

Veritas™ Cluster Server Installation Guide

Linux

5.0

Veritas Cluster Server Installation Guide

Copyright © 2006 Symantec Corporation. All rights reserved.

Veritas Cluster Server 5.0

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

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

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

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

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

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Linux is a registered trademark of Linus Torvalds.

Licensing and registration

Veritas Cluster Server is a licensed product. See the *Veritas Cluster Server Installation Guide* for license installation instructions.

Technical support

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

Contents

Chapter 1	Introducing VCS	
	About VCS	15
	VCS basics	15
	Multiple nodes	16
	Shared storage	16
	LLT and GAB	17
	Network channels for heartbeating	18
	Preexisting network partitions	18
	VCS seeding	18
Chapter 2	Preparing to install and configure VCS	
	About preparing to install VCS 5.0	21
	About VCS 5.0 features	22
	Symantec Product Authentication Service	22
	Veritas Cluster Management Console	25
	SMTP email notification for VCS events	27
	SNMP trap notification for VCS events	28
	Global clusters using VCS	28
	I/O fencing	28
	Preparing to install VCS 5.0	29
	Hardware requirements	31
	Required disk space	32
	Supported operating systems	32
	Supported software	33
	Supported software for the	
	Veritas Cluster Management Console	33
	Supported browsers for the Cluster Management	
	Console	34
	Installing root broker for Symantec Product	
	Authentication Service	35
	Creating encrypted files for Symantec Product	
	Authentication Service	36
	Installing the management server for the Veritas Cluster	
	Management Console	38
	Installing the management server on Solaris	39

Installing the management server on Windows 2003 systems	42
Copying the VCS installation guide and release notes to each node	44
Accessing Cluster Management Console information	44
Performing pre-installation tasks	45
Setting the PATH variable	45
Setting the MANPATH variable	46
Setting up the private network	46
Using network switches	47
Configuring SuSE network interfaces	47
Setting up shared storage	49
Setting up shared storage: SCSI	50
Setting up shared storage: Fibre Channel	51
Enabling communication between systems	52
Optimizing LLT media speed settings on private NICs	52
Guidelines for setting the media speed of the LLT interconnects	52
Setting up ssh on cluster systems	53
Configuring ssh	53
Obtaining VCS license keys	54
Mounting the product disc	55
Getting your VCS installation and configuration information ready	56
Optional VCS RPMs	60

Chapter 3 Installing and configuring VCS

About installing and configuring VCS	61
About the VCS installation program	62
Optional features of the installvcs program	62
Interacting with the installvcs program	63
Installing and configuring VCS 5.0	63
Overview of tasks	64
Checking the systems for installation	65
Starting the software installation	66
Specifying systems for installation	67
Licensing VCS	67
Choosing VCS RPMs	68
Choosing to install VCS RPMs or configure VCS	69
Configuring the cluster	69
Configuring the cluster in secure mode	70
Adding VCS users	72
Configuring cluster connector	72

Configuring the Cluster Management Console	73
Configuring SMTP email notification	74
Configuring SNMP trap notification	76
Configuring global clusters	77
Installing the VCS RPMs	78
Creating VCS configuration files	78
Verifying the NIC configuration	78
Starting VCS	79
Completing the installation	79
Copying the installation guide to each node	80
Setting up I/O fencing	80
Installing the VCS Java Console	80
Hardware requirements for the Java Console	81
Installing the Java Console on Linux	81
Installing the Java Console on a Windows workstation	82
Establishing cluster communication with	
the management server	82
Installing cluster connector	83
Installing the cluster connector on UNIX systems	83
Batch installation of cluster connector on UNIX systems	86
Installing cluster connector on Windows platforms	89
Accessing Cluster Management Console information	91
Verifying the cluster after installation	92
Installing VCS using installonly option	92
Configuring VCS using configure option	92
Performing VCS installation in a secure environment	92
Performing automated installations	94
Syntax used in response file	94
Example response file	95
Response file variable definitions	96
Checking licensing information on the system	100
Updating product licenses using vxlicinst	101
Replacing a VCS demo license with a permanent license	101
About installvcs command options	102
About the uninstallvcs program	104
Prerequisites	104
Uninstalling VCS 5.0	105
Removing VCS 5.0 RPMs	105
Running uninstallvcs from the VCS 5.0 disc	106
Uninstalling the Cluster Management Console	
management server	106
Uninstalling the management server from	
Solaris systems	106

Uninstalling the management server from Windows systems	107
Uninstalling the Cluster Management Console cluster connector	108
Uninstalling cluster connector from UNIX systems	108
Uninstalling cluster connector from Windows platforms	109

Chapter 4 Setting up I/O fencing

About I/O fencing	111
Preventing data corruption with I/O fencing	111
SCSI-3 persistent reservations	112
I/O fencing components	113
Data disks	113
Coordinator disks	113
I/O fencing operations	114
Preparing to configure I/O fencing	114
Checking shared disks for I/O fencing	114
Testing the shared disks for SCSI-3	115
Setting up I/O fencing for VCS	117
Initializing disks	118
Setting up coordinator disk groups	119
Requirements for coordinator disks	119
Creating the coordinator disk group and setting the coordinator attribute	119
Stopping VCS on all nodes	120
Configuring /etc/vxfendg disk group for I/O fencing	120
Updating /etc/vxfenmode file	121
Starting I/O fencing	122
Modifying VCS configuration to use I/O fencing	122
Verifying I/O fencing configuration	123
Removing permissions for communication	123
Additional I/O fencing information	124
vxfentsthdw options	124
Testing the coordinator disk group using vxfentsthdw -c	125
Using the -r option for non-destructive testing	126
Using the -m option	127
Using the -f option	127
Using the -g option	127
Testing a disk with existing keys	128
About VXFEN tunable parameters	128
Configuring the VXFEN parameters	129
How I/O fencing works in different event scenarios	130
About the vxfenadm utility	134

Registration key formatting	135
Troubleshooting I/O fencing	136
Node is unable to join cluster while another node is being ejected	136
vxfentsthdw fails when SCSI TEST UNIT READY command fails	136
Removing existing keys from disks	136
System panics to prevent potential data corruption	137
How vxfen driver checks for pre-existing split brain condition	137
Case 1: system 2 up, system 1 ejected (actual potential split brain)	138
Case 2: system 2 down, system 1 ejected (apparent potential split brain)	138
Clearing keys after split brain using vxfenclearpre command	139
Adding or removing coordinator disks	140
 Chapter 5	
Verifying the VCS installation	
About verifying the VCS installation	143
Verifying LLT and GAB configuration files	143
/etc/llthosts	143
/etc/llttab	144
/etc/gabtab	144
Verifying the main.cf file	145
Example main.cf, for clusters without the GCO option	146
Example main.cf, for clusters with the GCO option	148
Example main.cf for a centrally managed cluster using Cluster Management Console	148
Verifying LLT, GAB, and cluster operation	150
Verifying LLT	150
Using lltstat -n	150
Using lltstat -nvv	150
Verifying GAB	152
Verifying the cluster	152
hasys -display	153
Accessing the Veritas Cluster Management Console	155
Accessing the VCS documentation	156
 Chapter 6	
Upgrading to VCS 5.0	
About upgrading to VCS 5.0	157
Upgrading VCS using installvcs program	158
Upgrading VCS to 5.0	158

Removing deprecated resource types	158
Starting the upgrade	159
Checking upgrade requirements	160
Removing VCS RPMs from previous versions and installing VCS 5.0 RPMs	161
Starting VCS	162
Completing the upgrade	163
Using the halogin command for native OS accounts with VCS	164
Upgrading VCS in a secure environment	165
Upgrading VCS manually	165
Upgrading from VCS 4.0 MP2, 4.1, or 4.1 MP1	165
Licensing VCS	167
Starting LLT, GAB, and VCS	167
Upgrading to the VCS 5.0 Java Console	169
Upgrading from CommandCentral Availability 4.1 MP1	170
Upgrade order	170
Upgrading the management server on Solaris	170
Copying the VCS installation guide and release notes to each node	173
Accessing Cluster Management Console information	173
Upgrading the management server on Windows systems	173
Upgrading cluster monitor to cluster connector on UNIX systems	175
Upgrading cluster monitor to cluster connector on Windows platforms	177

Chapter 7 Adding and removing cluster nodes

About adding and removing nodes	179
Adding a node to a cluster	179
Setting up the hardware	180
Preparing for a manual installation	181
Installing VCS RPMs for a manual installation	181
Adding a license key	182
Checking licensing information on the system	182
Configuring LLT and GAB	182
Adding the node to the existing cluster	184
Starting VCS and verifying the cluster	184
Removing a node from a cluster	185
Verify the status of nodes and service groups	185
Deleting the leaving node from VCS configuration	186
Modifying configuration files on each remaining node	188
Unloading LLT and GAB and removing VCS on the leaving node	188

Chapter 8	Installing VCS on a single node	
	About installing VCS on a single node	191
	Creating a single-node cluster using the installer program	192
	Preparing for a single node installation	192
	Starting the installer for the single node cluster	192
	Creating a single-node cluster manually	193
	Setting the PATH variable	193
	Installing the VCS software manually	194
	Renaming the LLT and GAB startup files	194
	Modifying the startup files	194
	Configuring VCS	194
	main.cf file	194
	types.cf file	195
	Verifying single-node operation	196
	Adding a node to a single-node cluster	196
	Setting up a node to join the single-node cluster	197
	Installing VxVM, VxFS if necessary	198
	Installing and configuring Ethernet cards for private network	198
	Configuring the shared storage	198
	Bringing up the existing node	199
	Installing the VCS software manually	199
	Configuring LLT and GAB	199
	Configuring low latency transport (LLT)	200
	Configuring group membership and atomic broadcast (GAB)	202
	Starting LLT and GAB	202
	Reconfiguring VCS on the existing node	202
	Verifying configuration on both nodes	203
Appendix A	Advanced topics related to installing VCS	
	LLT over UDP	205
	When to use LLT over UDP	205
	Performance considerations	205
	Configuring LLT over UDP	206
	Broadcast address in the /etc/llttab file	206
	The link command in the /etc/llttab file	206
	The set-addr command in the /etc/llttab file	207
	Selecting UDP ports	208
	Configuring LLT on subnets	208
	Sample configuration: Direct-attached links	209
	Sample configuration: Links crossing IP routers	210
	Minimal downtime upgrade	212

Supported upgrades	212
Prerequisites for a minimal downtime upgrade	212
Planning for the minimal downtime upgrade	212
Minimal downtime upgrade limitations	212
Minimal downtime upgrade example	213
Minimal downtime example overview	213
Performing the minimal downtime example upgrade	214
Setting up a trust relationship between two authentication brokers	217

Appendix B Sample VCS installation and configuration output

About sample VCS installation and configuration	219
Installing the Root Broker	219
Installing the Cluster Management Console Management Server	221
Installing VCS 5.0	226
Start the product installer or the installvcs program	227
Installer performs initial system checks	228
License VCS	228
Installer checks for installed RPMs	228
Choose to install all VCS RPMs or required RPMs	228
Installer lists the RPMs	229
Configuring VCS 5.0	230
Configure the cluster	231
Configure the cluster in secure mode	231
Configuring security automatically	232
Configuring security semiautomatically using encrypted files	233
Configuring security semiautomatically answering prompts	234
Add VCS users	234
Configure cluster connector	235
Configure Cluster Management Console	236
Configure SMTP email notification	237
Configure SNMP trap notification	238
Configure the global cluster option	238
Installer installs the VCS RPMs	239
Installer creates VCS configuration files	239
Installer verifies the NIC configuration	240
Start VCS	240
Complete the installation	241
Uninstalling VCS 5.0	241

Appendix C	Deploying VCS in a Storage Foundation Volume Server environment	
	About VCS in a Storage Foundation VS environment	244
	Volume Serving mode using VCS	244
	Share Out mode using VCS	245
	About failover in a Storage Foundation VS environment	246
	About I/O fencing in a SFVS environment	247
	Deploying VCS in a Storage Foundation VS environment	247
	Prerequisites	247
	Installing VCS	248
	Configuring I/O fencing in Volume Serving mode	248
	Setting up the coordinator volume	248
	Editing the vxfenmode file	249
	Editing the main.cf file	250
	Configuring I/O fencing in the Share Out mode	251
	Creating coordinator volumes	251
	Editing the vxfenmode file	252
	Configuring application service groups	253
Appendix D	Configuring the Symantec License Inventory Agent	
	About the Symantec License Inventory Manager	258
	When the Symantec License Inventory Agent is installed	259
	When the server and access points are installed	259
	What you can do with the agent after it is installed	259
	How to remove the agent	260
	How to order the Symantec License Inventory	
	Manager license and media kit	261
Index		263

Introducing VCS

This chapter contains the following topics:

- [About VCS](#)
- [VCS basics](#)

About VCS

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.

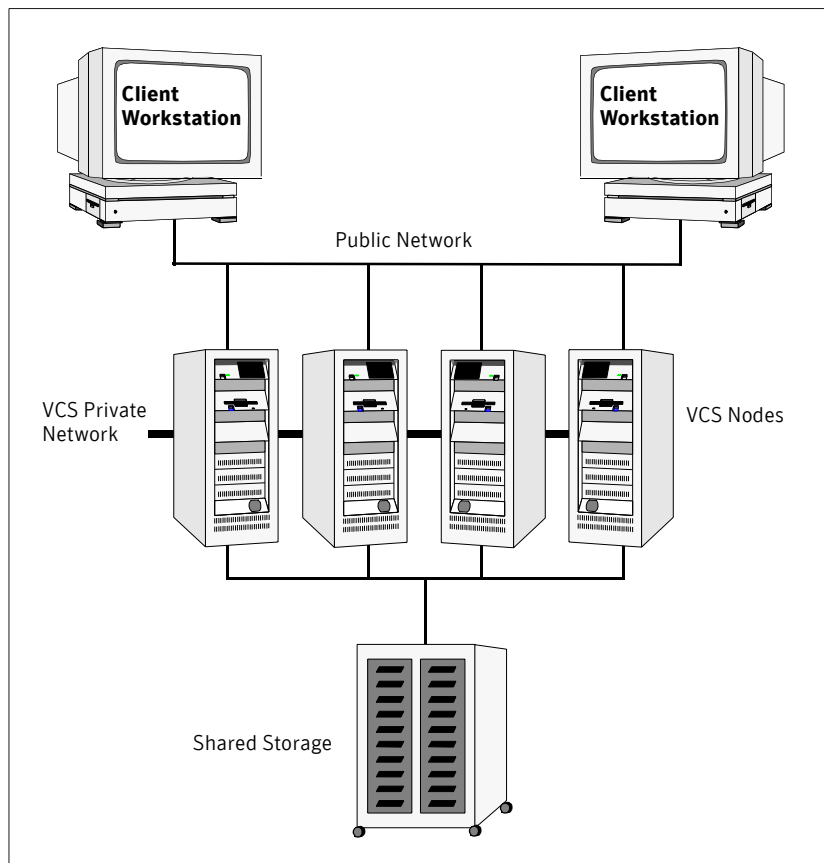
VCS basics

A single VCS cluster consists of multiple systems 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.

Client applications continue operation 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 connected to shared storage. 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.

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



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 and to recognize active nodes, nodes that are joining or leaving the cluster, and failed nodes. The private network requires two communication channels to guard against network partitions.

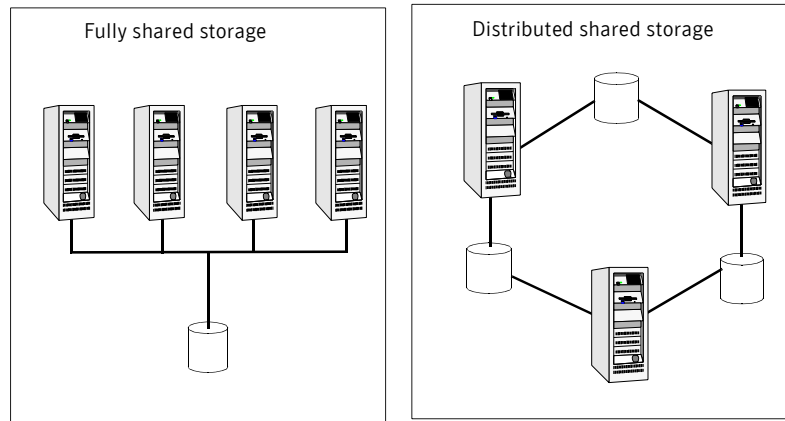
Shared storage

A VCS hardware configuration typically consists of multiple nodes connected to shared storage through I/O channels. Shared storage provides multiple systems

with an access path to the same data, and enables VCS to restart applications on alternate nodes when a node fails, which ensures high availability.

[Figure 1-2](#) illustrates the flexibility of VCS shared storage configurations. VCS nodes can only access physically-attached storage.

Figure 1-2 Two examples of shared storage configurations



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. The system administrator configures LLT by creating the configuration files `/etc/llthosts`, which lists all the nodes in the cluster, and `/etc/llttab`, which describes the local system's private network links to the other nodes in the cluster.
- GAB (Group Membership and Atomic Broadcast) provides the global message order required to maintain a synchronized state among the nodes. It monitors disk communications such as the VCS heartbeat utility. The system administrator configures the GAB driver by creating a configuration file (`/etc/gabtab`).

See "[Verifying LLT and GAB configuration files](#)" on page 143.


Network channels for heartbeating

For the VCS private network, two network channels must be available for heartbeating. These network connections are also used for transmitting 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. Refer to the *Veritas Cluster Server User's Guide* for more information on network partitioning.

[Figure 1-3](#) illustrates a two-node VCS cluster where sysA and sysB have two private network connections.

Figure 1-3 Two nodes connected by two ethernet connections



Preexisting network partitions

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

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 when:

- 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.

You need to perform a manual seed to run VCS from a cold start (all systems down) when one or more systems of the cluster are unavailable. VCS does not start service groups on a system until it has a seed.

Preparing to install and configure VCS

This chapter contains the following topics:

- [About preparing to install VCS 5.0](#)
- [About VCS 5.0 features](#)
- [Preparing to install VCS 5.0](#)
- [Performing pre-installation tasks](#)
- [Getting your VCS installation and configuration information ready](#)

About preparing to install VCS 5.0

Before you install any product, read the following Veritas Technical Support TechNote for the latest information on updates, patches, and software issues regarding this release:

<http://support.veritas.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>

About VCS 5.0 features

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:

- [Symantec Product Authentication Service](#)
- [Veritas Cluster Management Console](#)
- [SMTP email notification for VCS events](#)
- [SNMP trap notification for VCS events](#)
- [Global clusters using VCS](#)
- [I/O fencing](#)

Symantec Product Authentication Service

Symantec Product Authentication Service secures communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. For more information about the Authentication Service, see the *Veritas Cluster Server User's Guide*.

To configure the cluster in secure mode, VCS requires you to configure a system in your enterprise as root broker and all nodes in the cluster as authentication brokers.

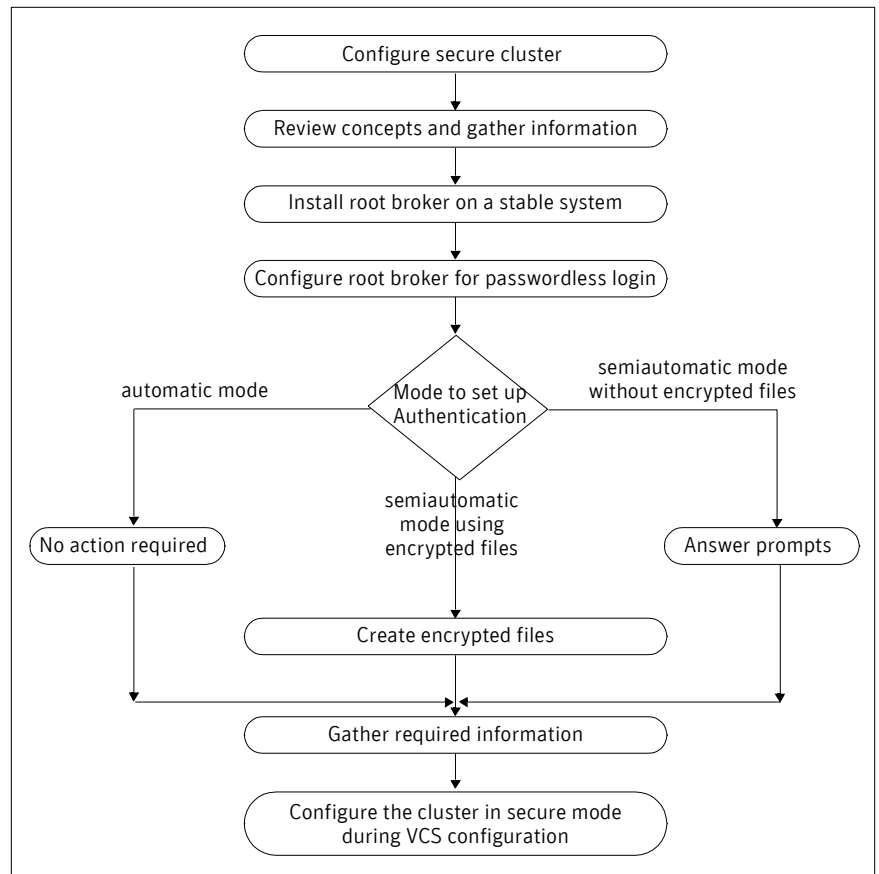
- **Root broker**
A root broker serves as the main registration and certification authority; it has a self-signed certificate and can authenticate other brokers. The root broker is only used during initial creation of an authentication broker.
- **Authentication brokers**
Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates that are signed by the root. Each node in VCS serves as an authentication broker.

You can set up Authentication Service for the cluster during the installation or after installation. Refer to the *Veritas Cluster Server User's Guide* to configure the cluster in secure mode after the installation and configuration process.

See [“Configuring the cluster in secure mode”](#) on page 70.

[Figure 2-4](#) depicts the flow of configuring VCS in secure mode.

Figure 2-4 Secure VCS cluster configuration flowchart



If you decide to enable Authentication Service, the root broker administrator must perform the following preparatory tasks:

- Install the root broker on another stable system.
 The root broker is the main registration and certification authority and can serve multiple clusters. Symantec recommends that you install a single root broker on a utility computer such as an email server or domain controller, which can be highly available.
 See [“Installing root broker for Symantec Product Authentication Service”](#) on page 35.
- Configure the root broker system for a passwordless login when you want to use the automatic mode.

The `installvcs` program provides the following modes to enable Symantec Product Authentication Service:

- In the automatic mode, the installer configures Authentication Service automatically without any user intervention.
You must provide the name of the root broker system.
- In the semiautomatic modes, the installer provides you an option to use encrypted files or answer the installer prompts to enable security. The semiautomatic mode requires the root broker administrator to set up the basic authentication environment and create principals for authentication brokers. You must complete the following preparatory tasks to configure security in the semiautomatic mode:

- | | |
|------------------------|--|
| With encrypted file | <ul style="list-style-type: none">■ The root broker administrator must create an encrypted file for each node in the cluster.
See "Creating encrypted files for Symantec Product Authentication Service" on page 36.■ You must fetch the encrypted files from the root broker administrator and copy the encrypted files to the installation node. Make a note of the path of these encrypted files. |
| Without encrypted file | <ul style="list-style-type: none">■ You must gather the following information from the root broker administrator:<ul style="list-style-type: none">- Root broker name- Root broker domain name- Root broker port (Default is 2821)- Authentication broker principal name for each node- Authentication broker password for each Authentication broker■ You must fetch the root_hash file from the root broker system and copy the root_hash file to a directory in the installation node. Make a note of the path of this root_hash file. |

Note: Make sure that the system clocks of the rook broker and authentication brokers are in sync.

Veritas Cluster Management Console

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

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

- If you want to use Cluster Management Console to manage multiple clusters, you must set up a management server.
- If you want to use the Cluster Management Console to manage a single cluster, choose the option to install the Cluster Management Console during VCS installation and configuration.

Operational mode

Configurational description

Local management of one cluster (single-cluster mode)

The Cluster Management Console is installed along with VCS on each node in the cluster and is configured for failover. It is integrated with VCS as part of the ClusterService service group. The Cluster Management Console offers robust cluster management capability and can be run from any supported Web browser on any system. See [“Configuring the Cluster Management Console”](#) on page 73.

Operational mode

Centralized, comprehensive, enterprise-wide administration of multiple clusters (multi-cluster mode)

Configurational description

One instance of the Cluster Management Console is installed outside all clusters on a standalone server. The console enables users to visually and intuitively input commands to the multi-cluster management engine, the *management server*. The management server initiates monitoring and management actions based upon those commands. The management server uses a database to store cluster configurations, cluster status, events, event policies, report jobs, report outputs, and more.

See [“Installing the management server for the Veritas Cluster Management Console”](#) on page 38.

If the management server and cluster nodes are separated by a firewall, a component called *cluster connector* is installed on each cluster node. Cluster connector enables communication with clusters through firewalls. Cluster connector also provides buffering for cluster data. If the console goes offline and then comes back online, it can retrieve data collected during the offline period from the cluster connector buffer.

See [“Configuring cluster connector”](#) on page 72.

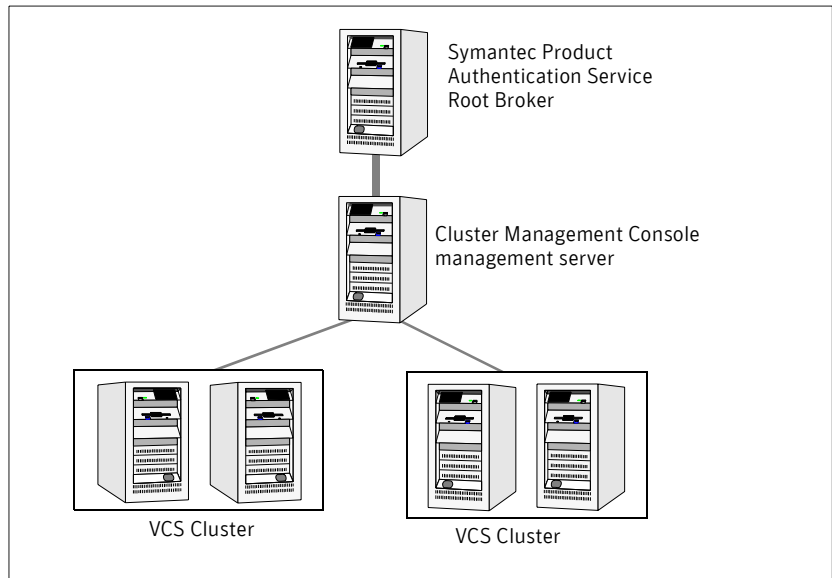
The console offers additional capability for administering users, reports, events, and notification. If the cluster environment includes licensed VCS global clusters, disaster recovery (DR) capability is also available.

The configurational differences between the operational modes mean that you cannot switch a single Cluster Management Console installation from one mode to the other. The modes are also incompatible on the same system.

Consequently, one system cannot offer both operational modes. However, the modes *can* co-exist in the same multi-cluster environment, with single-cluster-mode installations on VCS cluster nodes, and multi-cluster-mode installations on management server hosts. Such a deployment can be desirable if different IT administrators in your enterprise have different scopes of responsibility.

See *Veritas Cluster Server Centralized Management Guide* for more information.

Figure 2-5 Sample deployment for Veritas Cluster Management Console



SMTP email notification for VCS events

You have the option to configure SMTP email notification of VCS events by the VCS Notifier component. If you choose SMTP notification, be ready to answer prompts for the following information:

- The domain-based address of the SMTP server that is to send notification email about the events within the cluster, for example: `smtp.symantecexample.com`.
- The email address of each SMTP recipient to be notified, for example: `john@symantecexample.com`.
- The minimum severity of events for SMTP email notification. Events have four levels of severity: Information, Warning, Error, and SevereError.

Refer to the *Veritas Cluster Server User's Guide* for more information on SMTP notification.

SNMP trap notification for VCS events

You have the option to configure SNMP trap notification of VCS events by the VCS Notifier component. If you choose SNMP notification, be ready to answer prompts for the following information:

- The port number, 162 by default, for the SNMP trap daemon.
- The system name for each SNMP console.
- The minimum severity of events for SNMP trap notification. Events have four levels of severity: Information, Warning, Error, and SevereError.

Refer to the *Veritas Cluster Server User's Guide* for more information on SNMP notification.

Global clusters using VCS

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.

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

I/O fencing

I/O fencing protects the data on shared disks. When nodes in a cluster detect a change in cluster membership that could indicate a split brain condition, the fencing operation proceeds to determine which nodes are to retain access to the shared storage and which nodes are to be ejected from the cluster, thus preventing possible data corruption. The *Veritas Cluster Server User's Guide* describes I/O fencing concepts in detail. The `installvcs` program installs the VCS I/O fencing driver, `VRTSvxfen`.

Note: Symantec strongly recommends that you use VCS I/O fencing to deter potential split brain scenarios in your cluster.

See [“Setting up I/O fencing”](#) on page 111.

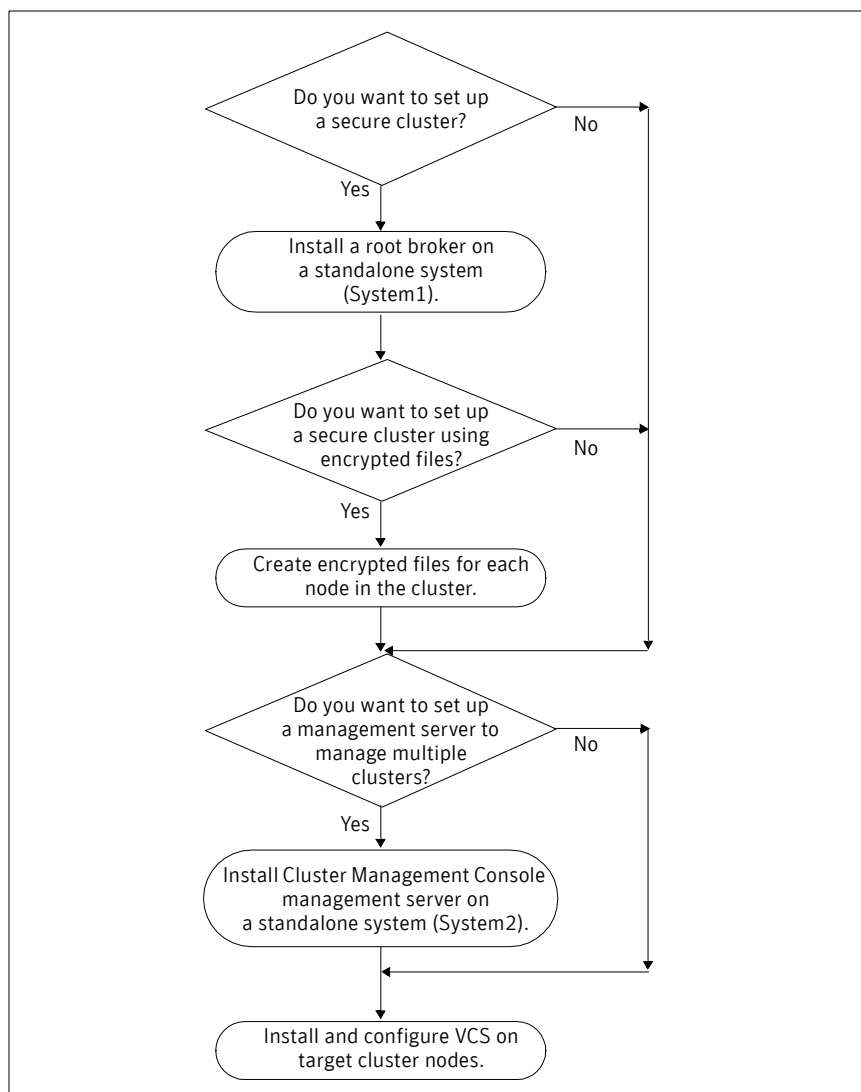
Preparing to install VCS 5.0

Each node on which you want to install VCS must meet the hardware and software requirements.

- [“Hardware requirements”](#) on page 31
- [“Supported operating systems”](#) on page 32
- [“Supported software”](#) on page 33

After planning the VCS features that you want to configure, you must prepare to configure these features. [Figure 2-6](#) represents the major tasks and decisions required to install and configure VCS.

Figure 2-6 Workflow for fresh install of VCS 5.0



Complete the following preparatory tasks based on the VCS features you want to configure:

- [“Installing root broker for Symantec Product Authentication Service”](#) on page 35
- [“Creating encrypted files for Symantec Product Authentication Service”](#) on page 36

- [“Installing the management server for the Veritas Cluster Management Console”](#) on page 38

Hardware requirements

Make sure that you meet the following requirements.

Table 2-1 Hardware requirements for a cluster

Item	Description
VCS systems	From 1 to 32 Linux systems running the supported Linux operating system version. See “Supported operating systems” on page 32.
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.
Disks	Typical VCS configurations require that shared disks support applications that migrate between systems in the cluster. The VCS I/O fencing feature requires that all disks used as data disks or as coordinator disks must support SCSI-3 Persistent Reservations (PR). The coordinator disk does not store data, so configure the disk as the smallest possible LUN on a disk array to avoid wasting space. See “Setting up I/O fencing” on page 111.
Disk space	See Table 2-2, “Disk space requirements and totals.”
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.
Fibre Channel or SCSI host bus adapters	Typical VCS configuration requires at least one SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.
RAM	Each VCS system requires at least 256 megabytes.

Required disk space

Confirm that your system has enough free disk space to install VCS. The following table 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	5 MB	143 MB	25 MB	169 MB	342 MB
Optional	3 MB	39 MB	1 MB	7 MB	50 MB
Required and optional total	8 MB	202 MB	26 MB	176 MB	412 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

Run VCS 5.0 on these operating systems at the suggested patch levels. Within a cluster, all nodes must use the same operating system version and patch level.

[Table 2-3](#) lists the architectures and operating systems on which VCS 5.0 operates. Symantec supports only those kernel binaries distributed by Red Hat and SUSE.

Table 2-3 Supported operating systems and architectures

Operating System	Kernel	Architecture
Red Hat Enterprise Linux 4 (RHEL 4) Update 3	2.6.9-34.EL	x86 (32-bit)
	2.6.9-34.smp	Intel Xeon (32-bit, 64-bit)
	2.6.9-34.hugemem	AMD Opteron (32-bit, 64-bit)
SUSE Linux Enterprise Server 9 (SLES 9) with SP3	2.6.5-7.244	x86 (32-bit)
	2.6.5-7.244-smp	Intel Xeon (32-bit, 64-bit)
	2.6.5-7.244-bigsmpp	AMD Opteron (32-bit, 64-bit)

Symantec products will 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* installing any Symantec product.

<http://support.veritas.com/docs/281993>

Supported software

- ext2, ext3, reiserfs, NFS, and bind on LVM2, Veritas Volume Manager (VxVM) 4.1 and 5.0, and raw disks.
- Veritas Volume Manager 4.1 with Veritas File System 4.1
- Veritas Volume Manager 5.0 with Veritas File System 5.0

Note: If you plan to set up VCS I/O fencing in a shared storage environment, Symantec recommends using VxVM versions 4.1 or 5.0.

Supported software for the Veritas Cluster Management Console

You can install the Cluster Management Console on a standalone system to manage multiple clusters or you can install the console on cluster nodes to manage a local cluster.

When you set up a management server to manage multiple clusters, you can connect to the clusters directly or install the cluster connector on cluster nodes to enable connection across firewalls.

Install Mode

Multi-cluster mode
 To manage multiple clusters.
 Installed on a standalone system designated the *management server*.

Single cluster mode
 To manage a single cluster.
 Installed on cluster nodes.

Supported software

- Solaris 8, 9, and 10, with patches indicated by Sun
- Windows 2000 Server, Advanced Server, and Datacenter, with SP4 or patches as indicated by Microsoft
- Windows Server 2003 Standard Edition, Datacenter Edition, Enterprise Edition, and Web Edition, with patches as indicated by Microsoft

Note: Windows Management Instrumentation (WMI) must be installed on Windows 2003 systems prior to installing Cluster Management Console.

Install Cluster Management Console in this mode only with VCS 5.0 in a fresh installation or upgrade scenario.

Install Mode

Cluster Connector

Installed on cluster nodes to enable a management server to manage a cluster across a firewall

Supported software

AIX

- VCS versions: 4.0, 4.0 MP1, 4.0 MP2, 4.0 MP3, and 5.0
- OS versions: AIX 5.2 ML6 (legacy) or later; AIX 5.3 TL4 with SP 3

Note: Cluster connector installations are not supported on clusters running AIX 5.1. Use direct connection to manage clusters running AIX 5.1.

HP-UX

- VCS versions: 4.1 and 5.0
- OS versions: HP-UX 11i v2

Linux

- VCS versions: 4.0, 4.0 MP1, 4.0 MP2, 4.1, 4.1 MP1, 4.1 MP2, and 5.0
- OS versions: RHEL 4 Update 3, SLES 9.

Note: Cluster connector installations are not supported on clusters running RHEL 3.0. Use direct connection to manage clusters running RHEL 3.0.

Solaris

- VCS versions: 4.0, 4.0 MP1, 4.0 MP2, 4.1, 4.1 MP1, and 5.0
- OS versions: 8, 9, and 10

Note: Cluster connector installations are not supported on clusters running Solaris 7. Use direct connection to manage clusters running Solaris 7.

Windows

- VCS versions: 4.1, 4.2, 4.2 RP1, 4.2 RP2, 4.3, 4.3 MP1
- OS versions: Windows 2000 Server, Advanced Server, and Datacenter, with SP4 or patches as indicated by Microsoft.
- Microsoft Windows Server 2003* Standard Edition, Datacenter Edition, Enterprise Edition, and Web Edition, with patches as indicated by Microsoft

* Windows Management Instrumentation (WMI) must be installed on Windows 2003 systems prior to installing Cluster Management Console.

Supported browsers for the Cluster Management Console

Veritas Cluster Management Console is supported on the following browsers:

- Microsoft Internet Explorer 6.0 or later

- Firefox 1.5 or later

Veritas Cluster Management requires the Macromedia Flash Plugin v8.0.

Installing root broker for Symantec Product Authentication Service

Install the root broker only if you plan on using Symantec Product Authentication Service. 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 *Veritas Cluster Server User's Guide* for more information. You can configure the Authentication Service during or after VCS installation.

See "[Symantec Product Authentication Service](#)" on page 22.

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: **east**
- 5 Review the output as the installer:
 - checks to make sure that the VCS supports the operating system
 - verifies that you are installing from the global zone (only on Solaris)
 - checks if the system already runs the security RPM
- 6 Review the output as the `installvcs` program checks for the installed RPMs on the system.
The `installvcs` program lists the RPMs that will be installed 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 Enter 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.
- 13 If you plan to configure the Authentication Service during VCS installation, choose to configure the cluster in secure mode when the installer prompts you.
See [“Installing and configuring VCS 5.0”](#) on page 63.

Creating encrypted files for Symantec Product Authentication Service

Create encrypted files only if you plan on choosing the semiautomatic mode that uses an encrypted file to configure the Authentication Service. The encrypted files must be created by the administrator on the root broker node. The administrator must create encrypted files for each node that would be a part of the cluster before you configure the Authentication Service for VCS. See *Veritas Cluster Server User's Guide* for more information. You can configure the Authentication Service during or after VCS installation.

See [“Symantec Product Authentication Service”](#) on page 22.

To create encrypted files

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

```
east> # vssat showalltrustedcreds
```

For example, the domain name would resemble “Domain Name: root@east.symantecexample.com” in the output.
- 2 For each node in the cluster, make sure that you have created an account on root broker system.
For example, to verify on node north:

```
east> # vssat showprpl --pdrtype root \  
--domain root@east.symantecexample.com --prplname north
```

 - 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:

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

- If the output displays an error similar to “Failed To Get Attributes For Principal,” then the account for given authentication broker is not created on this root broker. Proceed to [step 3](#).
- 3 Create a principal account for each authentication broker in the cluster. For example:

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

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

- 4 Make a note of the following information that is required for the input file for the encrypted file.
- hash - The root hash string that consists of 40 characters, as shown by the command:


```
east> # vssat showbrokerhash
```
 - identity - Authentication broker identity
 The value that you provide for `--prplname` in [step 3](#) (for example, north).
 - password - Authentication broker password
 The value that you provide for `--password` in [step 3](#).
 - root_domain - Domain name of the root broker system
 The value that you determined in [step 1](#).
 - broker_admin_password - Authentication broker password for Administrator account on the node
 Provide a password of at least five characters long.

- 5 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 north would resemble:

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

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

- 6 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.

- 7 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:

```
east> # vssat createpkg --in /tmp/north.blob.in \  
--out /tmp/north.blob.out --host_ctx north
```

Note that this command creates a encrypted file even if you provide wrong password for “password=” entry, but the encrypted file will fail to install on authentication broker node.

- 8 After you complete creating output files for the encrypted file, you must copy these files to the installer node.
- 9 If you plan to configure the Authentication Service during VCS installation, choose to configure the cluster in secure mode when the installer prompts you.
See “[Installing and configuring VCS 5.0](#)” on page 63.

Installing the management server for the Veritas Cluster Management Console

Install the Cluster Management Console management server only if you plan to centrally manage multiple clusters. Make sure you have a root broker in your domain. VCS clusters need not be secure to configure Cluster Management Console to manage multiple clusters.

See “[Veritas Cluster Management Console](#)” on page 25.

Install the Cluster Management Console management server and supporting components on a standalone system (outside any cluster but on the local network). Configure the management server to use a previously installed root broker or install and configure a root broker on the management server host.

You can install the management server on one of the following supported operating systems:

- [Installing the management server on Solaris](#)
- [Installing the management server on Windows 2003 systems](#)

Installing the management server on Solaris

You must install the management server on a system outside the cluster. This procedure follows a script of a successful installation. If at any step you experience a result other than the expected result that is documented here, you can click “n” to re-enter information. If you continue to have problems, click “q” to quit the installation and then verify the installation prerequisites.

To install the management server on Solaris

- 1 Insert the distribution media into the disc drive on the local system. At the command prompt, type the following command to run the setup program:


```
./installer -rsh
```

 The setup program (setup) presents copyright information followed by a menu titled, “Storage Foundation and High Availability Solutions 5.0”.
- 2 Enter **i** to specify a task.


```
Enter a Task: [I,C,L,P,U,D,Q,?] i
```

 Setup displays another menu that lists products that are available for installation.
- 3 Select the menu number that corresponds to Veritas Cluster Management Console.


```
Select a product to install: [1-13,b,q]
```

 Setup presents a description of the product.
- 4 Enter **1** to select a product component.


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

 Setup presents a message stating that it will install the management server.
- 5 Enter **y** to verify that the information up to this point is correct.


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

 Setup performs an initial system check of the local system and checks for installed packages on the local system. If these checks are satisfactory, setup lists the packages to be installed.


```
Storage Foundation and High Availability Solutions 5.0
installer will install the following CMC packages:
VRTSat          Symantec Product Authentication Service
VRTSperl        Veritas Perl 5.8.8 Redistribution
VRTSdbms3       Symantec Shared DBMS
VRTSjre15       Veritas Java Runtime Environment Redistribution
VRTSweb         Veritas Java Web Server
VRTScmcm        Veritas Cluster Management Console
VRTScmcdc       Veritas Cluster Management Console Documentation
Press [Return] to continue:
```
- 6 Press Enter.

You may install Cluster Management Console packages without performing configuration. The setup program gives you the option to configure Cluster Management Console now, and provides instructions for configuring Cluster Management Console later.

- 7 Enter **y** to configure Cluster Management Console.
Are you ready to configure CMC? [y,n,q] (y)
- 8 Enter a unique management server display name, such as:
Enter a unique management server display name: [?]
mgmtserver1_sol9
- 9 Enter the network address used by the management server, such as:
Enter the network address used by the management server [b,?]
mgmtserver1.symantec.com
- 10 When prompted, enter a location for the management server database.
Enter the desired location of the database to be used by the management server [b,?] (/opt/VRTScmc/db)
Setup repeats the management server display name, the management server network address, and the database location.
- 11 Enter **y** to verify that the information up to this point is correct.
Is this information correct? [y,n,q,b] (y)
Setup describes local user configuration and custom user configuration.
- 12 Configure a local user or a custom user as the initial management server administrator. This is the first user account that is enabled to log in to the Cluster Management Console.
Make your selection and then specify the following user authentication details:
 - For a local user, setup assumes that the domain name is the name of the local system and that the domain type is unixpwd, or UNIX password. When prompted for the initial management server user name, enter root or another administrator-level user for the local system.
 - For a custom user, you must explicitly specify the domain name and the domain type along with the user name. Follow the three separate prompts to enter this user information.

```
Local User:
Configure a user on the local machine as the initial admin user.
Custom User:
Configure a user manually.
1) Local User
2) Custom User
Enter '1' to enter the name of a local user, '2' to set up a
custom user:
[1-2,q] (1) 1
Storage Foundation and High Availability Solutions 5.0
Local admin user selection:
```


To log in to the CMC Management Server, enter the name of a local user to be set as the administrator. The domain and domain type will be automatically selected for you.

Enter the initial management server user name: [b,?] (root)
 Storage Foundation and High Availability Solutions 5.0
 Management Server admin user verification:
 Management Server User Name: root

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

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

Setup describes a particular management server service account, which the management server uses for secure internal communications with cluster connector. This account is named CMC_CC@CMC_SERVICES.

- 14** Enter a password for the management server service account and confirm it at the next prompt.

Enter a password for the CMC service account: **xxxxxx**
 Confirm the password you entered for the CMC service account: **xxxxxx**

When you install and configure cluster connector, you must provide this same password for the CMC_CC@CMC_SERVICES account.

- 15** Specify whether or not you want the management server to use a remote root broker for user authentication.

If you have already configured a root broker in your network, Symantec recommends that you enter **y** to use that existing root. Specify the additional details for that remote root broker exactly as specified.

If you do not have a currently-configured root broker, enter **n** to install and configure a root broker on the management server host.

After you enter **y** or **n**, setup installs an authentication broker on the management server and configures it to use whichever root broker you selected. When finished, setup presents:

- Installation progress percentages
- Status for writing the management server configuration file
- Status for creating secure internal service accounts

- 16** Enter **y** to start Veritas Cluster Management Console processes now.

Do you want to start Veritas Cluster Management Console processes now? [y,n,q,b] (y)

Setup presents startup progress percentages and, if successful, displays the following message:

Startup completed successfully on all systems.

- 17** Enter an encryption key of at least five characters.

Enter five or more characters to be used an encryption key: [b]
xxxxxx

```
This key must be retained in a secure file and referenced using  
the -enckeyfile option if the generated responsefile is to be  
used again.
```

```
Press [Return] to continue:
```

18 Press Enter to continue.

Record the location that setup provides for the installation log files, summary file, and response file. Also ensure that you record the initial admin user information. You *must* use this account to log in to the Cluster Management Console for the first time.

Installing the management server on Windows 2003 systems

You must install the management server on a system outside all clusters. Windows Management Instrumentation (WMI) is a prerequisite for installing and using the management server and cluster connector.

To install WMI

- 1 Log on as a user that has administrator privileges on the system on which you want to install WMI.
- 2 On the **Start** menu, click **Settings**, and then click **Control Panel**.
- 3 In the Control Panel window, double-click **Add or Remove Programs**.
- 4 In the task pane, click **Add/Remove Windows Components**.
- 5 Click **Management and Monitoring Tools**, then click **Details**.
- 6 Ensure that the WMI Windows Installer Provider is checked, and then click **OK**.
- 7 Click **Next**.
- 8 If prompted, insert the Windows CD and click **OK**.
- 9 After installation is complete, click **Finish**.
- 10 Restart your computer.

To install the management server on Windows

- 1 On the distribution disc, locate the **\installer** directory.
- 2 Double-click the **setup** file.

Depending upon the operating system, you may or may not receive the following warning message:

```
The publisher could not be verified. Are you sure you want to  
run this software?
```

If you receive this message, click **Run**.

- 3 In the Welcome to the Veritas Cluster Management Console Installation Manager dialog box, read the introduction and then click **Next**.
- 4 In the Installation and Configuration Options dialog box, click **Install a new management server on the local node**, and then click **Next**.
- 5 In the Management Server Installation Directory dialog box, leave the default installation path provided in the text box or click **Browse** to search for another installation location. Click **Next** to accept the path.
- 6 In the Management Server Information dialog box, enter the system name and IP address of the intended management server host.
 You cannot change the port specification, 14145, but it is provided to help you to prevent port conflicts when configuring other software. The other ports used by the Cluster Management Console are 8181 (HTTP), 8443 (HTTPS), and 2994 (DBMS; this port can be shared with other Symantec products)
- 7 In the Database File Path box, leave the default database path provided or click **Browse** to search for another location for the database. Click **Next** to accept the path.
- 8 In the Services Account Password dialog box, enter a password for the user account that cluster connector uses for management server communications, and then click **Next**.
 Record the password that you enter in a safe place. You must use it again whenever you install or configure cluster connector.
- 9 In the User Credential Confirmation dialog box, leave the automatically-detected user information provided or specify another user name, domain, and domain type.
 This user becomes the initial management server user. You must provide the credentials entered at this step when logging in to the management server for the first time.
- 10 In the Summary dialog box, review the information you have specified and, if satisfactory, click **Next** to accept it and start the installation.
 The Installing Veritas Cluster Management Console dialog box displays a progress bar and a status message window for the installation.
- 11 When you receive the following message, click **Next**:
 "Done deleting installation files from node..."
- 12 In the Completed the Symantec Veritas Cluster Management Console Installation Manager dialog box, review the information about how to connect to the management server and log in for the first time. Record this information in a safe place and then click **Finish**.

The installer creates log files at the following locations:

- **Installation logs** – C:\Documents and Settings\All Users\Application Data\VERITAS\Cluster Management Console. The file names are Install_GUI_0.log and Install_MSI_0.log.
- **Management server logs** – C:\Program Files\VERITAS\Cluster Management Console\log

Copying the VCS installation guide and release notes to each node

After you install the management server, copy the Install Guide and the Release Notes over to the management server. The documentation locations on the distribution disc are:

- **Install Guide**
cluster_server/docs/vcs_install.pdf
- **Release Notes**
cluster_server/release_notes/vcs_notes.pdf

Accessing Cluster Management Console information

Information about administering clusters in multi-cluster mode is available in the Veritas Cluster Server Centralized Management Guide. The online help includes information about administering clusters in both single-cluster and multi-cluster mode. If you want to access the information about managing a single cluster in printed documentation, you can install the documentation package to the desired system.

The documentation package name for each supported operating system is:

- AIX–VRTSvcs.doc
- HP-UX–VRTSvcsdc
Note that you can copy the documents from depot/VRTSvcsdc/VRTSvcsdc/opt/VRTS/docs.
- Linux–VRTSvcsdc
- Solaris–VRTSvcsdc

Performing pre-installation tasks

Table 2-4 lists the tasks you must perform before proceeding to install VCS.

Table 2-4 Pre-installation tasks

Task	Reference
Set the PATH and MANPATH variables.	“Setting the PATH variable” on page 45 “Setting the MANPATH variable” on page 46
Set up the private network.	“Setting up the private network” on page 46
Configure SuSE network interfaces	“Configuring SuSE network interfaces” on page 47
Set up shared storage for I/O fencing (optional)	“Setting up shared storage” on page 49
Enable communication between systems.	“Enabling communication between systems” on page 52
Review basic instructions to optimize LLT media speeds.	“Optimizing LLT media speed settings on private NICs” on page 52
Review guidelines to help you set the LLT interconnects.	“Guidelines for setting the media speed of the LLT interconnects” on page 52
Set up <code>ssh</code> on cluster systems.	“Setting up ssh on cluster systems” on page 53
Obtain license keys.	“Obtaining VCS license keys” on page 54
Mount the product disc	“Mounting the product disc” on page 55

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

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.

Setting up the private network

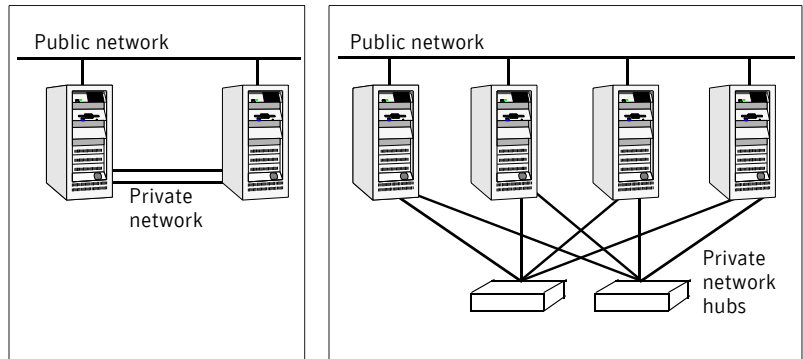
VCS requires you to setup a private network between the systems that will form part of a cluster.

To set up the private network

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

During the process of setting up heartbeat connections, note that a chance for data corruption exists if a failure removes all communications between the systems and still leaves the systems running and capable of accessing shared storage.

Figure 2-7 Private network setups: two-node and four-node clusters



- 4 Test network connections by temporarily assigning network addresses and use `telnet` or `ping` to verify communications. LLT uses its own protocol, and does not use TCP/IP. Therefore, to ensure the private network connections are used only for LLT communication and not for TCP/IP traffic, unplumb and unconfigure the temporary addresses after testing. The `installvcs` program configures the private network in the cluster during installation. See “[Installing and configuring VCS](#)” on page 61.

Using network switches

You can use network switches instead of hubs.

Configuring SuSE network interfaces

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

For 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.

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 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).

- Run the command:

```
ifconfig -a
```

- 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.

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.

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:

- 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 above steps for all interfaces displayed by `ifconfig -a` command.

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"
```

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

Setting up shared storage

The following sections describe setting up SCSI and Fibre 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 “[Setting up I/O fencing](#)” on page 111.

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

Setting up shared storage: SCSI

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 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.
 - 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 if the name of the disk group is dg, the name of the volume is vol, and the file system type is vxfs.

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

To set up shared storage for fibre channel

- 1 Connect the fibre channel disk to a cluster system.
- 2 Boot the system and change the settings of the fibre channel. Perform the 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 is detecting the fibre channel disks properly.
- 4 Create volumes. Format the shared disk and create required partitions on it.
 - 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 if the name of the disk group is `dg`, the name of the volume is `vol`, and the file system type is `vxfs`.

```
# mkfs -t vxfs /dev/vx/dsk/dg/vol01
```
- 5 Repeat [step 2](#) and [step 3](#) for all nodes in the clusters that require connections with fibre 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.

Enabling communication between systems

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

See “[Performing VCS installation in a secure environment](#)” on page 92.

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

Optimizing LLT media speed settings on private NICs

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

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, set the hub or switch port to the same setting as that used on the cards on each node.

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

Symantec does not recommend using dissimilar network cards for private links. Details for setting the media speeds for specific devices are outside of the scope of this manual. Consult the device's documentation for more information.

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. Verify that `ssh` is configured correctly before starting the installation process.

Secure Shell (`ssh`) is a program to log on to another computer over a network, to execute commands on a remote system, and to copy files from one system to another. The `ssh` 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.

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 prompted, 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 beginning with `ssh_dss` and ending 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 VCS is to be installed.

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 from which the installation is taking place:

```
# 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 prompted, enter your DSA passphrase.
You are ready to install VCS on several systems by running the `installvcs` program on any one of them or 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 will not be prompted for the passphrase. Refer to the Red Hat documentation for more information.
- 9 To verify that you can connect to the systems on which VCS is to be installed, 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.

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

Obtaining VCS license keys

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

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

<http://www.veritas.com/buy/vLicense/vLicenseHome.jhtml>

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

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

Even though other products are included on the enclosed software discs, you can only install the Symantec software products for which you have purchased a license.

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 connected by the network to the systems where you are installing VCS. The system that you are using to install VCS need not be part of the cluster.
- 2 Insert the product disc with the VCS software into a drive 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.

Depending on the OS distribution and architecture, type the command:

```
RHEL 4 (i686)      # cd /mnt/cdrom/rhel4_i686/cluster_server
```

```
RHEL 4 (x86_64)   # cd /mnt/cdrom/rhel4_x86_64/cluster_server
```

```
SLES 9 (i586)     # cd /mnt/cdrom/sles9_i586/cluster_server
```

```
SLES 9 (x86_64)  # cd /mnt/cdrom/sles9_x86_64/cluster_server
```

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: **north, south**

The required license keys

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 54.

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 60.

- To configure the Veritas Cluster Server you need:

The name of 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 used by the private networks among systems

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) | <ul style="list-style-type: none"> ■ The name of the Root Broker system
Example: east
See “Symantec Product Authentication Service” on page 22. ■ Access to the Root Broker system without use of a password. |
| For semiautomatic mode using encrypted files | <p>The path for the encrypted files that you get from the Root Broker administrator.</p> <p>See “Creating encrypted files for Symantec Product Authentication Service” on page 36.</p> |
| For semiautomatic mode without using encrypted files | <ul style="list-style-type: none"> ■ The fully-qualified hostname (FQDN) of the Root Broker . (e.g. east.symantecexample.com)
The above example given posits a system in the (DNS) domain symantecexample.com with the unqualified hostname east, which is designated as the Root Broker. ■ The root broker’s security domain (e.g. root@east.symantecexample.com) ■ The root broker’s port (e.g. 2821) ■ The path to the local root hash (e.g. /var/tmp/privatedir/root_hash) ■ The authentication broker’s principal name on the each cluster node (e.g. north.symantecexample.com and south.symantecexample.com) |

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

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

- To configure the Cluster Management Console to locally manage this cluster (optional), you need:

The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. Example: eth0
A virtual IP address of the NIC for the Cluster Management Console	This virtual IP address becomes a resource for use by the ClusterService group that includes the Cluster Management Console. The “Cluster Virtual IP address” can fail over to another cluster system, making the Web Console highly available. Example: 10.10.12.1
The netmask for the virtual IP address	The subnet used with the virtual address. Example: 255.255.240.0

- To configure the Cluster Management Console cluster connector (optional), you need:

The management server network address for Cluster Management Console	The Cluster Management Console cluster connector requires the management server network address. See “Veritas Cluster Management Console” on page 25. Example: mgmtserver1.symantecexample.com
A Cluster Management Console service account password	You must have set this account password while installing the management server.
The root hash of the management server	You can use <code>vssat showbrokerhash</code> command and copy the root hash of the management server.

- 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**

■ To configure I/O fencing:

The names of the three disks that form the coordinator disk group Example: **sda, sdb, sdc**

The DMP nodes names for each disk in the coordinator disk group (if using DMP) Example: **/dev/vx/dmp**

Optional VCS RPMs

The optional VCS RPMs include:

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

Installing and configuring VCS

This chapter contains the following topics:

- [About installing and configuring VCS](#)
- [About the VCS installation program](#)
- [Installing and configuring VCS 5.0](#)
- [Installing VCS using installonly option](#)
- [Configuring VCS using configure option](#)
- [Performing VCS installation in a secure environment](#)
- [Performing automated installations](#)
- [Checking licensing information on the system](#)
- [Updating product licenses using vxlicinst](#)
- [About installvcs command options](#)
- [About the uninstallvcs program](#)
- [Uninstalling VCS 5.0](#)

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 | Offers a high-level approach to installing multiple Veritas products. |
| installvcs program | Offers a direct approach to installing VCS. |

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

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, creating several detailed configuration files on each system
- Starting VCS processes

You can choose to configure the optional: Web-based Cluster Management Console, SNMP and SMTP notification features in the cluster, Symantec Product Authentication Services feature, or the wide area Global Cluster feature. Review the highlights of the information for which `installvcs` program prompts you as you proceed to configure.

See [“Preparing to install and configure VCS”](#) on page 21.

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

See [“About the uninstallvcs program”](#) on page 104.

Optional features of the `installvcs` program

[Table 3-5](#) specifies the optional actions that the `installvcs` program can perform.

Table 3-5 `installvcs` optional features

Optional action	Reference
Check the systems to verify that they meet the requirements to install VCS.	See “Checking the systems for installation” on page 65.
Upgrade VCS to version 5.0 if VCS currently runs on a cluster.	See “Upgrading VCS using <code>installvcs</code> program” on page 158.
Install VCS RPMs without configuring VCS.	See “Installing VCS using <code>installonly</code> option” on page 92.
Configure or reconfigure VCS when VCS RPMs are already installed.	See “Configuring VCS using <code>configure</code> option” on page 92.

Table 3-5 installvcs optional features

Optional action	Reference
Perform secure installations using values stored in a configuration file.	See “Performing VCS installation in a secure environment” on page 92.
Perform automated installations using values stored in a configuration file.	See “Performing automated installations” on page 94.

Interacting with the installvcs program

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

Note: Installation of VCS RPMs takes place only after you have confirmed the information. However, you must remove the partially installed VCS files before running the installvcs program again. See [“Uninstalling VCS 5.0”](#) on page 105.

At some points during the installation, the installer prompts you to type information and 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.

When the installer prompts you to answer a series of questions related to a configuration activity, 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 re-enter 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.

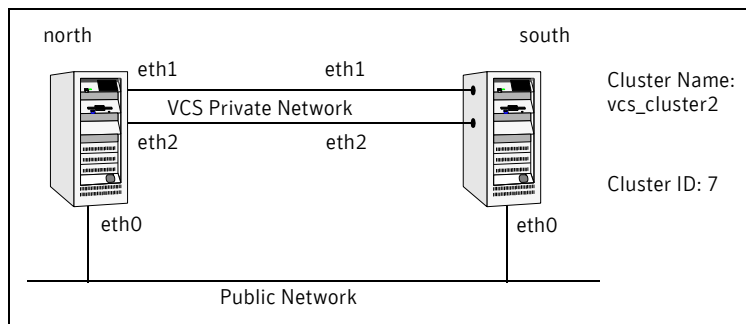
See [“Installing the VCS Java Console”](#) on page 80.

Installing and configuring VCS 5.0

[Figure 3-8](#) illustrates the systems on which you would install and run VCS. The example installation demonstrates how to install VCS on two systems: north and south. 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.

See “[Sample VCS installation and configuration output](#)” on page 219.

Figure 3-8 An example of a VCS installation on a two-node cluster



Overview of tasks

[Table 3-6](#) lists the installation and configuration tasks.

Table 3-6 Installation and configuration tasks

Task	Reference
Start the installation process and choose the installation	<ul style="list-style-type: none"> ■ “Checking the systems for installation” on page 65 (optional) ■ “Starting the software installation” on page 66 ■ “Specifying systems for installation” on page 67 ■ “Licensing VCS” on page 67 ■ “Choosing VCS RPMs” on page 68 ■ “Choosing to install VCS RPMs or configure VCS” on page 69

Table 3-6 Installation and configuration tasks

Task	Reference
Configure the cluster and optional features	<ul style="list-style-type: none"> ■ “Configuring the cluster” on page 69 ■ “Configuring the cluster in secure mode” on page 70 (optional) ■ “Adding VCS users” on page 72 (optional) ■ “Configuring cluster connector” on page 72 (optional) ■ “Configuring the Cluster Management Console” on page 73 (optional) ■ “Configuring SMTP email notification” on page 74 (optional) ■ “Configuring SNMP trap notification” on page 76 (optional) ■ “Configuring global clusters” on page 77 (optional)
Install the RPMs and create configuration files	<ul style="list-style-type: none"> ■ “Installing the VCS RPMs” on page 78 ■ “Creating VCS configuration files” on page 78
Start VCS and its components	<ul style="list-style-type: none"> ■ “Starting VCS” on page 79 ■ “Completing the installation” on page 79
Perform the post-installation tasks	<ul style="list-style-type: none"> ■ “Copying the installation guide to each node” on page 80 ■ “Setting up I/O fencing” on page 80 ■ “Installing the VCS Java Console” on page 80 ■ “Establishing cluster communication with the management server” on page 82 ■ “Installing cluster connector” on page 83
Verify the cluster	<ul style="list-style-type: none"> ■ “Verifying the cluster after installation” on page 92

Checking the systems for installation

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

```
installvcs -precheck system1 system2 ...
```

To check the systems

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

```
# cd /cdrom/cluster_server
```

- 2 Start the pre-installation check:

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

The program proceeds in a non-interactive mode, examining the systems for licenses, RPMs, disk space, and system-to-system communications.
- 3 Review the output as the program displays the results of the check and saves the results of the check in a log file.
See “[About installvcs command options](#)” on page 102.

Starting the software installation

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

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 containing 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: **north south**
For a single node installation, enter one name for the system.
See [“Starting the installer for the single node cluster”](#) on page 192.
- 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 5.0.
See [“Upgrading VCS using installvcs program”](#) on page 158.

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 100.

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 north: [?] **XXXX-XXXX-XXXX-XXXX-XXX**
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on north
VCS license registered on north
- 3 Enter keys for additional product features.
Do you want to enter another license key for north? [y,n,q,?]
(n) **y**

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

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

- 4 Review the output as the 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 south  
VCS license registered on south
```

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

Choosing VCS RPMs

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 to be installed.

```
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 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 cluster](#)” on page 69.
- Install RPMs on the systems and leave the cluster configuration steps for later.

To install VCS RPMs 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 RPMs.
- 3 Configure the cluster later.
See “[Configuring VCS using configure option](#)” on page 92.

Configuring the cluster

The installer provides you an option to configure VCS and its optional features.

Note: You can use `installvcs -configure` command to configure the cluster later and enter the system names where you want to configure VCS when the installer prompts you.

To configure the cluster

- 1 Enter **y** or press **Enter** at the prompt to configure VCS.
It is optional to configure VCS now. If you choose to configure VCS later, you can either do so manually or run the `installvcs -configure` command.
Are you ready to configure VCS?
[y,n,q] (y) **y**
- 2 Review the configuration requirements that the installer lists.
- 3 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**

- 4 Review the NICs available on the first system as the installer discovers and reports them.
- 5 Enter the details for the private heartbeat links.
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 north:
[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 north:
[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)
- 6 Choose whether to use the same NICs on all nodes.
 - If you want to use the same NICs for private heartbeat links on all nodes, make sure the same NICs are available on each system and enter **y**.
 - Enter **n** to use NICs with different device names on some of the nodes.
Are you using the same NICs for private heartbeat links on all systems? [y,n,q,b,?] (y)
- 7 Verify and confirm the information that the installer summarizes.

Configuring the cluster in secure mode

Before you configure a cluster in a secure mode, make sure to meet the requirements for automatic or semiautomatic mode of configuration. You can also enable Symantec Product Authentication Service later.

See [“Symantec Product Authentication Service”](#) on page 22.

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 have installed the root broker and enter **y**.
 - If you do not want to configure the cluster in secure mode and want to proceed to adding VCS users, enter **n**.
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,?]
Based on the mode of configuration you want to use, enter one of the following:

Option	Tasks
1. Automatic configuration	<p>Enter the name of the Root Broker system when prompted.</p> <p>Requires remote access to the Root Broker.</p> <p>Review the output as the installer verifies communication with the Root Broker system, checks vxatd process and version, and checks security domain.</p>
2. Semi-automatic using encrypted files	<p>Enter the path of the file for each node when prompted.</p>
3. Semi-automatic entering authentication information at installer prompts	<p>Enter the following Root Broker information as the installer prompts you:</p> <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> <p>Enter the following Authentication Broker information as the installer prompts you for each node:</p> <pre> Enter authentication broker principal name on north [b] (north.symantecexample.com) north.symantecexample.com Enter authentication broker password on north: Enter authentication broker principal name on south [b] (south.symantecexample.com) south.symantecexample.com Enter authentication broker password on south: </pre>
3	<p>After configuring the cluster in secure mode, proceed to configure the Cluster Management Console cluster connector.</p> <p>See “Configuring cluster connector” on page 72.</p>

Adding VCS users

If you have enabled Symantec Product Authentication Service, you do not need to add VCS users now.

Proceed to configure the Cluster Management Console cluster connector. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

See “[Configuring the cluster in secure mode](#)” on page 70.

See “[Configuring cluster connector](#)” on page 72.

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 cluster connector

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

See “[Veritas Cluster Management Console](#)” on page 25.

To configure cluster connector

- 1 Review the information to configure Cluster Management Console.
- 2 Choose whether to configure cluster connector or not. Do one of the following:
 - To configure cluster connector on the systems, press **Enter**.
Do you want this cluster to be managed by a management server? Enter 'y' if you have set up a management server.
[y,n,q] (y) **y**
 - To skip configuring cluster connector and advance to configuring Cluster Management Console for local cluster management, enter **n**.
See “[Configuring the Cluster Management Console](#)” on page 73.
- 3 Review the required information to configure cluster connector.
- 4 Enter the management server network address for the Cluster Management Console.
Enter the network address used by the management server [?]
(north) **mgmtserver1.symantecexample.com**
- 5 Verify and confirm the management server information.
- 6 Enter the following information that is required to securely communicate with the management server.
 - Password for the service account that is created during the management server installation
 - Hash of the Cluster Management Console management server's root broker
- 7 Verify and confirm the information.

Configuring the Cluster Management Console

If you want to locally manage this cluster, then you must configure the Cluster Management Console. Note that this cluster can also be a part of the clusters that are centrally managed by the management server.

See “[Veritas Cluster Management Console](#)” on page 25.

To configure the Cluster Management Console

- 1 Review the required information to configure the Cluster Management Console.
- 2 Choose whether to configure the Cluster Management Console or not. Do one of the following:
 - To configure the Cluster Management Console on the systems, press **Enter**.

```
Do you want to configure the Cluster Management Console  
[y,n,q] (y)
```

- To skip configuring the Cluster Management Console and advance to configuring SMTP, enter **n**.

See “[Configuring SMTP email notification](#)” on page 74.

- 3 Confirm whether you want to use the discovered public NIC on the first system. Do one of the following:

- If the discovered NIC is the one to use, press **Enter**.
- If you want to use a different NIC, type the name of a NIC to use and press **Enter**.

```
Active NIC devices discovered on north: eth0  
Enter the NIC for Cluster Management Console to use on north:  
[b,?] (eth0)
```

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

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

```
Is eth0 to be the public NIC used by all systems [y,n,q,b,?] (y)
```

- 5 Enter the virtual IP address for the Cluster Management Console.

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

- 6 Confirm the default netmask or enter another one:

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

- 7 Verify and confirm the Cluster Management Console information.

```
Cluster Management Console verification:
```

```
NIC: eth0  
IP: 10.10.12.1  
Netmask: 255.255.240.0
```

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

Configuring SMTP email notification

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

To configure SMTP email notification

- 1 Review the required information to configure the SMTP email notification.

- 2 Specify whether you want to configure the SMTP notification. Do one of the following:
 - To configure SMTP notification, press **Enter**.
Do you want to configure SMTP notification? [y,n,q] (y) **y**
 - To skip configuring SMTP notification and advance to configuring SNMP notification, enter **n**.
See “[Configuring SNMP trap notification](#)” on page 76.
- 3 Provide information to configure SMTP notification.
 - Enter the SMTP server’s host name.
Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,?] **smtp.example.com**
 - Enter the email address of each recipient.
Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,?] **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 it is also possible to configure notification after installation. Refer to the *Veritas Cluster Server User's Guide* for more information.

To configure the SNMP trap notification

- 1 Review the required information to configure the SNMP notification feature of VCS.
- 2 Specify whether you want to configure the SNMP notification. Do one of the following:
 - To configure SNMP notification, press **Enter**.
Do you want to configure SNMP notification? [y,n,q] (y)
 - To skip configuring SNMP notification and advance to configuring global clustering option, enter **n**.
See “[Configuring global clusters](#)” on page 77.
- 3 Provide information to configure SNMP trap notification.
 - Enter the SNMP trap daemon port.
Enter the SNMP trap daemon port: [b,?] (162)
 - Enter the SNMP console system name.
Enter the SNMP console system name: [b,?] **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

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

To configure the global cluster option

- 1 Review the required information to configure the global cluster option.
- 2 Specify whether you want to configure the global cluster option. Do one of the following:
 - To configure global cluster option, press **Enter**.

```
Do you want to configure the Global Cluster Option? [y,n,q]  
(y)
```
 - To skip configuring global cluster option and advance to installing VCS RPMs, enter **n**.
See [“Installing the VCS RPMs”](#) on page 78.
- 3 Provide information to configure the Global Cluster option.
If you configured Cluster Management Console to manage this cluster locally, the installer discovers and displays the virtual IP address and netmask used by the Cluster Management Console. You can use the same virtual IP address and netmask.
See [“Configuring the Cluster Management Console”](#) on page 73.
Do one of the following:
 - If you want to use the default values, press **Enter**.
 - If you do not want to use the default value, enter another IP address.
The installer prompts you for a NIC and value for the netmask.

```
Enter the Virtual IP address for Global Cluster Option:  
[b,?] (10.10.12.1)
```
- 4 Verify and confirm the configuration of the global cluster.

```
Global Cluster Option configuration verification:  
  
NIC: eth0  
IP: 10.10.12.1  
Netmask: 255.255.240.0  
  
Matching Cluster Management Console Virtual IP configuration  
  
Is this information correct? [y,n,q] (y)
```

Installing the 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 78.

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. Review the output as the installer uninstalls any previous versions and installs the VCS 5.0 RPMs.

Creating VCS configuration files

After installing the RPMs, the installer continues to create configuration files and copies them to each system:

```
Creating Cluster Server configuration files ..... Done
Copying configuration files to north..... Done
Copying configuration files to south..... 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 creates security principal or executes the encrypted file to create security principal on each node in the cluster. The installer creates the VxSS service group, creates Authentication Server credentials on each node in the cluster, and Web credentials for VCS users, and sets up trust with the root broker. Then, the installer proceeds to start VCS in secure mode.

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
north:
```

```
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 “[Configuring SuSE network interfaces](#)” on page 47.

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.

To start VCS

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

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

Completing the installation

After VCS 5.0 installation completes successfully, the installer creates summary, log, and response files. The files provide useful information that can assist you with the installation and can also assist future installations. [Table 3-7](#) specifies the files created at the end of the installation.

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

Table 3-7 File description

File	Description
summary file	<ul style="list-style-type: none">■ Lists RPMs installed on each system.■ Describes the cluster and its configured resources.■ Provides information for managing the cluster.
log file	Details the entire installation.
response file	Contains configuration information that can be used to perform secure or unattended installations on other systems. See “Example response file” on page 95.

Copying the installation guide to each node

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

Setting up I/O fencing

Symantec recommends you to set up the I/O fencing feature to prevent data corruption in the event of a communication breakdown in the cluster. Make sure that you do the following before you set up I/O fencing:

- Install a version of Veritas Volume Manager (VxVM) that supports SCSI-3 persistent reservations.
- Verify that the disks you intend to use for shared data storage and for coordinator disks support SCSI-3 PR (Persistent Reservations).

See [“Setting up I/O fencing”](#) on page 111.

Installing the VCS 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 or Windows 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; the latter enables each system in the cluster to be administered remotely.

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

Hardware requirements for the Java Console

The following items are the minimum hardware requirements for the Java Console:

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

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

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

Installing the Java Console on Linux

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 containing 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 VRTScscm-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 workstation or a Windows 2000 Professional Workstation to administer the cluster.

To install the Java Console on a Windows system

- 1 Insert the software disc with the VCS software into a drive on your Windows system.
- 2 Using Windows Explorer, select the disc drive.
- 3 Go to \windows\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.

Establishing cluster communication with the management server

Use the following list to prepare clusters for administration and management through the Cluster Management Console.

- Ensure that all clusters that you want to manage run a supported version of VCS.
[“Supported software for the Veritas Cluster Management Console”](#) on page 33.
- Decide which clusters are to use cluster connector to communicate with the management server, and then install cluster connector on each cluster. Cluster connector is a process agent. You must use cluster connector if a firewall exists between the management server and any clusters. You can install cluster connector when you install VCS 5.0. For other VCS versions, you must install cluster connector on each cluster node to establish a connection to the management server. Only certain versions of VCS can connect to the management server. See [“Installing cluster connector”](#) on page 83.
- Decide which clusters are to use a direct connection to communicate with the management server. If these clusters run a supported version of VCS, they require no further preparation. Refer to the *Veritas Cluster Server Centralized Management Guide* for more information on direct connection.

After you prepare the clusters for management server administration, start the Cluster Management Console and use it to configure the management server to connect to clusters using direct connection. Clusters using cluster connector connect to the management server automatically.

Refer to the *Veritas Cluster Server Centralized Management Guide* for more information on cluster connector.

Installing cluster connector

Cluster connector is a process agent, which you must use if a firewall exists between the management server and any clusters. You can install cluster connector on UNIX-based or Windows-based clusters. You can also use a batch feature to install cluster connector on UNIX-based clusters.

Note: You must stop all VCS Web consoles, VCS Java consoles, and agent wizards that are running on any cluster nodes before you install cluster connector.

For all cluster connector installations, cluster connector installs or upgrades Symantec Product Authentication Service (version 4.3) on target systems that do not already have it.

For Windows, you must install the authentication broker on the local system (the system from which you run the cluster connector installation). You can install cluster connector from the management server or from any other system that runs the authentication broker. For UNIX, one of the following two conditions must be true:

- You are installing cluster connector (locally or remotely) from a system running the authentication broker.
- You are installing cluster connector (locally or remotely) from a cluster node and that cluster node is in your install list.

For example, assume that you have nodes A, B, and C each in their own cluster; each have the authentication broker installed. You also have system X with no authentication broker. You cannot install cluster connector from X. You can install cluster connector from A to B and C to other nodes.

Installing the cluster connector on UNIX systems

Perform this procedure to use cluster connector for management server communications when the cluster is a supported VCS cluster. You can also use this procedure if you want to install or configure cluster connector after installing VCS 5.0 on a cluster.

To install cluster connector on a UNIX system

- 1 Insert the distribution disc into the drive on the local system. At the command prompt, type the following command to run the installer program:

```
./installer -rsh
```

The installer program presents copyright information followed by a menu titled, "Storage Foundation and High Availability Solutions 5.0".

- 2 Enter **i** to specify a task.

```
Enter a Task: [I,C,L,P,U,D,Q,?] i
```

The installer displays another menu that lists products that are available for installation.

- 3 Enter the menu number that corresponds to **Veritas Cluster Management Console**.

```
Select a product to install:nn
```

The installer presents a description of the product.

- 4 Enter **2** if you are prompted to select a product component.

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

The installer presents a message stating that it will install cluster connector. Note that option 1 only appears on Solaris systems.

- 5 Enter the name of one system in each cluster to be managed. Separate the system names with spaces.

```
Storage Foundation and High Availability Solutions 5.0
```

```
Enter the name of a system in each cluster that you want the  
management server to manage. Separate system names with spaces:
```

```
system1 system2 system3
```

The installer detects the systems that you enter, performs an initial check of those systems, and then checks for installed packages on those systems. If these checks are satisfactory, the installer lists the packages to be installed.

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

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

The installer performs an initial system check of the local system and checks for installed packages on the local system. If these checks are satisfactory, the installer program lists the packages to be installed.

```
Storage Foundation and High Availability Solutions 5.0
```

```
installer will install the following CMC packages:
```

```
VRTSat          Symantec Product Authentication Service  
VRTSperl        Veritas Perl 5.8.8 Redistribution  
VRTSjre15       Veritas Java Runtime Environment Redistribution  
VRTScmccc       Veritas Cluster Management Console Cluster  
Connector
```

```
Press [Return] to continue:
```

- 7 Press Enter.
You may install Cluster Management Console packages without performing configuration. The installer program gives you the option to configure Cluster Management Console now, and provides instructions for configuring Cluster Management Console later.
- 8 Enter **y** to configure Cluster Management Console.
Are you ready to configure CMC? [y,n,q] (y)
- 9 Enter the fully-qualified management server network address, such as:
Enter the network address used by the management server [?]
mgmtserver1.symantec.com
- 10 Enter **y** to verify that the information up to this point is correct.
Is this information correct? [y,n,q] (y)
- 11 Enter a password for the cluster connector service account.
The password is the password that was entered for the cluster connector service account during management server installation.
Enter the password for the CMC service account: ~~xxxxxx~~
- 12 Enter the root hash of the authentication broker installed on the management server, which you can get from the Root Broker administrator.
Enter the hash of the Management Server's root broker [?]
 - On Windows:
\program files\veritas\security\authentication\bin\vssat showbrokerhash
 - On UNIX systems:
/opt/VRTSat/bin/vssat showbrokerhashThe output of this command looks similar to the following:
Root Hash: 9dfde3d9aaebee084f8e35819c1fed7e6b01d2ae
Enter the alphanumeric string (the string you receive is different from the one shown).
- 13 Enter **y** to verify that the information up to this point is correct.
Is this information correct? [y,n,q] (y)
The installer presents:
 - Installation progress percentages
 - Authentication status messages
 - Cluster connector configuration status messages
- 14 Enter **y** to start Veritas Cluster Management Console processes.
Do you want to start Veritas Cluster Management Console processes now?
[y,n,q] (y)
The installer presents startup progress percentages and, if successful, displays the following message:

```
Startup completed successfully on all systems
```

- 15 Enter an encryption key of at least five characters.

This key is used to encrypt passwords in the response file. It must be referred to using the `-enckeyfile` option if the generated installation response file is to be used again.

```
A string of five or more characters is required to encrypt passwords in the responsefile
```

```
Enter five or more characters to be used an encryption key:
```

```
xxxxx
```

```
Press [Return] to continue:
```

- 16 Press Enter.

Record the location that the installer program provides for the installation log files, summary file, and response file.

Batch installation of cluster connector on UNIX systems

This process is available for a new installation on supported UNIX clusters and for the upgrade of UNIX clusters running CommandCentral Availability cluster monitor.

To perform a batch installation of cluster connector on UNIX

- 1 Insert the distribution disc into the drive on the local system. At the command prompt, type the following command to run the installer program:

```
./installer -rsh
```

The installer program presents copyright information followed by a menu titled, "Storage Foundation and High Availability Solutions 5.0".

- 2 Enter **i** to specify a task.

```
Enter a Task: [I,C,L,P,U,D,Q,?] i
```

The installer displays another menu that lists products that are available for installation.

- 3 Enter the menu number that corresponds to Veritas Cluster Management Console.

```
Select a product to install:nn
```

The installer presents a description of the product. The installer may also display the prompt:

- 4 Enter **2** if you are prompted to select a product component.

The installer presents a message stating that it will install cluster connector.

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

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

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

The installer performs an initial system check of the local system and checks for installed packages on the local system. If these checks are satisfactory, the installer program lists the packages to be installed.

Storage Foundation and High Availability Solutions 5.0

installer will install the following CMC packages:

VRTSat Symantec Product Authentication Service

VRTSperl Veritas Perl 5.8.8 Redistribution

VRTSjre15 Veritas Java Runtime Environment Redistribution

VRTScmccc Veritas Cluster Management Console Cluster

Connector

Press [Return] to continue:

- 6 Enter the name of one system in each cluster to be managed. Separate the system names with spaces.

Storage Foundation and High Availability Solutions 5.0

Enter the name of a system in each cluster that you want the management server to manage. Separate system names with spaces:

system1 system2 system3

The installer detects the systems that you enter, performs an initial check of those systems, and then checks for installed packages on those systems. If these checks are satisfactory, the installer lists the packages to be installed.

Storage Foundation and High Availability Solutions 5.0

The following CMC packages will be installed:

VRTSat Symantec Product Authentication Service

VRTSperl Veritas Perl 5.8.8 Redistribution

VRTSjre15 Veritas Java Runtime Environment Redistribution

VRTScmccc Veritas Cluster Management Console cluster

connector

Press [Return] to continue:

- 7 Press Enter.

You may install Cluster Management Console packages without performing configuration. The installer program gives you the option to configure Cluster Management Console now and provides instructions for configuring Cluster Management Console later.

- 8 Enter **y** to configure Cluster Management Console.

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

- 9 Enter the fully-qualified management server network address.

Enter the network address used by the management server [?]

mgmtserver1.symantec.com

The installation program repeats the management server address.

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

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

- 11 Enter a password for the cluster connector service account.
The password is the password that was entered for the cluster connector service account during management server installation.
Enter the password for the CMC service account: **xxxxxx**
- 12 Enter the root hash of the authentication broker installed on the management server.
Enter the hash of the Management Server's root broker [?]
To retrieve the root hash of the management server authentication broker, run the following command:
 - On Windows:
`\program files\veritas\security\authentication\bin\vssat showbrokerhash`
 - On UNIX systems:
`/opt/VRTSat/bin/vssat showbrokerhash`The output of this command looks similar to the following:
Root Hash: 9dfde3d9aaebec084f8e35819c1fed7e6b01d2ae
Enter the alphanumeric string (the string you receive is different from the one shown).
- 13 Enter **y** when prompted to verify that the information up to this point is correct.
Is this information correct? [y,n,q] (y)
The installer presents:
 - Installation progress percentages
 - Authentication status messages
 - Cluster connector configuration status messages
- 14 Enter **y** to start Veritas Cluster Management Console processes now.
Do you want to start Veritas Cluster Management Console processes now?
[y,n,q] (y)
The installer program presents startup progress percentages and, if successful, displays the following message:
Startup completed successfully on all systems
- 15 Enter an encryption key of at least five characters.
This key is used to encrypt passwords in the response file. It must be referred to using the `-enckeyfile` option if the generated installation response file is to be used again.
A string of five or more characters is required to encrypt passwords in the responsefile
Enter five or more characters to be used an encryption key:
xxxxx
Press [Return] to continue:

16 Press Enter.

Record the location that the installer program provides for the installation log files, summary file, and response file.

Installing cluster connector on Windows platforms

Cluster Management Console can manage VCS clusters that run on Windows. You can manage these clusters using cluster connector or direct connect just like the UNIX clusters.

Note: When installing cluster connector on 64-bit Windows platforms from a 32-bit system, the default install directory will show up as c:\Program Files. Symantec recommends that you change the 64-bit installation directory to C:\Program Files (x86).

To install cluster connector on Windows

- 1 Insert the distribution disc into the DVD drive on the local system.
- 2 On the distribution disc, locate the locate the \installer directory under \windows\cluster management console directory.
- 3 Double-click the setup.bat file.
Depending upon the operating system, you may or may not receive the following warning message:

```
The publisher could not be verified. Are you sure you want to run this software?
```


If you receive this message, click **Run**.
- 4 In the Welcome to the Veritas Cluster Management Console Installation Manager dialog box, read the introduction and then click **Next**.
- 5 In the Installation and Configuration Options dialog box, click **Add Clusters or clustered systems to a management server**, and then click **Next**.
- 6 In the Cluster Connector Cluster Selection dialog box, follow the dialog box instructions exactly as specified, and then click **Next**.
The installer performs a check for WMI on the specified nodes to ensure that they are ready for cluster connector installation.
- 7 When prompted, enter user account information for each cluster. If a cluster is secure, you are prompted for a domain name in addition to a user name and password that is valid for the cluster.
- 8 In the Cluster Connector Directory Selection dialog box, do one of the following:
 - Leave the default directories provided

- Click a directory and then click **Select** to specify another directory
 - Click **Reset** all to specify new directories on each node
- Click **Next** to accept the directories.

- 9 In the Management Server Information dialog box, provide the IP address for the management server to which cluster connector is intended to connect.

You cannot change the port specification, 14145, but it is provided to help you to prevent port conflicts when configuring other software. The other ports used by the Cluster Management Console are 8181 (HTTP), 8443 (HTTPS), and 2994 (DBMS; this port can be shared with other Symantec products)

- 10 In the Services Account Password dialog box:

- Enter a password for the user account that cluster connector uses for management server communications
- Enter the root hash of the authentication broker used by the authentication broker installed on the management server

The password is the password that was entered for the cluster connector service account during management server installation.

To retrieve the root hash of the management server authentication broker, run the following command:

- On Windows:
`\program files\veritas\security\authentication\bin\vssat showbrokerhash`
- On UNIX systems:
`/opt/VRTSat/bin/vssat showbrokerhash`

The output of this command looks similar to the following:

```
Root Hash:          9dfde3d9aaebee084f8e35819c1fed7e6b01d2ae
```

Enter or copy the alphanumeric string into the Root Hash text box (the string you receive is different from the one shown).

- 11 In the Summary dialog box, review the information you have specified and, if satisfactory, click **Next** to accept it and start the installation.

The Installing Veritas Cluster Management Console dialog box displays a progress bar and a status message window for the installation.

- 12 When you receive the following message, click **Next**:

```
"Done deleting installation files from node..."
```

- 13 In the Completed the Symantec Veritas Cluster Management Console Installation Manager dialog box, click **Finish**.

The installer creates log files at C:\Documents and Settings\All Users\Application Data\VERITAS\Cluster Management Console. The file names are Install_GUI_0.log and Install_MSI_0.log. The installer creates

Install_GUI_0.log on the system from which you run the cluster connector installation. The installer creates Install_MSI_0.log on the target systems.

Avoiding potential service group faults on secure Windows clusters

During cluster connector installation on secure Windows clusters, the CMC_CC@CMC_SERVICES service account fails to authenticate on the cluster nodes. After installation, the ClusterConnector resource faults on the node where the CMC_CC@CMC_SERVICES account is not authenticated. This result causes the CMC service group to fault. The ClusterConnector.log file contains the error message:

```
Can not get Cache Credential for CMC_CC
```

To avoid potential service group faults on secure Windows clusters

- 1 At a command prompt, change directory to the following directory:
Veritas\Security\Authentication\bin
This may be in
C:\Program Files or C:\Program Files\Common Files
- 2 Run the following command to verify if the CMC_CC credential was created:
vssat showcred | grep CMC_CC
- 3 Run the following command:
"vssat authenticate --domain vx:CMC_SERVICES --prplname CMC_CC
--password password_for_CMC_CC_user_created_during_MS_install
--broker MS_IPAddress:2821
Usage for this command is
vssat authenticate --domain <type:name> [--prplname <prplname>
[--password <password>]] [--broker <host:port>]
- 4 Repeat these steps on each node of the cluster.

Accessing Cluster Management Console information

Information about administering clusters in multi-cluster mode is available in the Veritas Cluster Server Centralized Management Guide. The online help includes information about administering clusters in both single-cluster and multi-cluster mode. If you want to access the information about managing a single cluster in printed documentation, you can install the documentation package to the desired system.

The documentation package name for each supported operating system is:

- AIX-VRTSvcs.doc
- HP-UX-VRTSvcsdc

Note that you can copy the documents from depot/VRTSvcsdc/ VRTSvcsdc/
opt/ VRTS/docs.

- Linux—VRTSvcsdc
- Solaris—VRTSvcsdc

Verifying the cluster after installation

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

See [“Verifying the VCS installation”](#) on page 143.

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 entered without creating any VCS configuration files.

Configuring VCS using configure option

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

See [“Configuring the cluster”](#) on page 69.

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

Performing VCS installation in a secure environment

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

See [“Example response file”](#) on page 95.

Note that a response file generated by the `installvcs` program contains descriptions and explanations of the variables and their values. By copying this

file to the other systems in the cluster and editing it to reflect the current local system, you can use the installation program with the `-responsefile` option to install and configure VCS identically on each system without being prompted.

To use `installvcs` in a secure environment

- 1 On one node in the cluster, start VCS installation using the `installvcs` program.
See “[Starting the software installation](#)” on page 66.
- 2 Review the output as the installer performs the initial system checks.
The installer detects the inability to communicate between systems.
- 3 Press Enter to install VCS on one system and create a response file with which you can install on other systems.

```
Would you like to install Cluster Server on systems north only  
and create a responsefile for systems south? [y,n,q] (y)
```
- 4 Enter all cluster information. Proceed with the installation and configuration tasks.
See “[Installing and configuring VCS 5.0](#)” on page 63.
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 response file within the file `/opt/VRTS/install/logs/installvcs-universaluniqueidentifier/installvcs-universaluniqueidentifier.response`.
- 6 If you start VCS before VCS is installed and started on all nodes in the cluster, you will see the output similar to:

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

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


For demo or site licenses, the license key need not be changed.
- 9 On the next system:

- Mount the product disc.
See “[Mounting the product disc](#)” on page 55.
- Start the software installation using the `installvcs -responsefile` option.

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

Where *uui* is the Universal Unique Identifier that the installer automatically assigned to the response file.
See “[Starting the software installation](#)” on page 66.

10 Repeat [step 7](#) through [step 9](#) until VCS has been installed on all nodes in the cluster.

Performing automated installations

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

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

To perform unattended installation

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

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

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

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

Syntax used in response file

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

For example,

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

or, in the case of an integer value:

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

or, in the case of a list:

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

Example response file

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

See [Table 3-8, "Response file variables."](#)

```
#
# installvcs configuration values:
#
$CPI::CFG{AT_ROOTDOMAIN}="root\@east.symantecexample.com";
$CPI::CFG{CMC_CC_CONFIGURED}=1;
$CPI::CFG{CMC_CLUSTERID}{east}=1146235600;
$CPI::CFG{CMC_MSADDR}{east}="mgmtserver1";
$CPI::CFG{CMC_MSADDR}{west}="mgmtserver1";
$CPI::CFG{CMC_MS_ROOT_HASH}="758a33dbd6fae751630058ace3dedb54e5
62fe98";
$CPI::CFG{CMC_SERVICE_PASSWORD}="U2FsdGVkX18vE5tn0hTSWwodThACc+
rX";
$CPI::CFG{ENCRYPTED}="U2FsdGVkX1+k2DHKVcnW7b6vrVghdh+zW4G0WFj5I
JA=";
$CPI::CFG{KEYS}{east}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{KEYS}{west}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{OBC_IGNOREWARNINGS}=0;
$CPI::CFG{OBC_MODE}="STANDALONE";
$CPI::CFG{OPT}{INSTALL}=1;
$CPI::CFG{OPT}{NOEXTRAPKGS}=1;
$CPI::CFG{OPT}{RSH}=1;
$CPI::CFG{SYSTEMS}=[ qw(east west) ];
$CPI::CFG{UPI}="VCS";
$CPI::CFG{VCS_ALLOWCOMMS}="Y";
$CPI::CFG{VCS_CLUSTERID}=13221;
$CPI::CFG{VCS_CLUSTERNAME}="vcs_cluster3";
$CPI::CFG{VCS_CSGNETMASK}="255.255.240.0";
$CPI::CFG{VCS_CSGNIC}{ALL}="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 3-8 lists the variables used in the response file and their definitions. Note that while some variables are labeled as required and others as optional, some of the optional variables, if used, make it necessary to define other optional variables. For example, all variables related to the cluster service group (CSGNIC, CSGVIP, and CSGNETMASK) must be defined if any are defined. The same is true for the SMTP notification (SMTPSERVER, SMTPRECP, and SMTPRSEV), SNMP trap notification (SNMPPORT, SNMPCONS, and SNMPCSEV), and the Global Cluster Option (CGONIC, GCOVIP, and GCONETMASK).

Table 3-8 Response file variables

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

Table 3-8 Response file variables

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

Table 3-8 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
\$CPI::CFG{OPT_LOGPATH}	Scalar	Opt'l	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.
\$CPI::CFG{CONFIGURE}	Scalar	Opt'l	Performs configuration if the RPMs are already installed using the <code>-installonly</code> option.
\$CPI::CFG{VCS_LLLINK#} {SYSTEM}	Scalar	Req'd	Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (LLTLINK1 and LLTLINK2). Up to four LLT links can be configured.
\$CPI::CFG{VCS_LLLINKLOWPRI} {SYSTEM}	Scalar	Opt'l	Defines a low priority heartbeat link. Typically, LLLINKLOWPRI is used on a public network link to provide an additional layer of communication.
\$CPI::CFG{VCS_CSGNIC}	Scalar	Opt'l	Defines the NIC for 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.
\$CPI::CFG{CSGVIP}	Scalar	Opt'l	Defines the virtual IP address to be used by the Cluster Management Console.
\$CPI::CFG{VCS_CSGNETMASK}	Scalar	Opt'l	Defines the Netmask of the virtual IP address to be used by the Cluster Management Console.
\$CPI::CFG{VCS_SMTPSERVER}	Scalar	Opt'l	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for web notification.

Table 3-8 Response file variables

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

Table 3-8 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
\$CPI::CFG{VCS_USERNAME}	List	Opt'l	List of names of users
\$CPI::CFG{VCS_USERPRIV}	List	Opt'l	List of privileges for users
\$CPI::CFG{OPT}{UNINSTALL}	Scalar	Opt'l	List of systems where VCS must be uninstalled.

Checking licensing information on the system

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

To check licensing information

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

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

- 2 Review the output to determine:

- The license key
- The type of license
- The product for which it applies
- Its expiration date, if any. Demo keys have expiration dates, permanent keys and site keys do not.

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

Features :=
Platform             = Linux
Version              = 5.0
Tier                 = 0
Reserved             = 0

Mode                 = VCS
```

Updating product licenses using vxlicinst

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

See “[Replacing a VCS demo license with a permanent license](#)” on page 101.

To update product licenses

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

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

Replacing a VCS demo license with a permanent license

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

To replace a demo key

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

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

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

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

```
# hstart
```

About installvcs command options

[Table 3-9](#) lists the `installvcs` command options. In addition to the `-precheck`, `-responsefile`, `-installonly`, and `-configure` options, the `installvcs` program has other useful options.

The `installvcs` command usage takes the following form:

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

Table 3-9 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 92.
<code>-enckeyfile</code> <i>encryption_key_file</i>	See the <code>-responsefile</code> and the <code>-encrypt</code> options.
<code>-encrypt password</code>	Encrypt <i>password</i> using the encryption key provided with the <code>-enckeyfile</code> option so that the encrypted password can be stored in response files.
<code>-installonly</code>	Install product RPMs on systems without configuring VCS. See “Installing VCS using installonly option” on page 92.
<code>-installpkgs</code>	Display VCS packages in correct installation order. Output can be used to create scripts for command line installs, or for installations over a network. See the <code>requiredpkgs</code> option.
<code>-keyfile</code> <i>ssh_key_file</i>	Specifies a key file for SSH. The option passes <code>-i ssh_key_file</code> with each SSH invocation.
<code>-license</code>	Register or update product licenses on the specified systems. Useful for replacing demo license.
<code>-logpath log_path</code>	Specifies that <i>log_path</i> , not <code>/opt/VRTS/install/logs</code> , is the location where <code>installvcs</code> log files, summary file, and response file are saved.
<code>-noextrapkgs</code>	Specifies that additional product RPMs such as VxVM and VxFS need not be installed. Note: VCS product upgrades in the future can be simplified if you do not install additional product RPMs.

Table 3-9 installvcs options

Option and Syntax	Description
-nolic	Install product RPMs on systems without licensing or configuration. License-based features or variants are not installed when using this option.
-nooptionalpkgs	Specifies that the optional product RPMs such as man pages and documentation need not be installed.
-nostart	Bypass starting VCS after completing installation and configuration.
-pkgpath <i>pkg_path</i>	Specifies that <i>pkg_path</i> contains all RPMs to be installed by installvcs program on all systems; <i>pkg_path</i> is the complete path of a directory, usually NFS mounted.
-precheck	Verify that systems meet the installation requirements before proceeding with VCS installation. Symantec recommends doing a precheck before installing VCS. See “Checking the systems for installation” on page 65.
-requiredpkgs	Displays all required VCS packages in correct installation order. Optional packages are not listed. Output can be used to create scripts for command line installs, or for installations over a network. See <code>installpkgs</code> option.
-responsefile <i>response_file</i> [-enckeyfile <i>encryption_key_file</i>]	Perform automated VCS installation using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. If not specified, the response file is automatically generated as <code>installerernumber.response</code> where <i>number</i> is random. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file. The <code>-enckeyfile</code> option and <i>encryption_key_file</i> name are required with the <code>-responsefile</code> option when the response file contains encrypted passwords. See “Performing VCS installation in a secure environment” on page 92. See “Performing automated installations” on page 94.

Table 3-9 `installvcs` options

Option and Syntax	Description
<code>-rsh</code>	Specifies that <code>rsh</code> and <code>rsh</code> are to be used for communication between systems instead of <code>ssh</code> and <code>scp</code> . This option requires that systems be pre-configured such that <code>rsh</code> commands between systems execute without prompting for passwords or confirmations
<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 “Symantec Product Authentication Service” on page 22.
<code>-tmp_path tmp_path</code>	Specifies that <code>tmp_path</code> , not <code>/var/tmp</code> , is the working directory for <code>installvcs</code> program. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.

About the uninstallecs program

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

If `uninstallecs` program does not remove an enterprise agent, see the documentation for the specific enterprise agent for instructions on removing it.

Prerequisites

- Before removing VCS from any node in the cluster, you must shut down applications such as Java Console or any VCS enterprise agents that depend on VCS.
- Before removing VCS from fewer than all nodes in a cluster, make sure that no service groups are running on the nodes from which VCS is uninstalled. You must also reconfigure VCS on the remaining nodes. See [“Adding and removing cluster nodes”](#) on page 179.

Uninstalling VCS 5.0

The example demonstrates how to uninstall VCS on two nodes: north and south. See “[Sample VCS installation and configuration output](#)” on page 219.

Removing VCS 5.0 RPMs

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

To uninstall VCS

- 1 Do one of the following to begin uninstalling:
 - If you can execute commands as superuser on the remote nodes in the cluster using `ssh` or `rsh` without supplying a password, run `uninstallvcs` program on one node to uninstall VCS on all nodes in the cluster.
 - If you cannot execute commands as superuser on remote nodes in the cluster using `ssh` or `rsh`, you must run `uninstallvcs` program on each node in the cluster.
- 2 Start `uninstallvcs` program.

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

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

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

Cluster Name: VCS_cluster2
Cluster ID Number: 7
Systems: north south
Service Groups: ClusterService groupA groupB
```
- 3 Answer the prompt to proceed with uninstalling the software.
 - To uninstall VCS on all nodes, press **Enter**.
 - To uninstall VCS only on specific nodes, enter **n**.
Note that if you enter **n** or if no VCS configuration files are found on the local node, the `uninstallvcs` program prompts you to enter a list of nodes from which you want to uninstall VCS.

```
Do you want to uninstall VCS from these systems? [y,n,q] (y)
```
- 4 Review the output as the `uninstallvcs` program continues to verify communication between systems and check the installations on each system to determine the RPMs to be uninstalled.

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

Running `uninstallvcs` from the VCS 5.0 disc

If you need to uninstall VCS after an incomplete installation, or if the `uninstallvcs` program is not available in `/opt/VRTS/install`, you may need to use the `uninstallvcs` program on the VCS 5.0 disc.

Uninstalling the Cluster Management Console management server

You must run the management server uninstallation on the management server host system.

Uninstalling the management server from Solaris systems

Use this procedure to remove the Cluster Management Console management server from the standalone management server host system.

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

To uninstall the management server from Solaris systems

- 1 Insert the product disc into the drive on the local system. At the command prompt, type the following command to run the installer program:
`./installer [-rsh]`
The installer program presents copyright information followed by a menu titled, "Storage Foundation and High Availability Solutions 5.0".
- 2 Enter **u** to specify uninstallation.
`Enter a Task: [I,C,L,P,U,D,Q,?] u`

The installer program displays another menu that lists products that are available for uninstallation.

- 3 Enter the menu number that corresponds to Veritas Cluster Management Console.
Select a product to uninstall: **nn**
The installer program presents a description of the product.
- 4 Enter **1** if you are prompted to select a product component.
Enter '1' to uninstall the Management Server, '2' to install the Cluster Connector: [1-2,q] (1) **1**
The installer program presents a message.
- 5 Enter **y** to verify that the information up to this point is correct.
Is this information correct? [y,n,q] (y)
It performs an initial system check of the local system and checks for installed packages on the local system. If these checks are satisfactory, the installer program lists the packages to be uninstalled.
- 6 Enter **y** to verify that you want to uninstall the management server.
Are you sure you want to uninstall CMC? [y,n,q] (y)
The installer program lists package dependencies and uninstallation progress percentages. If the uninstallation is successful, the installer program displays this message followed by the location of the uninstallation logs:
Uninstall completed successfully

Uninstalling the management server from Windows systems

Use this procedure to remove the Cluster Management Console management server from the standalone management server host system.

To uninstall the management server on Windows

- 1 On the Windows task bar, click **Start > Settings > Control Panel > Add or Remove Programs**.
- 2 In the Add or Remove Programs control panel, in the Currently installed programs list, click **Veritas Cluster Management Console Management Server**.
You may have to scroll through the list to find this entry.
- 3 On the right side of the selected entry, click **Remove**.
- 4 Follow the prompts in the uninstallation wizard, if any.

Uninstalling the Cluster Management Console cluster connector

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

Uninstalling cluster connector from UNIX systems

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

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

To uninstall cluster connector from UNIX systems

- 1 Insert the product disc into the drive on the local system. At the command prompt, type the following command to run the installer program:

```
./installer [-rsh]
```

The installer program presents copyright information followed by a menu titled, "Storage Foundation and High Availability Solutions 5.0".
- 2 Enter **u** to specify uninstallation.

```
Enter a Task: [I,C,L,P,U,D,Q,?] u
```

The installer program displays another menu that lists products that are available for uninstallation.
- 3 Enter the menu number that corresponds to Veritas Cluster Management Console.

```
Select a product to uninstall: nn
```

The installer program presents a description of the product.
- 4 Enter **2** if you are prompted to select a product component. Otherwise, proceed to [step 6](#).

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

The installer program presents a message stating that it will uninstall cluster connector.
- 5 The uninstall program prompts you for the name of at least one node in the cluster.

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

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

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

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

You must return and set up ssh.

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

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

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

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

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

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

```
Uninstall completed successfully
```

Uninstalling cluster connector from Windows platforms

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

To uninstall cluster connector on Windows

- 1 Insert the product disc into the DVD drive on the local system.
- 2 On the product disc, locate the \install directory for Cluster Management Console in the \windows folder.
- 3 Double-click the **setup.bat** file.
Depending upon the operating system, you may or may not receive the following warning message:

```
The publisher could not be verified. Are you sure you want to  
run this software?
```


If you receive this message, click **Run**.
- 4 In the Welcome to the Veritas Cluster Management Console Installation Manager dialog box, read the introduction and then click **Next**.
- 5 In the Installation and Configuration Options dialog box, click **Uninstall cluster connectors** and then click **Next**.
- 6 Follow the prompts in the uninstallation wizard. When available, click **Finish** to close the wizard.

Setting up I/O fencing

This chapter contains the following topics:

- [About I/O fencing](#)
- [Preparing to configure I/O fencing](#)
- [Setting up I/O fencing for VCS](#)
- [Additional I/O fencing information](#)
- [How I/O fencing works in different event scenarios](#)
- [About the vxfenadm utility](#)
- [Troubleshooting I/O fencing](#)

About I/O fencing

I/O Fencing is a feature that prevents data corruption in the event of a communication breakdown in a cluster.

Note: Symantec strongly recommends that you use VCS I/O fencing to deter potential split brain scenarios in your cluster.

Preventing data corruption with I/O fencing

To provide high availability, the cluster must be capable of taking corrective action when a node fails. In this situation, VCS configures its components to reflect the altered membership.

Problems arise when the mechanism that detects the failure breaks down because symptoms appear identical to those of a failed node. For example, if a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects and the remaining node takes corrective action. However,

the failure of private interconnects (instead of the actual nodes) would present identical symptoms and cause each node to determine its peer has departed. This situation typically results in data corruption because both nodes attempt to take control of data storage in an uncoordinated manner.

In addition to a broken set of private networks, other scenarios can generate this situation. If a system is so busy that it appears to stop responding or “hang,” the other nodes could declare it as dead. This declaration may also occur for nodes using 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 even though the system later returns and begins write operations.

VCS uses a technology called I/O fencing to remove the risk associated with split brain. I/O fencing allows write access for members of the active cluster and blocks access to storage from non-members; even a node that is alive is unable to cause damage.

SCSI-3 persistent reservations

SCSI-3 Persistent Reservations (SCSI-3 PR) are required for I/O fencing and resolve the issues of using SCSI reservations in a clustered SAN environment. SCSI-3 PR enables access for multiple nodes to a device and simultaneously blocks access for other nodes.

SCSI-3 reservations are persistent across SCSI bus resets and support multiple paths from a host to a disk. In contrast, only one host can use SCSI-2 reservations with one path. If the need arises to block access to a device because of data integrity concerns, only one host and one path remain active. The requirements for larger clusters, with multiple nodes reading and writing to storage in a controlled manner, make SCSI-2 reservations obsolete.

SCSI-3 PR uses a concept of registration and reservation. Each system registers its own “key” with a SCSI-3 device. Multiple systems registering keys form a membership and establish a reservation, typically set to “Write Exclusive Registrants Only.” The WERO setting enables only registered systems to perform write operations. For a given disk, only one reservation can exist amidst numerous registrations.

With SCSI-3 PR technology, blocking write access is as simple as removing a registration from a device. Only registered members can “eject” the registration of another member. A member wishing to eject another member issues a “preempt and abort” command. Ejecting a node is final and atomic; an ejected node cannot eject another node. In VCS, a node registers the same key for all paths to the device. A single preempt and abort command ejects a node from all paths to the storage device.

I/O fencing components

Fencing in VCS involves coordinator disks and data disks. Each component has a unique purpose and uses different physical disk devices. The fencing driver is `vxfen`.

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 added to a disk group are automatically fenced, as are new paths discovered to a device.

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 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 only register with them and do not need to reserve them.

These disks 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 coordinator disks before it can fence the peer from the data drives. This concept of racing for control of the coordinator disks to gain the ability to fence data disks is key to understanding prevention of split brain through fencing.

Dynamic Multipathing devices with I/O fencing

DMP allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. You can configure coordinator disks to use Veritas Volume Manager Dynamic Multipathing (DMP) feature.

For more information on using DMP, see the *Veritas Volume Manager Administrator's Guide*.

See “[Updating /etc/vxfenmode file](#)” on page 121

I/O fencing operations

I/O fencing, provided by the kernel-based fencing module (`vxfen`), performs identically on node failures and communications failures. When the fencing module on a node is informed of a change in cluster membership by the GAB module, it immediately begins the fencing operation. The node attempts to eject the key for departed nodes from the coordinator disks using the `preempt` and `abort` command. When the node successfully ejects the departed nodes from the coordinator disks, it ejects the departed nodes from the data disks. In a split brain scenario, both sides of the split would race for control of the coordinator disks. The side winning the majority of the coordinator disks wins the race and fences the loser. The loser then panics and reboots the system.

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 accompanying the Storage Foundation product that you are using.

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.

Checking shared disks for I/O fencing

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	Stores shared data
Coordinator disks	Act as a global lock during membership changes. Coordinator disks are small LUNs (typically three per cluster)

See [“Setting up shared storage”](#) on page 49.

Perform the following checks for I/O fencing disks:

- Identify three SCSI-3 PR compliant shared disks as coordinator disks. List the disks on each node and pick three disks as coordinator disks. For example, execute the following commands to list the disks:

```
# fdisk -l
```

- Test the shared disks using the `vxfcntlsthdw` script.
See [“Testing the shared disks for SCSI-3”](#) on page 115.

Testing the shared disks for SCSI-3

Use the `vxfcntlsthdw` utility to test the shared storage arrays support SCSI-3 persistent reservations and I/O fencing. Review the guidelines to run `vxfcntlsthdw` program, verify that the systems see the same disk, and proceed to test the disks. Make sure to test disks serving as coordinator disks.

See [“Setting up coordinator disk groups”](#) on page 119.

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

Review these guidelines for using `vxfcntlsthdw`

- 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.

- The two nodes must have `ssh` (default) or `rsh` communication. If you use `rsh`, launch the `vxfcntlsthdw` utility with the `-n` option.
See [“Enabling communication between systems”](#) on page 52.
After completing the testing process, remove permissions for communication and restore public network connections.
See [“Removing permissions for communication”](#) on page 123.
- To ensure both nodes are connected to the same disk during the testing, use the `vxfenadm -i diskpath` command to verify the disk serial number.
See [“Verifying the nodes see the same disk”](#) on page 115.

Verifying the nodes see the same disk

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 `vxfenadm` command with the `-i` option to verify that the same serial number for the LUN is returned on all paths to the LUN.

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

From node A, enter:

```
# vxfenadm -i /dev/sdr
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/sdt` path.

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

```
# vxfenadm -i /dev/sdc
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
Vendor id      : HITACHI
Product id     : OPEN-3
Revision       : 0117
Serial Number   : 0401EB6F0002
```

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

Testing the disks using `vxfentsthdw` script

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

If the utility does not show a message stating a disk is ready, verification has failed. Failure of verification can be the result of an improperly configured disk array. It can also be caused by 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/sdr is ready to be configured for I/O Fencing on node
north
```

See [“Adding or removing coordinator disks”](#) on page 140.

To test disks using `vxfentsthdw` script

- 1 Make sure system-to-system communication is functioning properly.
See [“Enabling communication between systems”](#) on page 52.
- 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 After reviewing the overview and warning that the tests overwrite data on the disks, 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: north  
Enter the second node of the cluster: south
```

- 4 Enter the names of the disks you are checking. For each node, the same disk may be known by a different name:

```
Enter the disk name to be checked for SCSI-3 PGR on node  
north in the format: /dev/sdx
```

```
/dev/sdx
```

```
Enter the disk name to be checked for SCSI-3 PGR on node  
south in the format: /dev/sdx
```

```
Make sure it's the same disk as seen by nodes north and south
```

```
/dev/sdx
```

If the disk names 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  
north
```

```
ALL tests on the disk /dev/sdx have PASSED
```

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

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

Setting up I/O fencing for VCS

Tasks involved in setting up I/O fencing include:

- [Initializing disks](#)
- [Setting up coordinator disk groups](#)
- [Stopping VCS on all nodes](#)
- [Configuring `/etc/vxfendg` disk group for I/O fencing](#)
- [Updating `/etc/vxfenmode` file](#)
- [Starting I/O fencing](#)
- [Modifying VCS configuration to use I/O fencing](#)
- [Verifying I/O fencing configuration](#)

Initializing 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 initialize them as VxVM disks and verify that all the nodes see the same disk. Use the example procedure; see the *Veritas Volume Manager Administrator's Guide* for more information on adding and configuring disks.

To initialize disks

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

```
# fdisk -l
```
- 2 If the Array Support Library (ASL) for the array you are adding is not installed, obtain and install it on each node before proceeding. The ASL for the supported storage device you are adding 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
=====
libvxCLARiION.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 adding and configuring 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.

Setting up coordinator disk groups

I/O fencing requires coordinator disks that are configured in a disk group and accessible to each node in the cluster. These disks enables the `vxfen` driver to resolve potential split brain conditions and prevent data corruption. Make sure to meet the requirements for coordinator disks and then create the coordinator disk group.

Requirements for coordinator disks

After adding and initializing disks for use as coordinator disks, make sure coordinator disks meet the following requirements:

- 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. See [“Testing the shared disks for SCSI-3”](#) on page 115.
- The coordinator disks must exist in a disk group (for example, `vxfencoordg`). See [“Creating the coordinator disk group and setting the coordinator attribute”](#) on page 119.
- Symantec recommends using hardware-based mirroring for coordinator disks.

Creating the coordinator disk group and setting the coordinator attribute

From one node, create a disk group named `vxfencoordg`. This group must contain three disks or LUNs.

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. You do this with a `vxddg set coordinator=on` command.

Refer to the *Veritas Volume Manager Administrator's Guide* for details on creating disk groups.

The example procedure assumes that the disks have the device names `/dev/sdz`, `/dev/sdaa`, and `sdab`.

To create the `vxencoorddg` disk group

- 1 On any node, create the disk group by specifying the device name of the disks:

```
# vxdbg -o coordinator=on init vxencoorddg sdz
```
- 2 Add the other two disks to the disk group:

```
# vxdbg -g vxencoorddg adddisk sdaa
# vxdbg -g vxencoorddg adddisk sdab
```

Stopping VCS on all nodes

Before configuring the coordinator disk for use, you must stop VCS on all nodes.

To stop VCS on all nodes

- ◆ On one node, enter:

```
# hastop -all
```

Configuring `/etc/vxfendg` disk group for I/O fencing

After setting up the coordinator disk group, configure it for use.

To configure the disk group for fencing

- 1 Deport the disk group:

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

```
# vxdbg -t import vxencoorddg
```
- 3 Deport the disk group. Deporting the disk group prevents the coordinator disks from serving other purposes:

```
# vxdbg deport vxencoorddg
```
- 4 On all nodes, type:

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

Do not use spaces between the quotes in the “`vxencoorddg`” text.

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

Based on the contents of the `/etc/vxfendg` and `/etc/vxfenmode` files, the `rc` script creates the `/etc/vxfentab` file for use by the `vxfen` driver when the system starts. The `rc` script also invokes the `vxfenconfig` command, which configures the `vxfen` driver to start and use the coordinator disks listed in

`/etc/vxfentab`. The `/etc/vxfentab` file is a generated file; do not modify this file.

Example `/etc/vxfentab` file

The `/etc/vxfentab` file gets created when you start the I/O fencing driver.

See “[Starting I/O fencing](#)” on page 122.

An example of the `/etc/vxfentab` file on one node resembles:

- Raw disk
 - `/dev/sdz`
 - `/dev/sdy`
 - `/dev/sdu`
- DMP disk
 - `/dev/vx/rdmp/sdz`
 - `/dev/vx/rdmp/sdy`
 - `/dev/vx/rdmp/sdu`

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

See “[Adding or removing coordinator disks](#)” on page 140.

Updating `/etc/vxfenmode` file

You must update the `/etc/vxfenmode` file to operate in SCSI-3 mode. 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, either `raw` or `dmp`, on all the nodes.

To update `/etc/vxfenmode` file

- ◆ On all cluster nodes, depending on the SCSI-3 mechanism you have chosen, type:
 - 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
```

Starting I/O fencing

You now need to start I/O fencing on each node. VxFEN, the I/O fencing driver, may already be running, so you need to restart the driver for the new configuration to take effect.

To stop I/O fencing on a node

- ◆ Stop the I/O fencing driver.
`# /etc/init.d/vxfen stop`

To start I/O fencing on a node

- ◆ Start the I/O fencing driver.
`# /etc/init.d/vxfen start`

Modifying VCS configuration to use I/O fencing

After adding coordinator disks and configuring 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. Modify the list of cluster attributes by adding the `UseFence` attribute and assigning its value of `SCSI3`.

```
cluster rac_cluster101
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, north) to the remaining cluster nodes.

For example, on each remaining node, enter:

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

- 8 On each node enter the following sequence of commands. These commands brings up VCS processes:

```
# /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: raw  
Cluster Members:
```

```
* 0 (north)  
1 (south)
```

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

Removing permissions for communication

After completing the installation of VCS and verification of disk support for I/O fencing, if you used `rsh`, remove the temporary `rsh` access permissions 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.

Additional I/O fencing information

Review additional information about I/O fencing, including an extended description of the `vxfsentsthdw` command, `vxfenadm` command, and a description of I/O fencing behavior to protect data in certain scenarios.

vxfsentsthdw options

[Table 4-10](#) describes the methods the utility provides to test storage devices.

Table 4-10 vxfsentsthdw options

vxfsentsthdw option	Description	When to use
-n	Utility uses <code>rsh</code> for communication.	Use when <code>rsh</code> is used for communication.
-r	Non-destructive testing. Testing of the disks for SCSI-3 persistent reservations occurs in a non-destructive way; that is, there is only testing for reads, not writes. May be used with <code>-m</code> , <code>-f</code> , or <code>-g</code> options.	Use during non-destructive testing.
-t	Testing of the return value of <code>SCSI TEST UNIT (TUR)</code> command under SCSI-3 reservations. A warning is printed on failure of TUR testing.	When you want to perform TUR testing.
-d	Use DMP devices. May be used with <code>-c</code> or <code>-g</code> options.	By default, the script picks up the OS paths for disks in the disk group. If you want the script to use the DMP path, use the <code>-d</code> option.
-c	Utility tests the coordinator disk group prompting for systems and devices, and reporting success or failure.	For testing disks in coordinator disk group.

Table 4-10 vxfcntlsthwdw options

vxfcntlsthwdw option	Description	When to use
-m	Utility runs manually, in interactive mode, prompting for systems and devices, and reporting success or failure. May be used with -r and -t options. -m is the default option.	For testing a few disks or for sampling disks in larger arrays.
-f <i>filename</i>	Utility tests system/device combinations listed in a text file. May be used with -r and -t options.	For testing several disks.
-g <i>disk_group</i>	Utility tests all disk devices in a specified disk group. May be used with -r and -t options.	For testing many disks and arrays of disks. Disk groups may be temporarily created for testing purposes and destroyed (ungrouped) after testing.

Testing the coordinator disk group using vxfcntlsthwdw -c

Use the vxfcntlsthwdw utility to verify disks are configured to support I/O fencing. In this procedure, the vxfcntlsthwdw utility tests the three disks one disk at a time from each node.

- From the node north, the disks are /dev/sdg, /dev/sdh, and /dev/sdi.
- From the node north, the disks are /dev/sdx, /dev/sdy, and /dev/sdz.

Note: To test the coordinator disk group using the vxfcntlsthwdw utility, the utility requires that the coordinator disk group, vxfcntlcoorddg, be accessible from two nodes.

To test the coordinator disk group using vxfcntlsthwdw -c

- 1 Use the vxfcntlsthwdw command with the -c option. For example:

```
# /opt/VRTSvcs/vxfen/bin/vxfcntlsthwdw -c vxfcntlcoorddg
```
- 2 Enter the nodes you are using to test the coordinator disks:
 Enter the first node of the cluster:
north
 Enter the second node of the cluster:

south

- 3 Review the output of the testing process for both nodes for all disks in the coordinator disk group. Each disk should display output that resembles:

```
ALL tests on the disk /dev/sdg have PASSED.  
The disk is now ready to be configured for I/O Fencing on node  
north as a COORDINATOR DISK.
```

```
ALL tests on the disk /dev/sdx have PASSED.  
The disk is now ready to be configured for I/O Fencing on node  
south as a COORDINATOR DISK.
```

- 4 After you test all disks in the disk group, the vxencoorddg disk group is ready for use.

Removing and replacing a failed disk

If a disk in the coordinator disk group fails verification, remove the failed disk or LUN from the vxencoorddg disk group, replace it with another, and retest the disk group.

If you need to replace a disk in an active coordinator disk group, refer to the troubleshooting procedure.

See [“Adding or removing coordinator disks”](#) on page 140.

To remove and replace a failed disk

- 1 Use the vxdiskadm utility to remove the failed disk from the disk group. Refer to the *Veritas Volume Manager Administrator’s Guide*.
- 2 Add a new disk to the node, initialize it, and add it to the coordinator disk group.
See [“Initializing disks”](#) on page 118.
See [“Setting up coordinator disk groups”](#) on page 119.
- 3 Retest the disk group.

Using the -r option for non-destructive testing

To test disk devices containing data you want to preserve, you can use the -r option with the -m, -f, or -g options, which are described in the following sections. For example, to use the -m option and the -r option, you can run the utility by entering:

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

When invoked with the -r option, the utility does not use tests that write to the disks. Therefore, it does not test the disks for all of the usual conditions of use.

Using the -m option

Review the procedure to test the shared disks. The utility uses the -m option. See “[Testing the shared disks for SCSI-3](#)” on page 115.

Using the -f option

Use the -f option to test disks that are listed in a text file. For example, you can create a file to test two disks shared by systems north and south that might resemble:

```
north /dev/sdz south /dev/sdy
north /dev/sdu south /dev/sdw
```

Where the first disk is listed in the first line and is seen by north as /dev/sdz and by south as /dev/sdy. The other disk, in the second line, is seen as /dev/sdu from north and /dev/sdw from south. Typically, the list of disks could be extensive.

Suppose you created the file named disks_blue. To test the disks, you would enter:

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw -f disks_blue
```

The utility reports the test results one disk at a time, just as for the -m option.

You can redirect the test results to a text file. Precede the command with “yes” to acknowledge that the testing destroys any data on the disks to be tested.

Caution: Be advised that by redirecting the command’s output to a file, a warning that the testing destroys data on the disks cannot be seen until the testing is done.

For example:

```
# yes | /opt/VRTSvcs/vxfen/bin/vxfentsthdw -f disks_blue >
blue_test.txt
```

Using the -g option

Use the -g option to test all disks within a disk group. For example, you create a temporary disk group consisting of all disks in a disk array and test the group.

Note: Do not import the test disk group as shared; that is, do not use the -s option.

The utility reports the test results one disk at a time. You can redirect the test results to a text file for review.

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw -g red_disks_dg >
redtest.txt
```

After testing, destroy the disk group and put the disks into disk groups as you need.

Testing a disk with existing keys

If the utility detects that a coordinator disk has existing keys, you see a message that resembles:

```
There are Veritas I/O Fencing keys on the disk. Please make sure
that I/O Fencing is shut down on all nodes of the cluster before
continuing.
```

```
***** WARNING!!!!!!!!!! *****
```

```
THIS SCRIPT CAN ONLY BE USED IF THERE ARE NO OTHER ACTIVE NODES
IN THE CLUSTER! VERIFY ALL OTHER NODES ARE POWERED OFF OR
INCAPABLE OF ACCESSING SHARED STORAGE.
```

```
If this is not the case, data corruption will result.
```

```
Do you still want to continue : [y/n] (default: n) y
```

The utility prompts you with a warning before proceeding. You may continue as long as I/O fencing is not yet configured.

About VXFEN tunable parameters

On each node, edit the file `/etc/sysconfig/vxfen` to change the value of the `vxfen` driver tunable global parameter, `vxfen_max_delay` and `vxfen_min_delay`. You must restart the system to put change into effect.

[Table 4-11](#) describes tunable parameters for the VXFEN driver.

Table 4-11 VXFEN tunable parameters

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
vxfen_debug_sz	Size of debug log in bytes ■ Values Default: 65536 Minimum: 65536 Maximum: 256K
vxfen_max_delay and vxfen_min_delay (See below.)	In the event of a network partition, the smaller cluster delays before racing for the coordinator disks. The time delayed allows a larger sub-cluster to win the race for the coordinator disks. The <code>vxfen_max_delay</code> and <code>vxfen_min_delay</code> parameters define the delay in seconds.

Table 4-11 VXFEN tunable parameters

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
vxfen_max_delay	<p>Specifies the maximum number of seconds that the smaller sub-cluster waits before racing with larger clusters for control of the coordinator disks.</p> <p>This value must be greater than the vxfen_min_delay value.</p> <ul style="list-style-type: none"> ■ Values <ul style="list-style-type: none"> Default: 60 Minimum: 0 Maximum: 600
vxfen_min_delay	<p>Specifies the minimum number of seconds that the smaller sub-cluster waits before racing with larger clusters for control of the coordinator disks. This value must be smaller than the vxfen_max_delay value.</p> <ul style="list-style-type: none"> ■ Values <ul style="list-style-type: none"> Default: 1 Minimum: 0 Maximum: 600

See [“Configuring the VXFEN parameters”](#) on page 129.

Configuring the VXFEN parameters

For the parameter changes to take effect, reconfigure the VXFEN module.

To reconfigure the VXFEN module

- 1 Unconfigure the VXFEN module.

```
# /sbin/vxfenconfig -U
```
- 2 Unload the module.

```
# /etc/init.d/vxfen stop
```
- 3 Edit the /etc/sysconfig/vxfen file.
 For example, change the entry from:

```
vxfen_min_delay=0
```

 to:

```
vxfen_min_delay=30
```
- 4 Start the VXFEN module.

```
# /etc/init.d/vxfen start
```
- 5 Start VCS.

```
# hastart
```
- 6 Bring the service groups online.

```
# hagrgr -online oragrgr -sys north
```

How I/O fencing works in different event scenarios

Table 4-12 describes how I/O fencing works to prevent data corruption in different failure event scenarios. For each event, corrective operator actions are indicated.

Table 4-12 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
Both private networks fail.	Node A races for majority of coordinator disks. If Node A wins race for coordinator disks, Node A ejects Node B from the shared disks and continues.	Node B races for majority of coordinator disks. If Node B loses the race for the coordinator disks, Node B removes itself from the cluster.	When Node B is ejected from cluster, repair the private networks before attempting to bring Node B back.
Both private networks function again after event above.	Node A continues to work.	Node B has crashed. It cannot start the database since it is unable to write to the data disks.	Restart Node B after private networks are restored.
One private network fails.	Node A prints message about an IOFENCE on the console but continues.	Node B prints message about an IOFENCE on the console but continues.	Repair private network. After network is repaired, both nodes automatically use it.

Table 4-12 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
Node A hangs.	<p>Node A is extremely busy for some reason or is in the kernel debugger.</p> <p>When Node A is no longer hung or in the kernel debugger, any queued writes to the data disks fail because Node A is ejected. When Node A receives message from GAB about being ejected, it removes itself from the cluster.</p>	<p>Node B loses heartbeats with Node A, and races for a majority of coordinator disks.</p> <p>Node B wins race for coordinator disks and ejects Node A from shared data disks.</p>	<p>Verify private networks function and restart Node A.</p>

Table 4-12 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
<p>Nodes A and B and private networks lose power. Coordinator and data disks retain power. Power returns to nodes and they restart, but private networks still have no power.</p>	<p>Node A restarts and I/O fencing driver (vxfen) detects Node B is registered with coordinator disks. The driver does not see Node B listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node A from joining the cluster. Node A console displays:</p> <p>Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.</p>	<p>Node B restarts and I/O fencing driver (vxfen) detects Node A is registered with coordinator disks. The driver does not see Node A listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node B from joining the cluster. Node B console displays:</p> <p>Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.</p>	<p>Resolve preexisting split brain condition.</p> <p>See “System panics to prevent potential data corruption” on page 137.</p>

Table 4-12 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
<p>Node A crashes while Node B is down. Node B comes up and Node A is still down.</p>	<p>Node A is crashed.</p>	<p>Node B restarts and detects Node A is registered with the coordinator disks. The driver does not see Node A listed as member of the cluster. The I/O fencing device driver prints message on console:</p> <pre>Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.</pre>	<p>Resolve preexisting split brain condition. See “System panics to prevent potential data corruption” on page 137.</p>

Table 4-12 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
The disk array containing two of the three coordinator disks is powered off.	Node A continues to operate as long as no nodes leave the cluster.	Node B continues to operate as long as no nodes leave the cluster.	
Node B leaves the cluster and the disk array is still powered off.	Node A races for a majority of coordinator disks. Node A fails because only one of three coordinator disks is available. Node A removes itself from the cluster.	Node B leaves the cluster.	Power on failed disk array and restart I/O fencing driver to enable Node A to register with all coordinator disks.

About the vxfenadm utility

Administrators can use the `vxfenadm` command to troubleshoot and test fencing configurations. The command's options for use by administrators are:

- g read and display keys
- i read SCSI inquiry information from device
- m register with disks
- n make a reservation with disks
- p remove registrations made by other systems
- r read reservations
- x remove registrations

Registration key formatting

The key defined by VxVM associated with a disk group consists of seven bytes maximum. This key becomes unique among the systems when the VxVM prefixes it with the ID of the system. The key used for I/O fencing, therefore, consists of eight bytes.

0							7
Node ID	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined

The keys currently assigned to disks can be displayed by using the `vxfenadm` command.

For example, from the system with node ID 1, display the key for the disk `/dev/sdy` by entering:

```
# vxfenadm -g /dev/sdy
Reading SCSI Registration Keys...
Device Name: /dev/sdy
SCSI ID => Host: 3 Channel: 0 Id: 9 Lun: 0
PRGeneration: -84
Total Number of Keys: 1
key[0]:
    Key Value [Numeric Format]: 65,45,45,45,45,45,45,45
    Key Value [Character Format]: A-----
```

The `-g` option of `vxfenadm` displays all eight bytes of a key value in two formats. In the numeric format, the first byte, representing the Node ID, contains the system ID plus 65. The remaining bytes contain the ASCII values of the letters of the key, in this case, “-----”. In the next line, the node ID 0 is expressed as “A;” node ID 1 would be “B.”

Troubleshooting I/O fencing

Headings indicate likely symptoms or procedures required for a solution.

Node is unable to join cluster while another node is being ejected

A cluster that is currently fencing out (ejecting) a node from the cluster prevents a new node from joining the cluster until the fencing operation is completed. The following are example messages that appear on the console for the new node:

```
...VCS FEN ERROR V-11-1-25 ... Unable to join running cluster
...VCS FEN ERROR V-11-1-25 ... since cluster is currently
fencing
...VCS FEN ERROR V-11-1-25 ... a node out of the cluster.

...VCS GAB.. Port b closed
```

If you see these messages when the new node is booting, the `vxfen` startup script on the node makes up to five attempts to join the cluster. If this is not sufficient to allow the node to join the cluster, restart the new node or attempt to restart `vxfen` driver with the command:

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

vxfersthdw fails when SCSI TEST UNIT READY command fails

If you see a message resembling:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node
FAILED.
Contact the storage provider to have the hardware configuration
fixed.
The disk array does not support returning success for a SCSI
TEST UNIT READY command when another host has the disk reserved
using SCSI-3 persistent reservations. This happens with Hitachi
Data Systems 99XX arrays if bit 186 of the system mode option is
not enabled.
```

Removing existing keys from disks

Review the procedure to remove the registration and reservation keys created by another node from a disk.

To remove the registration and reservation keys from disk

- 1 Create a file to contain the access names of the disks:

```
# vi /tmp/disklist
```


For example:

```
/dev/sdu
```

- 2 Read the existing keys:

```
# vxfenadm -g all -f /tmp/disklist
```

The output from this command displays the key:

```
Device Name: /dev/sdu
```

```
Total Number Of Keys: 1
```

```
key[0]:
```

```
Key Value [Numeric Format]: 65,49,45,45,45,45,45,45
```

```
Key Value [Character Format]: A1-----
```

- 3 If you know on which node the key was created, log in to that node and enter the following command:

```
# vxfenadm -x -kA1 -f /tmp/disklist
```

The key is removed.

- 4 If you do not know on which node the key was created, follow [step 5](#) through [step 7](#) to remove the key.

- 5 Register a second key “A2” temporarily with the disk:

```
# vxfenadm -m -k A2 -f /tmp/disklist
```

```
Registration completed for disk path /dev/sdu
```

- 6 Remove the first key from the disk by preempting it with the second key:

```
# vxfenadm -p -kA2 -f /tmp/disklist -vA1
```

```
key: A2----- preempted the key: A1----- on disk
```

```
/dev/sdu
```

- 7 Remove the temporary key assigned in [step 5](#).

```
# vxfenadm -x -kA2 -f /tmp/disklist
```

```
Deleted the key : [A2-----] from device /dev/sdu
```

No registration keys exist for the disk.

System panics to prevent potential data corruption

When a node experiences a split brain condition and is ejected from the cluster, it panics and displays the following console message:

```
VXFEN:vxfen_plat_panic: Local cluster node ejected from cluster  
to prevent potential data corruption.
```

How vxfen driver checks for pre-existing split brain condition

The vxfen driver functions to prevent an ejected node from rejoining the cluster after the failure of the private network links and before the private network links are repaired.

For example, suppose the cluster of system 1 and system 2 is functioning normally when the private network links are broken. Also suppose system 1 is

the ejected system. When system 1 restarts before the private network links are restored, its membership configuration does not show system 2; however, when it attempts to register with the coordinator disks, it discovers system 2 is registered with them. Given this conflicting information about system 2, system 1 does not join the cluster and returns an error from `vxfenconfig` that resembles:

```
vxfenconfig: ERROR: There exists the potential for a preexisting
split-brain. The coordinator disks list no nodes which are in
the current membership. However, they also list nodes which are
not in the current membership.
```

```
I/O Fencing Disabled!
```

Also, the following information is displayed on the console:

```
<date> <system name> vxfen: WARNING: Potentially a preexisting
<date> <system name> split-brain.
<date> <system name> Dropping out of cluster.
<date> <system name> Refer to user documentation for steps
<date> <system name> required to clear preexisting split-brain.
<date> <system name>
<date> <system name> I/O Fencing DISABLED!
<date> <system name>
<date> <system name> gab: GAB:20032: Port b closed
```

However, the same error can occur when the private network links are working and both systems go down, system 1 restarts, and system 2 fails to come back up. From the view of the cluster from system 1, system 2 may still have the registrations on the coordinator disks.

Case 1: system 2 up, system 1 ejected (actual potential split brain)

Determine if system1 is up or not. If it is up and running, shut it down and repair the private network links to remove the split brain condition. restart system 1.

Case 2: system 2 down, system 1 ejected (apparent potential split brain)

- 1 Physically verify that system 2 is down.
- 2 Verify the systems currently registered with the coordinator disks. Use the following command:


```
# vxfenadm -g all -f /etc/vxfentab
```

 The output of this command identifies the keys registered with the coordinator disks.
- 3 Clear the keys on the coordinator disks as well as the data disks using the command `/opt/VRTSvcs/rac/bin/vxfenclearpre`.

See [“Clearing keys after split brain using vxfenclearpre command”](#) on page 139.

- 4 Make any necessary repairs to system 2 and restart.

Clearing keys after split brain using vxfenclearpre command

When you have encountered a split brain condition, use the `vxfenclearpre` command to remove SCSI-3 registrations and reservations on the coordinator disks as well as on the data disks in all shared disk groups.

To clear keys after split brain

- 1 Shut down all other nodes in the cluster that have access to the shared storage. This prevents data corruption.
- 2 Start the script:


```
# cd /opt/VRTSvcs/vxfen/bin
# ./vxfenclearpre
```
- 3 Read the script's introduction and warning. Then, you can choose to let the script run.

Do you still want to continue: [y/n] (default : n) **y**

Informational messages resembling the following may appear on the console of one of the nodes in the cluster when a node is ejected from a disk/LUN:

```
<date> <system name> scsi: WARNING: /sbus@3,0/lpfs@0,0/
sd@0,1(sd91):
<date> <system name> Error for Command: <undecoded cmd 0x5f>
Error Level: Informational
<date> <system name> scsi: Requested Block: 0 Error Block 0
<date> <system name> scsi: Vendor: <vendor> Serial Number:
0400759B006E
<date> <system name> scsi: Sense Key: Unit Attention
<date> <system name> scsi: ASC: 0x2a (<vendor unique code
0x2a>), ASCQ: 0x4, FRU: 0x0
```

These informational messages may be ignored.

Cleaning up the coordinator disks...

Cleaning up the data disks for all shared disk groups...

Successfully removed SCSI-3 persistent registration and reservations from the coordinator disks as well as the shared data disks.

```
Reboot the server to proceed with normal cluster startup...
#
```

- 4 Restart all nodes in the cluster.

Adding or removing coordinator disks

Review the following information to:

- Replace coordinator disk in the coordinator disk group
- Destroy a coordinator disk group

Note: Adding or removing coordinator disks requires all services be shut down.

Note the following about the procedure:

- A coordinator disk group requires three disks/LUNs.
- When adding a disk, add the disk to the disk group `vx fencecoorddg` and retest the group for support of SCSI-3 persistent reservations.
- You can destroy the coordinator disk group such that no registration keys remain on the disks. The disks can then be used elsewhere.

To remove and replace a disk in the coordinator disk group

- 1 Log in as superuser on one of the cluster nodes.
- 2 If VCS is running, shut it down:

```
# hastop -all
```
- 3 Stop I/O fencing on all nodes:

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

This removes any registration keys on the disks.
- 4 Import the coordinator disk group. The file `/etc/vxfendg` includes the name of the disk group (typically, `vx fencecoorddg`) that contains the coordinator disks, so use the command:

```
# vxdg -tfc import `cat /etc/vxfendg`
```

where:

 - t specifies that the disk group is imported only until the node restarts.
 - f specifies that the import is to be done forcibly, which is necessary if one or more disks is not accessible.
 - C specifies that any import blocks are removed.
- 5 To remove disks from the disk group, use the VxVM disk administrator utility, `vx diskadm`.
You may also destroy the existing coordinator disk group. For example:

```
# vxdg destroy vx fencecoorddg
```
- 6 Add the new disk to the node, initialize it as a VxVM disk, and add it to the `vx fencecoorddg` disk group.
See [“Creating the coordinator disk group and setting the coordinator attribute”](#) on page 119.

- 7 Test the recreated disk group for SCSI-3 persistent reservations compliance. See [“Testing the coordinator disk group using vxfcntl -c”](#) on page 125.
- 8 After replacing disks in a coordinator disk group, deport the disk group:

```
# vxvg deport `cat /etc/vxvoldg`
```
- 9 On each node, start the I/O fencing driver:

```
# /etc/init.d/vxvob start
```
- 10 If necessary, restart VCS on each node:

```
# hstart
```


Verifying the VCS installation

This chapter contains the following topics:

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

About verifying the VCS installation

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

Verifying LLT and GAB configuration files

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

`/etc/llthosts`

The file `llthosts(4)` is a database, containing one entry per system, that links the LLT system ID (in the first column) with the LLT host name. This file is identical on each node in the cluster.

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

```
0 north
1 south
```

`/etc/llttab`

The file `llttab(1M)` contains 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 entries that resemble:

```
set-node north
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 entries that resemble:

```
set-node north
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, beginning with the `link` command, identify the two network cards used by the LLT protocol.

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

`/etc/gabtab`

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

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

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

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

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

Verifying the main.cf file

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

See [“Example main.cf, for clusters without the GCO option”](#) on page 146.

See [“Example main.cf, for clusters with the GCO option”](#) on page 148.

See [“Example main.cf for a centrally managed cluster using Cluster Management Console”](#) on page 148.

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 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 and includes the following:
 - The `ClusterService` service group includes the IP, NIC, and `VRTSWebApp` resources.
 - If you configured Cluster Management Console to manage this cluster locally, the `main.cf` includes the `VRTSWebApp` resource that includes `AppName = cmc` attribute.
 - If you configured Cluster Connector so that Cluster Management Console can centrally manage this cluster, the `main.cf` includes the CMC service group.

The CMC service group includes the `ClusterConnectorConfig` and `Process` resources.

- The service group also includes the notifier resource configuration, which is based on your input to `installvcs` program prompts about notification.

- The `installvcs` program also creates a resource dependency tree.
- If you installed VCS with the Global Cluster Option, the `ClusterService` service group contains an `Application` resource, `wac` (wide-area connector), whose attributes contain definitions for controlling the cluster in a Global Cluster environment.
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* and review the chapter on configuration concepts for descriptions and examples of `main.cf` and `types.cf` files for Linux systems.

Example main.cf, for clusters without the GCO option

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 north (
)

system south (
)

group ClusterService (
    SystemList = { north = 0, south = 1 }
    UserStrGlobal = "LocalCluster@https://10.182.2.76:8443;"
    AutoStartList = { north, south }
    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 }
    SntpServer = "smtp.example.com"
    SntpRecipients = { "ozzie@example.com" = Warning,
                      "harriet@example.com" = Error }
)

VRTSWebApp VCSweb (
    Critical = 0
    AppName = cmc
    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 = { north = 0, south = 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
// }
```

Example main.cf, for clusters with the GCO option

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

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

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

Example main.cf for a centrally managed cluster using Cluster Management Console

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

cluster vcs_cluster2 (
    UserNames = { "admin" = hqrJqlQnrMrrPzrLqo }
    Administrators = { "admin" }
    ClusterAddress = "10.10.12.1"
    CounterInterval = 5
)

system north (
)

system south (
```

```
)

group ClusterService (
    SystemList = { north, south }
    AutoStartList = { north, south }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

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

NIC csgnic (
    Device = eth0
)

VRTSWebApp VCSweb (
    Critical = 0
    AppName = cmc
    InstallDir = "/opt/VRTSweb/VERITAS"
    TimeForOnline = 5
    RestartLimit = 3
)

VCSweb requires webip
webip requires csgnic

group CMC (
    SystemList = { north, south }
    AutoStartList = { north, south }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

ClusterConnectorConfig CMC_ClusterConfig (
    MSAddress = "mgmtserver1.symantecexample.com"
    MSPort = 14145
    ClusterId = "1145613636"
    ClusterType = "vcs"
    ClusterPort = 14141
    VCSLoggingLevel = "TAG_A"
    Logging = "/opt/VRTScmccc/conf/cc_logging.properties"
    ClusterConnectorVersion = "5.0.1000.0"
)

Process CMC_ClusterConnector (
    PathName = "/bin/sh"
    Arguments = "/opt/VRTScmccc/bin/cluster_connector.sh"
)
```

CMC_ClusterConnector requires CMC_ClusterConfig

Verifying LLT, GAB, and cluster operation

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

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

Verifying LLT

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

Using `lltstat -n`

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

Node 1

```
# lltstat -n
```

Output resembles:

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

Node 2

```
# lltstat -n
```

Output resembles:

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

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

Using `lltstat -nvv`

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

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

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

Output resembles:

```
Node      State      Link      Status      Address
*0 north   OPEN
          eth1 UP      08:00:20:93:0E:34
          eth2 UP      08:00:20:93:0E:34
 1 south   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. It reports on the two nodes in the cluster, north and south, plus non-existent nodes. For each correctly configured node, the information should show a state of OPEN, a status for each link of UP, and an address for each link. However, the output in the example shows that for the node south the private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

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

Node 1

```
# lltstat -p
```

Output resembles:

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

```
31      gab      0x1F
      opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
      connects: 0 1
```

Verifying GAB

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

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

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

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

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

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

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

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

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

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

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

Verifying the cluster

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

```
# hastatus -summary
```

The output resembles:

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

A north           RUNNING       0
A south           RUNNING       0
```



```
-- GROUP STATE
-- Group          System      Probed  AutoDisabled  State

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

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

hasys -display

On one of the nodes, use the `hasys(1M)` command:

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

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

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

```
#System Attribute          Value
north  AgentsStopped          0
north  AvailableCapacity      100
north  CPUUsage                0
north  CPUUsageMonitoring      Enabled 0 ActionThreshold 0
                                ActionTimeLimit 0 Action NONE
                                NotifyThreshold 0 NotifyTimeLimit 0
north  Capacity                100
north  ConfigBlockCount        142
north  ConfigChecksum          4085
north  ConfigDiskState         CURRENT
north  ConfigFile               /etc/VRTSvcs/conf/config
north  ConfigInfoCnt           0
```

```
north ConfigModDate Fri May 26 17:22:48 2006
north ConnectorState Down
north CurrentLimits
north DiskHbStatus
north DynamicLoad 0
north EngineRestarted 0
north EngineVersion 5.0.00.0
north Frozen 0
north GUIIPAddr
north LLTNodeId 0
north LicenseType DEMO
north Limits
north LinkHbStatus eth1 UP eth2 UP
north LoadTimeCounter 0
north LoadTimeThreshold 600
north LoadWarningLevel 80
north NoAutoDisable 0
north NodeId 0
north OnGrpCnt 1
north ShutdownTimeout 120
north SourceFile ./main.cf
north SysInfo Linux:north,#1 Fri Apr 22 18:13:58 EDT
2005,2.6.9-34-default,i686
north SysName north
north SysState RUNNING
north SystemLocation
north SystemOwner
north TFrozen 0
north TRSE 0
```

north	UpDownState	Up
north	UserInt	0
north	UserStr	
north	VCSFeatures	DR
north	VCSMode	VCS

Accessing the Veritas Cluster Management Console

The VCS web-based Cluster Management Console enables you to monitor the cluster from any workstation on the public network. Supported browsers are Netscape Navigator 4.0 or later, or Internet Explorer 4.0 or later.

When VCS starts running in the cluster and the ClusterService Group comes up, the Web Console server starts.

To access the Web Console

- 1 From the browser, navigate to the Web Console by entering:

`http://hostname:8443/cmc`

or

`http://hostname:8181/cmc`

Where hostname is the system name or IP address.

For example:

`http://10.10.12.1:8443/cmc`

The IP address is the “Cluster virtual IP address” configured into the ClusterService Group.

- 2 On the Login screen, enter a valid user name and password. By default, the administrator of a new installation can log in as “admin” and use “password” as a password. For security, change your password at your earliest convenience.
- 3 Click Login to enter the Cluster Summary view.

Accessing the VCS documentation

If you had chosen to install the optional RPM `VRTSvcsdc`, then the directory `/opt/VRTS/docs` contains the documentation for VCS in Portable Document Format (PDF). The directory contains the following documents:

- `vcs_users.pdf`, *Veritas Cluster Server User's Guide*
- `vcs_bundled_agents.pdf`, *Veritas Cluster Server Bundled Agents Reference Guide*
- `vcs_agent_dev.pdf`, *Veritas Cluster Server Agent Developer's Guide*

Upgrading to VCS 5.0

This chapter includes the following topics:

- [About upgrading to VCS 5.0](#)
- [Upgrading VCS using installvcs program](#)
- [Upgrading VCS in a secure environment](#)
- [Upgrading VCS manually](#)
- [Upgrading to the VCS 5.0 Java Console](#)
- [Upgrading from CommandCentral Availability 4.1 MP1](#)

About upgrading to VCS 5.0

VCS 5.0 supports the following upgrade paths:

From	To	Task
VCS 4.1 MP2	VCS 5.0	Perform an automatic upgrade using the installvcs program.
VCS 4.0 MP1	VCS 5.0	Perform a manual upgrade.
VCS 4.1	VCS 5.0	Perform a manual upgrade.
VCS 4.1 MP1	VCS 5.0	Perform a manual upgrade.

See [“Upgrading VCS using installvcs program”](#) on page 158.

See [“Upgrading VCS manually”](#) on page 165.

For CommandCentral Availability upgrade from 4.1 MP1, see: [“Upgrading from CommandCentral Availability 4.1 MP1”](#) on page 170.

Upgrading VCS using installvcs program

If you are currently running a VCS cluster, including a VCS global cluster, you can run the installer to upgrade to VCS 5.0. The program detects the current configuration and prompts you to indicate whether you want to upgrade. While the installer is not able to verify the validity of the existing configuration, it is able to run extensive upgrade requirement checks before proceeding.

See “[About the VCS installation program](#)” on page 62.

If you want to upgrade from an earlier version of VCS to VCS 5.0 and use the Global Cluster option, you must first upgrade to standard VCS 5.0. After adding a license for the Global Cluster option, you can run the gcoconfig wizard. Note that the Global Cluster option requires a VCS HA/DR license. See the *Veritas Cluster Server User’s Guide* for instructions.

Upgrading VCS to 5.0

When you run `installvcs` on cluster systems that currently runs VCS 4.1 MP2, the program guides you through an upgrade procedure. The example upgrade demonstrates how to upgrade VCS 4.1 to VCS 5.0 on nodes north and south using `installvcs` program.

Upgrade tasks include:

- “[Removing deprecated resource types](#)” on page 158
- “[Starting the upgrade](#)” on page 159
- “[Checking upgrade requirements](#)” on page 160
- “[Removing VCS RPMs from previous versions and installing VCS 5.0 RPMs](#)” on page 161
- “[Starting VCS](#)” on page 162
- “[Completing the upgrade](#)” on page 163
- “[Using the halogin command for native OS accounts with VCS](#)” on page 164

Removing deprecated resource types

With VCS 5.0, 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. The list of resource types that are not used in VCS 5.0 are:

- `ServiceGroupHB`
- `CampusCluster`
- `ClusterMonitorConfig`

Note: The `ClusterConnectorConfig` resource type has replaced the `ClusterMonitorConfig` resource type.

If you are using the resource type `ServiceGroupHB`, Symantec recommends the use of I/O fencing.

Perform the following steps to remove the deprecated resource types.

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

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

Starting the upgrade

Start the `installvcs` program to automatically upgrade the cluster nodes.

To start the upgrade

- 1 Log in as superuser.
- 2 Mount the software disc.
See [“Mounting the product disc”](#) on page 55.
- 3 Ensure that HAD is running.

- 4 On the node where the disk was mounted, browse to the `cluster_server` directory and start the `installvcs` program:

```
# ./installvcs
```

See “[Starting the software installation](#)” on page 66.

- 5 After viewing a copyright notice, review the report after the program examines the configuration files and discovers the existing cluster configuration (including the `ClusterService` group, if it is defined):

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

```
Cluster Name: vcs_cluster2
Cluster ID Number: 7
Systems: north south
Service Groups: ClusterService XYZ_group
```

No checks have been made to ensure the validity of the configuration.

VCS version 4.1.10.00 is installed on this system

- 6 To upgrade to VCS 5.0, press **Enter**.

```
Do you want to upgrade to version 5.0 on these systems using the
current configuration? [y,n,q,?] (y) y
```

Checking upgrade requirements

The installer program verifies that your systems meet the requirements for upgrading.

To check upgrade requirements

- 1 Review the output as the program verifies that the upgrade can proceed on the systems.

The `installvcs` program checks the operating system level and system-to-system communication. It also creates logs for the VCS update.

`installvcs` requires that `ssh` commands used between systems execute without prompting for passwords or confirmations. If `installvcs` hangs or asks for a login password or hangs, stop `installvcs` and run it again with the `ssh` configured for passwordless logins, or configure `rsh` and use the `-rsh` option.

```
Checking ssh communication with north.....Linux 2.6.5-7.244-smp
Checking VCS installation on north..... 4.1.10.00
Checking VCS installation on south..... 4.1.10.00
```

- 2 Review the output as the installer checks the licenses that are currently in place on the cluster nodes. The installer also prompts to add additional license keys:

```
Checking system licensing
```



```
Installing licensing RPMs
VCS license registered on north
```

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

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

Removing VCS RPMs from previous versions and installing VCS 5.0 RPMs

The installer program is ready to remove the RPMs from VCS version 4.1 and install the 5.0 RPMs on each node after verifying the requirements.

To remove VCS RPMs from previous versions and add 5.0 RPMs

1 Review the output as the installer checks for existing RPMs.

```
Checking installed RPMs on north
Checking installed RPMs on south
```

2 Specify the RPMs to be installed on the cluster systems.

Additional RPMs are typically installed to simplify future upgrades.

- 1) Required Veritas Cluster Server RPMs - 397 MB required
- 2) All Veritas Cluster Server RPMs - 444 MB required

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

3 The installer lists the RPMs that will be installed or upgraded. Press **Enter** at the prompt.

```
installvcs will upgrade or install the following VCS RPMs:
```

```
VRTSvlic      Veritas Licensing
VRTSperl     Veritas Perl 5.8.8 Redistribution
VRTSicsco    Symantec Infrastructure Core Services Common
VRTSspbx     Symantec Private Branch Exchange
VRTSsmf      Symantec Service Management Framework
VRTSatClient Symantec Product Authentication Service Client
VRTSatServer Symantec Product Authentication Service
VRTSspt      Veritas Software Support Tools
SYMCLma     Symantec License Inventory Agent
VRTSllt     Veritas Low Latency Transport
VRTSgab     Veritas Group Membership and Atomic Broadcast
VRTSvxfen   Veritas I/O Fencing
VRTSvcs     Veritas Cluster Server
VRTSvcsmsg  Veritas Cluster Server Message Catalogs
VRTSacclib  Veritas ACC Library
VRTSvcsag   Veritas Cluster Server Bundled Agents
VRTSvcsdr   Veritas Cluster Server Disk Reservation Modules and
Utilities
VRTSvcsmn   Veritas Cluster Server Man Pages
VRTSvcsdc   Veritas Cluster Server Documentation
VRTSjre     Veritas Java Runtime Environment Redistribution
```

```

VRTSjre15    Veritas Java Runtime Environment Redistribution
VRTScscw    Veritas Cluster Server Configuration Wizards
VRTScscm    Veritas Cluster Server Cluster Manager
VRTScssim   Veritas Cluster Server Simulator
VRTSweb     Veritas Java Web Server
VRTScutil   Veritas Cluster Utilities
VRTScmcs    Veritas Cluster Management Console (Single Cluster
            Mode)
VRTScmccc   Veritas Cluster Management Console Cluster
            Connector

```

Press [Enter] to continue:

4 The installer is now ready to upgrade VCS. Press **Enter at the prompt.**

installvcs is now ready to upgrade VCS.

All VCS processes that are currently running must be stopped.

Are you sure you want to upgrade VCS? [y,n,q] (y) **y**

Updating gabtab on north..... Done

installvcs must now make configuration updates and stop the cluster before upgrading VCS RPMs

Are you ready to begin the Veritas Cluster Server upgrade at this time? [y,n,q] (y) **y**

5 View the output as the program backs up the types.cf and main.cf and freezes the service group configured in the cluster.

Backing up types.cf and main.cf..... Done

Freezing group XYZ_group..... Done

Updating types.cf file Done

Updating gabtab on north..... Done

Updating gabtab on south..... Done

6 The program attempts to stop VCS and uninstall the RPMs. Progress indicators show the status of these tasks.

7 The program now attempts to install VCS 5.0 RPMs. A progress indicator shows the status of the task.

8 If the cluster systems have SLES9 installed, the installer prompts you to make sure the NICs on the systems are boot-persistent. Refer to the operating system document for instructions.

Starting VCS

After the installvcs program upgrades the nodes, you can proceed to start VCS using the installvcs program.

To start VCS

1 The program prompts you to start VCS.

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

- If you have an NFS resource configured in the cluster, enter **n** and do not start VCS here. Edit the `main.cf` file to:
 - Remove the `IPResName` attribute.
 - Add the `Address` attribute, and specify the IP address as the value of this attribute.
 - Include details about the `NFSRestart` agent.
 Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for details about editing the `main.cf` file. After editing this file, attempt to start VCS.
- If you have configured resources for Oracle, Sybase, or DB2 databases, enter **n** and do not start VCS here.
 Refer to the respective Enterprise agent documentation for upgrade information. After upgrading the agents, attempt to start VCS.
- If you have an Apache resource configured in the cluster, enter **n** and do not start VCS here.
 Open the `main.cf` file and ensure that absolute path names are specified for the Apache resource attributes. Then attempt to start VCS.
- If you have none of the above resources configured on the node, enter **y** to start VCS.
 Starting VCS: 0%
 Starting VCS: 100%

2 The program attempts to unfreeze the service group and online the `ClusterService` group.

```
Unfreezing NFSGroup..... Done
Unfreezing myshare..... Done
Unfreezing testgrp..... Done
Onlining ClusterService Group on north..... Done
```

Completing the upgrade

After starting the cluster server processes, the upgrade is complete.

To complete upgrade

- 1 Note the locations of the summary and log files that the program creates.
 Upgrade log files and summary file are saved at:

```
/opt/VRTS/install/logs/installvcs-unique string/
```
- 2 Other files specific to the installer program are also created in the `/opt/VRTS/install/logs/installvcs-unique string/` directory.

- 3 Verify that the cluster is operating properly after the upgrade.
See “[Verifying the VCS installation](#)” on page 143.

Using the `halogin` command for native OS accounts with VCS

VCS has deprecated the `AllowNativeCliUsers` attribute. To use native OS accounts with VCS, use the `halogin` command. Once you run `halogin`, VCS encrypts and stores your VCS credentials in the your home directory for a specific time period. After running `halogin`, you do not need to 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.

Follow the instructions below if you used the `AllowNativeCliUsers` attribute.

To set up VCS authentication for clusters running in secure mode

Ensure that each native user running VCS commands must have a home directory on the system from which the user runs VCS commands.

- 1 Create VCS user accounts for all users and assign privileges to these users.
- 2 If you run VCS commands from a remote host, set these 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`.

- 3 Run the `halogin` command:

```
# halogin vcsusername password
```

Perform steps 2 and 3 for each VCS user.

To set up VCS authentication for clusters not running in secure mode

Ensure that each native user running VCS commands must have a home directory on the system from which the user runs VCS commands.

- 1 Create VCS user accounts for all users and assign privileges to these users.
- 2 For each VCS user, run the `halogin` command:

```
# halogin vcsusername password
```

Upgrading VCS in a secure environment

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

You can also upgrade VCS manually.

See [“Upgrading VCS manually”](#) on page 165.

Upgrading VCS manually

If you want to upgrade from VCS versions 4.0 MP2, 4.1, or 4.1 MP1 to VCS version 5.0, you must perform the following steps to do a manual upgrade:

- [“Upgrading from VCS 4.0 MP2, 4.1, or 4.1 MP1”](#) on page 165
- [“Licensing VCS”](#) on page 167
- [“Starting LLT, GAB, and VCS”](#) on page 167

If you used the `AllowNativeCliUsers` attribute in your earlier VCS version, then you must perform the post-upgrade task that allows you to use native OS accounts with VCS.

See [“Using the `halogin` command for native OS accounts with VCS”](#) on page 164.

Upgrading from VCS 4.0 MP2, 4.1, or 4.1 MP1

Ensure that you perform these steps on each node in the cluster.

To upgrade to VCS 5.0 manually

- 1 Freeze the node:

```
# hasys -freeze -persistent <nodename>
```
- 2 Stop HAD.
- 3 Remove the deprecated resource types.
See [“Removing deprecated resource types”](#) on page 158.
- 4 If you have NFS installed on the node, you must edit the `main.cf` file to include details about the NFSRestart agent. Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for details.
If you have configured Oracle, Sybase, or DB2 databases, refer to the respective Enterprise agent documentation for upgrade information.
- 5 Take a backup of the `main.cf` file.

- 6 Upgrade the operating system to:
 - Red Hat Enterprise Linux 4 (RHEL 4) Update 3
 - SUSE Linux Enterprise Server 9 (SLES 9) with SP3
- 7 Insert the software disc with the VCS software into a drive connected to the system.

The CD is automatically mounted. If the CD does not get automatically mounted, then enter:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```
- 8 Navigate to the folder containing the RPMs.

Depending on the OS distribution and architecture, enter this command:

 - On RHEL 4 (i686)

```
# cd /mnt/cdrom/rhel4_i686/rpms
```
 - On RHEL 4 (x86_64)

```
# cd /mnt/cdrom/rhel4_x86_64/rpms
```
 - On SLES 9 (i586)

```
# cd /mnt/cdrom/sles9_i586/rpms
```
 - On SLES 9 (x86_64)

```
# cd /mnt/cdrom/sles9_x86_64/rpms
```
- 9 Upgrade the RPMs:

```
rpm -U *
```
- 10 Make the following changes in the `main.cf` file as applicable. Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for details about editing the `main.cf` file.
 - If you have an NFS resource configured in the cluster, edit the `main.cf` file to:
 - Remove the `IPResName` attribute.
 - Add the `Address` attribute, and specify the IP address as the value of this attribute.
 - Include details about the `NFSRestart` resource.
 - If you have an Apache resource configured in the cluster, edit the `main.cf` file to:
 - Remove this include statement:

```
include "ApacheTypes.cf"
```
 - Ensure that absolute path names are specified for the Apache resource attributes.
 - If you have a Mount resource configured in the cluster, specify `vxfs` as the value of the `FSType` attribute in the `main.cf` file.

- If you have a Diskgroup resource configured in the cluster, and you want the same behavior as with the previous version, specify **0** as the values for the StartVolumes and StopVolumes attributes in the `main.cf` file.
- If you have configured resources for Oracle, Sybase, or DB2 databases, refer to the respective Enterprise agent documentation for information about upgrading the respective agents.

11 Start HAD.

12 Unfreeze the node.

Licensing VCS

If you are upgrading from versions 2.1 and higher, your existing VCS license is accepted and no action is required. You may also apply a new VCS 5.0 license key, if required.

See “[Obtaining VCS license keys](#)” on page 54.

To license VCS

- ◆ On each node, type the commands:

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

Starting LLT, GAB, and VCS

After installing VCS 5.0 using the `installvcs` program, the installer prompts you to start VCS and its components. You can also start VCS and its components manually.

To start LLT

- ◆ On each node, type:

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

If LLT is configured correctly on each node, the console output resembles:

```
Apr 5 14:46:18 north llc: LLT:10009: LLT Protocol available
```

See “[Verifying LLT](#)” on page 150.

To start GAB

- ◆ On each node, type:

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

If GAB is configured correctly on each node, the console output resembles:

```
Apr 5 14:46:29 north gab: GAB:20021: GAB available
```

Apr 5 14:51:50 north gab: GAB:20026: Port a registration waiting for seed port membership

See “[Verifying GAB](#)” on page 152.

To start VCS

- ◆ On each node, type:

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

If VCS is configured correctly on each node, the console output resembles:

```
VCS NOTICE V-16-1-10619 'HAD' starting on: north
VCS NOTICE V-16-1-10620 Waiting for local cluster configuration
status
VCS NOTICE V-16-1-10625 Local cluster configuration valid
VCS NOTICE V-16-1-11034 Registering for cluster membership
VCS NOTICE V-16-1-11035 Waiting for cluster membership
GAB INFO V-15-1-20036 Port h gen 265f06 membership ;1
GAB INFO V-15-1-20038 Port h gen 265f06 k_jeopardy 0
GAB INFO V-15-1-20040 Port h gen 265f06 visible 0
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10082 System (north) is in Regular Membership
- Membership: 0x2
VCS NOTICE V-16-1-10073 Building from local configuration
VCS NOTICE V-16-1-10066 Entering RUNNING state
GAB INFO V-15-1-20036 Port h gen 265f07 membership 01
VCS INFO V-16-1-10077 Received new cluster membership
VCS NOTICE V-16-1-10082 System (south) is in Regular Membership
- Membership: 0x3
```

See “[Verifying the cluster](#)” on page 152.

Upgrading to the VCS 5.0 Java Console

When you upgrade to VCS release 5.0, you must also upgrade the Java Console (GUI) that you installed on a separate system. VCS 5.0 supports earlier versions of the Java Console, but earlier VCS versions cannot support later versions of Java Console.

Note: The VCS 5.0 Java Console requires JRE version 1.5. If necessary, you can add it when you add the Java Console RPM.

Use one of the following applicable procedures:

To upgrade Java console on Linux

- 1 Log in as superuser on the node where you intend to install the RPM.
- 2 Remove the GUI from the previous installation.
For example, to remove a GUI installed with VCS 2.1, 2.1 P1, 2.1 P2, 2.2, 2.2 MP1, or 2.2 MP2, type:

```
# rpm -e VRTScscm
```
- 3 Insert the software disc with the VCS software into a drive connected to the system.
- 4 Install the VCS Java Console.
See “[Installing the Java Console on Linux](#)” on page 81.

To upgrade Java console on Windows 2000 Professional Workstations

- 1 Remove the Java-based cluster manager from previous installations.
 - From the Control Panel, double-click **Add/Remove Programs**.
 - Select **Symantec Cluster Manager**.
 - Click **Add/Remove**.
 - Follow the uninstall wizard instructions.
- 2 Add the new Java-based Cluster Manager.
See “[Installing the Java Console on a Windows workstation](#)” on page 82.

Upgrading from CommandCentral Availability 4.1 MP1

You can upgrade CommandCentral Availability 4.1 MP1 to the Cluster Management Console configured in multi-cluster mode.

This upgrade requires that you:

- Upgrade the management server in CommandCentral Availability to the management server in the Cluster Management Console.
You must perform this upgrade on a standalone server that is outside of all clusters.
- Upgrade cluster monitor, the cluster communications agent in CommandCentral Availability, to cluster connector, the cluster communications agent in the Cluster Management Console.
You must perform this upgrade on each cluster that uses cluster monitor.

Upgrade order

The management server in the Cluster Management Console cannot communicate with cluster monitor. To minimize the amount of management downtime and to maintain cluster histories in the database, Symantec recommends that you upgrade the management server first.

The upgraded management server can immediately begin to monitor and manage any direct-connection clusters in your enterprise. It cannot monitor or manage any clusters connected through cluster monitor until you upgrade each cluster monitor to cluster connector. From the time that you finish upgrading the management server until you upgrade cluster monitor, the status of the associated cluster is UNKNOWN.

Upgrading the management server on Solaris

You must perform this upgrade on a standalone server system that is outside all clusters and available on the local network. The system must currently host the CommandCentral Availability management server (predecessor to the Cluster Management Console management server). Symantec Product Authentication Service, a shared component, is installed during management server upgrade. If an older version of the service is installed, it is upgraded to the latest version.

This procedure follows a script of a successful upgrade. If at any step you experience a result other than the expected result that is documented here, you can click “n” to re-enter information. If you continue to have problems, click “q” to quit the installation and then verify the installation prerequisites.

To upgrade the management server on Solaris

- 1 Insert the disc into the drive on the local system. At the command prompt, type the following command to run the setup program:

```
./installer -rsh
```

The setup program presents copyright information followed by a menu titled, "Storage Foundation and High Availability Solutions 5.0".

- 2 Enter **i** to specify a task.

```
Enter a Task: [I,C,L,P,U,D,Q,?] i
```

The setup program displays another menu that lists products that are available for installation.

- 3 Select the menu number that corresponds to Veritas Cluster Management Console.

```
Select a product to install: [1-13,b,q]
```

The setup program presents a description of the product.

- 4 Enter **1** to select a product component.

```
Enter '1' to install the management server, '2' to install the cluster connector: [1-2,q] (1) 1
```

The setup program presents a message stating that it will install the management server.

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

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

The setup program performs an initial system check of the local system and checks for installed packages on the local system. If these checks are satisfactory, the setup program lists the packages to be installed.

Storage Foundation and High Availability Solutions 5.0 installer will upgrade or install the following Cluster Management Console packages:

```
VRTSat          Symantec Product Authentication Service
VRTSperl        Veritas Perl 5.8.8 Redistribution
VRTSjre15       Veritas Java Runtime Environment Redistribution
VRTSweb         Veritas Java Web Server
VRTSdbms3       Symantec Shared DBMS
VRTScmcm        Veritas Cluster Management Console for multiple cluster environments
VRTScmcdc       User Documentation for Veritas Cluster Management Console
Press [Return] to continue:
```

- 6 Press Enter.

The setup program displays the following message:

```
Installer is now ready to upgrade Cluster Management Console. All Cluster Management Console processes that are currently running must be stopped.
```

- 7 Enter **y** to confirm that you want to upgrade the Cluster Management Console.

```
Are you sure you want to upgrade Cluster Management Console?  
[y,n,q] (y)
```

- 8 Enter a password for the management server service account and confirm it at the next prompt.

```
Enter a password for the CMC service account:xxxxxx  
Confirm the password you entered for the CMC service  
account:xxxxxx
```

When you install and configure cluster connector, you must provide this same password for the CMC_CC@CMC_SERVICES account.

After you confirm the password for the service account, the setup program displays:

- A list of pre-upgrade tasks as it performs them
- Uninstallation progress percentages for prior components
- Installation progress percentages for the upgrade components

If the installation is successful, the setup program displays the following message:

```
Installation completed successfully on all systems. Upgrading  
management server configuration.
```

- 9 Enter **y** to start Veritas Cluster Management Console processes now.

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

The setup program displays progress percentages for starting the console.

After the console starts, the setup program displays the location of the upgrade log files and then logs in with the initial user account. The account information for this user is echoed to the screen so that you can see the initial user information.

The setup program reports the state of any managed clusters as EXITED. This is because the CommandCentral Availability cluster monitor cannot communicate with the Cluster Management Console management server. After you upgrade each cluster monitor to cluster connector, all clusters except for direct connection clusters show up in the Cluster Management Console. You must add direct connection clusters manually using the Administration tab in the Cluster Management Console.

The setup program eventually displays the following message:

```
End of upgrade
```

Copying the VCS installation guide and release notes to each node

After you install the management server, copy the Install Guide and the Release Notes over to the management server. The documentation locations on the distribution disc are:

- Install Guide
cluster_server/docs/vcs_install.pdf
- Release Notes
cluster_server/release_notes/vcs_notes.pdf

Accessing Cluster Management Console information

Information about administering clusters in multi-cluster mode is available in the Veritas Cluster Server Centralized Management Guide. The online help includes information about administering clusters in both single-cluster and multi-cluster mode. If you want to access the information about managing a single cluster in printed documentation, you can install the documentation package to the desired system.

The documentation package name for each supported operating system is:

- AIX–VRTSvcs.doc
- HP-UX–VRTSvcsdc
Note that you can copy the documents from depot/VRTSvcsdc/VRTSvcsdc/opt/VRTS/docs.
- Linux–VRTSvcsdc
- Solaris–VRTSvcsdc

Upgrading the management server on Windows systems

You must perform this upgrade on a standalone server system that is outside all clusters and available on the local network. The system must currently host the CommandCentral Availability management server (predecessor to the Cluster Management Console management server). Symantec Product Authentication Service, a shared component, is installed during management server upgrade. If an older version of the service is installed, it is upgraded to the latest version.

To upgrade the management server on Windows

- 1 Insert the disc into the drive on the local system.
- 2 On the distribution disc, locate the **\installer** directory.
- 3 Double-click the **setup** file.

Depending upon the operating system, you may or may not receive the following warning message:

The publisher could not be verified. Are you sure you want to run this software?

If you receive this message, click **Run**.

- 4 In the Welcome to the Veritas Cluster Management Console Installation Manager dialog box, read the introduction and then click **Next**.
- 5 In the Installation and Configuration Options dialog box, click **Upgrade the management server on the local node** and then click **Next**.
- 6 In the warning message box, read the information about:
 - Backing up your database using the Settings tab in your current (non-upgraded) console. This task can help to protect your cluster configurations, historical cluster data, and console settings if the upgrade fails or produces unexpected results.
 - Upgrading existing cluster connectors (CommandCentral Availability cluster monitors) to enable cluster management in the upgraded console. The Cluster Management Console cannot manage clusters that currently run cluster monitor until you perform this upgrade.Click **OK**.
- 7 When prompted, enter the password for the management server service account for cluster connector, `CC_CMC@CMC_SERVICES`.
Record the password that you enter in a safe place. You must use it again whenever you install or configure cluster connector.
- 8 In the Summary dialog box, read the summarized record of the upgrade and then click **Next**.
- 9 In the Upgrading to Cluster Management Console dialog box, when you see the following message, click **Next**:
`Restoration of databaes on node NodeName complete.`
where *NodeName* if the name of the system on which are performing the upgrade.
- 10 Click **Finish**.

Upgrading cluster monitor to cluster connector on UNIX systems

You must perform this upgrade on each cluster that formerly used CommandCentral Availability cluster monitor to communicate with the CommandCentral Availability management server. This procedure follows a script of a successful upgrade. If at any step you experience a result other than the expected result that is documented here, you can click “n” to re-enter information. If you continue to have problems, click “q” to quit the installation and then verify the installation prerequisites.

To upgrade cluster monitor to cluster connector

- 1** Insert the disc into the drive on the local system. At the command prompt, type the following command to run the setup program:

```
./installer -rsh
```

The setup program presents copyright information followed by a menu titled, “Storage Foundation and High Availability Solutions 5.0”.
- 2** Enter **i** to specify a task.

```
Enter a Task: [I,C,L,P,U,D,Q,?] i
```

The setup program displays another menu that lists products that are available for installation.
- 3** Select the menu number that corresponds to Veritas Cluster Management Console.

```
Select a product to install: [1-13,b,q]
```

The setup program presents a description of the product.
- 4** Enter **2** to select a product component.

```
Enter '1' to install the management server, '2' to install the cluster connector: [1-2,q] (1) 2
```

The setup program presents a message stating that it will install cluster connector.
- 5** Enter **y** to verify that the information up to this point is correct.

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

The setup program performs an initial system check of the local system and checks for installed packages on the local system. If these checks are satisfactory, the setup program lists the packages to be installed.

```
Storage Foundation and High Availability Solutions 5.0 installer will upgrade or install the following Cluster Management Console packages:
```

VRTSat	Symantec Product Authentication Service
VRTSperl	Veritas Perl 5.8.8 Redistribution
VRTSjre15	Veritas Java Runtime Environment Redistribution

```
VRTScmccc   Veritas Cluster Management Console Cluster
Connector
Press [Return] to continue:
```

6 Press Enter.

The setup program displays the following message:

```
Installer is now ready to upgrade Cluster Management Console.
All Cluster Management Console processes that are currently
running must be stopped.
```

7 Enter **y** to confirm that you want to upgrade the Cluster Management Console.

```
Are you sure you want to upgrade Cluster Management Console?
[y,n,q] (y)
```

8 Enter a password for the cluster connector service account.

The password is the password that was entered for the cluster connector service account during management server installation.

```
Enter the password for the CMC service account:xxxxxx
```

9 Enter the root hash of the authentication broker installed on the management server.

```
Enter the hash of the Management Server's root broker [?]
To retrieve the root hash of the management server authentication broker,
run the following command:
```

- On Windows:

```
&program files\veritas\security\authentication\bin\vssat
showbrokerhash
```

- On UNIX systems:

```
/opt/VRTSat/bin/vssat showbrokerhash
```

The output of this command looks similar to the following:

```
Root Hash:          9dfde3d9aaebec084f8e35819c1fed7e6b01d2ae
```

Enter the alphanumeric string (the string you receive is different from the one shown).

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

The setup program presents a list of:

- Prior component shutdown progress percentages
- Prior component uninstallation progress percentages
- Upgrade component installation progress percentages
- Cluster connector deployment messages
- Cluster connector configuration messages

11 Enter **y** to start Veritas Cluster Management Console processes now.

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


The setup program presents startup progress percentages and, if successful, displays the following message:

```
Installation completed successfully on all systems
```

After the console starts, the setup program displays the location of the upgrade log files and then displays:

- Information about each upgrade package that was installed.
- Status information for each system, service group, and resource in the cluster.
- A message that is an evaluation of the state of the upgrade (success or failure).

The setup program eventually displays the following message:

```
End of upgrade
```

Upgrading cluster monitor to cluster connector on Windows platforms

You must perform this upgrade on each cluster that formerly used CommandCentral Availability cluster monitor to communicate with the CommandCentral Availability management server.

To upgrade cluster monitor to cluster connector

- 1 In CommandCentral Availability, locate and record the value of the GUID attribute. This attribute is on the CCAvail_ClusterConfig resource.
- 2 Uninstall cluster monitor on each cluster node.
Refer to the CommandCentral Availability documentation for cluster monitor uninstallation instructions.
- 3 Insert the distribution disc into the DVD drive on the cluster node.
- 4 Obtain a command prompt and navigate to the **\installer** directory.
- 5 At the command line, enter the following command:

```
setup.bat -guid xxxxxx
```

where **xxxxxx** is the value of the GUID attribute you recorded in step 1. The setup program installs the Cluster Management Console cluster connector.

For information on the documentation that comes with this release, see:

- [“Copying the VCS installation guide and release notes to each node”](#) on page 173
- [“Accessing Cluster Management Console information”](#) on page 173

Adding and removing cluster nodes

This chapter contains the following topics:

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

About adding and removing nodes

After installing VCS and creating a cluster, you can add and remove nodes from the cluster. You can create a clusters of up to 32 nodes.

Adding a node to a cluster

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

See [“Preparing to install and configure VCS”](#) on page 21.

[Table 7-13](#) specifies the tasks involved in adding a cluster. The example demonstrates how to add a node east to already existing nodes, north and south.

Table 7-13 Tasks involved in adding a node to a cluster

Task	Reference
Set up the hardware.	“Setting up the hardware” on page 180
Install the software manually.	“Installing VCS software manually” on page 119


Table 7-13 Tasks involved in adding a node to a cluster

Task	Reference
Add a license key.	“Adding a license key” on page 182
Configure LLT and GAB.	“Configuring LLT and GAB” on page 182
Add the node to the existing cluster.	“Adding the node to the existing cluster” on page 184
Start VCS and verify the cluster.	“Starting VCS and verifying the cluster” on page 184

Setting up the hardware

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

Figure 7-9 Adding a node to a three-node cluster using two independent hubs



To set up the hardware

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

[Figure 7-9](#) 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

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 55.

To prepare for installation

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 selecting the optional RPMs, note:

- Symantec recommends that you install the RPMs for VCS manual pages (VRTSvcsmn) and VCS documentation (VRTSvcsdc). Install the documentation RPM on nodes where you want access to the documentation. See “[Setting up I/O fencing](#)” on page 111.
- The VCS configuration wizard (VRTSscsw) 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.

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.

```
# rpm -i VRTS11t-*.rpm
# rpm -i VRTSgab-*.rpm
# rpm -i VRTSvcs-*.rpm
# rpm -i VRTSvcsmg-*.rpm
# rpm -i VRTSvcsag-*.rpm
# rpm -i VRTScutil-*.rpm
# rpm -i VRTSscsw-*.rpm
# rpm -i VRTSvcsw-*.rpm
```

- 2 Install the optional RPMs, in the order shown. Omit those that you do not want to install.

```
# rpm -i VRTSvxfen-*.rpm
# rpm -i VRTSvcsmn-*.rpm
# rpm -i VRTSvcsdc-*.rpm
# rpm -i VRTSvcsdr-*.rpm
```

```
# rpm -i VRTScssim-*.rpm
# rpm -i VRTScscm-*.rpm
# rpm -i VRTSvcApache-*.rpm
```

Adding a license key

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

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

Checking licensing information on the system

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

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

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

Configuring LLT and GAB

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

To configure LLT

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

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

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

```
0 north
1 south
```

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

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

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

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

See “[/etc/llttab](#)” on page 144.

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
```

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

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

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

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

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

See “[/etc/gabtab](#)” on page 144.

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

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

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

To verify GAB

- 1 On the new node, run the command:

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

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

```
GAB Port Memberships
=====
Port a gen a3640003 membership 012
```

See “[Verifying GAB](#)” on page 152.

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

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

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

Adding the node to the existing cluster

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

To add the new node to the existing cluster

- 1 Enter the command:
`# haconf -makerw`
- 2 Add the new system to the cluster:
`# hasys -add east`
- 3 Stop VCS on the new node:
`# hastop -sys east`
- 4 Copy the main.cf file from an existing node to your new node:
`# rcp /etc/VRTSvcs/conf/config/main.cf east:/etc/VRTSvcs/conf/config/`
- 5 Start VCS on the new node:
`# hastart`
- 6 If necessary, modify any new system attributes.
- 7 Enter the command:
`# haconf -dump -makero`

Starting VCS and verifying the cluster

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

To start VCS and verify the cluster

- 1 From the new system, start VCS with the new system added to the cluster:
`# hastart`
- 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 7-14](#) specifies the tasks involved in removing a node from a cluster. In the example procedure, the cluster consists of nodes A, B, and C; node C is to leave the cluster.

Table 7-14 Tasks involved in removing a node

Task	Reference
<ul style="list-style-type: none"> ■ Back up the configuration file. ■ Check the status of the nodes and the service groups. 	“Verify the status of nodes and service groups” on page 185
<ul style="list-style-type: none"> ■ Switch or remove any VCS service groups on the node leaving the cluster. ■ Delete the node from VCS configuration. 	“Deleting the leaving node from VCS configuration” on page 186
Modify the lthosts and gabtab files to reflect the change.	“Modifying configuration files on each remaining node” on page 188
On the node leaving the cluster: <ul style="list-style-type: none"> ■ Modify startup scripts for LLT, GAB, and VCS to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. ■ Remove the VCS RPMs. 	“Unloading LLT and GAB and removing VCS on the leaving node” on page 188

Verify the status of nodes and service groups

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

To verify the status of the nodes and the service groups

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

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

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

```
# hastatus -summary

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

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

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

Deleting the leaving node from VCS configuration

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

To remove or switch service groups from the leaving node

- 1 Switch failover service groups from the leaving 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 leaving node; for example, `grp4` runs only on the leaving node.

```
# hagrps -dep
```
- 3 If the service group on the leaving node requires other service groups, that is, if it is a parent to service groups on other nodes, then unlink the service groups.

```
# haconf -makerw
# hagrps -unlink grp4 grp1
```

These commands enable you to edit the configuration and to remove the requirement `grp4` has for `grp1`.
- 4 Stop VCS on the leaving node:

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

```
# hastatus -summary
```

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

```

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

```

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

```

# hagr -modify grp3 SystemList -delete C
# hagr -modify grp4 SystemList -delete C

```

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

```

# hagr -resources grp4
    processx_grp4
    processy_grp4
# hares -delete processx_grp4
# hares -delete processy_grp4

```

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

```

# hagr -delete grp4

```

- 9 Check the status.

```

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

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

```

- 10 Delete the node from the cluster.

```

# hasys -delete C

```

- 11 Save the configuration, making it read only.

```

# haconf -dump -makero

```

Modifying configuration files on each remaining node

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

To modify the configuration files on a remaining node

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

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

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

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

For example, change:

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

to:

```
0 A
1 B
```

Unloading LLT and GAB and removing VCS on the leaving node

Perform the tasks on the node leaving 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 VRTScscm
# rpm -e VRTSvcs
# rpm -e VRTScscw
# rpm -e VRTScssim
```

```
# rpm -e VRTScutil
# rpm -e VRTSvcsdc
# rpm -e VRTSvcsmn
# rpm -e VRTSvcsag
# rpm -e VRTSvcsmg
# rpm -e VRTSvcs
# rpm -e VRTSvxfen
# rpm -e VRTSgab
# rpm -e VRTSllt
# rpm -e VRTSperl
```

- 4 Optionally, you can remove the infrastructure RPMs in the order shown:

```
# rpm -e VRTSvlic
# rpm -e VRTScpi
# rpm -e VRTSperl
# rpm -e VRTSweb
# rpm -e VRTSjre
```

- 5 Remove the LLT and GAB configuration files.

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


Installing VCS on a single node

This chapter contains the following topics:

- [About installing VCS on a single node](#)
- [Creating a single-node cluster using the 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 5.0 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 192.

See [“Creating a single-node cluster manually”](#) on page 193.

Creating a single-node cluster using the installer program

[Table 8-15](#) specifies the tasks involved in installing VCS on a single node using the installer program.

Table 8-15 Tasks to create a single-node cluster using the installer

Task	Reference
Prepare for installation.	“Preparing for a single node installation” on page 192
Install the VCS software on the system using the installer.	“Starting the installer for the single node cluster” on page 192

Preparing for a single node installation

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

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

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

For more information about LLT and GAB:

See [“LLT and GAB”](#) on page 17.

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 66.

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 standalone cluster.

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

Continue with the installation.

See “[Licensing VCS](#)” on page 67.

Creating a single-node cluster manually

[Table 8-16](#) specifies the tasks involved in installing VCS on a single node.

Table 8-16 Tasks to create a single-node cluster manually

Task	Reference
Set the PATH variable	“ Setting the PATH variable ” on page 193
Install the VCS software manually and adding a license key	“ Installing VCS software manually ” on page 119
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.	“ Renaming the LLT and GAB startup files ” on page 194
Modify the VCS startup file for single-node operation.	“ Modifying the startup files ” on page 194
Create and modify the VCS configuration files.	“ Configuring VCS ” on page 194
Start VCS and verify single-node operation.	“ Verifying single-node operation ” on page 196

Setting the PATH variable

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

To set the PATH variable

- ◆ Do one of the following:

- For the Bourne Shell (sh or ksh), type:

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

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

Installing the VCS software manually

Install the VCS 5.0 RPMs manually and install the license key.

See “[Installing VCS software manually](#)” on page 119.

See “[Adding a license key](#)” on page 182.

Renaming the LLT and GAB startup files

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

To rename the LLT and GAB startup files

- ◆ Rename the LLT and GAB startup files.

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

Modifying the startup files

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

Change the line:

```
ONENODE=no
```

To:

```
ONENODE=yes
```

Configuring VCS

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

main.cf file

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

- An “include” statement that specifies the file, `types.cf`, which defines the VCS bundled agent resources.

- The name of the cluster.
- The name of the systems that make up the cluster.

Editing the main.cf file

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

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

To edit the main.cf file

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

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

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

Example, main.cf

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

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

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

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

types.cf file

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

Verifying single-node operation

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

To verify single-node cluster

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

```
# hastart -onenode
```

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

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

Adding a node to a single-node cluster

[Table 8-17](#) specifies the activities involved in adding nodes to a single-node cluster. All nodes in the new cluster must run the same version of VCS. The example procedure refers to the existing single-node VCS node as Node A and the node that is to join Node A to form a multiple-node cluster as Node B.

Table 8-17 Tasks to add a node to a single-node cluster

Task	Reference
Set up Node B to be compatible with Node A	“Setting up a node to join the single-node cluster” on page 197
<ul style="list-style-type: none"> ■ Add ethernet cards for private heartbeat network for Node B ■ If necessary, add ethernet cards for private heartbeat network for Node A ■ Make the ethernet cable connections between the two nodes 	“Installing and configuring Ethernet cards for private network” on page 198
Connect both nodes to shared storage	“Configuring the shared storage” on page 198
<ul style="list-style-type: none"> ■ Bring up VCS on Node A ■ Edit the configuration file ■ Edit the startup scripts 	“Bringing up the existing node” on page 199

Table 8-17 Tasks to add a node to a single-node cluster

Task	Reference
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.	“Installing VCS software manually” on page 119
Edit the configuration files on Node B	“Configuring LLT and GAB” on page 199
Start LLT and GAB on Node B	“Starting LLT and GAB” on page 202
<ul style="list-style-type: none"> ■ Start LLT and GAB on Node A ■ Restart VCS on Node A ■ Modify service groups for two nodes 	“Reconfiguring VCS on the existing node” on page 202
<ul style="list-style-type: none"> ■ Start VCS on Node B ■ Verify the two-node cluster 	“Verifying configuration on both nodes” on page 203

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:
 - 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 removing the node from the cluster, remove the VCS RPMs and configuration files. See [“Removing a node from a cluster”](#) on page 185.
 - If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS. See [“Uninstalling VCS 5.0”](#) on page 105
 - If you renamed the LLT and GAB startup files, remove them. See [“Renaming the LLT and GAB startup files”](#) on page 194.
- 2 If necessary, install VxVM and VxFS. See [“Installing VxVM, VxFS if necessary”](#) on page 198.

Installing VxVM, VxFS if necessary

If VxVM with the cluster option or VxFS with the cluster option is installed on the existing node in the cluster, then the same versions must also be installed on the new node.

Refer to the appropriate documentation for VxVM and VxFS to verify the versions of the installed products and make sure the same version is running on all nodes that are to use any 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, two separate NICs on each node should be used, such that the failure of one NIC does not restrict heartbeating between the nodes.

See “[Setting up the private network](#)” on page 46.

To install and configure ethernet cards for private network

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

Configuring the shared storage

Make the connection to shared storage from Node B. Configure VxVM on Node B and reboot the node when you are prompted.

See “[Setting up shared storage](#)” on page 49.

See “[Setting up shared storage: Fibre Channel](#)” on page 51.

Bringing up the existing node

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

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

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

```
# hagr -freeze group -persistent
```

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

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

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

```
ONENODE=yes
```

To:

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

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

Installing the VCS software manually

Install the VCS 5.0 RPMs manually and install the license key.

See “[Installing VCS software manually](#)” on page 119.

See “[Adding a license key](#)” on page 182.

Configuring LLT and GAB

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

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

Configuring low latency transport (LLT)

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

Setting up `/etc/llthosts`

The file `llthosts(4)` is a database, containing one entry per system that links the LLT system ID (in the first column) with the LLT host name. You must create an identical file on each node in the cluster.

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

```
0 north
1 south
```

Setting Up `/etc/llttab`

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

See “[LLT directives](#)” on page 201.

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

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

The first line must identify the system on which the file exists. In the example above, the value for `set-node` could be `north`, `0`, or the file name `/etc/nodename`, provided the file contains the name of the system (`north` in this example). The next two lines, beginning with the `link` command, identify the two private network cards that the LLT protocol uses. The order of directives must be the same as in the sample file `/opt/VRTSllt/llttab`.

LLT directives

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

Table 8-18 LLT directives

Directive	Description
<code>set-node</code>	<p>Assigns the system ID or symbolic name. The system ID number must be unique for each system in the cluster, and must be in the range 0-31. The symbolic name corresponds to the system ID listed in <code>/etc/llthosts</code> file.</p> <p>Note that LLT fails to operate if any systems share the same ID.</p>
<code>link</code>	<p>Attaches LLT to a network interface. At least one link is required, and up to eight are supported. The first argument to <code>link</code> is a user-defined tag shown in the <code>lltstat(1M)</code> output to identify the link. It may also be used in <code>llttab</code> to set optional static MAC addresses.</p> <p>The second argument to <code>link</code> is the device name of the network interface. Its format is <code>device_name:device_instance_number</code>. The remaining four arguments to <code>link</code> are defaults; these arguments should be modified only in advanced configurations. There should be one link 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.</p>
<code>set-cluster</code>	<p>Assigns a unique cluster number. Use this directive when more than one cluster is configured on the same physical network connection. LLT uses a default cluster number of zero.</p>
<code>link-lowpri</code>	<p>Use this directive in place of <code>link</code> for public network interfaces. This directive prevents VCS communication on the public network until the network is the last link, and reduces the rate of heartbeat broadcasts. Note that LLT distributes network traffic evenly across all available network connections and, in addition to enabling VCS communication, broadcasts heartbeats to monitor each network connection.</p>

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

Additional considerations for LLT

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

Configuring group membership and atomic broadcast (GAB)

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

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

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

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

Starting LLT and GAB

On the new node, start LLT and GAB.

To start LLT and GAB

- 1 Start LLT on Node B.

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

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

Reconfiguring VCS on the existing node

- 1 On Node A, create the files `/etc/llttab`, `/etc/llthosts`, and `/etc/gabtab`. Use the files created on Node B as a guide, customizing the `/etc/llttab` for Node A.
- 2 Start LLT on Node A.

```
# /etc/init.d/llt start
```
- 3 Start GAB on Node A.

```
# /etc/init.d/gab start
```
- 4 Check the membership of the cluster.

```
# gabconfig -a
```
- 5 Start VCS on Node A.

```
# hastart
```
- 6 Make the VCS configuration writable.

```
# haconf -makerw
```
- 7 Add Node B to the cluster.

```
# hasys -add sysB
```

8 Add Node B to the system list of each service group.

■ List the service groups.

```
# hagrp -list
```

■ For each service group listed, add the node.

```
# hagrp -modify group SystemList -add sysB 1
```

Verifying configuration on both nodes

1 On Node B, check the cluster membership.

```
# gabconfig -a
```

2 Start the VCS on Node B.

```
# hastart
```

3 Verify that VCS is up on both nodes.

```
# hastatus
```

4 List the service groups.

```
# hagrp -list
```

5 Unfreeze the service groups.

```
# hagrp -unfreeze group -persistent
```

6 Implement the new two-node configuration.

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


Advanced topics related to installing VCS

This appendix contains the following topics:

- [LLT over UDP](#)
- [Minimal downtime upgrade](#)
- [Setting up a trust relationship between two authentication brokers](#)

LLT over UDP

VCS 5.0 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 when:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

Performance considerations

Because LLT over UDP is slower than LLT over Ethernet, LLT over UDP should only be used when the hardware configuration makes it necessary.

Configuring LLT over UDP

Following is a checklist for configuring LLT over UDP. Examples are provided in the sections that follow.

- 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 206. See the examples in the following sections.
- Make sure that each NIC has an IP address 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 non-well known UDP port. See “[Selecting UDP ports](#)” on page 208.
- Set the broadcast address correctly for direct-attached (non-routed) links.
- For 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 210.

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

[Table A-19](#) describes the fields of the `link` command shown in the `/etc/llttab` file examples.

See “[Sample configuration: Direct-attached links](#)” on page 209.

See “[Sample configuration: Links crossing IP routers](#)” on page 210.

Note that some of these fields differ from the command for standard LLT links.

Table A-19 Field description for link command in /etc/llttab

Field	Description
<tag-name>	A unique string that is used as a tag by LLT; for example link1, link2,....
<device>	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.
<node-range>	Nodes using the link. “-” indicates <i>all</i> cluster nodes are to be configured for this link.
<link-type>	Type of link; must be “udp” for LLT over UDP.
<udp-port>	Unique UDP port in the range of 49152-65535 for the link. See “ Selecting UDP ports ” on page 208.
<MTU>	“-” 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.
<IP address>	IP address of the link on the local node.
<bcast-address>	<ul style="list-style-type: none"> ■ For clusters having broadcasts enabled, 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. [Table A-20](#) describes the fields of the `set-addr` command.

See “[Sample configuration: Links crossing IP routers](#)” on page 210.

Table A-20 Field description for set-addr command in /etc/llttab

Field	Description
<node-id>	The ID of the cluster node; for example, 0.

Table A-20 Field description for set-addr command in /etc/l1ttab

Field	Description
<link tag-name>	The string used by LLT to identify the link; for example link1, link2,....
<address>	IP address assigned to the link for the peer node.

Selecting UDP ports

When selecting a UDP port, select an available 16-bit integer from the range described below.

- Use available ports (that is, ports that are not in use) in the private range 49152 to 65535
- Do not use:
 - 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; UDP ports 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 LLT on subnets

You need to make sure to properly configure the netmask and broadcast address when nodes reside on different subnets.

Configuring the netmask

If you have nodes on different subnets, set the netmask so that the nodes can access the subnets in use.

For example:

- 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

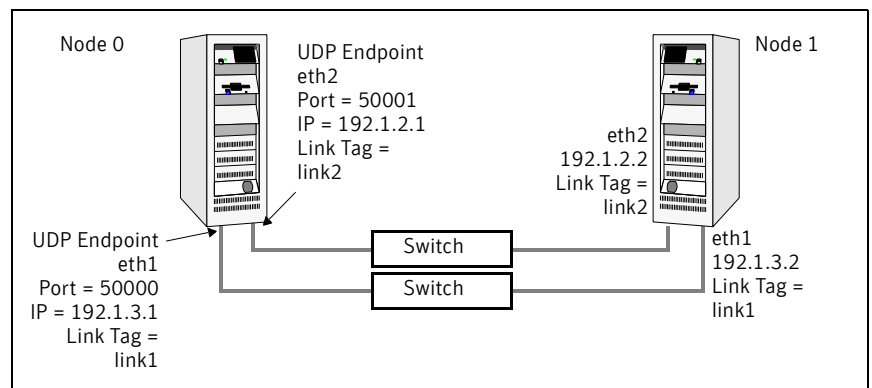
If you have nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.

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

The following illustration depicts a typical configuration of direct-attached links employing LLT over UDP.



The configuration represented by the following `/etc/llttab` file for Node 0 has directly attached crossover links or links connected through a hub or switch. These links do not cross routers.

Because LLT broadcasts requests peer nodes to discover their addresses, 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.

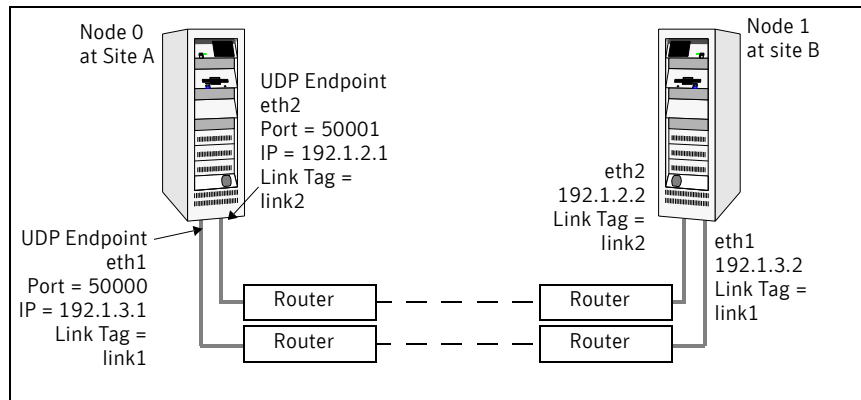
```
set-node Node0
set-cluster 1
#configure Links
#link <tag-name> <device> <node-range> <link-type> <udp port>
<MTU> <IP-address> <bcast-address>
link link1 udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 would resemble:

```
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

The following illustration depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows just two nodes of a four-node cluster.



The configuration represented by the following `/etc/llttab` file 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 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 would resemble:

```
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
```

Minimal downtime upgrade

Use a minimal downtime upgrade to upgrade VCS. This procedure minimizes downtime for the cluster that you want to upgrade. In situations where you can fail over all your service groups to running nodes, downtime equals the time that it takes to offline and online the service groups. In situations where you have a service group that you cannot fail over to a running node during upgrade, downtime for that service group equals the time that it takes to perform an upgrade and reboot the node.

Supported upgrades

Use this procedure to upgrade from VCS 4.0 or 4.1.

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 in order to minimize the downtime for any particular service group.
- Some rough guidelines are:
 - 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 tampering 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.

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

In your system list, you have each service group failing over to one other node.

- 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 A-10 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, it is best if sg3 is 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 `sg4` 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 machines by switching service groups.

Performing the minimal downtime example upgrade

This upgrade uses four nodes with four service groups—note that in this scenario the service groups cannot stay online for part of the upgrade. Remember to not add, remove, or change resources or service groups on any nodes during the upgrade as these changes are likely to get lost after the upgrade.

To establish running service groups

- 1 Establish where your service groups are online.

```
# hagrps -state
#Group      Attribute      System      Value
sg1         State         node01     |OFFLINE|
sg1         State         node02     |ONLINE|
sg2         State         node01     |OFFLINE|
sg2         State         node03     |ONLINE|
sg3         State         node01     |OFFLINE|
sg3         State         node04     |ONLINE|
sg4         State         node03     |ONLINE|
sg4         State         node04     |OFFLINE|
```

- 2 Switch the service groups from all the nodes that you are first upgrading (`node02` and `node03`) to the remaining nodes (`node01` and `node04`).

```
# hagrps -switch sg1 -to node01
# hagrps -switch sg2 -to node01
# hagrps -switch sg4 -to node04
```

- 3 Verify that your service groups are offline on the nodes targeted for upgrade.

```
# hagrps -state
#Group      Attribute      System      Value
sg1         State         node01     |ONLINE|
sg1         State         node02     |OFFLINE|
sg2         State         node01     |ONLINE|
sg2         State         node03     |OFFLINE|
sg3         State         node01     |OFFLINE|
sg3         State         node04     |ONLINE|
```

```

sg4          State          node03      |OFFLINE|
sg4          State          node04      |ONLINE|

```

During the next procedure, do not perform any configuration tasks. Do not start any modules.

To perform the minimum downtime upgrade on target nodes

- 1 On the target nodes, start the 5.0 installer for VCS.
- 2 Select the VCS installation.
- 3 Answer **n** when the installer asks:


```
Do you want to upgrade to version 5.0 on these systems using
the current configuration? [y,n,q,?] (y) n
```
- 4 Answer with the names of the nodes that you want to upgrade:


```
Enter the system names separated by spaces on which to install
VCS: node02 node03
```
- 5 Select either option **1** or **2** when the installer asks:


```
Select the packages to be installed on all systems? 2
```
- 6 Answer **n** when the installer completes and asks:


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

To edit the configuration and prepare for upgrade node01 and node04

- 1 When HAD is down on node02 and node03, you see this message:


```
Shutdown completed successfully on all systems.
```
- 2 After you see the above message, you can make the VCS configuration writable on node01 or node04. Note that you need to make the configuration writable because the installer froze the service groups during the upgrade.


```
# haconf -makerw
```
- 3 Unfreeze all service groups.


```
# hagrp -unfreeze sg1 -persistent
# hagrp -unfreeze sg2 -persistent
# hagrp -unfreeze sg3 -persistent
# hagrp -unfreeze sg4 -persistent
```
- 4 Dump the configuration and make it read-only.


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

To edit the configuration on node02 and node03

- 1 Open the main.cf file, and delete the `Frozen = 1` line for each service group as appropriate.
- 2 Save and close the file.

Minimal downtime upgrade

- 3 Reboot node02 and node03.
- 4 Wait for GAB to come up. In the console's output, look for a line that reads:
`Starting GAB is done.`

To upgrade and restart your clusters

- 1 On node01 and node04, take the service groups offline.

```
# hagrps -offline sg1 -sys node01
# hagrps -offline sg2 -sys node01
# hagrps -offline sg3 -sys node04
# hagrps -offline sg4 -sys node04
```
- 2 On node01 and node04, perform the upgrade.
See [“To perform the minimum downtime upgrade on target nodes”](#) on page 215.
- 3 When HAD is down on node01 and node04, you see this message:
`Shutdown completed successfully on all systems.`
- 4 Start vxfsconfig on node02 and node03.

```
# vxfsconfig -c
```
- 5 Start your cluster on node02 and node03.

```
# hstart
```
- 6 After the upgrade completes, reboot node01 and node04.

After you have rebooted the nodes, all four nodes now run the latest version of VCS.

In this example, you achieved minimal downtime because your service groups were down only from the point when you took them offline on node01 and node04, to the time VCS brought them online on node02 or node03 as appropriate.

Setting up a trust relationship between two authentication brokers

This procedure is a general prerequisite to add secure direct connection clusters to a management server or a peer management server.

To set up the trust relationship

- 1 Identify which two systems with authentication brokers are to participate in the trust relationship.

To set up a peer management server, these systems are:

- The local management server host
- The remote management server host

For adding secure direct connection clusters, these systems are:

- The system that hosts the authentication broker used by the cluster (can be a cluster node in the cluster)
- The local management server host

- 2 On one of the two systems, run the following command (the dashes preceding the command options are double dashes):

```
# vssat setuptrust --broker systemname:2821 --securitylevel low
```

Where systemname is the DNS server name or the IP address of the second (other) system.

- 3 On the second (other) system, obtain a command prompt and run the following command:

```
vssat setuptrust --broker systemname:2821 --securitylevel low
```

Where systemname is the DNS server name or the IP address of the first system.

If you are adding secure direct connection clusters, you are finished with this procedure.

If you are adding a peer management server, you must continue with the following steps:

- 4 On both systems, obtain a command prompt and run the following command (the dashes in steps 4–6 are single dashes):

```
# gaserver -list
```

This command returns the ID, system name, and address of the system that runs the command. Record these data for both systems.

- 5 On the first system, run the following command:

```
# gaserver -add ID systemname address [port]
```

Where ID, systemname, and address are the ID, system name, and IP address of the second (other) system, and port is the default Cluster Management Console port, 14145.

Setting up a trust relationship between two authentication brokers

- 6 On the second (other) system, obtain a command prompt and run the following command:

```
# gaserver -add ID systemname address [port]
```

Where ID, systemname, and address are the ID, system name, and IP address of the first system, and port is the default Cluster Management Console port, 14145.
- 7 To verify the peer management server, run the gaserver -list command on each management server host. If both command runs return data for both systems, the peer management server is configured successfully.

Sample VCS installation and configuration output

This appendix contains the following topics:

- [About sample VCS installation and configuration](#)
- [Installing the Root Broker](#)
- [Installing the Cluster Management Console Management Server](#)
- [Installing VCS 5.0](#)
- [Configuring VCS 5.0](#)
- [Uninstalling VCS 5.0](#)

About sample VCS installation and configuration

The sample installation involves installing VCS with all the optional features on two systems, north and south. Make sure that you performed the pre-installation tasks.

Installing the Root Broker

The Root Broker administrator must perform this task before you configure the cluster in secure mode or configure Cluster Connector for centralized management of clusters.

See [“Preparing to install VCS 5.0”](#) on page 29.

```
# cd cluster_server  
  
# ./installvcs -security
```

[3] Install Symantec Product Authentication Service Root Broker.

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

Enter the system name on which to install VxSS: **east**

Initial system check:

```
Checking rsh communication with east..... Linux 2.6.5-7.244-kmp
Checking VxSS installation on east ..... not installed
Checking kernel release on east ..... 244
Checking distribution match with east ..... OK
Checking architecture on east ..... OK
Checking rpm dist match with east..... OK
Checking for SE Linux on east ..... OK
```

Veritas Cluster Server 5.0 Installation Program

Checking installed rpms on east

Veritas Cluster Server 5.0 Installation Program

The following VxSS rpms will be installed:

```
VRTSperl      Veritas Perl 5.8.8 Redistribution
VRTSicsco     Symantec Infrastructure Core Services Common
VRTSspbx      Symantec Private Branch Exchange
VRTSatClient  Symantec Product Authentication Service Client
VRTSatServer  Symantec Product Authentication Service
```

Veritas Cluster Server 5.0 Installation Program

Installing VxSS: 0%

Installing VxSS: 100%

Veritas Cluster Server 5.0 Installation Program

Installation completed successfully on all systems

It is optional to configure VxSS now. If you choose to configure VxSS later, you can either do so manually or run the `installvcs -configure` command.

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

Enter password for Root Broker administrator for the Root Broker on host venus

Password must contain at least 5 characters.

Please enter a new password or <Control-C> to quit.

Enter password for Authentication Broker administrator for the
Authentication Broker on host venus

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

Veritas Cluster Server 5.0 Installation Program

Starting VxSS: 0%

Starting VxSS: 100%

Startup completed successfully on all systems

Installing the Cluster Management Console Management Server

The Cluster Management Console Management Server administrator must set up the Management Server before you configure this cluster for centralized management.

Installation Program

Copyright (c) 2006 Symantec Corporation. All rights reserved. Use of this product is subject to license terms. Federal Acquisitions: Commercial Software. Government Users Subject to Standard License Terms and Conditions.

Symantec, the Symantec Logo and all other Symantec product names and slogans are trademarks or registered trademarks of Symantec Corporation in the United States and certain other countries. Additional company and product names may be trademarks or registered trademarks of the individual companies and are respectfully acknowledged.

Logs for installer are being created in /var/tmp/installer-xxxxxx.

Storage Foundation and High Availability Solutions 5.0

Symantec Product Version Installed Licensed
=====

Symantec Licensing Utilities are not installed. Unable to determine product installation status.

Task Menu:

- I) Install/Upgrade a Product
- L) License a Product
- U) Uninstall a Product
- Q) Quit
- C) Configure an Installed Product
- P) Perform a PreinstallationCheck
- D) View a Product Description
- ?) Help

Enter a Task: [I,C,L,P,U,D,Q,?] **i**

Storage Foundation and High Availability Solutions 5.0

- 1) Veritas Cluster Server
- 2) Veritas File System
- 3) Veritas Volume Manager
- 4) Veritas Volume Replicator
- 5) Veritas Storage Foundation
- 6) Veritas Storage Foundation for Oracle
- 7) Veritas Storage Foundation for DB2
- 8) Veritas Storage Foundation for Sybase
- 9) Veritas Storage Foundation for Cluster File System
- 10) Veritas Storage Foundation for Oracle RAC
- 11) Veritas Cluster Management Console
- 12) Web Server for Storage Foundation Host Management
- 13) Symantec Product Authentication Service
- b) Back to previous menu

Select a product to install: [1-13,b,q] **11**

Storage Foundation and High Availability Solutions 5.0

Veritas Cluster Management Console component information and configuration:

Read the following component descriptions and choose the appropriate target.

Management Server:

The Management Server is the CMC component that manages your clusters.

Cluster Connector:

The Cluster Connector is the CMC component that acts as an agent to your clusters.

- 1) Management Server
- 2) Cluster Connector

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

Setup will install the Management Server.

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

Initial system check:

Checking CMC installation on venus not installed

Storage Foundation and High Availability Solutions 5.0

Checking installed packages on venus

.
.

Storage Foundation and High Availability Solutions 5.0

installer will install the following CMC packages:

VRTSat	Symantec Product Authentication Service
VRTSperl	Veritas Perl 5.8.8 Redistribution
VRTSdbms3	Symantec Shared DBMS
VRTSjre15	Veritas Java Runtime Environment Redistribution
VRTSweb	Veritas Java Web Server
VRTScmcm	Veritas Cluster Management Console
VRTScmcdc	Veritas Cluster Management Console Documentation

Press [Return] to continue:

It is possible to install CMC packages without performing configuration.

It is optional to configure CMC now. If you choose to configure CMC later, you can either do so manually or run the installcmc - configure command.

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

Storage Foundation and High Availability Solutions 5.0

The following information is required to configure the Management Server:

A unique Management Server display name
A unique Management Server address
The desired location of the Management Server database

The name of the user that will be the initial administrator

A service account password

Optionally, the address of a remote root broker

```
Enter a unique management server display name: [?] venus_cmc_ms  
Enter the network address used by the management server [b,?] (venus.symantecexample.com) venus.symantecexample.com  
Enter the desired location of the database to be used by the management server [b,?] (/opt/VRTScmc/db) /opt/VRTScmc/db
```

Storage Foundation and High Availability Solutions 5.0

Management Server configuration verification:

```
Management Server Display Name: venus_cmc_ms  
Management Server Address: venus.symantecexample.com  
Database location: /opt/VRTScmc/db
```

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

Storage Foundation and High Availability Solutions 5.0

Veritas Cluster Management Console initial admin user configuration:

Read the following user selection descriptions and choose the appropriate method.

Local User:

Configure a user on the local machine as the initial admin user.

Custom User:

Configure a user manually.

- 1) Local User
- 2) Custom User

Enter '1' to enter the name of a local user, '2' to set up a custom user: [1-2,q] (1) **1**

Storage Foundation and High Availability Solutions 5.0

Local admin user selection:

To log in to the CMC Management Server, enter the name of a local user to be set as the administrator. The domain and domain type will be automatically selected for you.

Enter the initial management server user name: [b,?] (root) **root**

Storage Foundation and High Availability Solutions 5.0

Management Server admin user verification:

Management Server User Name: root

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

Storage Foundation and High Availability Solutions 5.0

Set Veritas Cluster Management Console Management Server service account password:

The CMC Management Server creates a service account for secure communication between it and any managed clusters.

Enter a password for the CMC service account:

Confirm the password you entered for the CMC service account:

Storage Foundation and High Availability Solutions 5.0

Do you want to specify a remote root broker? [y,n,q,b,?] (n) **n**

Storage Foundation and High Availability Solutions 5.0

Installing CMC: 0%

.
.

Installing CMC: 100%

Installation completed successfully on all systems

Storage Foundation and High Availability Solutions 5.0

Creating Management Server configuration file Done

Creating CMC service account and domain for secure communication:

Creating service account private domain Done
Creating service account..... Done
Creating CLI service account..... Done
Authenticating the account CMC_CLI with the broker success

Do you want to start Veritas Cluster Management Console processes now? [y,n,q,b] (y)

```
Starting CMC: 0%
.
.
Starting CMC: 100%

Startup completed successfully on all systems

Storage Foundation and High Availability Solutions 5.0

A string of five or more characters is required to encrypt passwords
in the responsefile

Enter five or more characters to be used an encryption key: [b]
password

This key must be retained in a secure file and referenced using the
-enckeyfile option if the generated responsefile is to be used
again.

Press [Return] to continue:
Installation log files, summary file, and response file are saved
at:

/opt/VRTS/install/logs/installer-xxxxxx

Initial admin user info:
Management Server User Name: root
Management Server User Domain: venus.symantecexample.com
Management Server User Domain Type: unixpwd
```

Installing VCS 5.0

In this sample installation, the installation program installs VCS RPMs on two nodes, north and south.

- [Start the product installer or the installvcs program](#)
- [Installer performs initial system checks](#)
- [License VCS](#)
- [Installer checks for installed RPMs](#)
- [Choose to install all VCS RPMs or required RPMs](#)
- [Installer lists the RPMs](#)

Start the product installer or the installvcs program

To start the product installer

```
# ./installer
```

```
Install/Upgrade a product I
```

Veritas Cluster Server

```
Veritas Cluster Server 5.0 Installation Program  
Copyright (c) 2006 Symantec Corporation. All rights reserved.  
Symantec, the Symantec Logo 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 Licensed Software and Documentation are deemed to be  
"commercial computer software" and "commercial computer software  
documentation" as defined in FAR Sections 12.212 and DFARS  
Section 227.7202.
```

```
Logs for installvcs are being created in /var/tmp/installvcs-  
xxxxxx.
```

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

To start the installvcs program

```
# cd /cluster_server
```

```
# ./installvcs
```

```
Veritas Cluster Server 5.0 Installation Program
```

```
Copyright (c) 2006 Symantec Corporation. All rights reserved.  
Symantec, the Symantec Logo 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 Licensed Software and Documentation are deemed to be  
"commercial computer software" and "commercial computer software  
documentation" as defined in FAR Sections 12.212 and DFARS  
Section 227.7202.
```

```
Logs for installvcs are being created in /var/tmp/installvcs-  
xxxxxx.
```

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

Installer performs initial system checks

```
Initial system check:

Checking rsh communication with south.....Linux 2.6.9-34-default
Checking VCS installation on north ..... not installed
Checking VCS installation on south ..... not installed
```

License VCS

```
Checking system licensing

Installing licensing RPMs

VCS is not licensed on north

Enter a VCS license key for north: XXXX-XXXX-XXXX-XXXX-XXX
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on north
VCS license registered on north

Do you want to enter another license key for north? [y,n,q] (n)
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on south
VCS license registered on south

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

Installer checks for installed RPMs

```
Checking installed RPMs on north
Checking installed RPMs on south
```

Choose to install all VCS RPMs or required RPMs

```
VCS can be installed without optional RPMs to conserve
diskspace.

Additional rpms are typically installed to simplify future
upgrades.

1) Required Veritas Cluster Server rpms - 473 MB required
2) All Veritas Cluster Server rpms - 554 MB required
3) Storage Foundation Enterprise HA rpms - 1029 MB required

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

Installer lists the RPMs

```
The following VCS RPMs will be installed:
VRTSvlic Veritas Licensing
VRTSperl Veritas Perl 5.8.6 Redistribution
VRTSsicsco Symantec Common Infrastructure
VRTSspbx Symantec Private Branch Exchange
VRTSsmf Symantec Service Management Framework
VRTSat Symantec Product Authentication Service
VRTSspt Veritas Software Support Tools
SYMClma Symantec License Inventory Agent
VRTSllt Veritas Low Latency Transport
VRTSgab Veritas Group Membership and Atomic Broadcast
VRTSvxfen Veritas I/O Fencing
VRTSvcvcs Veritas Cluster Server
VRTSvcvcsmsg Veritas Cluster Server Message Catalogs
VRTSvcvcsag Veritas Cluster Server Bundled Agents
VRTSvcvcsmn Veritas Cluster Server Man Pages
VRTSvcvcsdc Veritas Cluster Server Documentation
VRTSjre Veritas Java Runtime Environment Redistribution
VRTSjre15 Veritas Java Runtime Environment Redistribution
VRTScutil Veritas Cluster Utilities
```

```
Press [Return] to continue:
...continued:
```

```
VRTSscsw Veritas Cluster Server Configuration Wizards
VRTSweb Veritas Java Web Server
VRTSscscm Veritas Cluster Server Cluster Manager
VRTSscsim Veritas Cluster Server Simulator V
RTSaclib Veritas ACC Library
VRTScmcs Veritas Cluster Management Console (Single Cluster
Mode)
VRTScmccc Veritas Cluster Management Console Cluster Connector
```

```
It is possible to install VCS rpms without performing
configuration.
```

Configuring VCS 5.0

You can proceed to configure VCS now or allow to install the VCS RPMs and then configure VCS at a later time.

It is optional to configure VCS now. If you choose to configure VCS later, you can either do so manually or run the `installvcs -configure` command.

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

To configure VCS, please answer the following sets of questions.

When a [b] is presented after a question, 'b' may be entered to go back to the first question of the configuration set.

When a [?] is presented after a question, '?' may be entered for help or additional information about the question.

Following each set of questions, the information you have entered will be presented for confirmation. To repeat the set of questions and correct any previous errors, enter 'n' at the confirmation prompt.

No configuration changes are made to the systems until all configuration questions are completed and VCS is installed successfully.

Perform the following tasks as the installer takes you through different screens:

- [Configure the cluster](#)
- [Configure the cluster in secure mode](#)
- [Add VCS users](#)
- [Configure Cluster Management Console](#)
- [Configure cluster connector](#)
- [Configure SMTP email notification](#)
- [Configure SNMP trap notification](#)
- [Configure the global cluster option](#)
- [Installer installs the VCS RPMs](#)
- [Installer creates VCS configuration files](#)
- [Installer verifies the NIC configuration](#)
- [Start VCS](#)
- [Complete the installation](#)

Configure the cluster

To configure VCS the following information is required:

A unique Cluster name
A unique Cluster ID number between 0-65535
Two or more NIC cards per system used for heartbeat links

One or more heartbeat links are configured as private links
One heartbeat link may be configured as a low priority link

All systems are being configured to create one cluster

Enter the unique cluster name: [?] **vcs_cluster2**
Enter the unique Cluster ID number between 0-65535: [b,?] **7**

Discovering NICs on north ...discovered eth0 eth1 eth2 eth3

Enter the NIC for the first private heartbeat NIC on north:
[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 north:
[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)

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

Checking Media Speed for eth2 on north 100Mb/s
Checking Media Speed for eth3 on north 100Mb/s
Checking Media Speed for eth2 on south 100Mb/s
Checking Media Speed for eth3 on south 100Mb/s
Cluster information verification:

Cluster Name: vcs_cluster2
Cluster ID Number: 7
Private Heartbeat NICs for north: link1=eth1 link2=eth2
Private Heartbeat NICs for south: link1=eth1 link2=eth2

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

Configure the cluster in secure mode

Veritas Cluster Server can be configured to utilize Symantec Security Services.

Running VCS in Secure Mode guarantees that all inter-system communication is encrypted and that users are verified with security credentials.

When running VCS in Secure Mode, NIS and system usernames and passwords are used to verify identity. VCS usernames and passwords are no longer utilized when a cluster is running in Secure Mode.

Before configuring a cluster to operate using Symantec Security Services, another system must already have Symantec Security Services installed and be operating as a Root Broker. Refer to the Veritas Cluster Server Installation Guide for more information on configuring a VxSS Root Broker.

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

If the VRTSat RPM is already installed, the installer provides you different modes to configure the cluster in secure mode.

Security can be configured completely automatically by the installer or it can also be configured semi automatically. In automatic mode, no user intervention is required. In the semi automatic mode, Authentication Broker related setup on Root Broker is expected to be performed by the Root Broker Administrator and CPI will ask for certain information that will be used to configure Authentication Brokers.

Security Menu

- 1) Configure security completely automatically
- 2) Provide AB credentials using BLOBs
- 3) Provide AB credentials without using BLOBs

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

Depending on the menu option you choose, the installer prompts you to select the configuration mode.

- [Configuring security automatically](#)
- [Configuring security semiautomatically using encrypted files](#)
- [Configuring security semiautomatically answering prompts](#)

Configuring security automatically

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

In order to enable Symantec Security Services on a VCS Cluster, Veritas Authentication Services (VRTSat rpm) must be installed on a system and operating as a Root Broker. Refer to the Veritas Cluster Server Installation Guide for more information on installing and configuring Veritas Authorization Services.

Enter the name of the VxSS Root Broker system: **east**

```
Checking ssh communication with venus ... Linux 2.6.9-34.default
Checking vxatd process ..... running
Checking vxatd version ..... 4.3.13.0
Checking security domain ..... root@venus.symantecexample.com
```

Systems will use root@east.symantecexample.com as its VxSS Domain

Configuring security semiautomatically using encrypted files

Make sure that you completed the pre-configuration tasks.

See [“Preparing to install VCS 5.0”](#) on page 29.

```
Select the Security option you would like to perform [1-3,q,?]
(1) 2
```

You need to create AB account for each cluster node on the Root Broker. Also you need to create a BLOB per cluster node. Please verify that the version of VRTSat installed on root broker supports BLOB creating. You need to use --in option to create BLOBs. Make sure that the input file is in the following format:

```
[setuptrust]
broker=<root_broker>
hash=<root_hash>
securitylevel=high

[configab]
identity=<ab_identity>
password=<ab_password>
root_domain=<root_FQDN>
root_broker=<root_broker>:<root_broker_port>
broker_admin_password=<root_broker_admin_password>
start_broker=false
enable_pbx=false
```

Refer to the VxSS Documentation for steps to create BLOBs. CPI needs a locally accessible path for BLOBs. You can either copy the BLOBs on north or mount the BLOBs using some removable media.

```
Do you want to continue? [y,n,q,b] (y)
```

```
Enter the path of BLOB for north: [b]/root/blob.north
```

```
Enter the path of BLOB for south: [b]/root/blob.south
```

Configuring security semiautomatically answering prompts

Make sure that you completed the pre-configuration tasks.

See [“Preparing to install VCS 5.0”](#) on page 29.

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

```
Veritas Cluster Server 5.0 Installation Program
```

```
You need to create authentication broker principal for each  
cluster node on the Root Broker. Refer to the VxSS Documentation  
for the configuration steps. Also make sure that the root_hash  
file is either copied to the installer node or it is locally  
accessible (via mounted file system or any other means). CPI  
will ask for the locally accessible path of root_hash file. You  
also need the AB principal passwords.
```

```
Press 'b' anytime (expect when prompted for passwords) to go to  
the previous menu.
```

```
Press [Return] to continue:
```

```
Enter root broker name: [b] east.symantecexample.com  
Enter root broker FQDN: [b] (symantecexample.com)  
Enter root broker domain: [b] (root@east.symantecexample.com)  
Enter root broker port: [b] (2821)  
Enter path to the locally accessible root hash [b]  
(/var/tmp/installvcs-1Lcljr/root_hash)  
Enter authentication broker principal name on north [b]  
(north.symantecexample.com)  
Enter authentication broker password on north:  
Enter authentication broker principal name on south [b]  
(south.symantecexample.com)  
Enter authentication broker password on south:
```

Proceed to configure Cluster Management Console.

See [“Configure Cluster Management Console”](#) on page 236.

Add VCS users

If you have enabled Symantec Product Authentication Service, you need not add VCS users now. Proceed to configure the Cluster Management Console.

See [“Configure cluster connector”](#) on page 235.

```
The following information is required to add VCS users:
```

```
A user name
A password for the user
User privileges (Administrator, Operator, or Guest)

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

Enter New Password:*****

Enter Again:*****

Do you want to add another user to the cluster? [y,n,q] (y) y

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

Enter Again:*****
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [?] a

Would you like to add another user? [y,n,q] (n) n

User: admin      Privilege: Administrators
User: smith     Privilege: Administrators
Passwords are not displayed

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

Configure cluster connector

Veritas Cluster Server 5.0 Configuration Program

Setup will now configure Veritas Cluster Management Console. If you have configured a management server, you can configure the cluster to be managed by the management server. You can also configure Veritas Cluster Management Console on cluster nodes to manage the local cluster.

```
Do you want this cluster to be managed by a management server?
Enter 'y' if you have set up a management server. [y,n,q] (y) y
```

To configure the CMC Cluster Connector, the following is required:

The Management Server address

```
Enter the network address used by the management server [?]
(north) mgmtserver1.symantecexample.com
```

Management Server Address: mgmtserver1.symantecexample.com

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

Veritas Cluster Server 5.0 Configuration Program

The following information is required to set up secure
communication with the CMC management server:

Password for the service account created during the management
server install (CMC_CC@CMC_SERVICES)

Hash of CMC management server's root broker. See the Veritas
Cluster Server Installation Guide for information on how to
retrieve the hash.

Enter the password for the CMC service account:

Enter the hash of the management server's root broker [?]
e96148d6ad9ea37b9b6a13f35512cb006a75be04

Root broker hash: e96148d6ad9ea37b9b6a13f35512cb006a75be04

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

Configure Cluster Management Console

```
The following information is required to configure the Cluster
Management Console:

A public NIC used by each system in the cluster
A Virtual IP address and netmask for the Cluster Management
Console

Do you want to configure the Cluster Management Console [y,n,q]
(y) y
Active NIC devices discovered on north: eth0
Enter the NIC for Cluster Manager (Web Console) to use on north:
[b,?] (eth0)

Is eth0 to be the public NIC used by all systems [y,n,q,b,?] (y)
y

Enter the Virtual IP address for Cluster Manager: [b,?]
10.10.12.1
Enter the netmask for IP 10.10.12.1: [b,?] (255.255.240.0)
Veritas Cluster Server 5.0 Configuration Program

Cluster Management Console verification:

NIC: eth0
IP: 10.10.12.1
Netmask: 255.255.240.0
```

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

Configure SMTP email notification

Veritas Cluster Server 5.0 Configuration Program

The following information is required to configure SMTP notification:

The domain-based hostname of the SMTP server
The email address of each SMTP recipient
A minimum severity level of messages to send to each recipient

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

Enter the domain-based hostname of the SMTP server
(example: smtp.yourcompany.com): [b,?] **smtp.symantecexample.com**
Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,?] **ozzie@symantecexample.com**
Enter the minimum severity of events for which mail should be
sent to ozzie@symantecexample.com [I=Information, W=Warning,
E=Error, S=SevereError]: [b,?] **w**

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@symantecexample.com**

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

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

Veritas Cluster Server 5.0 Configuration Program

SMTP email notification verification

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

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

Configure SNMP trap notification

```
Veritas Cluster Server 5.0 Configuration Program

System names of SNMP consoles to receive VCS trap messages
SNMP trap daemon port numbers for each console
A minimum severity level of messages to send to each console

Enter the SNMP trap daemon port: [b,?] (162)
Enter the SNMP console system name: [b,?] saturn
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

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

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

Veritas Cluster Server 5.0 Configuration Program

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)
```

Configure the global cluster option

```
Veritas Cluster Server 5.0 Configuration Program

The following is required to configure the Global Cluster
Option:

A public NIC used by each system in the cluster
A Virtual IP address and netmask

The Virtual IP address and NIC may be the same as those
configured for Cluster Management Console

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

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

```
Active NIC devices discovered on north: eth0
Enter the NIC for Global Cluster Option to use on north: [b,?]
(eth0)
Is eth0 to be the public NIC used by all systems [y,n,q,b,?] (y)
y
Enter the Virtual IP address for the Global Cluster Option:
[b,?] 10.10.12.1
Enter the netmask for IP 10.10.12.1: [b,?] (255.255.240.0)
```

Veritas Cluster Server 5.0 Configuration Program

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) **y**

Installer installs the VCS RPMs

Veritas Cluster Server 5.0 Configuration Program

Installing VCS: 0%

Veritas Cluster Server 5.0 Configuration Program

Installing VCS: 100%

Installation completed successfully on all systems

Installer creates VCS configuration files

```
Creating north security principal on venus ..... Done
Starting Veritas Security Services on north ..... Done
Creating south security principal on venus ..... Done
Starting Veritas Security Services on south ..... Done
```

```
Authenticating and deploying CMC service account from north:
Setting up trust with the broker venus.symantecexample.com
..... success
Authenticating the account CMC_CC with the broker... success
```

Deploying authentication package to south:

```
Creating package for south ..... Done
Copying package to south ..... Done
Extracting package on south ..... Done
```

Creating Cluster Server configuration files Done

```
Copying configuration files to north..... Done
Copying configuration files to south..... Done

Cluster Server configured successfully.
```

Installer verifies the NIC configuration

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

```
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 [“Configuring Network Interfaces” on page 30].

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

Start VCS

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

Veritas Cluster Server 5.0 Installation Program

Starting VCS: 0%

Veritas Cluster Server 5.0 Installation Program

Starting VCS: 100%

Startup completed successfully on all systems
```


Press [Return] to continue:

Complete the installation

Installation log files, summary file, and response file are saved at:

/opt/VRTS/install/logs/installvcs-xxxxxx

Uninstalling VCS 5.0

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

Veritas Cluster Server 5.0 Installation Program

Logs for uninstallvcs are being created in /var/tmp/
uninstallvcs-6flUca.

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

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

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

Veritas Cluster Server 5.0 Installation Program

Initial system check:

Checking rsh communication with south Linux 2.6.5-7.139-
default

Checking system licensing

Veritas Cluster Server 5.0 Installation Program

Checking installed RPMs on north

Checking installed RPMs on south

uninstallvcs is now ready to uninstall VCS.

All VCS processes that are currently running must be stopped.

Are you sure you want to uninstall VCS? [y,n,q] (y) **y**

Veritas Cluster Server 5.0 Installation Program

Uninstalling VCS: 100%

Uninstall completed successfully on all systems

Uninstallation log files, summary file, and response file are saved at:

/opt/VRTS/install/logs/uninstallvcs-7DAaUt

Deploying VCS in a Storage Foundation Volume Server environment

This appendix contains the following topics:

- [About VCS in a Storage Foundation VS environment](#)
- [Deploying VCS in a Storage Foundation VS environment](#)

Storage Foundation Volume Server (SF Volume Server) is a separately licensed feature of Veritas Storage Foundation™ by Symantec. An SF Volume Server license is currently available only through the Symantec customer access program. For information about participating in the access program and obtaining an SF Volume Server license, visit the following Symantec website:

<http://cap.symantec.com>

About VCS in a Storage Foundation VS environment

In the Storage Foundation Volume Server (Storage Foundation VS) environment, both Volume Clients (VCs) and Volume Servers (VSs) are candidates for high availability. Storage Foundation VS has two main modes that you can use:

- Share Out
- Volume Serving

Depending on whether you use Share Out or Volume Serving, different combinations of Volume Server and Volume Client clusters are available. VCS agents help ensure high availability for different mission-critical components of your SF Volume Server setup.

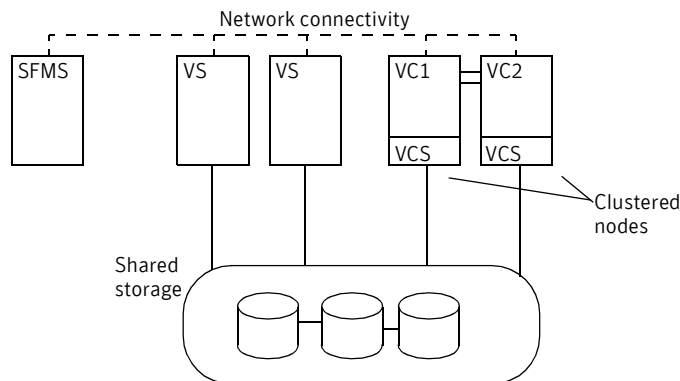
VCS agents help manage different resources. For Storage Foundation VS, you can use the SANVolume and DiskGroup agents.

- The SANVolume agent exists to provide management and monitoring functions for a SAN volume.
- The Disk Group agent allows you to manage and monitor your SAN disk groups on the Volume Server in the Share Out mode.

Volume Serving mode using VCS

In the Volume Serving mode, you install VCS on Volume Clients. You configure the application service group on the Volume Client systems. The SAN volume is the lowest resource in the application dependency graph for shared storage.

Figure C-11 Volume serving



Share Out mode using VCS

In Share Out mode, you can make the Volume Server or the Volume Client systems highly available using VCS. Multiple clustering combinations exist in this configuration.

The SAN disk group is the lowest resource in the application dependency graph for shared storage on the Volume Server systems. Similarly the SAN volume is the lowest resource in the application dependency graph for shared storage on the Volume Client systems.

For more information see:


[“Configuring application service groups”](#) on page 253.

Note: Symantec recommends that you configure the critical application service group on the Volume Server systems and the off-host processing applications in their service groups on the Volume Client systems.

Two VCS configuration options exist in the Share Out mode:

- One cluster with the Volume Client systems and the Volume Server systems.

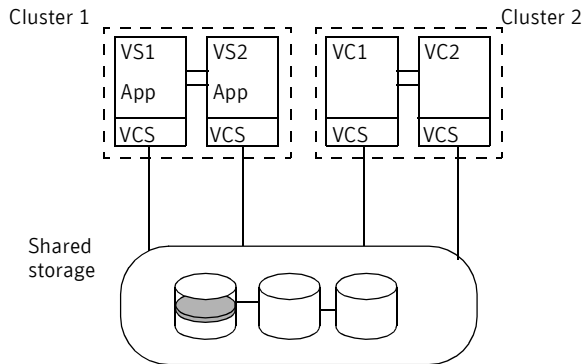
Figure C-12 A Volume Client and Volume Server cluster, note that this cluster does not support fencing



In the Share Out mode, when Volume Server and Volume Client systems are in the same cluster, maintain the start order of the SAN disk group resource and the SAN volume resources. For this, Symantec recommends the use of the service group dependencies. Make the service group containing the SAN disk group the child service group and the service group containing the SAN volume the parent service group. It is recommended that the link between the two service groups be an Online Global Firm type. For more information on service group dependencies see the section on service group dependencies in the *Veritas Cluster Server User's Guide*.

- Two clusters: The Volume Server systems and the Volume Client systems are in separate clusters.
Use the RemoteGroup agent to create a dependency between the Volume Client group and the Volume Server group.
See “[Resource dependencies for volume serving mode](#)” on page 255.

Figure C-13 Two clusters accessing shared storage



About failover in a Storage Foundation VS environment

You can use VCS to make Volume Server and Volume Client systems highly-available. The Volume Servers can be made highly-available only in the Share-Out mode and the Volume Clients can be made highly-available in both the Volume Serving and Share Out modes.

SAN Volumes on Volume Client systems are made highly available using the SANVolume agent. This agent achieves high-availability of the SAN Volume by attaching, monitoring, and unattaching a SAN volume. You can configure this agent as part of a VCS service group. The SAN volumes must reside on storage arrays that support SCSI-3 persistent reservations.

SAN disk groups on Volume Server systems are made highly-available using the DiskGroup agent. This agent achieves high-availability of the SAN disk groups by importing, monitoring, and deporting the SAN disk groups as required. You need to set the DiskGroupType attribute of this agent to SAN for a Storage Foundation VS environment.

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

About I/O fencing in a SFVS environment

In a Storage Foundation VS setup, use I/O fencing to protect data on shared storage on Volume Server and Volume Client systems.

I/O fencing on Volume Server systems is supported in the Share Out mode, while Volume Client systems can implement I/O fencing in both the Volume Serving and Share Out modes.

I/O fencing can be implemented on Volume Server systems in the Share Out mode using fencing with disk groups.

Implement I/O fencing on Volume Client systems in a different way compared to the Volume Server systems. In SCSI-3 fencing, nodes race to get control of coordinator disks. In comparison, the Volume Client systems race to get control of a coordinator volume, which you create on a dedicated disk group on the SAN. You create different disk groups for data.

Note: You cannot use coordinator volume and coordinator disks together for I/O fencing on the same system or in the same cluster.

See [“Setting up I/O fencing”](#) on page 111.

Deploying VCS in a Storage Foundation VS environment

- [“Prerequisites”](#) on page 247
- [“Installing VCS”](#) on page 248
- [“Configuring I/O fencing in Volume Serving mode”](#) on page 248
- [“Configuring I/O fencing in the Share Out mode”](#) on page 251
- [“Configuring application service groups”](#) on page 253

Prerequisites

- Connectivity to both storage and to the network
 - Connect the Volume Server and the Volume Client to the same storage
 - The Volume Server, Volume Client, and Storage Foundation Management Server should be able to communicate over the network.
- SCSI-3-capable storage

Installing VCS

Depending on your requirements, you can install VCS on either Volume Client systems or Volume Server systems. In the Volume Serving mode, you can only install VCS on Volume Client systems. In the Share Out mode, you can install VCS on Volume Server and Volume Client systems.

Configuring I/O fencing in Volume Serving mode

In Volume Serving mode, you must configure a coordinator volume for I/O fencing on the Volume Clients.

See “[About I/O fencing in a SFVS environment](#)” on page 247.

Setting up the coordinator volume

This procedure describes how to set up a coordinator volume.

To set up the coordinator volume

- 1 Create the SANVM domain.
You can use the same domain as the one that the data disk groups use. You do not have to create a special domain for the coordinator volume.
- 2 Create the Volume Server set in the domain.
You can use the same volume server set as the one used for data disk groups. You do not have to create a special volume server set for the coordinator volume.
- 3 Put the Volume Server in a Volume Server set.
- 4 Create a SAN disk group in the Volume Server set.

```
vxadm -m sfms.example.com diskgroup create sample_cvool_dg  
Disk1_19 vsset=vsset1 dgtype=san
```

Where `sfms.example.com` is the SFMS hostname; `sample_cvool_dg` is the disk group name; `vsset1` is the volume server set’s name; and the disk group is of a SAN type.

- 5 Create a SAN volume in the SAN disk group. Create the coordinator volume in the coordinator disk group. Use the following command:

```
vxadm -m sfms.example.com -g sample_cvool_dg volume create  
sample_cvolsample_cvool 10m
```

Where `sfms.example.com` is the SFMS hostname; `sample_cvool_dg` is the disk group name; `sample_cvool` is the SAN volume, and `10m` is the size of the volume.

- 6 Export the SAN volumes to all volume clients in the cluster with read-write permissions.

```
vxadm -m sfms.example.com -g sample_cvol_dg volume export  
sample_cvol sysA.veritas.com exportperms=RDWR
```

Where `sfms.example.com` is the SFMS hostname; `sample_cvol_dg` is the disk group name; `sample_cvol` is the SAN volume; and `sysA` is the Volume Client system.

Once you have set up the coordinator volume, you need to initialize it using the `/opt/VRTSvcs/vxfen/bin/vxfencvol` utility. Attach the SAN volume, and then initialize it with the following commands:

To attach the SAN volume and initialize the volume

- 1 Attach the coordinator volume on a Volume Client with read-write permissions.

```
vxadm -m sfms.example.com -g sample_cvol_dg volume attach  
sample_cvol sysA.example.com access_policy=RDWR  
io_access=ANY
```

- 2 Initialize the coordinator volume.

```
/opt/VRTSvcs/vxfen/bin/vxfencvol -init -cvol /dev/vx/rdisk/  
testdom1/sample_cvol_dg/sample_cvol
```

- 3 After initialization, unattach the coordinator volume.

```
vxadm -m sfms.example.com -g sample_cvol_dg volume unattach  
sample_cvol sysA.example.com.
```

Editing the vxfenmode file

Before you start to modify the `vxfenmode` file, you must stop fencing. You then need to modify the `vxfenmode` file. The file is in `/etc/vxfenmode`.

To edit the vxfenmode file

The following code is a sample configuration for SANVM in the `vxfenmode` file:

```
vxfen_mode=customized  
vxfen_mechanism=sanvm  
disk_group=sample_cvol_dg  
cvol_name=sample_cvol  
vs_host=vsA.example.com  
sanvm_domain_name=domain1
```

Where the:

- `disk_group` is the name of the coordinator SAN disk group.
- `cvol_name` is the name of the coordinator SAN volume.

- `vs_host` is the name of the volume server. Symantec recommends that this be the name associated with a virtual IP address residing on the volume server.
- `sanvm_domain_name` is the name of the SANVM domain that the coordinator disk group is a part of. This is optional in the Share Out mode.

Editing the main.cf file

Edit the `main.cf` file to configure VCS for fencing.

To edit the VCS configuration to add the UseFence attribute

- 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. Modify the list of cluster attributes by adding the attribute, `UseFence`, and assign it a value of `SCSI3`. For example, with the attribute added this portion of the file resembles:

```
cluster vcs_cluster2 (
    UserNames = { admin = "cDRpdxPmHpzS." }
    Administrators = { admin }
    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 .
```
- 7 Using `rcp`, or some other available utility, copy the VCS configuration file to the other nodes. For example, on each node:

```
# rcp north:/etc/VRTSvcs/conf/config/main.cf
      /etc/VRTSvcs/conf/config
```
- 8 With the configuration file in place on each system, shut down and then restart each system.

```
# reboot
```

Note: To ensure that I/O fencing is shut down properly, use the `shutdown` command instead of the `reboot` command.

Configuring I/O fencing in the Share Out mode

In Share Out mode, Volume Client systems use I/O fencing with coordinator volumes as described above. Volume Server systems, however, use I/O fencing with coordinator disks.

Note: You cannot use coordinator volume and coordinator disks together for I/O fencing on the same system or in the same cluster.

See “[Setting up I/O fencing](#)” on page 111.

See “[About I/O fencing in a SFVS environment](#)” on page 247.

Creating coordinator volumes

The following steps are to be executed on a volume server system:

- 1 Create a SAN disk group.

```
vxdbg -o dgtype=san init sample_cvol_dg newdg01=Disk_1
```

Where `sample_cvol_dg` is the disk group name; and the disk group is of a SAN type.

- 2 Create the coordinator volume in the coordinator disk group. Use the following command:

```
vxassist -g sample_cvol_dg make sample_cvol 10m
```

Where `sample_cvol_dg` is the disk group name; `sample_cvol` is the SAN volume, and 10m is the size of the volume.

- 3 Export the SAN volumes to all volume clients in the cluster with read-write permissions.

```
vxadm -g sample_cvol_dg volume export sample_cvol  
vcssysA.veritas.com exportperms=RDWR
```

Where `vcssysA.veritas.com` is a Volume Client system that the SAN volume `sample_cvol` is being exported to.

Once you have set up the coordinator volume, you need to initialize it using the `/opt/VRTSvcs/vxfen/bin/vxfencvol` utility. Attach the SAN volume, and then initialize it with the following commands:

To attach the SAN volume and initialize the volume

- 1 Attach the coordinator volume on a Volume Client with read-write permissions.

```
vxadm -g sample_cvol_dg volume attach sample_cvol  
sysA.example.com access_policy=RDWR io_access=ANY
```

- 2 Initialize the coordinator volume.

```
/opt/VRTSvcs/vxfen/bin/vxfencvol -init -cvol /dev/vx/rdisk/
vol_server/sample_cvol_dg/sample_cvol
```

- 3 After initialization, unattach the coordinator volume.

```
vxadm -g sample_cvol_dg volume unattach sample_cvol
sysA.example.com.
```

Editing the vxfenmode file

Before you start to modify the vxfenmode file, you must stop fencing. You then need to modify the vxfenmode file. The file is in `/etc/vxfenmode`.

To edit the vxfenmode file

The following code is a sample configuration for SANVM in the vxfenmode file:

```
vxfen_mode=customized
vxfen_mechanism=sanvm
disk_group=example_cvol_dg
cvol_name=example_cvol
vs_host=vol_server.example.com
```

Where the:

- `disk_group` is the name of the SAN disk group that you want to use when configuring I/O fencing for a SANVM environment.
- `cvol_name` is the name of the coordinator SAN volume.
- `vs_host` is the name of the volume server.

To edit the VCS configuration to add the UseFence attribute

- 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. Modify the list of cluster attributes by adding the attribute, `UseFence`, and assign it a value of `SCSI3`. For example, with the attribute added this portion of the file resembles:

```
cluster vcs_cluster2 (
    UserNames = { admin = "cDRpdxPmHpzS." }
    Administrators = { admin }
    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 .
```
- 7 Using `rcp`, or some other available utility, copy the VCS configuration file to the other nodes. For example, on each node:

```
# rcp north:/etc/VRTSvcs/conf/config/main.cf
      /etc/VRTSvcs/conf/config
```
- 8 With the configuration file in place on each system, shut down and then restart each system.

```
# reboot
```

Note: To ensure that I/O fencing is shut down properly, use the `shutdown` command instead of the `reboot` command.

Configuring application service groups

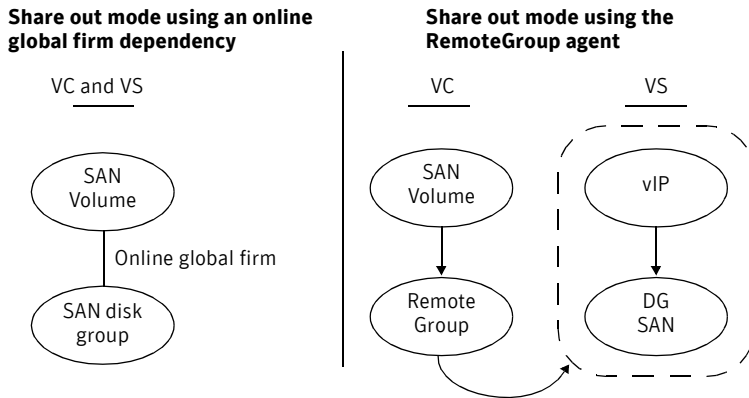
In Share Out mode, if the Volume Server and Volume Client nodes are configured in different clusters, the SAN volumes on Volume Clients depend on the SAN disk groups on the Volume Server. In order to maintain the start order to satisfy this dependency, you should use the `RemoteGroup` agent when you have two clusters.

If the Volume Server and Volume Client nodes are configured in one cluster, use service group dependencies when you have the Volume Server and the Volume Client systems in one cluster.

If the Volume Server and Volume Client nodes are configured in the same cluster, make the service group containing the SAN disk group the child service group and the service group containing the SAN volume the parent service group. The link between the two service groups should be an `Online Global Firm` type.

See the *Veritas Cluster Server User's Guide* and the *Veritas Cluster Server Bundled Agents Reference Guide* on information about configuring and using the `RemoteGroup` agent and about service group dependencies.

Figure C-14 Resource dependencies for share out mode



SANVolume resource in Share Out mode in the main.cf file

```
SANVolume svol (
  SANDiskGroup = sample_cvool_dg
  SANVolume = sample_cvool
  VolumeServer = "vol_server.example.com"
)
```


DiskGroup Resource in Share Out mode in the main.cf file

```
DiskGroup only_sandg_vNFS_DiskGroup (
  DiskGroup = sample_cvool_dg
  MonitorReservation = 1
  DiskGroupType = SAN
)
```

In Volume Serving mode, the SAN volume is the lowest resource in the application dependency graph for shared storage.

Figure C-15 Resource dependencies for volume serving mode

Volume serving mode



SANVolume resource in Volume Serving mode in the main.cf file

```
SANVolume vNFS_SANVolume (  
  Domain = testdom1  
  SANDiskGroup = sample_cvool_dg  
  SANVolume = sample_cvool  
  VolumeServer = "vol_server.example.com"  
)
```


Configuring the Symantec License Inventory Agent

This appendix includes the following topics:

- [About the Symantec License Inventory Manager](#)
- [When the Symantec License Inventory Agent is installed](#)
- [When the server and access points are installed](#)
- [What you can do with the agent after it is installed](#)
- [How to remove the agent](#)
- [How to order the Symantec License Inventory Manager license and media kit](#)

The Symantec License Inventory Manager installation disc is available separately. For information on how to order the full product, see “[How to order the Symantec License Inventory Manager license and media kit](#)” on page 261. The installation media provides online documentation with details on all topics discussed in this appendix.

Read the following Technical Support TechNote for the latest information on updates, patches, and software issues regarding this product:

<http://support.veritas.com/docs/282183>

You can also download the *Symantec License Inventory Agent 4.1 Release Notes*, from this website.

About the Symantec License Inventory Manager

The Symantec License Inventory Manager (license inventory manager) is an enterprise asset management tracking tool that inventories Symantec Information Availability products in your network and consolidates critical information on the deployment of these products to facilitate license management and compliance tracking. Using the information provided by the license inventory manager, you can:

- Determine all the Symantec software products and licenses being used in your enterprise
- Achieve easier license self-compliance management
- Know your Enterprise License Agreement deployment status
- Reduce administrative overhead for managing license compliance
- Renew support and maintenance based on the licenses you have deployed
- Gain more control over your Symantec software usage
- Manage department chargebacks based on actual software usage
- Use more flexible licensing and pricing models
- Exploit detailed deployment data to perform return on investment analyses for purchased software

The license inventory manager is a three-tiered system that consists of a server tier, access point tier, and an agent tier. The server tier is the Symantec License Inventory Server, which consolidates and stores information that it gathers from the agents and access points.

The optional access point tier includes Symantec License Inventory Access Points and serves as a consolidation layer between the agents and server.

The agent tier includes Symantec License Inventory Agents, which are deployed on individual hosts in a network. Each agent gathers product information on the supported Symantec products that are installed on the agent's host, then sends the information to an access point or the server.

When the Symantec License Inventory Agent is installed

The Symantec product installer installs or upgrades the agent on the host with the Symantec product. The agent is installed in the following directory:

`/opt/SYMC1ma`

The agent is installed with a default configuration that minimizes its impact on a running system. The minimum configuration prevents remote communication with the agent to keep its data and interfaces secure.

When the server and access points are installed

The server and access points are not installed automatically. If you want to use the Symantec License Inventory Manager, you must manually install the server and, optionally, the access points. After you install the server and access points, the agents can gather information and you can create inventory reports.

You can install the server and access points from the Symantec License Inventory Manager installation disc.

What you can do with the agent after it is installed

If you are already participating in a Symantec sales program that requires the use of the agent, or if you want to order and deploy the Symantec License Inventory Manager, you can use the agent to track Symantec products on the systems on which it was installed. To use the agent, however, you must manually configure it to enable remote communication between the agent and its server or access point.

Complete instructions for reconfiguring the agent are provided in the *Symantec License Inventory Manager 4.1 Release Notes*. You can download this document from the following website:

<http://support.veritas.com/docs/282183>

How to remove the agent

If you do not want to use the Symantec License Inventory Manager, you can remove the agent using the operating system package removal commands to remove the agent packages, which include SYMClma and VRTSsmf.

The server and access point also use the VRTSsmf package. If the server or access point is installed on this host with the agent, you can remove the SYMClma package, but not the VRTSsmf package. If neither the server nor the access point is installed on this host, you can remove both the SYMClma and VRTSsmf packages.

If you remove both packages, remove the SYMClma package first.

[Table D-1](#) lists the commands required to remove these packages on the supported platforms.

Table D-1 Package removal commands required to remove the agent

Platform	Package removal command
AIX	<code>installp -u VRTSlma</code> <code>installp -u VRTSsmf</code>
HP-UX	<code>swremove SYMClma</code> <code>swremove VRTSsmf</code>
Linux	<code>rpm evv SYMClma</code> <code>rpm evv VRTSsmf</code>
Solaris	<code>pkgrm VRTSlma</code> <code>pkgrm VRTSsmf</code>

Later, you can reinstall the agent with the Symantec License Inventory Manager installation disc. This disc is available in the Symantec License Inventory Manager kit.

How to order the Symantec License Inventory Manager license and media kit

To order a Symantec License Inventory Manager license and media kit, contact your Symantec sales representative.

The installation media provides online documentation for the Symantec License Inventory Manager. You can contact your sales representative to order printed copies of the documentation. The documents you can order include:

- *Symantec License Inventory Manager Installation and Configuration Guide*
- *Symantec License Inventory Manager Administrator's Guide*
- *Symantec License Inventory Manager User's Guide*

Index

A

- accessing
 - documentation 44, 173
 - installation guide 44, 173
 - release notes 44, 173
- adding
 - users 72, 234
- adding node
 - to a cluster 179
 - to a one-node cluster 196
- attributes
 - UseFence 122

B

- bundled agents
 - types.cf file 194

C

- cables
 - cross-over Ethernet 180
- cables, crossover Ethernet 46
- centralized cluster management 72, 235
- checking systems 228
- cluster
 - creating a single-node cluster, installer 192
 - creating a single-node cluster, manual 193
 - four-node configuration 16
 - removing a node from 185
 - verifying 92
 - verifying operation 152
- cluster connector
 - uninstalling, UNIX 108
 - uninstalling, Windows 109
- cluster management 73, 236
- Cluster Management Console 25
 - accessing Web Console 155
 - Management Server 221
- Cluster Management Console, documents 44, 173
- Cluster Manager
 - installing Java Console 80

- cold start, running VCS 18
- CommandCentral Availability
 - upgrading 170
- command-line options 102
- commands
 - gabconfig 152, 202
 - hastart 184
 - hastatus 152
 - hasys 153
 - licensevcs, vxlicinst 167
 - lltconfig 144
 - lltstat 150
 - vxdisksetup (initializing disks) 118
 - vxfen start 122
 - vxfenadm 134
 - vxfclearpre 139
 - vxlicinst 101, 182
 - vxlicrep 100, 182
- communication channels 18
- communication disk 18
- configuration files
 - main.cf 145
 - types.cf 145, 195
- configuring
 - GAB 202
 - hardware 31
 - LLT, manual 200
 - private network 46
 - ssh 53
 - switches 47
- configuring security
 - automatically 232
 - semiautomatically 233
 - semiautomatically, answering prompts 234
- configuring VCS 69, 230
 - adding users 72, 234
 - basic cluster 231
 - Cluster Connector 72, 235
 - Cluster Management Console 72, 73, 235, 236
 - event notification 74, 76, 237, 238
 - global clusters 77, 238
 - overview 64

- secure mode 70, 231
- coordinator disks
 - for I/O fencing 113
 - setting up 119
- crossover cables 46

D

- data disks
 - for I/O fencing 113
- deploying VCS
 - Storage Foundation VS 247
- directives, LLT 201
- disk space
 - directories 31
- disk space, required 31
- disks
 - adding and initializing 118
 - coordinator 119
 - testing with vxfcntlshdw 115
 - verifying node access 115
- documentation
 - accessing 156

E

- eeprom, parameters 47
- encrypted files, security 233
- Ethernet controllers 180

F

- fibre channel 31

G

- GAB
 - description 17
 - manual configuration 202
 - port membership information 152
 - starting 167
 - verifying 152
- gabconfig command 152, 202
 - a (verifying GAB) 152
 - in gabtab file 144
- gabtab file
 - creating 202
 - verifying after installation 144
- Global Cluster option 28
- global clusters 28
- global clusters, configuration 77, 238

H

- hardware
 - configuration 16
 - configuring network and storage 31
- hastart 184
- hastatus -summary command 152
- hasys -display command 153
- hubs 46
- hubs, independent 180

I

- I/O fencing
 - checking disks 114
 - event scenarios 130
 - operations 114
 - setting up 117
 - shared storage 114
 - starting 122
 - testing and scenarios 130
- installation
 - required disk space 32
- Installing 38
- installing
 - Management Server 221
 - Windows 173
 - management server 38
 - management server, Solaris 39
 - management server, Windows 42
 - required disk space 31
 - Root Broker 35, 219
 - using installvcs program 62
- installing and configuring VCS
 - overview 64
- installing VCS
 - checking systems 65
 - choosing RPMs 68, 228
 - licensing 67, 228
 - overview 64
 - required information 56
 - RPMs list 229
 - starting 66
 - starting installer 227
 - system check 228
 - utilities 61
- installing VCS, example 63
- installvcs 62
 - options 62
- installvcs prompts

b 63
n 63
y 63

J

Java Console
installing 80
installing on UNIX 80
installing on Windows workstation 82
upgrading on UNIX 169
upgrading on Windows workstation 169
upgrading VCS 169

L

license keys
adding with vxlicinst 101, 167, 182
obtaining 54
replacing demo key 101
licenses, information about 100
licenses, showing information 182
licensevcs 167
licensing commands
vxlicinst 54
vxlicrep 54
vxlictest 54
licensing VCS 67, 228
links, private network 46, 144
LLT
description 17
directives 201
interconnects 52
manual configuration 200
starting 167
verifying 150
LLT directives
link 201
link-lowpri 201
set-cluster 201
set-node 201
lltconfig command 144
llthosts file, verifying after installation 143
lltstat command 150
llttab file, verifying after installation 144

M

MAC addresses 47
main.cf file 145

contents after installation 146
example 145
Management Console
installing 38
Management Server 221
installing
Windows 173
management server
trust between 217
uninstalling 106
uninstalling, UNIX 106
upgrading on Solaris 170
managing cluster, locally 236
managing clusters, centrally 72, 235
MANPATH variable, setting 46
manual installation
preparing 181
media speed 52
optimizing 52
membership information 152
Microsoft Windows
installing Management Server on 173
minimal downtime upgrade 212
example 213
mounting, software disc 55

N

network partition
preexisting 18
protecting against 16
Network partitions
protecting against 18
network switches 47
NFS 15

O

operating system
supported 32
optimizing
media speed 52
overview, VCS 15

P

parameters, eeprom 47
PATH variable
setting 45, 193
VCS commands 150

- peers 217
- persistent reservations, SCSI-3 49
- port a
 - membership 152
- port h
 - membership 152
- port membership information 152
- preparing
 - manual installation 181
- private network, configuring 46

R

- RAM, installation requirement 31
- registrations
 - key formatting 135
- removing a system from a cluster 185
- requirements
 - Ethernet controllers 31
 - fibre channel 31
 - hardware 31
 - RAM Ethernet controllers 31
 - SCSI host bus adapter 31
- reservations
 - description 112
- Root Broker 22
 - installing 35, 219
- rpm -e command 188
- rsh 52, 67, 92

S

- sample output 219
- SCSI host bus adapter 31
- SCSI-3
 - persistent reservations 49
- SCSI-3 persistent reservations
 - verifying 117
- seeding 18
 - automatic 18
 - manual 18
- setting
 - MANPATH variable 46
 - PATH variable 45, 193
- setting up, shared storage 49
- Share Out mode
 - Storage Foundation VS 245
- shared storage
 - setting up 49
- single-node cluster

- adding a node to 196
- single-system cluster
 - creating 192, 193
 - modifying startup files 194
- SMTP email notification 74, 237
- SMTP notifications 27
- SNMP notifications 27
- SNMP trap notification 76, 238
- ssh 52, 67, 92
 - configuring 53
- starting installation
 - installvcs program 66, 227
 - Veritas product installer 66, 227
- starting VCS 79, 240
- starting VCS after manual upgrade 167
- starting VCS after rpm -i 168
- storage
 - fully shared vs. distributed 16
 - shared 16
- Storage Foundation Volume Server 244
 - modes 244
- Storage Foundation VS 244
 - Share Out 245
 - Volume Serving 244
- switches 47
- Symantec Product Authentication Service 22, 35, 70, 219, 231
- system communication using rsh, ssh 52
- system state attribute value 153

T

- trust relationship 217
- types.cf 194
 - bundled agents 194
- types.cf file 195
 - included in main.cf 145

U

- uninstalling
 - cluster connector, UNIX 108
 - cluster connector, Windows 109
 - management server 106
 - management server, UNIX 106
- uninstalling, VCS 104, 241
- uninstallvcs 104, 241
- upgrade
 - minimal downtime 212
- upgrading

Command Central Availability 170
 management server, Solaris 170
 minimal downtime 212

V

variables

MANPATH 46
 PATH 45, 193

VCS

basics 15
 command directory path variable 150
 configuration files
 main.cf 145
 types.cf 145
 coordinator disks 119
 documentation 156
 example installation 63
 global clusters 28
 installation example 63
 installing 63
 installing using program 62
 replicated states on each system 16
 starting 167, 168
 supported Linux OS 32

VCS I/O fencing

shared storage 49

verifying

cluster 92
 NIC configuration 78, 240

Veritas 2

volume server 244

Volume Serving mode

Storage Foundation VS 244

vxdisksetup command 118

VXFEN

tunable parameters 128

vxfen command 122

vxfenadm command 134

vxfenclearpre command 139

VxFS, supported version 33

vxlicinst 54

vxlicinst command 101, 167, 182

vxlicrep 54

vxlicrep command 100, 182

vxlictest 54

VxVM, supported version 33

W

Web Console

accessing after installation 155

Windows

installing Management Server on 173

