasks	Who performs this task
Decide one of the following the configuration modes to set up a cluster in secure mode: ■ Automatic mode ■ Semi-automatic mode ■ Manual mode	VCS administrator
Install the root broker on a stable system in the enterprise. See “Installing the root broker for the security infrastructure” on page 31.	AT administrator
On the root broker system, create authentication broker accounts for each node in the cluster. See “Creating authentication broker accounts on root broker system” on page 32. AT administrator requires the following information from the VCS administrator: ■ Node names that are designated to serve as authentication brokers ■ Password for each authentication broker	AT administrator
To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator. See “Creating encrypted files for the security infrastructure” on page 33. AT administrator requires the following additional information from the VCS administrator: ■ Administrator password for each authentication broker Typically, the password is the same for all nodes.	AT administrator
To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.	AT administrator
Copy the files that are required to configure a cluster in secure mode to the system from where you plan to install and configure VCS. See “Preparing the installation system for the security infrastructure” on page 35.	VCS administrator

Installing the root broker for the security infrastructure

Install the root broker only if you plan to use AT to configure the cluster in secure mode. The root broker administrator must install and configure the root broker before you configure the Authentication Service for VCS. Symantec recommends that you install the root broker on a stable system that is outside the cluster. You can install the root broker on an AIX, HP-UX, Linux, or Solaris system.

See Symantec Product Authentication Service documentation for more information.

See [“About Symantec Product Authentication Service \(AT\)”](#) on page 19.

To install the root broker

- 1 Change to the directory where you can start the installvcs program:

```
# cd cluster_server
```

- 2 Start the Root Broker installation program:

```
# ./installvcs -security
```

- 3 Select to install the Root Broker from the three choices that the installer presents:

```
3 Install Symantec Security Services Root Broker
```

- 4 Enter the name of the system where you want to install the Root Broker.

```
Enter the system name on which to install VxSS: venus
```

- 5 Review the output as the installer does the following:

- Checks to make sure that the VCS supports the operating system
- Verifies that you install from the global zone (only on Solaris)
- Checks if the system is already configured for security

- 6 Review the output as the installvcs program checks for the installed RPMs on the system.

The installvcs program lists the RPMs that the program is about to install on the system. Press Enter to continue.

- 7 Review the output as the installer installs the root broker on the system.
- 8 Enter **y** when the installer prompts you to configure the Symantec Product Authentication Service.

- 9 Enter a password for the root broker. Make sure the password contains a minimum of five characters.
- 10 Enter a password for the authentication broker. Make sure the password contains a minimum of five characters.
- 11 Press the Enter key to start the Authentication Server processes.

```
Do you want to start Symantec Product Authentication Service  
processes now? [y,n,q] y
```

- 12 Review the output as the installer starts the Authentication Service.

Creating authentication broker accounts on root broker system

On the root broker system, the administrator must create an authentication broker (AB) account for each node in the cluster.

To create authentication broker accounts on root broker system

- 1 Determine the root broker domain name. Enter the following command on the root broker system:

```
venus> # vssat showalltrustedcreds
```

For example, the domain name resembles "Domain Name:
root@venus.symantecexample.com" in the output.

- 2 For each node in the cluster, verify whether an account exists on the root broker system.

For example, to verify that an account exists for node galaxy:

```
venus> # vssat showprpl --pdrtype root \  
--domain root@venus.symantecexample.com --prplname galaxy
```

- If the output displays the principal account on root broker for the authentication broker on the node, then delete the existing principal accounts. For example:

```
venus> # vssat deleteprpl --pdrtype root \  
--domain root@venus.symantecexample.com \  
--prplname galaxy --silent
```

- If the output displays the following error, then the account for the given authentication broker is not created on this root broker:

```
"Failed To Get Attributes For Principal"
```

Proceed to step 3.

- 3 Create a principal account for each authentication broker in the cluster. For example:

```
venus> # vssat addprpl --pdrtype root --domain \  
root@venus.symantecexample.com --prplname galaxy \  
--password password --prpltype service
```

You must use this password that you create in the input file for the encrypted file.

Creating encrypted files for the security infrastructure

Create encrypted files (BLOB files) only if you plan to choose the semiautomatic mode that uses an encrypted file to configure the Authentication Service. The administrator must create the encrypted files on the root broker node. The administrator must create encrypted files for each node that is going to be a part of the cluster before you configure the Authentication Service for VCS.

To create encrypted files

- 1 Make a note of the following root broker information. This information is required for the input file for the encrypted file:

hash	The value of the root hash string, which consists of 40 characters. Execute the following command to find this value:
------	---

```
venus> # vssat showbrokerhash
```

root_domain	The value for the domain name of the root broker system. Execute the following command to find this value:
-------------	--

```
venus> # vssat showalltrustedcreds
```

- 2 Make a note of the following authentication broker information for each node. This information is required for the input file for the encrypted file:

identity	The value for the authentication broker identity, which you provided to create authentication broker principal on the root broker system. This is the value for the --prplname option of the <code>addprpl</code> command. See “Creating authentication broker accounts on root broker system” on page 32.
password	The value for the authentication broker password, which you provided to create authentication broker principal on the root broker system. This is the value for the --password option of the <code>addprpl</code> command. See “Creating authentication broker accounts on root broker system” on page 32.
broker_admin_password	The value for the authentication broker password for Administrator account on the node. This password must be at least five characters.

3 For each node in the cluster, create the input file for the encrypted file.

The installer presents the format of the input file for the encrypted file when you proceed to configure the Authentication Service using encrypted file. For example, the input file for authentication broker on galaxy resembles:

```
[setuptrust]
broker=venus.symantecexample.com
hash=758a33dbd6fae751630058ace3dedb54e562fe98
securitylevel=high

[configab]
identity=galaxy
password=password
root_domain=vx:root@venus.symantecexample.com
root_broker=venus.symantecexample.com:2821
broker_admin_password=ab_admin_password
start_broker=false
enable_pbx=false
```

4 Back up these input files that you created for the authentication broker on each node in the cluster.

Note that for security purposes, the command to create the output file for the encrypted file deletes the input file.

- 5 For each node in the cluster, create the output file for the encrypted file from the root broker system using the following command.

```
RootBroker> # vssat createpkg \
--in /path/to/blob/input/file.txt \
--out /path/to/encrypted/blob/file.txt \
--host_ctx AB-hostname
```

For example:

```
venus> # vssat createpkg --in /tmp/galaxy.blob.in \
--out /tmp/galaxy.blob.out --host_ctx galaxy
```

Note that this command creates an encrypted file even if you provide wrong password for "password=" entry. But such an encrypted file with wrong password fails to install on authentication broker node.

- 6 After you complete creating the output files for the encrypted file, you must copy these files to the installer node.

Preparing the installation system for the security infrastructure

The VCS administrator must gather the required information and prepare the installation system to configure a cluster in secure mode.

To prepare the installation system for the security infrastructure

- ◆ Depending on the configuration mode you decided to use, do one of the following:

Automatic mode Do the following:

- Gather the root broker system name from the AT administrator.
- During VCS configuration, choose the configuration option 1 when the installvcs program prompts.

Semi-automatic mode Do the following:

- Copy the encrypted files (BLOB files) to the system from where you plan to install VCS.
Note the path of these files that you copied to the installation system.
- During VCS configuration, choose the configuration option 2 when the installvcs program prompts.

- Manual mode
- Do the following:
- Copy the root_hash file that you fetched to the system from where you plan to install VCS.
Note the path of the root hash file that you copied to the installation system.
 - Gather the root broker information such as name, fully qualified domain name, domain, and port from the AT administrator.
 - Note the principal name and password information for each authentication broker that you provided to the AT administrator to create the authentication broker accounts.
 - During VCS configuration, choose the configuration option 3 when the installvcs program prompts.

Performing preinstallation tasks

Table 3-2 lists the tasks you must perform before proceeding to install VCS.

Table 3-2 Preinstallation tasks

Task	Reference
Obtain license keys.	See “Obtaining VCS license keys” on page 37.
Set up the private network.	See “Setting up the private network” on page 38.
Configure SuSE network interfaces	See “Configuring SuSE network interfaces” on page 39.
Enable communication between systems.	See “Setting up inter-system communication” on page 41.
Set up ssh on cluster systems.	See “Setting up ssh on cluster systems” on page 42.
Set up shared storage for I/O fencing (optional)	See “Setting up shared storage” on page 43.
Set the PATH and the MANPATH variables.	See “Setting the PATH variable” on page 46. See “Setting the MANPATH variable” on page 47.
Review basic instructions to optimize LLT media speeds.	See “Optimizing LLT media speed settings on private NICs” on page 47.

Table 3-2 Preinstallation tasks (*continued*)

Task	Reference
Review guidelines to help you set the LLT interconnects.	See “ Guidelines for setting the media speed of the LLT interconnects ” on page 47.
Mount the product disc	See “ Mounting the product disc ” on page 48.
Verify the systems before installation	See “ Performing automated pre-installation check ” on page 48.

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.

To register and receive a software license key, go to the Symantec Licensing Portal at the following location:

<https://licensing.symantec.com>

Make sure you have your Software Product License document. You need information in this document to retrieve and manage license keys for your Symantec product. After you receive the license key, you can install the product.

Click the Help link at this site to access the *License Portal User Guide* and FAQ.

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 the features and their descriptions that are encoded in a license key

You can only install the Symantec software products for which you have purchased a license. The enclosed software discs might include other products for which you have not purchased a license.

Setting up the private network


VCS requires you to set up a private network between the systems that form a cluster. You can use either NICs or aggregated interfaces to set up private network.

You can use network switches instead of hubs.

Refer to the *Veritas Cluster Server User's Guide* to review VCS performance considerations.

Figure 3-2 shows two private networks for use with VCS.

Figure 3-2 Private network setups: two-node and four-node clusters



To set up the private network

- 1 Install the required network interface cards (NICs).
Create aggregated interfaces if you want to use these to set up private network.
- 2 Connect the VCS private NICs on each system.
- 3 Use crossover Ethernet cables, switches, or independent hubs for each VCS communication network. Note that the crossover Ethernet cables are supported only on two systems.

Ensure that you meet the following requirements:

- The power to the switches or hubs must come from separate sources.
- On each system, you must use two independent network cards to provide redundancy.
- The network interface card to set up private interface is not part of any aggregated interface.

During the process of setting up heartbeat connections, consider a case where a failure removes all communications between the systems.

Note that a chance for data corruption exists under the following conditions:

- The systems still run, and
 - The systems are capable to access shared storage.
- 4 Test the network connections. Temporarily assign network addresses and use `telnet` or `ping` to verify communications.

LLT uses its own protocol, and does not use TCP/IP. So, you must ensure that the private network connections are used only for LLT communication and not for TCP/IP traffic. To verify this requirement, unplumb and unconfigure the temporary addresses.

The `installvcs` program configures the private network in the cluster during installation.

See [“About installing and configuring VCS”](#) on page 51.

Configuring SuSE network interfaces

You must perform additional network configuration on SuSE. You need not perform this procedure for the systems that run SLES 10 or later. By default, SLES 10 uses `udev` to achieve persistent interface names. Refer to the OS documentation for information on configuring persistent interfaces on SLES 10.

In rare cases where RedHat does not automatically configure the network interfaces, RedHat users may also have to perform the network configuration.

Review the following tasks that allow VCS to function properly:

- VCS must be able to find the same network interface names across reboots.
- VCS must have network interfaces up before LLT starts to run.

Symantec suggests the following steps for configuring network interfaces on SUSE.

Note: You must not reboot the system between configuring the persistent interface names and configuring the interfaces to be up before starting LLT.

Note: The MAC address in the `ifcfg-eth-id-mac` file can be in uppercase or lowercase. SUSE, and therefore the Veritas product installer, ignores the file with lowercase MAC address if the file with uppercase MAC address is present.

To configure persistent interface names for network devices

- 1 Navigate to the hotplug file in the `/etc/sysconfig` directory:

```
# cd /etc/sysconfig
```

- 2 Open the hotplug file in an editor.
- 3 Set `HOTPLUG_PCI_QUEUE_NIC_EVENTS` to `yes`:

```
HOTPLUG_PCI_QUEUE_NIC_EVENTS=yes
```

- 4 Run the command:

```
ifconfig -a
```

- 5 Make sure that the interface name to MAC address mapping remains same across the reboots.

Symantec recommends adding the `PERSISTENT_NAME` entries to the configuration files for all the network interfaces (including the network interfaces that are not used).

For each ethernet interface displayed, do the following:

- If a file named `/etc/sysconfig/network/ifcfg-eth-id-mac`, where `mac` is the hardware address of that interface, does not exist, then do the following: Create the file.

If a file exists for the same network interface with the name `/etc/sysconfig/network/ifcfg-ethX`, then copy the contents of that file into the newly created file. The variable `ethX` represents the interface name.

- Add the following line at the end of the file `/etc/sysconfig/network/ifcfg-eth-id-mac`.

```
PERSISTENT_NAME=ethX
```

where `ethX` is the interface name.

For example:

```
# ifconfig -a
eth0      Link encap:Ethernet  HWaddr 00:02:B3:DB:38:FE
          inet addr:10.212.99.30  Bcast:10.212.99.255
          Mask:255.255.254.0
          inet6 addr: fe80::202:b3ff:fedb:38fe/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:453500 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8131 errors:0 dropped:0 overruns:0 carrier:0
```

```
collisions:0 txqueuelen:1000
RX bytes:35401016 (33.7 Mb) TX bytes:999899 (976.4 Kb)
Base address:0xdce0 Memory:fcf20000-fcf40000
```

If a file named `etc/sysconfig/network/ifcfg-eth-id-00:02:B3:DB:38:FE` does not exist, do the following task:

- Create the file.
- If the file `/etc/sysconfig/network/ifcfg-eth0` exists, then copy the contents of this file into `etc/sysconfig/network/ifcfg-eth-id-00:02:B3:DB:38:FE`.

Add the following to the end of the file named `etc/sysconfig/network/ifcfg-eth-id-00:02:B3:DB:38:FE`,

```
PERSISTENT_NAME=eth0
```

Perform the procedure for all the interfaces that the `ifconfig -a` command displays.

To configure interfaces to be up before starting LLT

- 1 For each network interface that you want LLT to use, find its MAC address by running the `ifconfig` command:

```
# ifconfig eth0
eth0      Link encap:Ethernet  HWaddr 00:0C:0D:08:C4:32
```

Where `eth0` is the sample network interface name. The output displays `00:0C:0D:08:C4:32` as the interface's MAC address.

- 2 Navigate to the config file in the `/etc/sysconfig/network` directory:

```
# cd /etc/sysconfig/network
```

- 3 Open the config file in an editor.
- 4 Append the string `eth-id-macaddress` to the `MANDATORY_DEVICES` list in the config file. Separate each address with a space, for example:

```
MANDATORY_DEVICES="eth-id-00:0C:0D:08:C4:31
eth-id-00:0C:0D:08:C4:32"
```

Setting up inter-system communication

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 installer uses `ssh`. You must grant root privileges for the system

where you run `installvcs` program. This privilege facilitates to issue `ssh` or `rsh` commands 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. Similarly, `rsh` must be configured in such a way to not prompt for passwords.

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

Warning: The `rsh` and `ssh` commands to the remote systems, where VCS is to be installed, must not print any extraneous characters.

Setting up ssh on cluster systems

Use the Secure Shell (`ssh`) to install VCS on all systems in a cluster from a system outside of the cluster. Before you start the installation process, verify that `ssh` is configured correctly.

Use Secure Shell (`ssh`) to do the following:

- Log on to another system over a network
- Execute commands on a remote system
- Copy files from one system to another

The `ssh` shell provides strong authentication and secure communications over channels. It is intended to replace `rlogin`, `rsh`, and `rcp`.

The Remote Shell (`rsh`) is disabled by default to provide better security. Use `ssh` for remote command execution.

Configuring ssh

The procedure to configure `ssh` uses OpenSSH example file names and commands.

Note: You can configure `ssh` in other ways. Regardless of how `ssh` is configured, complete the last step in the example to verify the configuration.

To configure ssh

- 1 Log on to the system from which you want to install VCS.
- 2 Generate a DSA key pair on this system by running the following command:

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

- 3 Accept the default location of `~/.ssh/id_dsa`.

4 When the command prompts, enter a passphrase and confirm it.

5 Change the permissions of the `.ssh` directory by typing:

```
# chmod 755 ~/.ssh
```

6 The file `~/.ssh/id_dsa.pub` contains a line that begins with `ssh_dss` and ends with the name of the system on which it was created. Copy this line to the `/root/.ssh/authorized_keys2` file on all systems where you plan to install VCS.

If the local system is part of the cluster, make sure to edit the `authorized_keys2` file on that system.

7 Run the following commands on the system where you are installing:

```
# exec /usr/bin/ssh-agent $SHELL
# ssh-add
```

This step is shell-specific and is valid for the duration the shell is alive.

8 When the command prompts, enter your DSA passphrase.

You are ready to install VCS on several systems in one of the following ways:

- Run the `installvcs` program on any one of the systems
- Run the `installvcs` program on an independent system outside the cluster

To avoid running the `ssh-agent` on each shell, run the X-Window system and configure it so that you are not prompted for the passphrase. Refer to the Red Hat documentation for more information.

9 To verify that you can connect to the systems where you plan to install VCS, type:

```
# ssh -x -l root north ls
# ssh -x -l root south ifconfig
```

The commands should execute on the remote system without having to enter a passphrase or password.

Setting up shared storage

The following sections describe how to set up the SCSI and the Fiber Channel devices that the cluster systems share. For VCS I/O fencing, the data disks must support SCSI-3 persistent reservations. You need to configure a coordinator disk group that supports SCSI-3 PR and verify that it works.

See [“About setting up I/O fencing”](#) on page 87.

See also the *Veritas Cluster Server User's Guide* for a description of I/O fencing.

Setting up shared storage: SCSI

Perform the following steps to set up shared storage.

To set up shared storage

- 1 Connect the disk to the first cluster system.
- 2 Power on the disk.
- 3 Connect a terminator to the other port of the disk.
- 4 Boot the system. The disk is detected while the system boots.
- 5 Press CTRL+A to bring up the SCSI BIOS settings for that disk.

Set the following:

- Set Host adapter SCSI ID = 7, or to an appropriate value for your configuration.
- Set Host Adapter BIOS in Advanced Configuration Options to Disabled.

- 6 Format the shared disk and create required partitions on it.

Perform the following:

- Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is /dev/sdc. Identify whether the shared disk is sdc, sdb, and so on.
- Type the following command:

```
# fdisk /dev/shareddiskname
```

For example, if your shared disk is sdc, type:

```
# fdisk /dev/sdc
```

- Create disk groups and volumes using Volume Manager utilities.
- To apply a file system on the volumes, type:

```
# mkfs -t fs-type /dev/vx/dsk/disk-group/volume
```

For example, enter the following command:

```
# mkfs -t vxfs /dev/vx/dsk/dg/vol01
```

Where the name of the disk group is dg, the name of the volume is vol01, and the file system type is vxfs.

- 7 Power off the disk.

- 8 Remove the terminator from the disk and connect the disk to the other cluster system.
- 9 Power on the disk.
- 10 Boot the second system. The system can now detect the disk.
- 11 Press Ctrl+A to bring up the SCSI BIOS settings for the disk.
Set the following:
 - Set Host adapter SCSI ID = 6, or to an appropriate value for your configuration. Note that the SCSI ID should be different from the one configured on the first cluster system.
 - Set Host Adapter BIOS in Advanced Configuration Options to Disabled.
- 12 Verify that you can view the shared disk using the `fdisk` command.

Setting up shared storage: Fiber Channel

Perform the following steps to set up fiber channel.

To set up shared storage for fiber channel

- 1 Connect the fiber channel disk to a cluster system.
- 2 Boot the system and change the settings of the fiber channel. Perform the following tasks for all QLogic adapters in the system:
 - Press Alt+Q to bring up the QLogic adapter settings menu.
 - Choose **Configuration Settings**.
 - Click Enter.
 - Choose **Advanced Adapter Settings**.
 - Click Enter.
 - Set the Enable Target Reset option to **Yes** (the default value).
 - Save the configuration.
 - Reboot the system.
- 3 Verify that the system detects the fiber channel disks properly.
- 4 Create volumes. Format the shared disk and create required partitions on it and perform the following:
 - Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is `/dev/sdc`.
Identify whether the shared disk is `sdc`, `sdb`, and so on.

- Type the following command:

```
# fdisk /dev/shareddiskname
```

For example, if your shared disk is sdc, type:

```
# fdisk /dev/sdc
```

- Create disk groups and volumes using Volume Manager utilities.
- To apply a file system on the volumes, type:

```
# mkfs -t fs-type /dev/vx/dsk/disk-group/volume
```

For example, enter the following command:

```
# mkfs -t vxfs /dev/vx/dsk/dg/vol01
```

Where the name of the disk group is dg, the name of the volume is vol01, and the file system type is vxfs.

- 5 Repeat step 2 and step 3 for all nodes in the clusters that require connections with fiber channel.
- 6 Power off this cluster system.
- 7 Connect the same disks to the next cluster system.
- 8 Turn on the power for the second system.
- 9 Verify that the second system can see the disk names correctly—the disk names should be the same.

See [“Verifying that the nodes have access to the same disk”](#) on page 93.

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
```

Setting the MANPATH variable

Set the MANPATH variable to view the manual pages.

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
```

See incident 82099 on the Red Hat support web site for more information.

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 the switches or the hubs that are used for the LLT interconnections must match that of the interface cards. Incorrect settings can cause poor network performance or even network failure.

Guidelines for setting the media speed of the LLT interconnects

Review the following guidelines for setting the media speed of the LLT interconnects:

■ If you have hubs or switches 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 hubs or switches for LLT interconnects and you do not use the Auto_Negotiation media speed setting, then do the following:
 Set the hub or switch port to the same setting as 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.

Mounting the product disc

You must have superuser (root) privileges to load the VCS software.

To mount the product disc

- 1 Log in as superuser on a system where you want to install VCS.
 The system from which you install VCS need not be part of the cluster. The systems must be in the same subnet.
- 2 Insert the product disc with the VCS software into a drive that is connected to the system.
 The disc is automatically mounted.
- 3 If the disc does not automatically mount, then enter:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

- 4 Navigate to the location of the RPMs.

```
# cd /mnt/cdrom/dist_arch/cluster_server
```

Where *dist* is rhel4, rhel5, sles9, or sles10, and *arch* is i686 or x86_64 for RHEL and i586 or x86_64 for SLES.

Performing automated pre-installation check

Before you begin 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 system1 system2 ...
```

You can also use the Veritas Installation Assessment Service utility for a detailed assessment of your setup.

See [“Veritas Installation Assessment Service”](#) on page 17.

To check the systems

- 1 Navigate to the folder that contains the `installvcs` program.

See [“Mounting the product disc”](#) on page 48.

- 2 Start the pre-installation check:

```
# ./installvcs -precheck galaxy nebula
```

The program proceeds in a noninteractive mode to examine 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 of the check in a log file.

See [“About installvcs program command options”](#) on page 57.

Installing and configuring VCS

This chapter includes the following topics:

- [About installing and configuring VCS](#)
- [Getting your VCS installation and configuration information ready](#)
- [About the VCS installation program](#)
- [Installing and configuring VCS 5.0MP3](#)
- [Verifying and updating licenses on the system](#)
- [Accessing the VCS documentation](#)

About installing and configuring VCS

You can install Veritas Cluster Server on clusters of up to 32 systems. You can install VCS using one of the following:

Veritas product installer	Use the product installer to install multiple Veritas products.
installvcs program	Use this to install just VCS.

The Veritas product installer and the installvcs program use ssh to install by default. Refer to the *Getting Started Guide* for more information.

Getting your VCS installation and configuration information ready

The VCS installation and configuration program prompts you for information about certain VCS components.

When you perform the installation, prepare the following information:

- To install VCS RPMs you need:

The system names where you plan to install VCS

Example: **galaxy, nebula**

The required license keys

Depending on the type of installation, keys include:

- A valid site license key
- A valid demo license key
- A valid license key for VCS global clusters

See [“Obtaining VCS license keys”](#) on page 37.

To decide whether to install:

- the required VCS RPMs
- all the VCS RPMs

Install only the required RPMs if you do not want to configure any optional components or features.

The default option is to install all RPMs.

See [“Optional VCS RPMs”](#) on page 55.

- To configure Veritas Cluster Server you need:

A name for the cluster

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 "_".

Example: **vcs_cluster27**

A unique ID number for the cluster

A number in the range of 0-65535. Within the site that contains the cluster, each cluster must have a unique ID.

Example: **7**

The device names of the NICs that the private networks use among systems

A network interface card that is not part of any aggregated interface, or an aggregated interface.

Do not use the network interface card that is used for the public network, which is typically eth0.

Example: **eth1, eth2**

- To configure VCS clusters in secure mode (optional), you need:

For automatic mode (default)	■ The name of the Root Broker system Example: <code>east</code> See “About Symantec Product Authentication Service (AT)” on page 19. ■ Access to the Root Broker system without use of a password.
For semiautomatic mode using encrypted files	The path for the encrypted files that you get from the Root Broker administrator. See “Creating encrypted files for the security infrastructure” on page 33.
For semiautomatic mode without using encrypted files	■ The fully-qualified hostname (FQDN) of the Root Broker . (e.g. <code>east.symantecexample.com</code>) The given example puts a system in the (DNS) domain <code>symantecexample.com</code> with the unqualified hostname <code>east</code>, which is designated as the Root Broker. ■ The root broker’s security domain (e.g. <code>root@east.symantecexample.com</code>) ■ The root broker’s port (e.g. <code>2821</code>) ■ The path to the local root hash (e.g. <code>/var/tmp/privatedir/root_hash</code>) ■ The authentication broker’s principal name on each cluster node (e.g. <code>galaxy.symantecexample.com</code> and <code>nebula.symantecexample.com</code>)

- To add VCS users, which is not required if you configure your cluster in secure mode, you need:

User names	Example: <code>smith</code>
User passwords	Enter the password at the prompt.
To decide user privileges	Users have three levels of privileges: A=Administrator, O=Operator, or G=Guest. Example: <code>A</code>

- To configure SMTP email notification (optional), you need:

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), you need:

The port number for the SNMP trap daemon	The default port number 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), you need:

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: **eth0**

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 the following packages:

- VRTScmccc – Veritas Cluster Management Console Cluster Connector
- VRTScmcs – Veritas Cluster Management Console for Single Cluster Mode
- VRTScssim – VCS Simulator
- VRTScscm – Veritas Cluster Server Cluster Manager
- VRTSvcsmn – Manual pages for VCS commands

About the VCS installation program

You can access the `installvcs` program from the command line or through the Veritas product installer.

The VCS installation program is interactive and manages the following tasks:

- Licensing VCS
- Installing VCS RPMs on multiple cluster systems
- Configuring VCS, by creating several detailed configuration files on each system
- Starting VCS processes

You can choose to configure different optional features, such as the following:

- SNMP and SMTP notification
- The Symantec Product Authentication Services feature
- The wide area Global Cluster feature

Review the highlights of the information for which `installvcs` program prompts you as you proceed to configure.

See [“About preparing to install VCS”](#) on page 27.

The `uninstallvcs` program, a companion to `installvcs` program, uninstalls VCS RPMs.

See [“About the uninstallvcs program”](#) on page 169.

Optional features of the `installvcs` program

[Table 4-1](#) specifies the optional actions that the `installvcs` program can perform.

Table 4-1 installvcs optional features

Optional action	Reference
Check the systems to verify that they meet the requirements to install VCS.	See “ Performing automated pre-installation check ” on page 48.
Upgrade VCS to version 5.0 MP3 if VCS currently runs on a cluster.	See “ Upgrading to VCS 5.0 MP3 ” on page 129.
Install VCS RPMs without configuring VCS.	See “ Installing VCS using installonly option ” on page 60.
Configure or reconfigure VCS when VCS RPMs are already installed.	See “ Configuring VCS using configure option ” on page 60.
Perform secure installations using the values that are stored in a configuration file.	See “ Installing VCS with a response file where ssh or rsh are disabled ” on page 187.
Perform automated installations using the values that are stored in a configuration file.	See “ Performing automated VCS installations ” on page 180.

Interacting with the installvcs program

As you run the program, you are prompted to answer yes or no questions. A set of responses that resemble **[y, n, q, ?] (y)** typically follow these questions. The response within parentheses is the default, which you can select by pressing the Enter key. Enter the **?** character to get help to answer the prompt. Enter **q** to quit the installation.

Installation of VCS RPMs takes place only after you have confirmed the information. However, you must remove the partially installed VCS files before you run the installvcs program again.

See “[About the uninstallvcs program](#)” on page 169.

During the installation, the installer prompts you to type information. The installer expects your responses to be within a certain range or in a specific format. The installer provides examples. If you are prompted to enter an item from a list, enter your selection exactly as it is shown in the list.

The installer also prompts you to answer a series of questions that are related to a configuration activity. For such questions, you can enter the **b** character to return to the first prompt in the series. When the installer displays a set of information items you have entered, you are prompted to confirm it. If you answer **n**, the program lets you reenter all of the information for the set.

You can install the VCS Java Console on a single system, which is not required to be part of the cluster. Note that the `installvcs` program does not install the VCS Java Console.

See [“Installing the Java Console”](#) on page 80.

About `installvcs` program 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 [ system1 system2... ] [ options ]
```

[Table 4-2](#) lists the `installvcs` command options.

Table 4-2 `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 60.
<code>-enckeyfile</code> <i>encryption_key_file</i>	See the <code>-responsefile</code> and the <code>-encrypt</code> options.
<code>-encrypt</code> <i>password</i>	Encrypt password using the encryption key that is provided with the <code>-enckeyfile</code> option so that the encrypted password can be stored in response files.
<code>-hostfile</code>	Specifies the location of a file that contains the system names for the installer.
<code>-installonly</code>	Install product RPMs on systems without configuring VCS. See “Installing VCS using installonly option” on page 60.
<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.

Table 4-2 installvcs options (*continued*)

Option and Syntax	Description
<code>-logpath log_path</code>	Specifies that <code>log_path</code> , 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.
<code>-nolic</code>	Install product RPMs on systems without licensing or configuration. License-based features or variants are not installed when using this option.
<code>-nooptionalpkgs</code>	Specifies that the optional product RPMs such as man pages and documentation need not be installed.
<code>-nostart</code>	Bypass starting VCS after completing installation and configuration.
<code>-pkgpath pkg_path</code>	Specifies that <code>pkg_path</code> contains all RPMs that the <code>installvcs</code> program is about to install on all systems. The <code>pkg_path</code> is the complete path of a directory, usually NFS mounted.
<code>-precheck</code>	Verify that systems meet the installation requirements before proceeding with VCS installation. Symantec recommends doing a precheck before installing VCS. See “Performing automated pre-installation check” on page 48.
<code>-requiredpkgs</code>	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.

Table 4-2 installvcs options (*continued*)

Option and Syntax	Description
<code>-responsefile</code> <code>response_file</code> <code>[-enckeyfile</code> <code>encryption_key_file]</code>	Perform automated VCS installation using the system and the configuration information that is stored in a specified file instead of prompting for information. The <code>response_file</code> must be a full path name. If not specified, the response file is automatically generated as <code>installerernumber.response</code> where number 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 <code>encryption_key_file</code> name are required with the <code>-responsefile</code> option when the response file contains encrypted passwords. See “Installing VCS with a response file where ssh or rsh are disabled” on page 187. See “Performing automated VCS installations” on page 180.
<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 preconfigured such that <code>rsh</code> commands between systems execute without prompting for passwords or confirmations
<code>-security</code>	Enable or disable Symantec Product Authentication Service in a VCS cluster that is running. Install and configure Root Broker for Symantec Product Authentication Service. See “About Symantec Product Authentication Service (AT)” on page 19.
<code>-serial</code>	Performs the installation, uninstallation, start, and stop operations on the systems in a serial fashion. By default, the installer performs these operations simultaneously on all the systems.
<code>-timeout</code>	Specifies the timeout value (in seconds) for each command that the installer issues during the installation. The default timeout value is set to 600 seconds.
<code>-tmppath</code> <code>tmp_path</code>	Specifies that <code>tmp_path</code> is the working directory for <code>installvcs</code> program. This path is different from the <code>/var/tmp</code> path. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.

Table 4-2 installvcs options (continued)

Option and Syntax	Description
<code>-verbose</code>	Displays the details when the installer installs the RPMs. By default, the installer displays only a progress bar during the RPMs installation.

Installing VCS using installonly option

In certain situations, users may choose to install the VCS RPMs on a system before they are ready for cluster configuration. During such situations, the `installvcs -installonly` option can be used. The installation program licenses and installs VCS RPMs on the systems that you enter 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. You can configure VCS when you are ready for cluster configuration. The `installvcs` program prompts for cluster information, and creates VCS configuration files without performing installation.

See [“Configuring the basic cluster”](#) on page 68.

The `-configure` option can be used to reconfigure a VCS cluster. VCS must not be running on systems when this reconfiguration is performed.

If you manually edited the `main.cf` file, you need to reformat the `main.cf` file.

See [“Reformatting VCS configuration files on a stopped cluster”](#) on page 60.

Reformatting VCS configuration files on a stopped cluster

When you manually edit VCS configuration files (for example, the `main.cf` or `types.cf` file) you can potentially create formatting issues that may cause the installer to interpret the cluster configuration information incorrectly.

If you have manually edited any of the configuration files, you need to perform one of the following before you run the installation program:

- On a running cluster, perform an `haconf -dump` command. This command saves the configuration files and ensures that they do not have formatting errors before you run the installer.
- On cluster that is not running, perform the `haconf -cftocmd` and then the `haconf -cmdtoconf` commands to format the configuration files.

Note: Remember to make back up copies of the configuration files before you edit them.

You also need to use this procedure if you have manually changed the configuration files before you perform the following actions using the installer:

- Upgrade VCS
- Uninstall VCS

For more information about the `main.cf` and `types.cf` files, refer to the *Veritas Cluster Server User's Guide*.

To display the configuration files in the correct format on a running cluster

- ◆ Run the following commands to display the configuration files in the correct format:

```
# haconf -dump
```

To display the configuration files in the correct format on a stopped cluster

- ◆ Run the following commands to display the configuration files in the correct format:

```
# hacf -cftocmd config
```


```
# hacf -cmdtoctf config
```

Installing and configuring VCS 5.0MP3

The example installation demonstrates how to install VCS on two systems: galaxy and nebula. The example installation chooses to install all VCS RPMs and configures all optional features. For this example, the cluster's name is `vcs_cluster2` and the cluster's ID is 7.

[Figure 4-1](#) illustrates the systems on which you would install and run VCS.

Figure 4-1 An example of a VCS installation on a two-node cluster



Overview of tasks

[Table 4-3](#) lists the installation and the configuration tasks.

Table 4-3 Installation and configuration tasks

Task	Reference
License and install VCS	■ See “Starting the software installation” on page 63. ■ See “Specifying systems for installation” on page 64. ■ See “Licensing VCS” on page 65. ■ See “Choosing VCS RPMs for installation” on page 65. ■ See “Choosing to install VCS RPMs or configure VCS” on page 66. ■ See “Installing VCS RPMs” on page 76.
Configure the cluster and its features	■ See “Starting the software configuration” on page 67. ■ See “Specifying systems for configuration” on page 68. ■ See “Configuring the basic cluster” on page 68. ■ See “Adding VCS users” on page 72. (optional) ■ See “Configuring SMTP email notification” on page 72. (optional) ■ See “Configuring SNMP trap notification” on page 74. (optional) ■ See “Configuring global clusters” on page 75. (optional)
Create configuration files	See “Creating VCS configuration files” on page 77.

Table 4-3 Installation and configuration tasks (*continued*)

Task	Reference
Start VCS and its components	■ See “Starting VCS” on page 78. ■ See “Completing the installation” on page 79.
For clusters that run in secure mode, enable LDAP authentication plug-in if VCS users belong to LDAP domain.	■ See “Enabling LDAP authentication for clusters that run in secure mode” on page 79.
Perform the post-installation tasks	■ See “About configuring VCS clusters for data integrity” on page 85. ■ See “Installing the Java Console” on page 80.
Verify the cluster	See “Verifying the cluster after installation” on page 82.

Starting the software installation

You can install VCS using the Veritas product installer or the `installvcs` program.

Note: The system from where you install VCS must run the same Linux distribution as the target systems.

To install VCS using the product installer

- 1 Confirm that you are logged in as the superuser and mounted the product disc.
- 2 Start the installer.

```
# ./installer
```

The installer starts the product installation program with a copyright message and specifies the directory where the logs are created.

- 3 From the opening Selection Menu, choose: `I` for "Install/Upgrade a Product."
- 4 From the displayed list of products to install, choose: Veritas Cluster Server.

To install VCS using the `installvcs` program

- 1 Confirm that you are logged in as the superuser and mounted the product disc.

- 2 Navigate to the folder that contains the `installvcs` program.

```
# cd /cluster_server
```

- 3 Start the `installvcs` program.

```
# ./installvcs
```

The installer begins with a copyright message and specifies the directory where the logs are created.

Specifying systems for installation

The installer prompts for the system names on which you want to install and then performs an initial system check.

To specify system names for installation

- 1 Enter the names of the systems where you want to install VCS.

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

For a single node installation, enter one name for the system.

See [“Creating a single-node cluster using the installer program”](#) on page 155.

- 2 Review the output as the installer verifies the systems you specify.

The installer does the following:

- Checks that the local node running the installer can communicate with remote nodes
If the installer finds `ssh` binaries, it confirms that `ssh` can operate without requests for passwords or passphrases.
- Makes sure the systems use the proper operating system
- Checks whether a previous version of VCS is installed
If a previous version of VCS is installed, the installer provides an option to upgrade to VCS 5.0MP3.

Licensing VCS

The installer checks whether VCS license keys are currently in place on each system. If license keys are not installed, the installer prompts you for the license keys.

See [“Checking licensing information on the system”](#) on page 82.

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 installer prompts for each node.

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

- 3 Enter keys for additional product features.

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

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

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

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

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

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

Choosing VCS RPMs for installation

The installer verifies for any previously installed RPMs and then based on your choice installs all the VCS RPMs or only the required RPMs.

To install VCS RPMs

- 1 Review the output as the installer checks the RPMs that are already installed.
- 2 Choose the VCS RPMs that you want to install.

```
Select the RPMs to be installed on all systems? [1-3,q,?]  
(3) 2
```

Based on what RPMs you want to install, enter one of the following:

- 1 Installs only the required VCS RPMs.
 - 2 Installs all the VCS RPMs.

You must choose this option to configure any optional VCS feature. Note that this option is the default if you already installed the SF HA RPMs.
 - 3 Installs all the VCS and the SF HA RPMs. (default option)

If you already installed the SF HA RPMs, the installer does not list this option.
- 3 View the list of RPMs that the installer would install on each node.

If the current version of a RPM is on a system, the installer removes it from the RPM installation list for the system.

Choosing to install VCS RPMs or configure VCS

While you must configure VCS before you can use VCS, you can do one of the following:

- Choose to install and configure VCS now.
See [“Configuring the basic cluster”](#) on page 68.
- Install packages on the systems and leave the cluster configuration steps for later.

To install VCS packages now and configure VCS later

- 1 If you do not want to configure VCS now, enter n at the prompt.

```
Are you ready to configure VCS? [y,n,q] (y) n
```

The utility checks for the required file system space and makes sure that any processes that are running do not conflict with the installation. If requirements for installation are not met, the utility stops and indicates the actions required to proceed with the process.

- 2 Review the output as the installer uninstalls any previous versions and installs the VCS 5.0 MP3 packages.
- 3 Configure the cluster later.

See [“Configuring VCS using configure option”](#) on page 60.

Starting the software configuration

You can configure VCS using the Veritas product installer or the `installvcs` program.

To configure VCS using the product installer

- 1 Confirm that you are logged in as the superuser and mounted the product disc.
- 2 Start the installer.

```
# ./installer
```

The installer starts the product installation program with a copyright message and specifies the directory where the logs are created.

- 3 From the opening Selection Menu, choose: `c` for "Configure an Installed Product."
- 4 From the displayed list of products to configure, choose: Veritas Cluster Server.

To configure VCS using the installvcs program

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

```
# cd /cluster_server
```

- 3 Start the installvcs program.

```
# ./installvcs -configure
```

The installer begins with a copyright message and specifies the directory where the logs are created.

Specifying systems for configuration

The installer prompts for the system names on which you want to configure VCS. The installer performs an initial check on the systems that you specify.

To specify system names for installation

- 1 Enter the names of the systems where you want to configure VCS.

```
Enter the system names separated by spaces on which to configure  
VCS: galaxy nebula
```

- 2 Review the output as the installer verifies the systems you specify.

The installer does the following tasks:

- Checks that the local node running the installer can communicate with remote nodes
If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases.
- Makes sure the systems use the proper operating system
- Checks whether VCS is installed
- Exits if VCS 5.0MP3 is not installed

Configuring the basic cluster

Enter the cluster information when the installer prompts you.

To configure the cluster

- 1 Review the configuration instructions that the installer presents.
- 2 Enter the unique cluster name and cluster ID.

```
Enter the unique cluster name: [?] vcs_cluster2
```

```
Enter the unique Cluster ID number between 0-65535: [b,?] 7
```

- 3 Review the NICs available on the first system as the installer discovers and reports them.

The private heartbeats can either use NIC or aggregated interfaces. To use aggregated interfaces for private heartbeat, enter the name of the aggregated interface. To use a NIC for private heartbeat, enter a NIC which is not part of an aggregated interface.

- 4 Enter the network interface card details for the private heartbeat links.

You must choose the network interface cards or the aggregated interfaces that the installer discovers and reports. If your nodes run SLES, then you must choose the network interface cards which are not part of any aggregated interface instead of the aggregated interface at this time.

Later when the installer prompts to start VCS after product configuration, you must manually edit the `/etc/litab` file for the following cases:

- You want to use aggregated interfaces for nodes that run SLES.
- You want to use aggregated interfaces that the installer has not discovered.

See [“Starting VCS”](#) on page 78.

You must not enter the network interface card that is used for the public network (typically `eth0`.)

```
Enter the NIC for the first private heartbeat NIC on galaxy:
```

```
[b,?] eth1
```

```
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 galaxy:
```

```
[b,?] eth2
```

```
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)
```

- 5 Choose whether to use the same NIC details to configure private heartbeat links on other systems.

Are you using the same NICs for private heartbeat links on all systems? [y,n,q,b,?] (y)

If you want to use the NIC details that you entered for galaxy, make sure the same NICs are available on each system. Then, enter **y** at the prompt.

If the NIC device names are different on some of the systems, enter **n**. Provide the NIC details for each system as the program prompts.

- 6 Verify and confirm the information that the installer summarizes.

Configuring the cluster in secure mode

If you want to configure the cluster in secure mode, make sure that you meet the prerequisites for secure cluster configuration.

The `installvcs` program provides different configuration modes to configure a secure cluster. Make sure that you completed the pre-configuration tasks for the configuration mode that you want to choose.

See [“Preparing to configure the clusters in secure mode”](#) on page 27.

To configure the cluster in secure mode

- 1 Choose whether to configure VCS to use Symantec Product Authentication Service.

Would you like to configure VCS to use Symantec Security Services? [y,n,q] (n) **y**

- If you want to configure the cluster in secure mode, make sure you meet the prerequisites and enter **y**.
- If you do not want to configure the cluster in secure mode, enter **n**.
You must add VCS users when the configuration program prompts.
See [“Adding VCS users”](#) on page 72.

- 2 Select one of the options to enable security.

Select the Security option you would like to perform [1-3,q,?]

Review the following configuration modes. Based on the configuration that you want to use, enter one of the following values:

Option 1. Automatic configuration	Enter the name of the Root Broker system when prompted. Requires a remote access to the Root Broker. Review the output as the installer verifies communication with the Root Broker system, checks vxatd process and version, and checks security domain.
Option 2. Semiautomatic configuration	Enter the path of the encrypted file (BLOB file) for each node when prompted.
Option 3. Manual configuration	Enter the following Root Broker information as the installer prompts you: <pre>Enter root Broker name: east.symantecexample.com Enter root broker FQDN: [b] (symantecexample.com) symantecexample.com Enter root broker domain: [b] (root@east.symantecexample.com) root@east.symantecexample.com Enter root broker port: [b] (2821) 2821 Enter path to the locally accessible root hash [b] (/var/tmp/ installvcs-1Lcljr/root_hash) /root/root_hash</pre> Enter the following Authentication Broker information as the installer prompts you for each node: <pre>Enter authentication broker principal name on galaxy [b] (galaxy.symantecexample.com) galaxy.symantecexample.com Enter authentication broker password on galaxy: Enter authentication broker principal name on nebula [b] (nebula.symantecexample.com) nebula.symantecexample.com Enter authentication broker password on nebula:</pre>

- 3 After you provide the required information to configure the cluster in secure mode, the program prompts you to add VCS users.

Adding VCS users

If you have enabled Symantec Product Authentication Service, you do not need to add VCS users now. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

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.

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the 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 you want to configure SMTP notification? [y,n,q] (y) y
```

If you do not want to configure the SMTP notification, you can skip to the next configuration option.

See [“Configuring SNMP trap notification”](#) on page 74.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- 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,?] ozzie@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,?] harriet@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 you can also configure the 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 you want to configure SNMP notification? [y,n,q] (y)

If you skip this option and if you had installed a valid HA/DR license, the installer presents you with an option to configure this cluster as global cluster. If you did not install an HA/DR license, the installer proceeds to configure VCS based on the configuration details you provided.

See [“Configuring global clusters”](#) on page 75.

- 3 Provide information to configure SNMP trap notification.

Provide the following information:

- 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,?] **saturn**

- 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

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster. If not, the installer proceeds to configure VCS based on the configuration details you provided.

You can configure global clusters to link clusters at separate locations and enable wide-area failover and disaster recovery. The installer adds basic global cluster information to the VCS configuration file. You must perform additional configuration tasks to set up a global cluster. Note that you can also run the `gcoconfig` utility in each cluster later to update the VCS configuration file for global cluster.

See *Veritas Cluster Server User's Guide* for instructions to set up VCS global clusters.

Note: If you installed a HA/DR license to set up replicated data cluster or campus cluster, skip this installer option.

To configure the global cluster option

- 1 Review the required information to configure the global cluster option.
- 2 Specify whether you want to configure the global cluster option.

```
Do you want to configure the Global Cluster Option? [y,n,q] (y)
```

If you skip this option, the installer proceeds to configure VCS based on the configuration details you provided.

- 3 Provide information to configure this cluster as global cluster.

The installer prompts you for a NIC, a virtual IP address, and value for the netmask.

- 4 Verify and confirm the configuration of the global cluster.

```
Global Cluster Option configuration verification:
```

```
NIC: eth0
```

```
IP: 10.10.12.1
```

```
Netmask: 255.255.240.0
```

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

Installing VCS RPMs

After the installer gathers all the configuration information, the installer installs the RPMs on the cluster systems. If you already installed the RPMs and chose to configure or reconfigure the cluster, the installer proceeds to create the configuration files.

See [“Creating VCS configuration files”](#) on page 77.

The utility checks for the required file system space and makes sure that any processes that are running do not conflict with the installation. If requirements for installation are not met, the utility stops and indicates the actions that are required to proceed with the process. Review the output as the installer uninstalls any previous versions and installs the VCS 5.0MP3 RPMs.

Creating VCS configuration files

After you install the RPMs and provide the configuration information, the installer continues to create configuration files and copies them to each system:

```
Creating Cluster Server configuration files ..... Done
Copying configuration files to galaxy..... Done
Copying configuration files to nebula..... Done
Cluster Server configured successfully.
```

If you chose to configure the cluster in secure mode, the installer also configures the Symantec Product Authentication Service.

Depending on the mode you chose to set up Authentication Service, the installer does one of the following:

- Creates the security principal
- Executes the encrypted file to create security principal on each node in the cluster

The installer then does the following before the installer starts VCS in secure mode:

- Creates the VxSS service group
- Creates the Authentication Server credentials on each node in the cluster
- Creates the Web credentials for VCS users
- Sets up trust with the root broker

Verifying the NIC configuration

The installer verifies on all the nodes if all NICs have PERSISTENT_NAME set correctly.

If the persistent interface names are not configured correctly for the network devices, the installer displays the following warnings:

```
Verifying that all NICs have PERSISTENT_NAME set correctly on
galaxy:
```

For VCS to run correctly, the names of the NIC cards must be boot persistent

```
CPI WARNING V-9-122-1021
No PERSISTENT_NAME set for NIC with MAC address
00:11:43:33:17:28 (present name eth0), though config file exists!
CPI WARNING V-9-122-1022
No config file for NIC with MAC address 00:11:43:33:17:29
```

```
(present name eth1) found!  
CPI WARNING V-9-122-1022  
No config file for NIC with MAC address 00:04:23:ac:25:1f  
(present name eth3) found!
```

PERSISTENT_NAME is not set for all the NICs.
You need to set them manually before the next reboot.

Set the PERSISTENT_NAME for all the NICs.

See [“Performing preinstallation tasks”](#) on page 36.

Warning: If the installer finds the network interface name to be different from the name in the configuration file, then the installer exits.

Starting VCS

You can now start VCS and its components on each system. If you chose to configure the cluster in secure mode, the installer also starts the Authentication Service processes on each node in the cluster.

Note: Do not opt to start VCS now if you want to configure private heartbeats to use aggregated interfaces that the installer has not discovered or to use aggregated interfaces on nodes that run SLES..

See [“Configuring the basic cluster”](#) on page 68.

To start VCS

- 1 Confirm to start VCS and its components on each node.

If you want to use aggregated interfaces on the nodes that run SLES or the aggregated interfaces that the installer has not discovered for private heartbeats, enter **n**. Skip to step 2.

```
Do you want to start Veritas Cluster Server processes now?  
[y,n,q] (y) n
```

- 2 Do the following to use aggregated interfaces for private heartbeats:
 - Edit the `/etc/llttab` file to replace the names of NICs with the names of the aggregated interfaces.
 - Reboot the system for the configuration changes to take effect.

Completing the installation

After VCS 5.0 MP3 installation completes successfully, the installer creates summary, log, and response files. The files provide the useful information that can assist you with the installation and can also assist future installations.

Review the location of the installation log files, summary file, and response file that the installer displays.

[Table 4-4](#) specifies the files that are created at the end of the installation.

Table 4-4 File description

File	Description
summary file	■ Lists the RPMs that are installed on each system. ■ Describes the cluster and its configured resources. ■ Provides the information for managing the cluster.
log file	Details the entire installation.
response file	Contains the configuration information that can be used to perform secure or unattended installations on other systems. See “Example response file” on page 181.

Enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as Sun Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

See the *Symantec Product Authentication Service Administrator's Guide*.

The LDAP schema and syntax for LDAP commands (such as, ldapadd, ldapmodify, and ldapsearch) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
 - UserObjectClass (the default is posixAccount)
 - UserObject Attribute (the default is uid)
 - User Group Attribute (the default is gidNumber)
 - Group Object Class (the default is posixGroup)
 - GroupObject Attribute (the default is cn)

- Group GID Attribute (the default is gidNumber)
- Group Membership Attribute (the default is memberId)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)
- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain. To enable LDAP authentication plug-in, you must verify the LDAP environment, add the LDAP domain in AT, and then verify LDAP authentication. The AT component packaged with VCS requires you to manually edit the VRTSlocal.conf file to enable LDAP authentication.

Refer to the *Symantec Product Authentication Service Administrator's Guide* for instructions.

If you have not already added VCS users during installation, you can add the users later.

See *Veritas Cluster Server User's Guide* for instructions to add VCS users.

Installing the Java Console

You can administer VCS using the VCS Java-based graphical user interface, Java Console. After VCS has been installed, install the Java Console on a Windows NT/2000 Professional system, or Linux system. The system from which you run the Java Console can be a system in the cluster or a remote workstation. A remote workstation enables each system in the cluster to be administered remotely.

Review the information about using the Cluster Manager and the Configuration Editor components of the Java Console. For more information, refer to the *Veritas Cluster Server User's Guide*.

Hardware requirements for the Java Console

The minimum hardware requirements for the Java Console follow:

- Pentium II 300 megahertz
- 256 megabytes of RAM
- 800x600 display resolution
- 8-bit color depth of the monitor
- A graphics card that is capable of 2D images

Note: Symantec recommends using Pentium III, 400MHz, 256MB RAM, and 800x600 display resolution.

The version of the Java™ 2 Runtime Environment (JRE) requires 32 megabytes of RAM. This version is supported on the Intel Pentium platforms that run the Linux kernel v 2.2.12 and glibc v2.1.2-11 (or later).

Symantec recommends using the following hardware:

- 48 megabytes of RAM
- 16-bit color mode
- The KDE and the KWM window managers that are used with displays set to local hosts

Installing the Java Console on Linux

Review the procedure to install the Java console.

To install Java console on Linux

- 1 Insert the VCS software disc into a drive on the system.

The software automatically mounts the disc on /mnt/cdrom.

- 2 If the disc does not get automatically mounted, then enter:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

- 3 Navigate to the folder that contains the RPMs.

```
# cd /mnt/cdrom/dist_arch/cluster_server/rpms
```

Where dist is the Linux distribution, rhel4 or sles9 and arch is the corresponding architecture.

- For RHEL4, replace arch with i686, ia64, or x86_64.
 - For SLES9, replace arch with i586, ia64, or x86_64.
- 4 Install the RPM using rpm -i command.

```
# rpm -i VRTScsm-5.0.00.0-GA_GENERIC.noarch.rpm
```

Installing the Java Console on a Windows workstation

You can install the VCS Java Console (Cluster Manager) on a Windows NT/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\VCSWindowsInstallers\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

When you have used installvcs program and chosen to configure and start VCS, VCS and all components are properly configured and can start correctly. You must verify that your cluster operates properly after the installation.

See [“About verifying the VCS installation”](#) on page 101.

Verifying and updating licenses on the system

After you install VCS, you can verify the licensing information using the vxlicrep program. You can replace the demo licenses with a permanent license.

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 following output to determine the following information:

- 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 have expiration dates.

```

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

Features :=
Platform              = Linux
Version               = 5.0 MP3
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 use a demo license, you can replace the demo license.

See [“Replacing a VCS demo license with a permanent license”](#) on page 83.

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 command does not shut down any running applications.

- 3 Enter the permanent license key using the following command on each node:

```
# cd /opt/VRTS/bin  
# ./vxllicinst -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:

```
# hastart
```

Accessing the VCS documentation

The software disc contains the documentation for VCS in Portable Document Format (PDF) in the cluster_server/docs directory. After you install VCS, Symantec recommends that you copy the PDF version of the documents to the /opt/VRTS/docs directory on each node to make it available for reference.

To access the VCS documentation

- ◆ Copy the PDF from the software disc (cluster_server/docs/) to the directory /opt/VRTS/docs.

Configuring VCS clusters for data integrity

This chapter includes the following topics:

- [About configuring VCS clusters for data integrity](#)
- [About I/O fencing components](#)
- [About setting up I/O fencing](#)
- [Preparing to configure I/O fencing](#)
- [Setting up I/O fencing](#)

About configuring VCS clusters for data integrity

When a node fails, VCS takes corrective action and configures its components to reflect the altered membership. If an actual node failure did not occur and if the symptoms were identical to those of a failed node, then such corrective action would cause a split-brain situation.

Some example scenarios that can cause such split-brain situations are as follows:

- **Broken set of private networks**
If a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects. The remaining node then takes corrective action. The failure of the private interconnects, instead of the actual nodes, presents identical symptoms and causes each node to determine its peer has departed. This situation typically results in data corruption because both nodes try to take control of data storage in an uncoordinated manner
- **System that appears to have a system-hang**

If a system is so busy that it appears to stop responding, the other nodes could declare it as dead. This declaration may also occur for the nodes that use the hardware that supports a "break" and "resume" function. When a node drops to PROM level with a break and subsequently resumes operations, the other nodes may declare the system dead. They can declare it dead even if the system later returns and begins write operations.

I/O fencing is a feature that prevents data corruption in the event of a communication breakdown in a cluster. VCS uses I/O fencing to remove the risk that is associated with split brain. I/O fencing allows write access for members of the active cluster. It blocks access to storage from non-members so that even a node that is alive is unable to cause damage.

After you install and configure VCS, you must configure I/O fencing in VCS to ensure data integrity.

About I/O fencing components

The shared storage for VCS must support SCSI-3 persistent reservations to enable I/O fencing. VCS involves two types of shared storage:

Data disks	Store shared data
Coordination points	Act as a global lock during membership changes

I/O fencing in VCS involves coordination points and data disks. Each component has a unique purpose and uses different physical disk devices. The fencing driver, known as vxfen, directs VxVM as necessary to carry out actual fencing operations at the disk group level.

About data disks

Data disks are standard disk devices for data storage and are either physical disks or RAID Logical Units (LUNs). These disks must support SCSI-3 PR and are part of standard VxVM or CVM disk groups.

CVM is responsible for fencing data disks on a disk group basis. Disks that are added to a disk group and new paths that are discovered for a device are automatically fenced.

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. Racing for

control of the coordination points to fence data disks is the key to understand how fencing prevents split brain.

Disks that act as coordination points are called coordinator disks. Coordinator disks are three standard disks or LUNs set aside for I/O fencing during cluster reconfiguration. Coordinator disks do not serve any other storage purpose in the VCS configuration. Users cannot store data on these disks or include the disks in a disk group for user data. The coordinator disks can be any three disks that support SCSI-3 PR. Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.

Symantec recommends using the smallest possible LUNs for coordinator disks. Because coordinator disks do not store any data, cluster nodes need to only register with them and do not need to reserve them.

You can configure coordinator disks to use Veritas Volume Manager Dynamic Multipathing (DMP) feature. Dynamic Multipathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. So, you can configure I/O fencing to use either DMP devices or the underlying raw character devices. I/O fencing uses SCSI-3 disk policy that is either raw or dmp based on the disk device that you use. The disk policy is raw by default.

See the *Veritas Volume Manager Administrator's Guide*.

You can use iSCSI devices as coordinator disks for I/O fencing. However, I/O fencing supports iSCSI devices only when you use DMP disk policy. If you use iSCSI devices as coordinator disks, make sure that the `/etc/vxfsnmode` file has the disk policy set to DMP.

For the latest information on supported hardware visit the following URL:


<http://entsupport.symantec.com/docs/283161>

About setting up I/O fencing


Figure 5-1 illustrates the tasks involved to configure I/O fencing.

Figure 5-1 Workflow to configure I/O fencing

Preparing to set up I/O fencing



Setting up I/O fencing



See [“Setting up I/O fencing”](#) on page 95.

I/O fencing requires the coordinator disks be configured in a disk group. The coordinator disks must be accessible to each node in the cluster. These disks enable the vxfen driver to resolve potential split-brain conditions and prevent data corruption.

Review the following requirements for coordinator disks:

- You must have three coordinator disks.
- Each of the coordinator disks must use a physically separate disk or LUN.
- Each of the coordinator disks should exist on a different disk array, if possible.

- You must initialize each disk as a VxVM disk.
- The coordinator disks must support SCSI-3 persistent reservations.
- The coordinator disks must exist in a disk group (for example, vxencoorddg).
- Symantec recommends using hardware-based mirroring for coordinator disks.

The I/O fencing configuration files include:

/etc/vxfendg	You must create this file to include the coordinator disk group information.
/etc/vxfenmode	You must set the I/O fencing mode to SCSI-3. You can configure the vxfen module to use either DMP devices or the underlying raw character devices. Note that you must use the same SCSI-3 disk policy on all the nodes. The SCSI-3 disk policy can either be raw or dmp. The policy is raw by default.
/etc/vxfentab	When you run the vxfen startup file to start I/O fencing, the script creates this /etc/vxfentab file on each node with a list of all paths to each coordinator disk. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Thus any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all paths to the coordinator disks. Note: The /etc/vxfentab file is a generated file; do not modify this file. An example of the /etc/vxfentab file on one node resembles as follows: ■ Raw disk: <pre>/dev/sdx /dev/sdy /dev/sdz</pre> ■ DMP disk: <pre>/dev/vx/rdmp/sdx /dev/vx/rdmp/sdy /dev/vx/rdmp/sdz</pre>

In some cases you must remove disks from or add disks to an existing coordinator disk group.

Warning: If you remove disks from an existing coordinator disk group, then be sure to remove the registration and reservation keys from these disks before you add the disks to another disk group.

Preparing to configure I/O fencing

Make sure you performed the following tasks before configuring I/O fencing for VCS:

- Install the correct operating system.
- Install the VRTSvxfen RPM when you installed VCS.
- Install a version of Veritas Volume Manager (VxVM) that supports SCSI-3 persistent reservations (SCSI-3 PR).
Refer to the installation guide that comes with the Storage Foundation product that you use.

The shared storage that you add for use with VCS software must support SCSI-3 persistent reservations, a functionality that enables the use of I/O fencing.

Perform the following preparatory tasks to configure I/O fencing:

Initialize disks as VxVM disks	See “Initializing disks as VxVM disks” on page 90.
Identify disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 92.
Check shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 92.
The tasks involved in checking the shared disks for I/O fencing are as follows:	
■ Verify that the nodes have access to the same disk	
■ Test the disks using the vxfsentsthdw utility	

Initializing disks as VxVM disks

Install the driver and HBA card. Refer to the documentation from the vendor for instructions.

After you physically add shared disks to the nodes, you must do the following:

- Initialize them as VxVM disks
- Verify that all the nodes see the same disk

See the *Veritas Volume Manager Administrator's Guide* for more information on how to add and configure disks.

To initialize disks as VxVM disks

- 1 Make the new disks recognizable. On each node, enter:

```
# fdisk -l
```

- 2 If the Array Support Library (ASL) for the array that you add is not installed, obtain and install it on each node before proceeding.

The ASL for the supported storage device that you add is available from the disk array vendor or Symantec technical support.

- 3 Verify that the ASL for the disk array is installed on each of the nodes. Run the following command on each node and examine the output to verify the installation of ASL.

The following output is a sample:

```
# vxddladm listsupport all
LIBNAME                                VID
=====
libvxCLARiON.so                        DGC
libvxcscovrts.so                       CSCOVRTS
libvxemc.so                            EMC
```

- 4 Scan all disk drives and their attributes, update the VxVM device list, and reconfigure DMP with the new devices. Type:

```
# vxdisk scandisks
```

See the Veritas Volume Manager documentation for details on how to add and configure disks.

- 5 To initialize the disks as VxVM disks, use one of the following methods:
 - Use the interactive `vxdiskadm` utility to initialize the disks as VxVM disks. For more information see the *Veritas Volume Managers Administrator's Guide*.
 - Use the `vxdisksetup` command to initialize a disk as a VxVM disk.

```
vxdisksetup -i device_name format=cdsdisk
```

The example specifies the CDS format:

```
# vxdisksetup -i sdr format=cdsdisk
```

Repeat this command for each disk you intend to use as a coordinator disk.

Identifying disks to use as coordinator disks

After you add and initialize disks, identify disks to use as coordinator disks.

To identify the coordinator disks

- 1 List the disks on each node.

For example, execute the following commands to list the disks:

```
# fdisk -l
```

- 2 Pick three SCSI-3 PR compliant shared disks as coordinator disks.

Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure VCS meets the I/O fencing requirements. You can test the shared disks using the `vxfsentsthdw` utility. The two nodes must have `ssh` (default) or `rsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfsenadm` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

Make sure to test the disks that serve as coordinator disks.

The `vxfsentsthdw` utility has additional options suitable for testing many disks. Review the options for testing the disk groups (`-g`) and the disks that are listed in a file (`-f`). You can also test disks without destroying data using the `-r` option.

See *Veritas Cluster Server User's Guide*.

Checking that disks support SCSI-3 involves the following tasks:

- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 93.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfsentsthdw utility”](#) on page 94.

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the `vxfcntlsthdw` utility, you must verify that the systems see the same disk.

To verify that the nodes have access to the same disk

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed VCS.

Warning: The tests overwrite and destroy data on the disks unless you use the `-r` option.

- 2 Make sure system-to-system communication functions properly.

See [“Setting up inter-system communication”](#) on page 41.

After you complete the testing process, remove permissions for communication and restore public network connections.

See [“Removing permissions for communication”](#) on page 100.

- 3 Ensure that both nodes are connected to the same disk during the testing. Use the `vxfenadm` command to verify the disk serial number. If you use `rsh`, launch the `vxfenadm` utility with the `-n` option.

```
vxfenadm -i diskpath
```

Refer to the `vxfenadm` (1M) manual page.

For example, an EMC disk is accessible by the `/dev/sdx` path on node A and the `/dev/sdy` path on node B.

From node A, enter:

```
# vxfenadm -i /dev/sdx

SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E

Vendor id : EMC
Product id : SYMMETRIX
Revision : 5567
Serial Number : 42031000a
```

The same serial number information should appear when you enter the equivalent command on node B using the `/dev/sdy` path.

On a disk from another manufacturer, Hitachi Data Systems, the output is different and may resemble:

```
# vxfenadm -i /dev/sdz

SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E

Vendor id      : HITACHI
Product id     : OPEN-3
Revision      : 0117
Serial Number  : 0401EB6F0002
```

Testing the disks using vxfentsthdw utility

This procedure uses the `/dev/sdx` disk in the steps.

If the utility does not show a message that states a disk is ready, the verification has failed. Failure of verification can be the result of an improperly configured disk array. The failure can also be due to a bad disk.

If the failure is due to a bad disk, remove and replace it. The `vxfentsthdw` utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/sdx is ready to be configured for I/O Fencing on
node galaxy
```

For more information on how to replace coordinator disks, refer to the *Veritas Cluster Server User's Guide*.

To test the disks using vxfentsthdw utility

- 1 Make sure system-to-system communication functions properly.
- 2 From one node, start the utility.

Do one of the following:

- If you use `ssh` for communication:

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw
```

- If you use `rsh` for communication:

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw -n
```

- 3 The script warns that the tests overwrite data on the disks. After you review the overview and the warning, confirm to continue the process and enter the node names.

```
***** WARNING!!!!!!!!!! *****
THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!

Do you still want to continue : [y/n] (default: n) y
Enter the first node of the cluster: galaxy
Enter the second node of the cluster: nebula
```

- 4 Enter the names of the disks that you want to check. Each node may know the same disk by a different name:

```
Enter the disk name to be checked for SCSI-3 PGR on node
galaxy in the format: /dev/sdx
/dev/sdx
Enter the disk name to be checked for SCSI-3 PGR on node
nebula in the format: /dev/sdx
Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/sdx
```

If the serial numbers of the disks are not identical, then the test terminates.

- 5 Review the output as the utility performs the checks and report its activities.
- 6 If a disk is ready for I/O fencing on each node, the utility reports success:

```
The disk is now ready to be configured for I/O Fencing on node
galaxy

ALL tests on the disk /dev/sdx have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

- 7 Run the vxfsentsthdw utility for each disk you intend to verify.

Setting up I/O fencing

Make sure you completed the preparatory tasks before you set up I/O fencing.

Tasks that are involved in setting up I/O fencing include:

Table 5-1 Tasks to set up I/O fencing

Action	Description
Setting up coordinator disk groups	See “ Setting up coordinator disk groups ” on page 96.
Configuring I/O fencing	See “ Configuring I/O fencing ” on page 96.
Modifying VCS configuration to use I/O fencing	See “ Modifying VCS configuration to use I/O fencing ” on page 98.
Verifying I/O fencing configuration	See “ Verifying I/O fencing configuration ” on page 99.

Setting up coordinator disk groups

From one node, create a disk group named `vxfencoorddg`. This group must contain three disks or LUNs. If you use VxVM 5.0 or later, you must also set the coordinator attribute for the coordinator disk group. VxVM uses this attribute to prevent the reassignment of coordinator disks to other disk groups.

Note that if you create a coordinator disk group as a regular disk group, you can turn on the coordinator attribute in Volume Manager.

Refer to the *Veritas Volume Manager Administrator’s Guide* for details on how to create disk groups.

The following example procedure assumes that the disks have the device names `sdx`, `sdz`, and `sdz`.

To create the `vxfencoorddg` disk group

- 1 On any node, create the disk group by specifying the device names:

```
# vxdg init vxfencoorddg sdx sdz sdz
```

- 2 If you use VxVM 5.0 or later, set the coordinator attribute value as "on" for the coordinator disk group.

```
# vxdg -g vxfencoorddg set coordinator=on
```

Configuring I/O fencing

After you set up the coordinator disk group, you must do the following to configure I/O fencing:

- Create the I/O fencing configuration file `/etc/vxfendg`

- Update the I/O fencing configuration file `/etc/vxfenmode`
- Start I/O fencing

To update the I/O fencing files and start I/O fencing

- 1 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

- 2 Import the disk group with the `-t` option to avoid automatically importing it when the nodes restart:

```
# vxdg -t import vxfencoorddg
```

- 3 Deport the disk group. Deporting the disk group prevents the coordinator disks from serving other purposes:

```
# vxdg deport vxfencoorddg
```

- 4 On each nodes, type:

```
# echo "vxfencoorddg" > /etc/vxfendg
```

Do not use spaces between the quotes in the "vxfencoorddg" text.

This command creates the `/etc/vxfendg` file, which includes the name of the coordinator disk group.

- 5 On all cluster nodes depending on the SCSI-3 mechanism, type one of the following selections:

- For DMP configuration:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
```

- For raw device configuration:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_raw /etc/vxfenmode
```

- 6 To check the updated `/etc/vxfenmode` configuration, enter the following command on one of the nodes. For example:

```
# more /etc/vxfenmode galaxy
```

- 7 If the I/O fencing driver `vxfen` is already running, stop the I/O fencing driver. You need to restart the driver for the new configuration to take effect.

```
# /etc/init.d/vxfen stop
```

- 8 Start the I/O fencing driver.

The `vxfen` startup script also invokes the `vxfenconfig` command, which configures the `vxfen` driver to start and use the coordinator disks that are listed in `/etc/vxfentab`.

```
# /etc/init.d/vxfen start
```

Modifying VCS configuration to use I/O fencing

After you add coordinator disks and configure I/O fencing, add the `UseFence = SCSI3` cluster attribute to the VCS configuration file `/etc/VRTSvcs/conf/config/main.cf`. If you reset this attribute to `UseFence = None`, VCS does not make use of I/O fencing abilities while failing over service groups. However, I/O fencing needs to be disabled separately.

To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

- 3 Make a backup copy of the `main.cf` file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

- 4 On one node, use vi or another text editor to edit the main.cf file. To modify the list of cluster attributes, add the UseFence attribute and assign its value as SCSI3.

```
cluster clus1
UserNames = { admin = "CDRpdxPmHpzS." }
Administrators = { admin }
HacliUserLevel = COMMANDROOT
CounterInterval = 5
UseFence = SCSI3
)
```

- 5 Save and close the file.
- 6 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

- 7 Using rcp or another utility, copy the VCS configuration file from a node (for example, galaxy) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp galaxy:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

- 8 On each node enter the following command.

```
# /opt/VRTS/bin/hastart
```

Verifying I/O fencing configuration

Verify from the vxfenadm output that the SCSI-3 disk policy reflects the configuration in the /etc/vxfenmode file.

To verify I/O fencing configuration

- ◆ On one of the nodes, type:

```
# vxfenadm -d

I/O Fencing Cluster Information:
=====

Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:

    * 0 (galaxy)
    1 (nebula)

RFSM State Information:
    node 0 in state 8 (running)
    node 1 in state 8 (running)
```

Removing permissions for communication

Make sure you completed the installation of VCS and the verification of disk support for I/O fencing. If you used `rsh`, remove the temporary `rsh` access permissions that you set for the nodes and restore the connections to the public network.

If the nodes use `ssh` for secure communications, and you temporarily removed the connections to the public network, restore the connections.

Verifying the VCS installation

This chapter includes the following topics:

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

About verifying the VCS installation

After you install and configure VCS, you can inspect the contents of the key VCS configuration files that you have installed and modified during the process. These files reflect the configuration that is based on the information you supplied. You can also run VCS commands to verify the status of LLT, GAB, and the cluster.

About the 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.

The information that these LLT and GAB configuration files contain is as follows:

- The `/etc/llthosts` file

The file `llthosts` is a database that contains one entry per system. This file 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 the entries that resemble:

```
0      galaxy
1      nebula
```

■ The `/etc/llttab` file

The file `llttab` contains the information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the private network links that correspond to the specific system.

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

```
set-node galaxy
set-cluster 2
link eth1 eth1 - ether - -
link eth2 eth2 - ether - -
```

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

```
set-node galaxy
set-cluster 2
link eth1 eth-00:04:23:AC:12:C4 - ether - -
link eth2 eth-00:04:23:AC:12:C5 - ether - -
```

The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the `link` command. These lines identify the two network cards that the LLT protocol uses.

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.

■ The `/etc/gabtab` file

After you install VCS, the file `/etc/gabtab` contains a `gabconfig(1)` command that configures the GAB driver for use.

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

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

The `-c` option configures the driver for use. The `-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. The Gigabit Ethernet controller does not support the use of `-c -x`.

About the VCS configuration file main.cf

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

See [“Sample main.cf file for VCS clusters”](#) on page 104.

See [“Sample main.cf file for global clusters”](#) on page 106.

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 also located 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 definition 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.
- If you set up the optional I/O fencing feature for VCS, then the `UseFence = SCSI3` attribute that you added is present.
- If you configured the cluster in secure mode, the `main.cf` includes the `VxSS` service group and "`SecureClus = 1`" cluster attribute.
- The `installvcs` program creates the `ClusterService` service group. The group includes the IP, NIC, and `VRTSWebApp` resources.

The service group also has the following characteristics:

- 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 set up global clusters, the `ClusterService` service group contains an `Application` resource, `wac` (wide-area connector). This resource's attributes contain definitions for controlling the cluster in a global cluster environment.

Refer to the *Veritas Cluster Server User's Guide* for information about managing VCS global clusters.

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

Sample main.cf file for VCS clusters

The following sample main.cf is for a secure cluster that is managed locally by Cluster Management Console.

```
include "types.cf"

cluster vcs_cluster2 (
    UserNames = { admin = cDRpdxPmHpzS, smith = dKLhKJkHLh }
    ClusterAddress = "10.10.12.1"
    Administrators = { admin, smith }
    CounterInterval = 5
    SecureClus = 1
)

system galaxy (

system nebula (

group ClusterService (
    SystemList = { galaxy = 0, nebula = 1 }
    UserStrGlobal = "LocalCluster@https://10.182.2.76:8443;"
    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

IP webip (
    Device = eth0
    Address = "10.10.12.1"
    NetMask = "255.255.240.0"
)

NIC csgnic (
    Device = eth0
)
```

```
NotifierMngr ntfr (  
    SnmpConsoles = { "saturn" = Error, "jupiter" = SevereError }  
    SmtplibServer = "smtp.example.com"  
    SmtplibRecipients = { "ozzie@example.com" = Warning,  
                          "harriet@example.com" = Error }  
)  
  
VRTSWebApp VCSweb (  
    Critical = 0  
    InstallDir = "/opt/VRTSweb/VERITAS"  
    TimeForOnline = 5  
    RestartLimit = 3  
)  
  
VCSweb requires webip  
ntfr requires csgnic  
webip requires csgnic  
  
// resource dependency tree  
//  
// group ClusterService  
// {  
//     VRTSWebApp VCSweb  
//     {  
//         IP webip  
//         {  
//             NIC csgnic  
//         }  
//     }  
//     NotifierMngr ntfr  
//     {  
//         NIC csgnic  
//     }  
// }  
  
group VxSS (  
    SystemList = { galaxy = 0, nebula = 1 }  
    Parallel = 1  
    OnlineRetryLimit = 3  
    OnlineRetryInterval = 120  
)
```

```

Phantom phantom_vxss (
)

ProcessOnOnly vxatd (
    IgnoreArgs = 1
    PathName = "/opt/VRTSat/bin/vxatd"
)

// resource dependency tree
//
// group VxSS
// {
//     Phantom phantom_vxss
//     ProcessOnOnly vxatd
// }

```

Sample main.cf file for global clusters

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

```

.
.
group ClusterService (
    SystemList = { galaxy = 0, nebula = 1 }

    UserStrGlobal = "LocalCluster@https://10.182.2.78:8443;"

    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

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

```

Verifying the LLT, GAB, and VCS configuration files

Make sure that the LLT, GAB, and VCS configuration files contain the information you provided during VCS installation and configuration.

To verify the LLT, GAB, and VCS configuration files

- 1 Navigate to the location of the configuration files:
 - LLT
/etc/llthosts
/etc/llttab
 - GAB
/etc/gabtab
 - VCS
/etc/VRTSvcs/conf/config/main.cf
- 2 Verify the content of the configuration files.
 - See [“About the LLT and GAB configuration files”](#) on page 101.
 - See [“About the VCS configuration file main.cf”](#) on page 103.

Verifying LLT, GAB, and cluster operation

Verify the operation of LLT, GAB, and the cluster using the VCS commands.

To verify LLT, GAB, and cluster operation

- 1 Log in to any node in the cluster as superuser.
- 2 Make sure that the PATH environment variable is set to run the VCS commands.
 - See [“Setting the PATH variable”](#) on page 46.
- 3 Verify LLT operation.
 - See [“Verifying LLT”](#) on page 108.
- 4 Verify GAB operation.
 - See [“Verifying GAB”](#) on page 110.
- 5 Verify the cluster operation.
 - See [“Verifying the cluster”](#) on page 111.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also 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.

To verify LLT

- 1
- Log in as superuser on the node galaxy.
- 2
- Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node           State           Links
*0 galaxy      OPEN            2
 1 nebula      OPEN            2
```

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

- 3
- Log in as superuser on the node nebula.
- 4
- Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
Node           State           Links
 0 galaxy      OPEN            2
*1 nebula      OPEN            2
```

- 5
- To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv | more
```

The output on galaxy resembles:

Node	State	Link	Status	Address
*0 galaxy	OPEN	eth1	UP	08:00:20:93:0E:34
		eth2	UP	08:00:20:93:0E:34
1 nebula	OPEN	eth1	UP	08:00:20:8F:D1:F2
		eth2	DOWN	
2	CONNWAIT	eth1	DOWN	
		eth2	DOWN	
3	CONNWAIT	eth1	DOWN	
		eth2	DOWN	
.				
.				
.				
31	CONNWAIT	eth1	DOWN	
		eth2	DOWN	

Note that the output lists 32 nodes. The command reports the status on the two nodes in the cluster, galaxy and nebula, along with the details for the non-existent nodes.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- An address for each link

However, the output in the example shows different details for the node nebula. The private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

- 6
- To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node galaxy in a two-node cluster:

```
lltstat -p
```

The 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

Verify the GAB operation using the `gabconfig -a` command. This command returns the GAB port membership information.

The ports indicate the following:

- Port a
- Nodes have GAB communication
 - gen a36e0003 is a randomly generated number
 - membership 01 indicates that nodes 0 and 1 are connected
- Port h
- VCS is started
 - gen fd570002 is a randomly generated number
 - membership 01 indicates that nodes 0 and 1 are both running VCS

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

To verify GAB

- 1 To verify that GAB operates, type the following command on each node:

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

- 2 Review the output of the command:

- If GAB operates, the following GAB port membership information is returned:

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

- If GAB does not operate, 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
```

Verifying the cluster

Verify the status of the cluster using the `hastatus` command. This command returns the system state and the group state.

Refer to the `hastatus(1M)` manual page.

Refer to the *Veritas Cluster Server User's Guide* for a description of system states and the transitions between them.

To verify the cluster

- 1 To verify the status of the cluster, type the following command:

```
hastatus -summary
```

The output resembles:

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

A galaxy              RUNNING              0
A nebula              RUNNING              0

-- GROUP STATE
-- Group              System              Probed  AutoDisabled  State

B ClusterService galaxy              Y        N              ONLINE
B ClusterService nebula              Y        N              OFFLINE
```

- 2 Review the command output for the following information:
- The system state
If the value of the system state is RUNNING, VCS is successfully installed and started.
 - The ClusterService group state
In the sample output, the group state lists the ClusterService group, which is ONLINE on galaxy and OFFLINE on nebula.

Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

Refer to the `hasys(1M)` manual page.

Refer to the *Veritas Cluster Server User's Guide* for information about the system attributes for VCS.

To verify the cluster nodes

- ◆ On one of the nodes, type the `hasys -display` command:

```
hasys -display
```

The example shows the output when the command is run on the node galaxy. The list continues with similar information for nebula (not shown) and any other nodes in the cluster.

#System	Attribute	Value
galaxy	AgentsStopped	0
galaxy	AvailableCapacity	100
galaxy	CPUUsage	0
galaxy	CPUUsageMonitoring	Enabled 0 ActionThreshold 0 ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
galaxy	Capacity	100
galaxy	ConfigBlockCount	142
galaxy	ConfigCheckSum	4085
galaxy	ConfigDiskState	CURRENT
galaxy	ConfigFile	/etc/VRTSvcS/conf/config
galaxy	ConfigInfoCnt	0
galaxy	ConfigModDate	Fri May 26 17:22:48 2006
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	
galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	5.0.30.0
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	LLTNodeId	0
galaxy	LicenseType	DEMO

#System	Attribute	Value
galaxy	Limits	
galaxy	LinkHbStatus	<i>eth1</i> UP <i>eth2</i> UP
galaxy	LoadTimeCounter	0
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	ShutdownTimeout	120
galaxy	SourceFile	./main.cf
galaxy	SysInfo	Linux:galaxy, #1 Fri Apr 22 18:13:58 EDT 2005, 2.6.9-34-default, i686
galaxy	SysName	galaxy
galaxy	SysState	RUNNING
galaxy	SystemLocation	
galaxy	SystemOwner	
galaxy	TFrozen	0
galaxy	TRSE	0
galaxy	UpDownState	Up
galaxy	UserInt	0
galaxy	UserStr	
galaxy	VCSFeatures	DR
galaxy	VCSMode	VCS

Upgrading VCS

This chapter includes the following topics:

- [About VCS 5.0 MP3 upgrade](#)
- [VCS supported upgrade paths](#)
- [Upgrading VCS in secure enterprise environments](#)
- [About minimal downtime upgrade](#)
- [About changes to VCS bundled agents](#)
- [Upgrading to VCS 5.0 MP3](#)

About VCS 5.0 MP3 upgrade

If you want to upgrade VCS to 5.0 MP3, you can use one of the following programs depending on the VCS version installed on the nodes:

Veritas installer or installvcs program	Upgrade VCS from VCS 4.1 or VCS 4.1 Maintenance Pack versions to VCS 5.0 MP3
installmp program	Upgrade VCS 5.0 or VCS 5.0 Maintenance Pack versions to VCS 5.0 MP3

See [“VCS supported upgrade paths”](#) on page 115.

VCS supported upgrade paths

If you are currently running a cluster with any earlier VCS versions that is supported for upgrade, you can run the installer to upgrade to VCS 5.0 MP3.

Review the supported upgrade path tables for VCS clusters on RHEL, SLES, and OEL operating systems. Depending on the upgrade path, the tables provide information on whether to use the `installvcs` or the `installmp` program.

The following variations apply to the upgrade paths:

- To upgrade VCS 4.1MP4 on RHEL5:
 - Upgrade to RHEL5U1 or RHEL5U2.
 - Upgrade to 4.1MP4RP2.
 - Upgrade to VCS 5.0MP3.
- To upgrade to VCS5.0MP3 on SLES10SP1:
 - Upgrade to SLES10SP1.
 - Upgrade VCS to 4.1MP4RP2.
 - Upgrade to VCS5.0MP3.
- To upgrade to VCS5.0MP3 on SLES9SP4:
 - Upgrade to SLES9SP4.
 - Upgrade VCS to 4.1MP4RP2.
 - Upgrade to VCS5.0MP3.

[Figure 7-1](#) lists the supported upgrade paths for Red Hat Enterprise Linux.

Figure 7-1 Supported upgrade paths for RHEL

Upgrade scenarios	From		To		Upgrade program to use
	VCS	RHEL	VCS	RHEL	
VCS upgrade and RHEL upgrade	4.1	RHEL4U1	5.0MP3	RHEL4U3	installvcs
	4.1MP1	RHEL4U2		RHEL4U4	
	4.1MP2			RHEL4U5	
	4.1MP3			RHEL4U6	
	4.1MP2	RHEL4U3	5.0MP3	RHEL4U4	
	4.1MP3			RHEL4U5	
				RHEL4U6	
	4.1MP4	RHEL4U4	5.0MP3	RHEL4U5	
				RHEL4U6	
	4.1MP4	RHEL5	5.0MP3	RHEL5U1	
VCS upgrade				RHEL5U2	
	4.1MP2	RHEL4U3	5.0MP3	RHEL4U3	
	4.1MP3				
	4.1MP4	RHEL4U4	5.0MP3	RHEL4U4	
VCS upgrade and RHEL upgrade	5.0	RHEL4U3	5.0MP3	RHEL4U4	installmp
	5.0MP1			RHEL4U5	
	5.0MP2			RHEL4U6	
VCS upgrade	5.0	RHEL4U3	5.0MP3	RHEL4U3	
	5.0MP1				
	5.0MP2				

Figure 7-2 lists the supported upgrade paths for SUSE Linux Enterprise Server.

Figure 7-2 Supported upgrade paths for SLES

Upgrade scenarios	From		To		Upgrade program to use
	VCS	SLES	VCS	SLES	
VCS upgrade and SLES upgrade	4.1	SLES9SP1	5.0MP3	SLES9SP3	installvcs
	4.1MP1	SLES9SP2		SLES9SP4	
	4.1MP2				
	4.1MP3				
	4.1	SLES9SP3	5.0MP3	SLES9SP4	
	4.1MP1				
	4.1MP2				
	4.1MP3				
	4.1MP3	SLES10	5.0MP3	SLES10SP1	
	4.1MP4				
VCS upgrade	4.1	SLES9SP3	5.0MP3	SLES9SP3	
	4.1MP1				
	4.1MP2				
	4.1MP3				
VCS upgrade and SLES upgrade	5.0	SLES9SP3	5.0MP3	SLES9SP4	installmp
	5.0MP1				
	5.0MP2				
VCS upgrade	5.0MP2	SLES9SP3	5.0MP3	SLES9SP3	

[Figure 7-3](#) lists the supported upgrade paths for Oracle Enterprise Linux.

Figure 7-3 Supported upgrade paths for OEL

Upgrade scenarios	From		To		Upgrade program to use
	VCS	OEL	VCS	OEL	
VCS upgrade and OEL upgrade	5.0MP2	OEL4 U4	5.0MP3	OEL4 U4	installmp
				OEL4 U5	

Upgrading VCS in secure enterprise environments

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). So, run the `installvcs` program on each node to upgrade a cluster to VCS 5.0MP3. On the first node, the program updates the configuration and stops the cluster before you upgrade the system. On the other nodes, it uninstalls the previous version and installs VCS 5.0MP3. After the last node is upgraded and started, the upgrade is complete.

About minimal downtime upgrade

Use a minimal downtime upgrade to upgrade VCS. This procedure minimizes downtime for the cluster that you want to upgrade. Depending on the situation, you can calculate the approximate downtime as follows:

You can fail over all your service groups to the nodes that are up.	Downtime equals the time that is taken to offline and online the service groups.
You have a service group that you cannot fail over to a node that runs during upgrade.	Downtime for that service group equals the time that is taken to perform an upgrade and reboot the node.

Prerequisites for a minimal downtime upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade.

Planning for the minimal downtime upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group.

Some rough guidelines follow:

- Split the cluster in half. If the cluster has an odd number of nodes, calculate $(n+1)/2$, and start the upgrade with the even number of nodes.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.

Minimal downtime upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the minimum downtime upgrade:

- While you perform the upgrades, do not choose any configuration options.
- While you perform the upgrades, do not start any modules.
- When you start the installer, only select VCS.
- While you perform the upgrades, do not add or remove service groups to any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.

Minimal downtime upgrade example


In this example, you have four nodes: node01, node02, node03, and node04. You also have four service groups: sg1, sg2, sg3, and sg4. Each service group is running on one node as follows:

- node01 runs sg2.
- node02 runs sg1.
- node03 runs sg4.
- node04 runs sg3.

In your system list, you have each service group that fails over to one other node as follows:

- sg1 can fail over between node01 and node02.
- sg2 can fail over between node01 and node03.
- sg3 can fail over between node01 and node04.
- sg4 can fail over between node03 and node04.

[Figure 7-4](#) shows four nodes, four service groups, and their failover paths.

Figure 7-4 Four nodes, four service groups, and their failover paths

Minimal downtime example overview

This example presumes that you have at least one service group (in this case sg3), that cannot stay online on both nodes during the upgrade. In this situation, sg3 must be a low-priority service group. The cluster is split with node02 and node03 together for the first upgrade, and node01 and node04 together for the next upgrade.

You switch sg1 to run on node01. Switch sg2 to run on node04. You then perform the upgrade on node02 and node03. When you finish the upgrade on node02 and node03, you need to upgrade node01 and node04.

Your cluster is down when you stop HAD on node01 and node04, but have not yet started node02 and node03.

You have to take your service groups offline manually on node01 and node04. When you start node02 and node03, the service groups come online. Reboot node01 and node04 when the upgrade completes. They then rejoin the cluster and you can balance the load on systems by switching service groups.

About changes to VCS bundled agents

Review the changes to VCS bundled agents if you upgrade to VCS 5.0 MP3.

Deprecated agents

The following agents are no longer supported:

- CampusCluster
- ClusterMonitorConfig

- SANVolume (deprecated since 5.0 MP1)
- Service group heartbeat (ServiceGroupHB)—VCS does not support service group heartbeats in this release. Symantec recommends using I/O fencing.

Removing deprecated resource types

With VCS 5.0MP3, certain resource type definitions are no longer used. Before you start the upgrade process, you must remove the resources of the deprecated resource types from your cluster configuration.

Review the changes to VCS agents in version 5.0 MP3.

See [“About changes to VCS bundled agents”](#) on page 121.

Perform the following steps to remove the deprecated resource types.

Note: Make sure you start VCS on the local node before starting on the other nodes. This standard ensures that HAD reads the configuration from the local node and updates it on the remaining nodes.

To remove the deprecated resource types

- 1 Save the VCS configuration and stop the VCS engine.

```
# haconf -dump -makero  
# hastop -all -force
```

- 2 Back up the configuration file, main.cf to a location on the cluster node.
- 3 Edit the main.cf located under /etc/VRTSvcs/conf/config.

Perform the following instructions:

- Remove the resource of the deprecated resource types.
You must modify the resource dependencies to ensure that the configuration works properly.
- Save the main.cf.
- Reformat the main.cf file.

```
# hacf -cftocmd config  
# hacf -cmdtocf config
```

- 4 Verify the configuration.

```
# cd /etc/VRTSvcs/conf/config  
# hacf -verify config
```

- 5 Start VCS on the local node.
- 6 Start VCS on other nodes.

New agents

The following new agents were added in the 5.0 release:

- NFSRestart—Provides high availability for NFS record locks.
- ProcessOnOnly—Starts and monitors a user-specified process.
- RemoteGroup—Monitors and manages a service group on another system.
- SANVolume—Monitors volumes in a SAN environment managed using Storage Foundation Volume Server.

However, this agent is deprecated since VCS 5.0 MP1.

The following new agent is added in the 5.0 MP3 release:

- DiskGroupSnap—Verifies the configuration and data integrity in a campus cluster environment.

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on these new agents.

New and modified attributes for VCS 5.0 MP3 agents

[Table 7-1](#) lists the attributes that VCS adds or modifies when you upgrade from VCS 4.1 to VCS 5.0 MP3.

Table 7-1

New and modified attributes for VCS 5.0MP3 agents for upgrades from VCS 4.1

Agent	New and modified attributes	Default Value
Apache		
New attributes		
	EnableSSL	INFO
	EnvFile	
	IntentionalOffline	
	PidFile	
	ResLogLevel	30
	SecondLevelMonitor	

Table 7-1 New and modified attributes for VCS 5.0MP3 agents for upgrades from VCS 4.1 (*continued*)

Agent	New and modified attributes	Default Value
	SecondLevelTimeout	
	SharedObjDir	
	User	0
Modified attributes		
	Address is changed to HostName	
	Postdirective is changed to DirectiveAfter	
	Predirective is changed to DirectiveBefore	
	ServerRoot is changed to httpdDir	""
	ConfigFile	""
Application		
Modified attributes		
	SupportedActions	{ "program.vfd", "user.vfd", "cksum.vfd", getcksum }
DiskGroup		
New attributes		
	DiskGroupType	private
	UmountVolumes	0
Modified attributes		
	StopVolumes	1
	StartVolumes	1

Table 7-1 New and modified attributes for VCS 5.0MP3 agents for upgrades from VCS 4.1 (*continued*)

Agent	New and modified attributes	Default Value
	SupportedActions	{ "license.vfd", "disk.vfd", "udid.vfd", "verifyplex.vfd", checkudid, numdisks, campusplex, joindg, splitdg, getvxvminfo, volinuse }
DiskGroupSnap		
New attributes		
	ActionTimeOut	120
	MonitorInterval	300
	NumThreads	1
	ArgList	{ TargetResName, FDSiteName }
	TargetResName	
	FDSiteName	
DNS		
New attributes		
	SupportedActions	{ "dig.vfd", "keyfile.vfd", "master.vfd" }
	ResRecord	
	CreatePTR	
	OffDelRR	
IP		
Modified attribute		
	SupportedActions	{ "device.vfd" "route.vfd" }
LVMVolumeGroup		
New attributes		
	SupportedActions	{ volinuse }

Table 7-1 New and modified attributes for VCS 5.0MP3 agents for upgrades from VCS 4.1 (*continued*)

Agent	New and modified attributes	Default Value
Mount		
New attributes		
	RegList	{ VxFSMountLock }
	SecondLevelMonitor	
	SecondLevelTimeout	30
	VxFSMountLock	0
Modified attributes		
	SupportedActions	{ "mountpoint.vfd", "mounted.vfd", "vxfslic.vfd", "chgmtntlock", "mountentry.vfd" }
NFS		
New attributes		
	NFSSecurity	
	NFSv4Support	
	LockFileTimeout	180
Modified attributes		
	IPResName: Renamed Address	
	LockRecovery: Replaced by NFSLockFailover attribute in NFSRestart agent	
	Operations	OnOnly
	RestartLimit	1
NIC		
Modified attribute		

Table 7-1 New and modified attributes for VCS 5.0MP3 agents for upgrades from VCS 4.1 (*continued*)

Agent	New and modified attributes	Default Value
Process	SupportedActions	{"device.vfd"}
	Modified attribute	
Share	SupportedActions	{ "program.vfd", getcksum }
	New attributes	
	SupportedActions	{ "direxists.vfd" }

[Table 7-2](#) lists the attributes that VCS adds or modifies when you upgrade from VCS 5.0 or later to VCS 5.0 MP3.

Table 7-2 New and modified attributes for VCS 5.0MP3 agents for upgrades from VCS 5.0

Agent	New and modified attributes	Default Value
Apache		
New attributes		
	PidFile	
DiskGroup	IntentionalOffline	0
	New attributes	
Modified attributes	UmountVolumes	0
	SupportedActions	{ "license.vfd", "disk.vfd", "udid.vfd", "verifyplex.vfd", checkudid, numdisks, campusplex, joindg, splitdg, getvxvminfo, volinuse }
DNS		

Table 7-2 New and modified attributes for VCS 5.0MP3 agents for upgrades from VCS 5.0 (*continued*)

Agent	New and modified attributes	Default Value
New attributes		
	SupportedActions	{ "dig.vfd", "keyfile.vfd", "master.vfd" }
	ResRecord	
	CreatePTR	
	OffDelRR	
LVMVolumeGroup		
New attributes		
	SupportedActions	{ volinuse }
Mount		
New attributes		
	RegList	{ VxFSMountLock }
	VxFSMountLock	0
Modified attributes		
	SupportedActions	{ "mountpoint.vfd", "mounted.vfd", "vxfslic.vfd", "chgmntlock", "mountentry.vfd" }
NFSRestart		
New attributes		
	SupportedActions	{ "lockdir.vfd", "nfsconf.vfd" }
Share		
New attributes		
	SupportedActions	{ "direxists.vfd" }

Upgrading to VCS 5.0 MP3

You must upgrade the following to version 5.0 MP3:

VCS	Depending on your upgrade path, use <code>installvcs</code> or <code>installmp</code> program to upgrade VCS. See “Upgrading VCS to version 5.0 MP3” on page 129.
VCS agents	See “Upgrading the VCS agents” on page 136.
VCS Cluster Manager	See “Upgrading the Cluster Manager (Java Console)” on page 136.
VCS Simulator	See “Upgrading the VCS Simulator” on page 136.

Upgrading VCS to version 5.0 MP3

Make sure that you removed the deprecated resource types before you start the upgrade.

See [“Removing deprecated resource types”](#) on page 122.

To reduce the downtime of the cluster, you can upgrade earlier versions of VCS to version 5.0 MP3 in the following phases:

Select one or a group of nodes to upgrade first.	Leave the other nodes with the VCS and the applications running.
Upgrade the selected nodes until you reach the last node or group of nodes.	Perform the following procedures: ■ “Performing the pre-upgrade tasks” on page 130. ■ “Upgrading VCS” on page 131. ■ “Performing the post-upgrade tasks” on page 132.
Upgrade the last node or group of nodes and complete the upgrade.	Perform the following procedure: “Completing the upgrade on all the nodes in the cluster” on page 133.

If you used native OS accounts with the earlier versions of VCS, you must create new VCS accounts.

See [“Creating new VCS accounts if you used native OS accounts”](#) on page 134.

Review the strategies to determine which nodes to target, and which to leave running.

See [“About minimal downtime upgrade”](#) on page 119.

Note: If you did not configure I/O fencing for your cluster, ignore all commands related to I/O fencing in the procedure.

Performing the pre-upgrade tasks

Perform the following procedures before you upgrade VCS.

To perform pre-upgrade tasks

- 1 Select a node (or a group of nodes) in the cluster to upgrade.
- 2 Log in as superuser on one of the nodes where you want to upgrade VCS.
- 3 Mount the software disc.
See [“Mounting the product disc”](#) on page 48.
- 4 Verify that `/opt/VRTS/bin` is set in your `PATH` environment variable to execute all product commands.
- 5 Switch the service groups to the nodes where you plan to upgrade VCS later.

```
hagrp -switch service_group -to nodename
```

If the node you want to upgrade is part of a parallel group, then offline the group on that node.

```
hagrp -offline service_group -sys local_node
```

- 6 Make the VCS configuration writable. On a node that you want to upgrade, type:

```
haconf -makerw
```

- 7 Freeze the service groups. On each node that you selected to upgrade, type:

```
hasys -freeze -persistent nodename
```

- 8 Save the VCS configuration.

```
haconf -dump -makero
```

To stop VCS and its components

- 1 Stop any application agents that are installed on the VxVM disk (example, NBU agent). Perform the following steps to stop the application agents.

- Stop the application agents that are installed on VxVM disk on all the systems.

```
haagent -stop agentname -sys sysname
```

- Ensure that the agent processes are not running.

```
ps -ef | grep Agent
```

This command does not list any processes in the VxVM installation directory.

2 Stop VCS and its components.

```
/etc/init.d/vcs stop
/etc/init.d/vxfen stop
/etc/init.d/gab stop
/etc/init.d/llt stop
```

To upgrade Linux operating system

- ◆ If necessary, upgrade your operating system on the nodes:
 - Upgrade your operating system to one of the supported kernel versions. See [“Supported operating systems”](#) on page 23. See [“VCS supported upgrade paths”](#) on page 115.
 - Shut down and reboot the nodes.

Upgrading VCS

Perform the following procedure to upgrade VCS to version 5.0 MP3. Depending on the upgrade path, you must either use the `installvcs` program or the `installmp` program.

Note: If you use the `installvcs` program to upgrade VCS, make sure that VCS is running on the nodes where you perform the upgrade task.

See [“VCS supported upgrade paths”](#) on page 115.

To upgrade VCS on the target nodes

- 1 Insert the VCS 5.0 MP3 software disc into the disc drive of one of the nodes.
- 2 Mount the disc on a suitable mount point.
- 3 Depending on the upgrade path, navigate to the folder that contains either the `installvcs` program or the `installmp` program.

4 Upgrade to VCS 5.0 MP3.

Warning: Make sure to enter the node names where you want to upgrade VCS. Otherwise, the upgrade program chooses all the nodes in the cluster.

If you upgrade VCS from 4.1 or 4.1 MP versions, then enter the following command:

```
./installmp [-rsh] node1 node2
```

If you upgrade VCS from 5.0 or 5.0 MP versions, then enter the following command:

```
./installvcs [-rsh] node1 node2
```

- 5 If you had added custom type definitions in the original types.cf file, you must add them to the new types.cf file.
- 6 After the initial system checks and the requirements checks are complete, press Return to start upgrading the packages.
- 7 When the installation is complete, note the locations of the summary, log, and response files indicated by the installer.

Do not start the GAB and LLT processes. Do not start any VCS processes at this time.
- 8 Edit the main.cf file to configure any new attributes.
- 9 Change the cluster ID in the /etc/llttab file.

Find the line containing “set-cluster” and change the cluster ID following this keyword. Make sure that the new cluster ID is unique within the LAN.

Performing the post-upgrade tasks

After upgrading VCS, perform the following procedure.

To perform the post-upgrade tasks

- 1 If the main.cf file does not contain “Frozen=1” for all group definitions, then edit the main.cf file to freeze all the groups.

Add the “Frozen = 1” line to all group definitions.

Example: If original group definition is

```
Group oracle_sg (
SystemList = { galaxy = 0, nebula = 1 }
AutoStartList = { galaxy, nebula }
```

The new group definition, after adding “Frozen = 1”, should be:

```
Group oracle_sg (
SystemList = { galaxy = 0, nebula = 1 }
AutoStartList = { galaxy, nebula }
Frozen = 1
```

- 2 Start VCS and its components. On each node, enter:

```
/etc/init.d/llt start
/etc/init.d/gab start
/sbin/gabconfig -cx
/etc/init.d/vxfen start
```

If you chose to upgrade a single node, you must edit the /etc/sysconfig/vcs file to set ONENODE=no before you start VCS.

```
/etc/init.d/vcs start
```

Completing the upgrade on all the nodes in the cluster

Perform the following procedure only when you reach the last node (or group of nodes) in the cluster.

To complete the upgrade on all the nodes in the cluster

- 1 On the last node (or group of nodes) that you want to upgrade, offline all the groups that are online:

```
hagr -offline service_group -sys nodename
```

- 2 On one of the nodes in the upgraded cluster, do the following:

- Make the configuration writable.

```
haconf -makerw
```

- Unfreeze all the services groups.

```
hagrp -unfreeze service_group -persistent
```

- Bring all the service groups online. For each service group, run the following command:

```
hagrp -online service_group -sys nodename
```

- Save the configuration:

```
haconf -dump -makero
```

3 Upgrade the last node (or group of nodes). Perform the following procedures:

- Perform the pre-upgrade on these nodes.
See [“Performing the pre-upgrade tasks”](#) on page 130.
- Upgrade VCS on these nodes.
See [“Upgrading VCS”](#) on page 131.
- Modify the `/etc/llttab` file and provide the cluster ID for the new cluster.
- Start all VCS components on the last nodes that were upgraded.

```
/etc/init.d/llt start
/etc/init.d/gab start
/sbin/gabconfig -cx
/etc/init.d/vxfen start
```

If you chose to upgrade a single node, you must edit the `/etc/sysconfig/vcs` file to set `ONENODE=no` before you start VCS.

```
/etc/init.d/vcs start
```

Creating new VCS accounts if you used native OS accounts

VCS has deprecated the `AllowNativeCliUsers` attribute. To use native OS accounts with VCS, use the `halogin` command. After you run the `halogin` command, VCS encrypts and stores your VCS credentials in your home directory for a specific time period. After you run the `halogin` command, you need not authenticate yourself every time you run a VCS command. In secure clusters, the command also sets up a trust relationship and retrieves a certificate from an authentication broker.

See the *Veritas Cluster Server User's Guide* for information on assigning user privileges to OS user groups for clusters running in secure mode and clusters not running in secure mode.

Perform the following procedure if you used the AllowNativeCliUsers attribute.

Ensure that each native user running VCS commands has a home directory on the system from which the user runs VCS commands.

To set up VCS authentication for clusters running in secure mode

- 1 Set the configuration (main.cf) mode to read/write.

```
# haconf -makerw
```

- 2 Assign proper privileges to the OS users or user groups.

Each OS user must perform steps 3 and 4.

- 3 If the user executes VCS commands from a remote host, set the following environment variables:

- VCS_HOST—Name of the VCS node on which you run commands. You may specify the virtual IP address associated with the cluster.
- VCS_DOMAIN—Name of the VxSS domain to which the user belongs.
- VCS_DOMAINTYPE—Type of VxSS domain: unixpwd, nt, nis, nisplus, or vx.

- 4 Run the `halogin` command:

```
$ halogin vcsusername password
```

To set up VCS authentication for clusters not running in secure mode

- 1 Set the configuration (main.cf) mode to read/write.

```
# haconf -makerw
```

- 2 Create VCS user accounts for all users and assign privileges to these users.

- 3 Each VCS user must run the `halogin` command:

```
$ halogin vcsusername password
```

Upgrading the VCS agents

The `installvcs` program does not upgrade the VCS agents for DB2, Oracle, and Sybase. If previous versions of these agents are installed on your cluster, you must upgrade these agents manually.

See the Installation and Configuration Guide for the agent that you want to upgrade.

The VCS agents are backward compatible. The agents that work with VCS 4.1 also work with VCS 5.0 MP3.

See *Veritas Cluster Server Release Notes* for supported versions of the agents with VCS 5.0 MP3.

Upgrading the Cluster Manager (Java Console)

This release includes updates for Cluster Manager (Java Console).

To upgrade the Java Console on a Windows client

- 1 Stop Cluster Manager if it is running.
- 2 Remove Cluster Manager from the system.
- 3 Insert the software disc into a drive on your Windows system.
- 4 Start the installer from the following path:
`\windows\VCSWindowsInstallers\ClusterManager\EN\setup.exe`
- 5 Follow the wizard instructions to complete the installation.

Upgrading the VCS Simulator

This release includes updates for VCS Simulator.

To upgrade VCS Simulator on a Windows client

- 1 Stop all instances of VCS Simulator.
- 2 Stop VCS Simulator, if it is running.
- 3 Remove VCS Simulator from the system.
- 4 Insert the software disc into a drive on your Windows system.
- 5 Start the installer from the following path:
`\windows\VCSWindowsInstallers\Simulator\EN\vrtsvcssim.msi`
- 6 Follow the wizard instructions to complete the installation.

Adding and removing cluster nodes

This chapter includes 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 you install VCS and create a cluster, you can add and remove nodes from the cluster. You can create a cluster of up to 32 nodes.

Adding a node to a cluster

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

See [“Hardware requirements”](#) on page 21.

[Table 8-1](#) specifies the tasks that are involved in adding a cluster. The example demonstrates how to add a node east to already existing nodes, north and south.

Table 8-1 Tasks that are involved in adding a node to a cluster

Task	Reference
Set up the hardware.	See “Setting up the hardware” on page 138.


Table 8-1 Tasks that are involved in adding a node to a cluster *(continued)*

Task	Reference
Install the software manually.	See “Preparing for a manual installation when adding a node” on page 140. See “Installing VCS RPMs for a manual installation” on page 140.
Add a license key.	See “Adding a license key” on page 142.
For a cluster that is running in secure mode, verify the existing security setup on the node.	See “Verifying the existing security setup on the node” on page 143.
Configure LLT and GAB.	See “Configuring LLT and GAB” on page 145.
Add the node to the existing cluster.	See “Adding the node to the existing cluster” on page 146.
Start VCS and verify the cluster.	See “Starting VCS and verifying the cluster” on page 147.

Setting up the hardware

[Figure 8-1](#) shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

Figure 8-1 Adding a node to a three-node cluster using two independent hubs



To set up the hardware

1 Connect the VCS private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a two-node cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a two-node cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

Figure 8-1 illustrates a new node being added to an existing three-node cluster using two independent hubs.

2 Connect the system to the shared storage, if required.

Preparing for a manual installation when adding a node

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

See [“Mounting the product disc”](#) on page 48.

To prepare for installation

```
◆ # cd /mnt/cdrom/dist_arch/cluster_server/rpms
```

Installing VCS RPMs for a manual installation

VCS has both required and optional RPMs. Install the required RPMs first. All RPMs are installed in the /opt directory.

When you select the optional RPMs, review the following information:

- Symantec recommends that you install the RPMs for VCS manual pages (VRTSvcsmn).
- The I/O fencing RPM (VCSvxfen) can be used only with the shared disks that support SCSI-3 Persistent Reservations (PR). See the *Veritas Cluster Server User's Guide* for a conceptual description of I/O fencing. You need to test shared storage for SCSI-3 PR and to implement I/O fencing.
See [“About setting up I/O fencing”](#) on page 87.
- The VCS configuration wizard (VRTScscw) RPM includes wizards for the installation and configuration of Veritas products that require VCS configuration.
- To use the Java Console with VCS Simulator, you must install the VRTScssim and VRTScscm RPMs.

Use this procedure if you install VCS for the first time. Make sure the system does not have any of the VCS RPMs already installed. If VCS is already installed, either remove the RPMs before you perform this procedure or upgrade VCS on the new node.

See [“About VCS 5.0 MP3 upgrade”](#) on page 115.

Perform the steps to install VCS RPMs on each node in the cluster.

To install VCS RPMs on a node

- 1 Install the required VCS RPMs in the order shown. Do not install any RPMs already installed on the system. Pay special attention to operating system distribution and architecture. The following commands use variables such as *dist* for the supported Linux distribution, *arch* for supported Linux architecture, and *version* for the RPM version of a specific distribution.

■ RHEL, required RPMs:

```
# rpm -i VRTSatClient-4.3.34.4-4.i386.rpm
# rpm -i VRTSatServer-4.3.34.4-4.i386.rpm
# rpm -i VRTSicsco-1.3.28.0-0.i386.rpm
# rpm -i VRTSspb-1.3.28.0-0.i386.rpm
# rpm -i VRTSperl-5.8.8.0-dist.arch.rpm
# rpm -i VRTSspt-5.0.00.2-GA.noarch.rpm
# rpm -i VRTSvlic-3.02.33.5500-0.arch.rpm
# rpm -i VRTSllt-5.0.30.00-MP3_dist.arch.rpm
# rpm -i VRTSgab-5.0.30.00-MP3_dist.arch.rpm
# rpm -i VRTSvxfen-5.0.30.00-MP3_dist.arch.rpm
# rpm -i VRTSvcs-5.0.30.00-MP3_dist.i686.rpm
# rpm -i VRTSvcsmg-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTSacclib-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTSvcsag-5.0.30.00-MP3_dist.i686.rpm
# rpm -i VRTSvcsdr-5.0.30.00-MP3_dist.arch.rpm
# rpm -i VRTSjre15-1.5.3.5-5.i386.rpm
# rpm -i VRTScscw-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTSweb-5.0.1-GA4_GENERIC.noarch.rpm
# rpm -i VRTScutil-version_GENERIC.noarch.rpm
```

Where *dist* is RHEL4 or RHEL5 and *arch* is i686 or x86_64.

The value of the variable *version* for RHEL4 is 5.0-GA and for RHEL5 is 5.0-MP3.

■ SLES, required RPMs:

```
# rpm -i VRTSatClient-4.3.34.4-4.i386.rpm
# rpm -i VRTSatServer-4.3.34.4-4.i386.rpm
# rpm -i VRTSicsco-1.3.28.0-0.i386.rpm
# rpm -i VRTSspb-1.3.28.0-0.i386.rpm
# rpm -i VRTSperl-5.8.8.0-dist.arch.rpm
# rpm -i VRTSspt-5.0.00.2-GA.noarch.rpm
# rpm -i VRTSvlic-3.02.33.5500-0.arch.rpm
# rpm -i VRTSllt-5.0.30.00-MP3_dist.arch.rpm
# rpm -i VRTSgab-5.0.30.00-MP3_dist.arch.rpm
# rpm -i VRTSvxfen-5.0.30.00-MP3_dist.arch.rpm
# rpm -i VRTSvcs-5.0.30.00-MP3_dist.i586.rpm
# rpm -i VRTSvcsmg-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTSacclib-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTSvcsag-5.0.30.00-MP3_dist.i586.rpm
# rpm -i VRTSvcsdr-5.0.30.00-MP3_dist.arch.rpm
```

```
# rpm -i VRTSjre15-1.5.3.5-5.i386.rpm
# rpm -i VRTScscw-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTSweb-5.0.1-GA4_GENERIC.noarch.rpm
# rpm -i VRTScutil-version_GENERIC.noarch.rpm
```

Where *dist* is SLES9 or SLES10 and *arch* is i586 or x86_64.

The value of the variable *version* for SLES9 is 5.0-GA and for SLES10 is 5.0-MP3.

- 2 Install the optional RPMs, in the order shown. Omit those that you do not want to install.

■ RHEL, optional RPMs:

```
# rpm -i VRTSvcsmn-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTScscm-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTScssim-5.0.30.00-MP3_dist.i686.rpm
# rpm -i VRTScmcs-5.0.30.00-50MP3_dist.i686.rpm
# rpm -i VRTScmccc-5.0.30.00-50MP3_dist.i686.rpm
```

Where *dist* is RHEL4 or RHEL5.

■ SLES, optional RPMs:

```
# rpm -i VRTSvcsmn-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTScscm-5.0.30.00-MP3_GENERIC.noarch.rpm
# rpm -i VRTScssim-5.0.30.00-MP3_dist.i586.rpm
# rpm -i VRTScmcs-5.0.30.00-50MP3_dist.i586.rpm
# rpm -i VRTScmccc-5.0.30.00-50MP3_dist.i586.rpm
```

Where *dist* is SLES9 or SLES10.

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.

Verifying the existing security setup on the node

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

See [“Configuring LLT and GAB”](#) on page 145.

To verify the existing security setup on the node

- 1 If node east is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Else, proceed to configuring the authentication broker on node east.

See [“Configuring the authentication broker on node east”](#) on page 144.

- 2 Find out the root broker to which the node east belongs using the following command.

```
# vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
-k "BrokerName"
```

- 3 If the node east already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.

See [“Setting up VCS related security configuration”](#) on page 144.

- 4 If the node east belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node east.

- Kill `/opt/VRTSat/bin/vxatd` process.
- Remove the credential that RB2 has given to AB on node east.

```
# vssat deletetcred --domain type:domainname \
--prplname prplname
```

Configuring the authentication broker on node east

Configure a new authentication broker (AB) on node east. This AB belongs to root broker RB1.

To configure the authentication broker on node east

- 1 Create a principal for node east on root broker RB1. Execute the following command on root broker RB1.

```
# vssat addprpl --pdrtype root --domain domainname \
--prplname prplname --password password \
--prpltype service
```

- 2 Ensure that there is no clock skew between the times on node east and RB1.
- 3 Copy the /opt/VRTSat/bin/root_hash file from RB1 to node east.
- 4 Configure AB on node east to talk to RB1.

```
# vxatd -o -a -n prplname -p password -x vx -y domainname -q \
rootbroker -z 2821 -h roothash_file_path
```

- 5 Verify that AB is configured properly.

```
# vssat showbrokermode
```

The command should return 1, indicating the mode to be AB.

Setting up VCS related security configuration

Perform the following steps to configure VCS related security settings.

Setting up VCS related security configuration

- 1 Start /opt/VRTSat/bin/vxatd process.
- 2 Create HA_SERVICES domain for VCS.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

- 3 Add VCS and webserver principal to AB on node east.

```
# vssat addprpl --pdrtype ab --domain HA_SERVICES --prplname
webserver_VCS_prplname --password new_password --prpltype
service --can_proxy
```

- 4 Create /etc/VRTSvcs/conf/config/.secure file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

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 add 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.

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

```
set-node east
set-cluster 2
link eth1 eth1 - ether - -
link eth2 eth2 - 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
```

The file on the new node should be the same. Symantec recommends that you 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
```

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
```

The `-n` flag indicates to VCS the number of nodes that must 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 110.

- 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:

```
# hstop -sys east
```

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

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

- 5 Start VCS on the new node:

```
# hstart
```

- 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:

```
# hstart
```

- 2 Run the GAB configuration command on each node to verify that port a and port h include the new node in the membership:

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

Removing a node from a cluster

Table 8-2 specifies the tasks that are 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 8-2 Tasks that are involved in removing a node

Task	Reference
■ Back up the configuration file. ■ Check the status of the nodes and the service groups.	See “ Verifying the status of nodes and service groups ” on page 148.
■ Switch or remove any VCS service groups on the node departing the cluster. ■ Delete the node from VCS configuration.	See “ Deleting the departing node from VCS configuration ” on page 149.
Modify the llthosts and gabtab files to reflect the change.	See “ Modifying configuration files on each remaining node ” on page 152.
For a cluster that is running in a secure mode, remove the security credentials from the leaving node.	See “ Removing security credentials from the leaving node ” on page 152.
On the node departing the cluster: ■ 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.	See “ Unloading LLT and GAB and removing VCS on the departing node ” on page 153.

Verifying 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 departing node from VCS configuration

Before you remove a node from the cluster you need to identify the service groups that run on the node.

You then need to perform the following actions:

- Remove the service groups that other service groups depend on, or
- Switch the service groups to another node that other service groups depend on.

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—if it is a parent to service groups on other nodes—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. Make sure that any service group that you want to fail over is 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 the service groups that run only on the departing node, delete the resources from the group before you delete the group.

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

- 8 Delete the service group that is 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`. 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, make sure that *N* is not greater than the actual number of nodes in the cluster. When *N* is greater than the number of nodes, 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/llhosts` 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
```

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node C. Perform the following steps.

To remove the security credentials

- 1 Kill `/opt/VRTSat/bin/vxatd` process.
- 2 Remove the root credentials on node C.

```
# vssat deletecred --domain type:domainname --prplname prplname
```

Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node that is 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. Start by removing the following RPMs, which may have been optionally installed, in the order shown:

```
# rpm -e VRTScmccc
# rpm -e VRTScmcs
# rpm -e VRTScssim
# rpm -e VRTScscm
# rpm -e VRTSvcsmn
# rpm -e VRTScutil
# rpm -e VRTSweb
# rpm -e VRTScscw
# rpm -e VRTSjre15
# rpm -e VRTSjre
# rpm -e VRTSvcldr
# rpm -e VRTSvcsg
# rpm -e VRTSaclib
# rpm -e VRTSvcsmg
# rpm -e VRTSvc
# rpm -e VRTSvxfen
# rpm -e VRTSgab
# rpm -e VRTSllt
# rpm -e VRTSvlic
# rpm -e VRTSspt
# rpm -e VRTSsmf
# rpm -e VRTSperl
# rpm -e VRTSspb
# rpm -e VRTSicsco
# rpm -e VRTSatServer
```

```
# rpm -e VRTSatClient  
# rpm -e SYMCIma
```

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 includes the following topics:

- [About installing VCS on a single node](#)
- [Creating a single-node cluster using the installer 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 5.0MP3 on a single node. You can subsequently add another node to the single-node cluster to form a multinode 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 installer program or you can add it manually.

See [“Creating a single-node cluster using the installer program”](#) on page 155.

See [“Creating a single-node cluster manually”](#) on page 157.

Creating a single-node cluster using the installer program

[Table 9-1](#) specifies the tasks that are involved to install VCS on a single node using the installer program.

Table 9-1 Tasks to create a single-node cluster using the installer

Task	Reference
Prepare for installation.	See “Preparing for a single node installation” on page 156.
Install the VCS software on the system using the installer.	See “Starting the installer for the single node cluster” on page 156.

Preparing for a single node installation

You can use the installer program to install a cluster on a single system for either of the two following purposes:

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

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

For more information about LLT and GAB:

See [“About LLT and GAB”](#) on page 15.

Starting the installer for the single node cluster

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

See [“Starting the software installation”](#) on page 63.

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

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

Enter a single system name. The installer 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 stand-alone 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 65.

Creating a single-node cluster manually

[Table 9-2](#) specifies the tasks that you need to perform to install VCS on a single node.

Table 9-2 Tasks to create a single-node cluster manually

Task	Reference
Set the PATH variable	See “Setting the path variable for a manual single node installation” on page 157.
Install the VCS software manually and add a license key	See “Installing the VCS software manually on a single node” on page 157.
Remove any LLT or GAB configuration files and rename LLT and GAB startup files. A single-node cluster does not require the node-to-node communication service, LLT, or the membership communication service, GAB.	See “Renaming the LLT and GAB startup files” on page 158.
Modify the VCS startup file for single-node operation.	See “Modifying the startup files” on page 158.
Create and modify the VCS configuration files.	
Start VCS and verify single-node operation.	See “Verifying single-node operation” on page 158.

Setting the path variable for a manual single node installation

Set the path variable.

See [“Setting the PATH variable”](#) on page 46.

Installing the VCS software manually on a single node

Install the VCS 5.0MP3 RPMs manually and install the license key.

Refer to the following sections:

- See [“Preparing for a manual installation when adding a node”](#) on page 140.
- See [“Installing VCS RPMs for a manual installation”](#) on page 140.
- See [“Adding a license key”](#) on page 142.

Renaming the LLT and GAB startup files

You may need the LLT and GAB startup files 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
```

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

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. The node that is to join Node A to form a multiple-node cluster is Node B.

Table 9-3 specifies the activities that you need to perform to add nodes to a single-node cluster.

Table 9-3 Tasks to add a node to a single-node cluster

Task	Reference
Set up Node B to be compatible with Node A.	See “Setting up a node to join the single-node cluster” on page 160.
■ 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.	See “Installing and configuring Ethernet cards for private network” on page 160.
Connect both nodes to shared storage.	See “Configuring the shared storage” on page 161.
■ Bring up VCS on Node A. ■ Edit the configuration file. ■ Edit the startup scripts.	See “Bringing up the existing node” on page 161.
If necessary, install VCS on Node B and add a license key. Make sure Node B is running the same version of VCS as the version on Node A.	See “Installing the VCS software manually when adding a node to a single node cluster” on page 162.
Edit the configuration files on Node B.	See “Configuring LLT and GAB” on page 145.
Start LLT and GAB on Node B.	See “Starting LLT and GAB” on page 165.
■ Start LLT and GAB on Node A. ■ Restart VCS on Node A. ■ Modify service groups for two nodes.	See “Reconfiguring VCS on the existing node” on page 165.
■ Start VCS on Node B. ■ Verify the two-node cluster.	See “Verifying configuration on both nodes” on page 166.

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

- 1 Do one of the following tasks:
 - If VCS is not currently running on Node B, proceed to step [2](#).
 - If the node you plan to add as Node B is currently part of an existing cluster, remove the node from the cluster. After you remove the node from the cluster, remove the VCS RPMs and configuration files.
See [“Removing a node from a cluster”](#) on page 148.
 - If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS.
See [“Uninstalling VCS 5.0MP3”](#) on page 170.
 - If you renamed the LLT and GAB startup files, remove them.
See [“Renaming the LLT and GAB startup files”](#) on page 158.
- 2 If necessary, install VxVM and VxFS.
See [“Installing VxVM or VxFS if necessary”](#) on page 160.

Installing VxVM or VxFS if necessary

If you have either VxVM or VxFS with the cluster option installed on the existing node, install the same version on the new node.

Refer to the appropriate documentation for VxVM and VxFS to verify the versions of the installed products. Make sure the same version runs on all nodes where you want to use shared storage.

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, use two separate NICs on each node. The two NICs provide redundancy for heartbeating.

See [“Setting up the private network”](#) on page 38.

To install and configure Ethernet cards for private network

- 1 Shut down VCS on Node A.

```
# hstop -local
```

- 2 Shut down the node to get to the OK prompt:

```
# sync;sync;init 0
```

- 3 Install the Ethernet card on Node A.

If you want to use aggregated interface to set up private network, configure aggregated interface.

- 4 Install the Ethernet card on Node B.

If you want to use aggregated interface to set up private network, configure aggregated interface.

- 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. Configure VxVM on Node B and reboot the node when you are prompted.

See [“Configuring the shared storage”](#) on page 161.

See [“Setting up shared storage: Fiber Channel”](#) on page 45.

Bringing up the existing node

Bring up the node.

To bring up the 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 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 software manually when adding a node to a single node cluster

Install the VCS 5.0MP3 RPMs manually and install the license key.

Refer to the following sections:

- See [“Preparing for a manual installation when adding a node”](#) on page 140.
- See [“Installing VCS RPMs for a manual installation”](#) on page 140.
- See [“Adding a license key”](#) on page 142.

Configuring LLT

VCS uses the Low Latency Transport (LLT) protocol for all cluster communications as a high-performance, low-latency replacement for the IP stack. LLT has two major functions.

It handles the following tasks:

- Traffic distribution
- Heartbeat traffic

Configured as described in the following sections.

Setting up /etc/llthosts

The file `llthosts(4M)` is a database. This file contains 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 the 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 164.

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

```
set-node north
set-cluster 2
link eth1 eth1 - ether - -
link eth2 eth2 - ether - -
```

The first line must identify the system where the file exists. In the preceding example, the value for `set-node` can be: `north`, `0`, or the file name `/etc/nodename`. The file needs to contain the name of the system (`north` in this example) to use these choices. 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 9-4](#) describes the LLT directives for LLT setup.

Table 9-4 LLT directives

Directive	Description
<code>set-node</code>	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 in the <code>/etc/llthosts</code> file. Note that LLT fails to operate if any systems share the same ID.
<code>link</code>	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(1M)</code> output to identify the link. It may also be used in <code>llttab</code> to set optional static MAC addresses. 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 0xCAFE. 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.
<code>set-cluster</code>	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.
<code>link-lowpri</code>	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. It also enables VCS communication, and broadcasts heartbeats to monitor each network connection.

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 GAB when adding a node to a single node cluster

VCS uses the Group Membership Services/Atomic Broadcast (GAB) protocol for cluster membership and reliable cluster communications. GAB has two major functions.

It handles the following tasks:

- Cluster membership
- Cluster communications

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 an `/etc/gabtab` file:

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

The `-c` option configures the driver for use. The `-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

Reconfigure VCS on the existing nodes.

To reconfigure VCS on existing nodes

- 1 On Node A, create the files `/etc/llttab`, `/etc/llthosts`, and `/etc/gabtab`. Use the files that are 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.

```
# hagr -list
```

- For each service group that is listed, add the node.

```
# hagr -modify group SystemList -add sysB 1
```

Verifying configuration on both nodes

Verify the configuration for the nodes.

To verify the nodes' configuration

- 1 On Node B, check the cluster membership.

```
# gabconfig -a
```

- 2 Start the VCS on Node B.

```
# hstart
```

- 3 Verify that VCS is up on both nodes.

```
# hastatus
```

- 4 List the service groups.

```
# hagrps -list
```

- 5 Unfreeze the service groups.

```
# hagrps -unfreeze group -persistent
```

- 6 Implement the new two-node configuration.

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


Uninstalling VCS

This chapter includes the following topics:

- [About the uninstallvcs program](#)
- [Prerequisites for using the uninstallvcs program](#)
- [Uninstalling VCS 5.0MP3](#)

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 enterprise agents, but offers uninstallation if proper RPMs dependencies on `VRTSvcs` are found.

If `uninstallvcs` program does not remove an enterprise agent, see the documentation for the specific enterprise agent for instructions on how to remove it.

Prerequisites for using the uninstallvcs program

Review the following prerequisites before you uninstall VCS:

- Before you remove VCS from any node in the cluster, shut down the applications that depend on VCS. For example, applications such as Java Console or any high availability agents for VCS.
- Before you remove VCS from fewer than all nodes in a cluster, stop the service groups on the nodes from which you uninstall VCS. You must also reconfigure VCS on the remaining nodes.

See [“About adding and removing nodes”](#) on page 137.

- If you have manually edited any of the VCS configuration files, you need to reformat them.
See [“Reformatting VCS configuration files on a stopped cluster”](#) on page 60.

Uninstalling VCS 5.0MP3

You must meet the following conditions to use the `uninstallvcs` program to uninstall VCS on all nodes in the cluster at one time:

- Make sure that the communication exists between systems. By default, the uninstaller uses `ssh`.
- Make sure you can execute `ssh` or `rsh` commands as superuser on all nodes in the cluster.
- Make sure that the `ssh` or `rsh` is configured to operate without requests for passwords or passphrases.

If you cannot meet the prerequisites, then you must run the `uninstallvcs` program on each node in the cluster.

The example demonstrates how to uninstall VCS using the `uninstallvcs` program. The `uninstallvcs` program uninstalls VCS on two nodes: north and south. The example procedure uninstalls VCS from all nodes in the cluster.

Removing VCS 5.0MP3 RPMs

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

To uninstall VCS

- 1 Log in as superuser from the node where you want to uninstall VCS.
- 2 Start `uninstallvcs` program.

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

The program specifies the directory where the logs are created. The program displays a copyright notice and 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: galaxy nebula
Service Groups: ClusterService groupA groupB
```

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

Select one of the following:

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

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

- 4 If the `uninstallvcs` program prompts, enter a list of nodes from which you want to uninstall VCS.

The `uninstallvcs` program prompts this information in one of the following conditions:

- You enter `n`.
- The program finds no VCS configuration files on the local node.

- 5 Review the output as the `uninstallvcs` program continues to do the following:

- Verifies the communication between systems
- Checks the installations on each system to determine the RPMs to be uninstalled

- 6 If RPMs, such as enterprise 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.

- 7 Review the uninstaller report after the verification.
- 8 Press Enter to uninstall the VCS RPMs.


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

Running `uninstallvcs` from the VCS 5.0MP3 disc

You may need to use the `uninstallvcs` program on the VCS 5.0 MP3 disc in one of the following cases:

- You need to uninstall VCS after an incomplete installation.
- The `uninstallvcs` program is not available in `/opt/VRTS/install`.

Advanced VCS installation topics

This appendix includes the following topics:

- [Using the UDP layer for LLT](#)
- [Performing automated VCS installations](#)
- [Installing VCS with a response file where ssh or rsh are disabled](#)

Using the UDP layer for LLT

VCS 5.0MP3 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

Note: LLT over UDP is not supported on IPv6.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

LLT over UDP is slower than LLT over Ethernet. Only use LLT over UDP when the hardware configuration makes it necessary.

Configuring LLT over UDP

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on different physical networks.
If the LLT private links are not on different physical networks, then make sure that the links are on separate subnets. Set the broadcast address in `/etc/llttab` explicitly depending on the subnet for each link.
See [“Broadcast address in the `/etc/llttab` file”](#) on page 174.
- Make sure that each NIC has an IP address that is configured before configuring LLT.
- Make sure the IP addresses in the `/etc/llttab` files are consistent with the IP addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
See [“Selecting UDP ports”](#) on page 176.
- Set the broadcast address correctly for direct-attached (non-routed) links.
- For the links that cross an IP router, disable broadcast features and specify the IP address of each link manually in the `/etc/llttab` file.
See [“Sample configuration: links crossing IP routers”](#) on page 179.

Broadcast address in the `/etc/llttab` file

The broadcast address is set explicitly for each link in the following example.

```
# cat /etc/llttab
set-node Node0
set-cluster 1

link link1 udp - udp 50000 - 10.20.30.1 10.20.30.255
link link2 udp - udp 50001 - 10.20.31.1 10.20.31.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

```
# ifconfig
eth2 Link encap:Ethernet HWaddr 00:04:23:AC:2B:E4
    inet addr:10.20.30.1 Bcast:10.20.30.255 Mask:255.255.255.0
eth3 Link encap:Ethernet HWaddr 00:04:23:AC:2B:E5
    inet addr:10.20.31.1 Bcast:10.20.31.255 Mask:255.255.255.0
```

The link command in the `/etc/llttab` file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See [“Sample configuration: direct-attached links”](#) on page 177.

- See “[Sample configuration: links crossing IP routers](#)” on page 179.

Note that some of the fields in [Table A-1](#) on page 175, differ from the command for standard LLT links.

[Table A-1](#) describes the fields of the link command that are shown in the /etc/llttab file examples.

Table A-1 Field description for link command in /etc/llttab

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example udp. A place holder string. On other unix platforms like Solaris or HP, this entry points to a device file (for example, /dev/udp). Linux does not have devices for protocols. So this field is ignored.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “ Selecting UDP ports ” on page 176.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command displays the current value.
<i>IP address</i>	IP address of the link on the local node.
<i>bcast-address</i>	■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the /etc/llttab file

The `set-addr` command in the /etc/llttab file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 179.

[Table A-2](#) describes the fields of the set-addr command.

Table A-2 Field description for set-addr command in /etc/llttab

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.
<i>link tag-name</i>	The string that LLT uses to identify the link; for example link1, link2,....
<i>address</i>	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address      State
udp        0      0 *:32768          *: *
udp        0      0 *:956            *: *
udp        0      0 *:tftp           *: *
udp        0      0 *:sunrpc         *: *
udp        0      0 *:ipp            *: *
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the /etc/services file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use.

For example, with the following interfaces:

- For first network interface

```
IP address=192.168.30.1, Broadcast address=192.168.30.255,  
Netmask=255.255.255.0
```

■ For second network interface

```
IP address=192.168.31.1, Broadcast address=192.168.31.255,  
Netmask=Mask:255.255.255.0
```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.


An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100  
  
link link1 udp - udp 50000 - 192.168.30.1 192.168.30.255  
link link2 udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: direct-attached links

[Figure A-1](#) depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure A-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 resembles:


```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
```

```
link link1 udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

Figure A-2 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure A-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1

link link1 udp - udp 50000 - 192.1.3.1 -
link link2 udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr      0 link1 192.1.1.1
set-addr      0 link2 192.1.2.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
```

```
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3
```

```
#disable LLT broadcasts
set-bcasthb   0
set-arp       0
```

The `/etc/llttab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 udp - udp 50000 - 192.1.1.1 -
link link2 udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr      1 link1 192.1.3.1
set-addr      1 link2 192.1.4.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb   0
set-arp       0
```

Performing automated VCS installations

Using `installvcs` program with the `-responsefile` option is useful not only for installing and configuring VCS within a secure environment. This option is also useful 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.

When the systems are set up and meet the requirements for installation, you can perform an unattended installation. You perform the 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 in the response file

The syntax of the Perl statements that are included in the response file varies. It can depend on whether the variables require scalar or list values.

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 that installvcs creates after the example VCS installation. The file 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 that are required for installation.

See [“Response file variable definitions”](#) on page 182.

```
#
# 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}="758a33dbd6fae716...3deb54e562fe98";
```

```
$CPI::CFG{CMC_SERVICE_PASSWORD}="U2FsdVkXl8v...n0hTSWwodThc+rX";
$CPI::CFG{ENCRYPTED}="U2FsdGVkXl+k2DHcnW7b6...ghdh+zw4G0WFIJA=";
$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}="eth0";
$CPI::CFG{VCS_CSGVIP}="10.10.12.1";
$CPI::CFG{VCS_LLTLINK1}{east}="eth1";
$CPI::CFG{VCS_LLTLINK1}{west}="eth1";
$CPI::CFG{VCS_LLTLINK2}{east}="eth2";
$CPI::CFG{VCS_LLTLINK2}{west}="eth2";

$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 A-3 Response file variables

Variable	Description
\$CPI::CFG{OPT}{INSTALL}	Installs and configures VCS. List or scalar: scalar Optional or required: required
\$CPI::CFG{OPT}{INSTALLONLY}	Installs VCS RPMs. Configuration can be performed at a later time using the <code>-configure</code> option. List or scalar: scalar Optional or required: optional

Table A-3 Response file variables (*continued*)

Variable	Description
\$CPI::CFG{SYSTEMS}	List of systems on which the product is to be installed, uninstalled, or configured. List or scalar: list Optional or required: required
\$CPI::CFG{SYSTEMSCFG}	List of systems to be recognized in configuration if secure environment prevents all systems from being installed at once. List or scalar: list Optional or required: optional
\$CPI::CFG{UPI}	Defines the product to be installed, uninstalled, or configured. List or scalar: scalar Optional or required: required
\$CPI::CFG{OPT}{KEYFILE}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
\$CPI::CFG{OPT}{LICENSE}	Licenses VCS only. List or scalar: scalar Optional or required: optional
\$CPI::CFG{OPT}{NOLIC}	Installs the product without any license. List or scalar: scalar Optional or required: optional
\$CPI::CFG{AT_ROOTDOMAIN}	Defines the name of the system where the root broker is installed. List or scalar: list Optional or required: optional

Table A-3 Response file variables (*continued*)

Variable	Description
\$CPI::CFG{OPT}{PKGPATH}	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. List or scalar: scalar Optional or required: optional
\$CPI::CFG{OPT}{TMPPATH}	Defines the location where a working directory is created to store temporary files and the depots that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
\$CPI::CFG{OPT}{RSH}	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. List or scalar: scalar Optional or required: optional
\$CPI::CFG{DONOTINSTALL} {RPM}	Instructs the installation to not install the optional RPMs in the list. List or scalar: list Optional or required: optional
\$CPI::CFG{DONOTREMOVE} {RPM}	Instructs the uninstallation to not remove the optional RPMs in the list. List or scalar: list Optional or required: optional
\$CPI::CFG{VCS_CLUSTERNAME}	Defines the name of the cluster. List or scalar: scalar Optional or required: required
\$CPI::CFG{VCS_CLUSTERID}	An integer between 0 and 65535 that uniquely identifies the cluster. List or scalar: scalar Optional or required: required

Table A-3 Response file variables (*continued*)

Variable	Description
\$CPI::CFG{KEYS} {SYSTEM}	List of keys to be registered on the system. List or scalar: scalar Optional or required: optional
\$CPI::CFG{OPT_LOGPATH}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
\$CPI::CFG{CONFIGURE}	Performs the configuration if the RPMs are already installed using the <code>-installonly</code> option. List or scalar: scalar Optional or required: optional
\$CPI::CFG{VCS_LLTLINK#} {SYSTEM}	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. List or scalar: scalar Optional or required: required
\$CPI::CFG{VCS_LLTLINKLOWPRI} {SYSTEM}	Defines a low priority heartbeat link. Typically, LLTLINKLOWPRI is used on a public network link to provide an additional layer of communication. List or scalar: scalar Optional or required: optional
\$CPI::CFG{VCS_CSGNIC}	Defines the NIC for Cluster Management Console to use on a system. 'ALL' can be entered as a system value if the same NIC is used on all systems. List or scalar: scalar Optional or required: optional
\$CPI::CFG{CSGVIP}	Defines the virtual IP address that the Cluster Management Console uses. List or scalar: scalar Optional or required: optional

Table A-3 Response file variables (*continued*)

Variable	Description
\$CPI::CFG{VCS_CSGNETMASK}	Defines the Netmask of the virtual IP address that the Cluster Management Console uses. List or scalar: scalar Optional or required: optional
\$CPI::CFG{VCS_SMTPSERVER}	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. List or scalar: scalar Optional or required: optional
\$CPI::CFG{VCS_SMTPRECP}	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. List or scalar: list Optional or required: optional
\$CPI::CFG{VCS_SMTPRSEV}	Defines the 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. List or scalar: list Optional or required: optional
\$CPI::CFG{VCS_SNMPPORT}	Defines the SNMP trap daemon port (default=162). List or scalar: scalar Optional or required: optional
\$CPI::CFG{VCS_SNMPCONS}	List of SNMP console system names List or scalar: list Optional or required: optional

Table A-3 Response file variables (continued)

Variable	Description
\$CPI::CFG{VCS_SNMPSEV}	Defines the 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. List or scalar: list Optional or required: optional
\$CPI::CFG{VCS_USERENPW}	List of encoded passwords for users List or scalar: list Optional or required: optional
\$CPI::CFG{VCS_USERNAME}	List of names of users List or scalar: list Optional or required: optional
\$CPI::CFG{VCS_USERPRIV}	List of privileges for users List or scalar: list Optional or required: optional
\$CPI::CFG{OPT}{UNINSTALL}	List of systems where VCS must be uninstalled. List or scalar: scalar Optional or required: optional

Installing VCS with a response file where ssh or rsh are disabled

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, VCS creates a response file.

See [“Example response file”](#) on page 181.

The response file that the installvcs program generates contains descriptions and explanations of the variables and their values. You copy this file to the other systems in the cluster, and edit 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 63.

- 2 Review the output as the installer performs the initial system checks.
The installer 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 galaxy only  
and create a responsefile for systems nebula? [y,n,q] (y)
```

- 4 Enter all cluster information. Proceed with the installation and configuration tasks.

See [“Installing and configuring VCS 5.0MP3”](#) on page 61.

The `installvcs` program installs and configures VCS on systems where communication is possible.

- 5 After the installation is complete, review the installer report.

The installer stores the `installvcs-universaluniqueidentifier` response file in the `/opt/VRTS/install/logs/installvcs-universaluniqueidentifier/.response` directory where `universaluniqueidentifier` is a variable to uniquely identify the file.

- 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 Use 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, edit the response file.

For the variables 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, perform the following:

- Mount the product disc.

See [“Mounting the product disc”](#) on page 48.

- Start the software installation using the `installvcs -responsefile` option.

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

Where uui is the Universal Unique Identifier that the installer automatically assigned to the response file.

See [“Starting the software installation”](#) on page 63.

- 10 Repeat step 7 through step 9 until VCS has been installed on all nodes in the cluster.

Index

A

- about
 - global clusters 17
- adding
 - users 72
- adding node
 - to a one-node cluster 159
- attributes
 - UseFence 98

C

- cables
 - cross-over Ethernet 139
- cluster
 - creating a single-node cluster
 - installer 155
 - manual 157
 - four-node configuration 14
 - removing a node from 148
 - verifying 82
 - verifying operation 111
- Cluster Management Console 20
- Cluster Manager
 - installing Java Console 80
- cold start
 - running VCS 16
- commands
 - gabconfig 110, 165
 - hastart 147
 - hastatus 111
 - hasys 112
 - lltconfig 101
 - lltstat 108
 - vxdisksetup (initializing disks) 90
 - vxlicinst 83, 142
 - vxlicrep 82, 143
- communication channels 15
- communication disk 15
- configuring
 - GAB 165
 - hardware 21

- configuring (*continued*)
 - LLT
 - manual 163
 - private network 38
 - ssh 42
 - switches 38
- configuring VCS
 - adding users 72
 - event notification 72, 74
 - global clusters 75
 - overview 62
 - secure mode 70
 - starting 67
- controllers
 - private Ethernet 38
- coordinator disks
 - DMP devices 86
 - for I/O fencing 86
 - setting up 96

D

- data disks
 - for I/O fencing 86
- directives
 - LLT 164
- disk space
 - directories 21
 - language pack 21
 - required 21
- disks
 - adding and initializing 90
 - coordinator 96
 - testing with vxfcntlshdw 92
 - verifying node access 93
- documentation
 - accessing 84

E

- eprom
 - parameters 38
- Ethernet controllers 38, 139

F

fibre channel 21

G

GAB

- description 15
- manual configuration 165
- port membership information 110
- verifying 110

gabconfig command 110, 165

- a (verifying GAB) 110

gabtab file

- creating 165
- verifying after installation 101

global clusters 17

- configuration 75

H

hardware

- configuration 14
- configuring network and storage 21

hastart 147

hastatus -summary command 111

hasys -display command 112

hubs 38

- independent 139

I

I/O fencing

- checking disks 92
- setting up 95
- shared storage 92

installation

- required disk space 22

installing

- post 79
- required disk space 21
- Root Broker 31

installing and configuring VCS

- overview 62

installing VCS

- choosing depots 65
- choosing filesets 65
- choosing packages 65
- choosing RPMs 65
- licensing 65
- overview 62
- required information 52

installing VCS (*continued*)

- starting 63
- utilities 51

installvcs

- options 55

installvcs prompts

- b 56
- n 56
- y 56

J

Java Console

- installing 80
- installing on UNIX 80
- installing on Windows workstation 81

L

language packages

- disk space 21

license keys

- adding with vxlicinst 83, 142
- obtaining 37
- replacing demo key 83

licenses

- information about 82
- showing information 143

licensing commands

- vxlicinst 37
- vxlicrep 37
- vxlictest 37

licensing VCS 65

links

- private network 101

LLT

- description 15
- directives 164
- interconnects 47
- manual configuration 163
- verifying 108

LLT directives

- link 164
- link-lowpri 164
- set-cluster 164
- set-node 164

lltconfig command 101

llthosts file

- verifying after installation 101

lltstat command 108

- lfttab file
 - verifying after installation 101

M

- MAC addresses 38
- main.cf file
 - contents after installation 104
- MANPATH variable
 - setting 47
- manual installation
 - preparing 140
- media speed 47
 - optimizing 47
- membership information 110
- minimal downtime upgrade 119
 - example 120
- mounting
 - software disc 48

N

- network partition
 - preexisting 16
 - protecting against 14
- Network partitions
 - protecting against 15
- network switches 38
- NFS 13

O

- optimizing
 - media speed 47
- overview
 - VCS 13

P

- parameters
 - EEPROM 38
- PATH variable
 - setting 46
 - VCS commands 107
- persistent reservations
 - SCSI-3 43
- port a
 - membership 110
- port h
 - membership 110
- port membership information 110

- preparing
 - manual installation 140
- prerequisites
 - uninstalling 169
- private network
 - configuring 38

R

- RAM
 - installation requirement 21
- removing a system from a cluster 148
- remsh 64, 68
- requirements
 - Ethernet controllers 21
 - fibre channel 21
 - hardware 21
 - RAM Ethernet controllers 21
 - SCSI host bus adapter 21
- Root Broker 19
 - installing 31
- rpm -e command 153
- rsh 41, 64, 68

S

- SCSI host bus adapter 21
- SCSI-3
 - persistent reservations 43
- SCSI-3 persistent reservations
 - verifying 95
- seeding 16
 - automatic 16
 - manual 16
- setting
 - MANPATH variable 47
 - PATH variable 46
- single-node cluster
 - adding a node to 159
- single-system cluster
 - creating 155, 157
 - modifying startup files 158
- SMTP email notification 72
- SNMP trap notification 74
- ssh 41, 64, 68
 - configuring 42
- starting configuration
 - installvcs program 68
 - Veritas product installer 67

- starting installation
 - installvcs program 64
 - Veritas product installer 63
- starting VCS 78
- storage
 - fully shared vs. distributed 14
 - shared 14
- switches 38
- Symantec Product Authentication Service 19, 31, 70
- system communication using rsh
 - ssh 41
- system state attribute value 111

U

- uninstalling
 - prerequisites 169
 - VCS 169
- uninstallvcs 169
- upgrade
 - minimal downtime 119
- upgrading
 - minimal downtime 119

V

- variables
 - MANPATH 47
 - PATH 46
- VCS
 - basics 13
 - command directory path variable 107
 - configuration files
 - main.cf 103
 - coordinator disks 96
 - documentation 84
 - replicated states on each system 14
- verifying
 - cluster 82
 - NIC configuration 77
- vxdisksetup command 90
- vxlicinst 37
- vxlicinst command 83, 142
- vxlicrep 37
- vxlicrep command 82, 143
- vxlictest 37