

Veritas™ Cluster Server Installation Guide

Linux

5.1

Veritas Cluster Server Installation Guide

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

Product version: 5.1

Document version: 5.1.1

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's maintenance offerings include the following:

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

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

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

Contacting Technical Support

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

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

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

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

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

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

Licensing and registration

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

customercare.symantec.com

Customer service

Customer service information is available at the following URL:

www.symantec.com/customercare

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

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

Documentation feedback

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

sfha_docs@symantec.com

Maintenance agreement resources

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

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

Additional enterprise services

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

Enterprise services that are available include the following:

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

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

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4
Section 1 Installation overview and planning	19
Chapter 1 Introducing VCS	21
About Veritas Cluster Server	21
About VCS basics	21
About multiple nodes	22
About shared storage	23
About LLT and GAB	23
About network channels for heartbeating	24
About preexisting network partitions	24
About VCS seeding	24
About Veritas product licensing	25
About VCS features	26
About Veritas Operations Services	27
About VCS notifications	27
About global clusters	28
About I/O fencing	28
About VCS optional components	29
About Symantec Product Authentication Service (AT)	29
About Cluster Manager (Java Console)	30
About Veritas Cluster Server Management Console	30
About VCS Simulator	30
Chapter 2 Planning to install VCS	33
VCS installation requirements	33
Hardware requirements	33
Supported Linux operating systems	35
Required Linux RPMs for VCS	36
Required Solaris patches for VCS	37
Supported software	37
I/O fencing requirements	38
VCS installation methods	40

	About the VCS installation program	41
	About the Web-based installer	43
	About response files	44
	Typical VCS cluster setup models	46
Section 2	Preinstallation tasks	49
Chapter 3	Preparing to install VCS	51
	About preparing to install VCS	51
	Performing preinstallation tasks	51
	Obtaining VCS license keys	52
	Setting up the private network	53
	Configuring persistent network interfaces	55
	Setting up inter-system communication	55
	Setting up shared storage	57
	Setting the PATH variable	60
	Setting the MANPATH variable	60
	Setting the kernel.panic tunable	61
	Optimizing LLT media speed settings on private NICs	61
	Guidelines for setting the media speed of the LLT interconnects	61
	Mounting the product disc	62
	Performing automated preinstallation check	62
	Reformatting VCS configuration files on a stopped cluster	63
	Getting your VCS installation and configuration information ready	64
Section 3	Installation using the script-based installer	69
Chapter 4	Installing VCS	71
	Installing VCS using the installer	71
Chapter 5	Preparing to configure VCS	77
	Preparing to configure the clusters in secure mode	77
	Installing the root broker for the security infrastructure	81
	Creating authentication broker accounts on root broker system	82
	Creating encrypted files for the security infrastructure	83

	Preparing the installation system for the security infrastructure	85
	About configuring VCS clusters for data integrity	86
	About I/O fencing components	87
	About I/O fencing configuration files	88
	About planning to configure I/O fencing	91
	Setting up the CP server	96
	Installing the CP server using the installer	97
	Configuring security on the CP server	98
	Setting up shared storage for the CP server database	99
	Configuring the CP server using the configuration utility	99
	Configuring the CP server manually	106
	Verifying the CP server configuration	108
Chapter 6	Configuring VCS	109
	Overview of tasks for VCS configuration using installvcs program	109
	Starting the software configuration	110
	Specifying systems for configuration	111
	Configuring the basic cluster	111
	Configuring the virtual IP of the cluster	113
	Configuring the cluster in secure mode	115
	Adding VCS users	118
	Configuring SMTP email notification	118
	Configuring SNMP trap notification	120
	Configuring global clusters	122
	Completing the VCS configuration	123
	Verifying the NIC configuration	124
	Verifying and updating licenses on the system	125
	Checking licensing information on the system	125
	Updating product licenses using vxlicinst	126
Chapter 7	Configuring VCS clusters for data integrity	129
	Setting up disk-based I/O fencing using installvcs program	129
	Initializing disks as VxVM disks	129
	Configuring disk-based I/O fencing using installvcs program	130
	Checking shared disks for I/O fencing	132
	Setting up server-based I/O fencing using installvcs program	136
	Verifying security configuration on VCS cluster to use CP server coordination point	137
	Configuring server-based I/O fencing	139

Section 4	Installation using the Web-based installer	149
Chapter 8	Installing VCS	151
	Before using the Veritas Web-based installer	151
	Starting the Veritas Web-based installer	152
	Obtaining a security exception on Mozilla Firefox	152
	Performing a pre-installation check with the Veritas Web-based installer	152
	Installing VCS with the Veritas Web-based installer	153
Chapter 9	Configuring VCS	155
	Configuring VCS using the web-based installer	155
Section 5	Installation using response files	159
Chapter 10	Performing automated VCS installation	161
	Installing VCS using response files	161
	Response file variables to install VCS	162
	Sample response file for installing VCS	164
Chapter 11	Performing automated VCS configuration	167
	Configuring VCS using response files	167
	Response file variables to configure VCS	168
	Sample response file for configuring VCS	174
Chapter 12	Performing automated I/O fencing configuration for VCS	177
	Configuring I/O fencing using response files	177
	Response file variables to configure disk-based I/O fencing	178
	Sample response file for configuring disk-based I/O fencing	179
	Response file variables to configure server-based I/O fencing	180
	Sample response file for configuring server-based I/O fencing	182

Section 6	Manual installation	185
Chapter 13	Performing preinstallation tasks	187
	Preparing for a manual installation	187
	Requirements for installing VCS	187
Chapter 14	Manually installing VCS	189
	About VCS manual installation	189
	Installing VCS software manually	189
	Viewing the list of VCS RPMs	190
	Installing VCS RPMs for a manual installation	191
	Adding a license key for a manual installation	193
	Copying the installation guide to each node	195
	Installing VCS using Kickstart	195
	About Kickstart	195
	Summary of Kickstart installation tasks	196
	Setting up an NFS exported location for Veritas product RPMs	196
	Copying the rpms to the install/NFS server	196
	Creating the Veritas Kickstart script file	197
	Creating and modifying the Kickstart configuration file	198
	Completing the Kickstart process	198
Chapter 15	Manually configuring VCS	201
	Configuring LLT for a manual installation	201
	Setting up /etc/llthosts for a manual installation	202
	Setting up /etc/llttab for a manual installation	202
	LLT directives for a manual installation	202
	Additional considerations for LLT for a manual installation	203
	Configuring GAB for a manual installation	204
	Configuring VCS	204
	Configuring the cluster UUID when creating a cluster manually	205
	Starting LLT, GAB, and VCS for a manual installation	206
	Modifying the VCS configuration	207
	Configuring the ClusterService group	207

Chapter 16	Manually configuring the clusters for data integrity	209
	Setting up disk-based I/O fencing manually	209
	Identifying disks to use as coordinator disks	210
	Setting up coordinator disk groups	210
	Creating I/O fencing configuration files	211
	Modifying VCS configuration to use I/O fencing	212
	Verifying I/O fencing configuration	213
	Setting up server-based I/O fencing manually	214
	Preparing the CP servers manually for use by the VCS cluster	215
	Configuring server-based fencing on the VCS cluster manually	218
	Configuring Coordination Point agent to monitor coordination points	222
	Verifying server-based I/O fencing configuration	224
Section 7	Upgrading VCS	225
Chapter 17	Planning to upgrade VCS	227
	About upgrading to VCS 5.1	227
	VCS supported upgrade paths	228
	Upgrading VCS in secure enterprise environments	229
	About phased upgrade	229
	Prerequisites for a phased upgrade	229
	Planning for a phased upgrade	230
	Phased upgrade limitations	230
	Phased upgrade example	230
	Phased upgrade example overview	231
Chapter 18	Performing a typical VCS upgrade using the installer	233
	Before upgrading from 4.x using the script-based or Web-based installer	233
	Upgrading VCS using the script-based installer	234
	Upgrading VCS with the Veritas Web-based installer	235
Chapter 19	Performing a phased upgrade	237
	Performing a phased upgrade from VCS 5.0 MP3	237
	Moving the service groups to the second subcluster	237

	Upgrading the operating system on the first subcluster	241
	Upgrading the first subcluster	241
	Preparing the second subcluster	243
	Activating the first subcluster	247
	Upgrading the operating system on the second subcluster	248
	Upgrading the second subcluster	249
	Finishing the phased upgrade	250
Chapter 20	Performing an automated VCS upgrade using response files	253
	Upgrading VCS using response files	253
	Response file variables to upgrade VCS	254
	Sample response file for upgrading VCS	256
Section 8	Post-installation tasks	257
Chapter 21	Performing post-installation tasks	259
	About enabling LDAP authentication for clusters that run in secure mode	259
	Enabling LDAP authentication for clusters that run in secure mode	261
	Accessing the VCS documentation	267
	Removing permissions for communication	267
Chapter 22	Installing or upgrading VCS components	269
	Installing the Java Console	269
	Software requirements for the Java Console	269
	Hardware requirements for the Java Console	270
	Installing the Java Console on Linux	270
	Installing the Java Console on a Windows system	271
	Upgrading the Java Console	271
	Installing VCS Simulator	272
	Software requirements for VCS Simulator	272
	Installing VCS Simulator on Windows systems	272
	Reviewing the installation	273
	Upgrading VCS Simulator	273
	Upgrading the VCS agents	274

Chapter 23	Verifying the VCS installation	275
	About verifying the VCS installation	275
	About the LLT and GAB configuration files	275
	About the cluster UUID	278
	About the VCS configuration files	278
	Sample main.cf file for VCS clusters	279
	Sample main.cf file for global clusters	282
	Verifying the LLT, GAB, and VCS configuration files	284
	Verifying LLT, GAB, and cluster operation	285
	Verifying LLT	285
	Verifying GAB	288
	Verifying the cluster	290
	Verifying the cluster nodes	290
Section 9	Uninstalling VCS	295
Chapter 24	Uninstalling VCS using the installer	297
	Preparing to uninstall VCS	297
	Uninstalling VCS 5.1 using the script-based installer	298
	Removing VCS 5.1 RPMs	298
	Running <code>uninstallvcs</code> from the VCS 5.1 disc	299
	Uninstalling VCS with the Veritas Web-based installer	299
	Removing the CP server configuration using the removal script	300
Chapter 25	Uninstalling VCS using response files	305
	Uninstalling VCS using response files	305
	Response file variables to uninstall VCS	306
	Sample response file for uninstalling VCS	307
Section 10	Adding and removing nodes	309
Chapter 26	Adding and removing cluster nodes	311
	About adding and removing nodes	311
	Adding nodes using the VCS installer	311
	Manually adding a node to a cluster	314
	Setting up the hardware	315
	Installing the VCS software manually when adding a node	316
	Setting up the node to run in secure mode	317
	Configuring LLT and GAB	319
	Configuring I/O fencing on the new node	321

	Adding the node to the existing cluster	326
	Starting VCS and verifying the cluster	327
	Removing a node from a cluster	327
	Verifying the status of nodes and service groups	328
	Deleting the departing node from VCS configuration	329
	Modifying configuration files on each remaining node	332
	Removing the node configuration from the CP server	332
	Removing security credentials from the leaving node	333
	Unloading LLT and GAB and removing VCS on the departing node	333
Chapter 27	Adding a node to a single-node cluster	335
	Adding a node to a single-node cluster	335
	Setting up a node to join the single-node cluster	336
	Installing and configuring Ethernet cards for private network	337
	Configuring the shared storage	338
	Bringing up the existing node	338
	Installing the VCS software manually when adding a node to a single node cluster	339
	Creating configuration files	339
	Starting LLT and GAB	339
	Reconfiguring VCS on the existing node	340
	Verifying configuration on both nodes	341
Section 11	Installation reference	343
Appendix A	VCS installation packages	345
	Veritas Cluster Server installation packages	345
Appendix B	Installation command options	349
	Command options for installvcs program	349
	Command options for uninstallvcs program	354
Appendix C	Changes to bundled agents in VCS 5.1	357
	Deprecated agents	357
	New agents	357
	New and modified attributes for 5.1 agents	358
	Manually removing deprecated resource types and modifying attributes	365

	Creating new VCS accounts if you used native operating system accounts	366
Appendix D	Sample main.cf files	369
	Sample configuration files for CP server	369
	CP server hosted on a single node main.cf file	369
	CP server hosted on an SFHA cluster main.cf file	371
Appendix E	Installing VCS on a single node	377
	About installing VCS on a single node	377
	Creating a single-node cluster using the installer program	378
	Preparing for a single node installation	378
	Starting the installer for the single node cluster	378
	Creating a single-node cluster manually	379
	Setting the path variable for a manual single node installation	379
	Installing VCS software manually on a single node	380
	Renaming the LLT and GAB startup files	380
	Modifying the startup files	380
	Configuring VCS	380
	Verifying single-node operation	381
Appendix F	Configuring LLT over UDP using IPv4	383
	Using the UDP layer for LLT	383
	When to use LLT over UDP	383
	Configuring LLT over UDP	383
	Broadcast address in the /etc/llttab file	384
	The link command in the /etc/llttab file	385
	The set-addr command in the /etc/llttab file	385
	Selecting UDP ports	386
	Configuring the netmask for LLT	386
	Configuring the broadcast address for LLT	387
	Sample configuration: direct-attached links	387
	Sample configuration: links crossing IP routers	389
Appendix G	Configuring LLT over UDP using IPv6	391
	Using the UDP layer of IPv6 for LLT	391
	When to use LLT over UDP	391
	Configuring LLT over UDP using IPv6	391
	The link command in the /etc/llttab file	392
	The set-addr command in the /etc/llttab file	393
	Selecting UDP ports	393

	Sample configuration: direct-attached links	394
	Sample configuration: links crossing IP routers	395
Appendix H	Troubleshooting VCS installation	397
	What to do if you see a licensing reminder	397
	Restarting the installer after a failed connection	398
	Starting and stopping processes for the Veritas products	398
	Installer cannot create UUID for the cluster	399
	Script-based installer cannot start LLT, GAB, VCS, or VxFEN on a newly added node	399
	LLT startup script displays errors	400
	The vxfsth utility fails when SCSI TEST UNIT READY command fails	400
	Issues during server-based fencing start up on VCS cluster node	401
	cpsadm command on the VCS cluster node gives connection error	401
	Authentication failure	401
	Authorization failure	402
	Preexisting split-brain	402
	Adding a node to the secure cluster whose root broker system has failed	403
Appendix I	Sample VCS cluster setup diagrams for CP server-based I/O fencing	405
	Configuration diagrams for setting up server-based I/O fencing	405
	Two unique client clusters served by 3 CP servers	406
	Client cluster served by highly available CPS and 2 SCSI-3 disks	407
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	408
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	410
Index	413

Installation overview and planning

- [Chapter 1. Introducing VCS](#)
- [Chapter 2. Planning to install VCS](#)

Introducing VCS

This chapter includes the following topics:

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

About Veritas Cluster Server

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

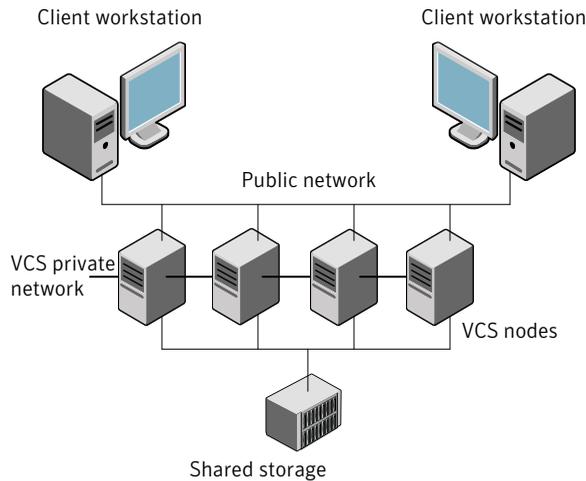
About VCS basics

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

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

[Figure 1-1](#) illustrates a typical VCS configuration of four nodes that are connected to shared storage.

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



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

VCS configuration files are as follows:

- `main.cf`—Defines the cluster, including services groups and resources.
- `types.cf`—Defines the resource types.
- `/etc/sysconfig/vcs`—Defines the startup and the shutdown behavior of VCS during system reboot and shutdown.

The `main.cf` and `types.cf` files reside in the `/etc/VRTSvcs/conf/config` folder by default.

See [“About the VCS configuration files”](#) on page 278.

About multiple nodes

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

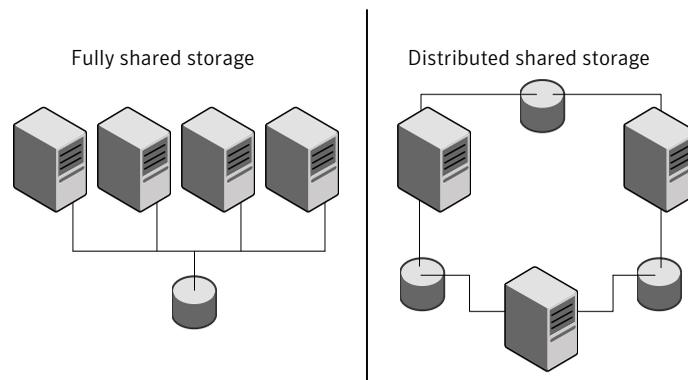
About shared storage

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

VCS nodes can only access physically-attached storage.

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

Figure 1-2 Two examples of shared storage configurations



About LLT and GAB

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

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

LLT configuration files are as follows:

- `/etc/llthosts`—lists all the nodes in the cluster
- `/etc/llttab`—describes the local system's private network links to the other nodes in the cluster

GAB (Group Membership and Atomic Broadcast) provides the global message order that is required to maintain a synchronized state among the nodes. It monitors disk communications such as the VCS heartbeat utility. The `/etc/gabtab` file is the GAB configuration file.

LLT and GAB initialization configuration files include:

- /etc/sysconfig/llt
- /etc/sysconfig/gab

See “[About the LLT and GAB configuration files](#)” on page 275.

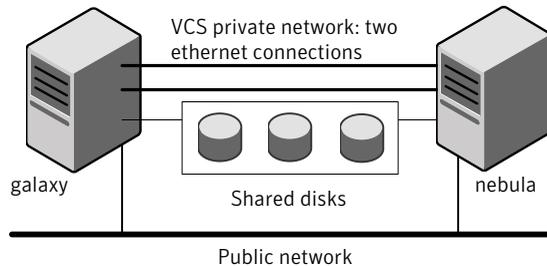
About network channels for heartbeating

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

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

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

Figure 1-3 Two Ethernet connections connecting two nodes



About preexisting network partitions

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

About VCS seeding

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

GAB automatically seeds nodes under the following conditions:

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

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

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

About Veritas product licensing

This release of the Veritas products introduces the option to install without a license key. The keyless license strategy does not eliminate the need to obtain a license. A software license is a legal instrument governing the usage or redistribution of copyright protected software. The administrator and company representatives must ensure that a server or cluster is entitled to the license level for the products installed. Symantec reserves the right to ensure entitlement and compliance through auditing.

If you encounter problems while licensing this product, visit the Symantec licensing support website.

The Veritas product installer prompts you to select one of the following licensing methods:

- Install a license key for the product and features that you want to install.
When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
- Continue to install without a license key.
The installer prompts for the product modes and options that you want to install, and then sets the required product level.
Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server. If you do not comply with the above terms, continuing to use the Veritas product is a violation of your end user license agreement, and results in warning messages. For more information about keyless licensing, see the following URL:
<http://go.symantec.com/sfhakeyless>

If you upgrade to this release from a prior release of the Veritas software, the product installer does not change the license keys that are already installed. The existing license keys may not activate new features in this release.

If you upgrade with the product installer, or if you install or upgrade with a method other than the product installer, you must do one of the following to license the products:

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.
See [“Setting or changing the product level for keyless licensing”](#) on page 193.
See the `vxkeyless(1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the 5.1 products you have purchased.
See the `vxlicinst(1m)` manual page.

You can also use the above options to change the product levels to another level that you are authorized to use. For example, you can add the replication option to the installed product. You must ensure that you have the appropriate license for the product level and options in use.

Note: In order to change from one product stack to another, additional steps may be required.

We recommend updating to keyless licensing for the following reasons:

- enables 5.1 functionality.
- allows you to change the product level easily.

About VCS features

You can use the Veritas Operation Services to assess your setup for VCS installation.

See [“About Veritas Operations Services”](#) on page 27.

VCS offers the following features that you can configure during VCS configuration:

VCS notifications	See “About VCS notifications” on page 27.
VCS global clusters	See “About global clusters” on page 28.
I/O fencing	See “About I/O fencing” on page 28.

About Veritas Operations Services

Veritas Operations Services (VOS) is a Web-based application that is designed specifically for Veritas CommandCentral and Veritas Storage Foundation and High Availability products. VOS increases operational efficiency and helps improve application availability.

VOS automates and simplifies administrator tasks, including:

- Determining if systems are ready to install or upgrade Veritas products
- Gathering deployment and usage information on Veritas products
- Receiving notifications about the latest updates for:
 - Patches
 - Hardware Compatibility Lists (HCLs)
 - Array Support Libraries (ASLs)
 - Array Policy Modules (APMs)
- Determining whether your Veritas product configurations conform to best practices
- Managing server and environmental configuration data from a single Website
- Interpreting Unified Message Identifier (UMI) codes and their solutions
- Identifying and downloading patches for Veritas products

To access VOS, go to:

<http://vos.symantec.com/>

About VCS notifications

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

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

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

About global clusters

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

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

About I/O fencing

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

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

The fencing operation determines the following:

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

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

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

You can configure I/O fencing to use one or both of the following components as coordination points:

Coordinator disk	I/O fencing that uses coordinator disks is referred to as disk-based I/O fencing. Disk-based I/O fencing ensures data integrity in a single cluster.
Coordination point server (CP server)	I/O fencing that uses at least one CP server system is referred to as server-based I/O fencing. Server-based I/O fencing ensures data integrity in multiple clusters.

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

About VCS optional components

You can add the following optional components to VCS:

Symantec Product Authentication Service	See “About Symantec Product Authentication Service (AT)” on page 29.
Veritas Cluster Server Management Console	See “About Veritas Cluster Server Management Console” on page 30.
Cluster Manager (Java console)	See “About Cluster Manager (Java Console)” on page 30.
VCS Simulator	See “About VCS Simulator” on page 30.

To configure most optional components, install all RPMs when the installation program prompts you. Note that most consoles are separate downloads.

About Symantec Product Authentication Service (AT)

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

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

- **Root broker**

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

A root broker on a stable external system can serve multiple clusters. Symantec recommends that you install a single root broker on a utility system. The utility system, such as an email server or domain controller, can be highly available. You can also configure one of the nodes in the VCS cluster to serve as a root and an authentication broker.

- **Authentication brokers**

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

See Symantec Product Authentication Service documentation for more information.

See “[Preparing to configure the clusters in secure mode](#)” on page 77.

About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers complete administration capabilities for your cluster. Use the different views in the Java Console to monitor clusters and VCS objects, including service groups, systems, resources, and resource types.

You can download the console from <http://go.symantec.com/vcsmc>.

Symantec also offers the Veritas Cluster Server (VCS) Management Console to manage clusters.

See “[About Veritas Cluster Server Management Console](#)” on page 30.

See *Veritas Cluster Server Administrator's Guide*.

About Veritas Cluster Server Management Console

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

You can configure Veritas Cluster Server Management Console to manage multiple clusters.

Refer to the *Veritas Cluster Server Management Console Implementation Guide* for installation, upgrade, and configuration instructions.

For information on updates and patches for VCS Management Console, see <http://seer.entsupport.symantec.com/docs/308405.htm>.

To download the most current version of VCS Management Console, go to <http://www.symantec.com/business/cluster-server> and click **Utilities**.

About VCS Simulator

VCS Simulator enables you to simulate and test cluster configurations. Use VCS Simulator to view and modify service group and resource configurations and test failover behavior. VCS Simulator can be run on a stand-alone system and does not require any additional hardware. VCS Simulator can be installed only a Windows operating system.

VCS Simulator runs an identical version of the VCS High Availability Daemon (HAD) as in a cluster, ensuring that failover decisions are identical to those in an actual cluster.

You can test configurations from different operating systems using VCS Simulator. For example, you can run VCS Simulator on a Windows system and test VCS configurations for Windows, Linux, and Solaris clusters. VCS Simulator also enables creating and testing global clusters.

You can administer VCS Simulator from the Java Console or from the command line.

To download VCS Simulator, go to <http://www.symantec.com/business/cluster-server> and click **Utilities**.

Planning to install VCS

This chapter includes the following topics:

- [VCS installation requirements](#)
- [VCS installation methods](#)
- [Typical VCS cluster setup models](#)

VCS installation requirements

You can install VCS on clusters of up to 32 systems. Every node where you want to install VCS must meet the hardware and the software requirements.

This is document version 5.1.0. Before you continue, make sure that you are using the current version of this guide. It is online at:

http://sfdoccentral.symantec.com/sf/5.1/linux/pdf/vcs_install.pdf

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/335001>

The hardware compatibility list contains information about supported hardware and is updated regularly. For the latest information on supported hardware visit the following URL:

<http://entsupport.symantec.com/docs/330441>

Before installing or upgrading Veritas Cluster Server, review the current compatibility list to confirm the compatibility of your hardware and software.

Hardware requirements

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

Table 2-1 Hardware requirements for a VCS cluster

Item	Description
VCS nodes	From 1 to 32 Linux systems that run the supported Linux operating system version. See “ Supported Linux operating systems ” on page 35.
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.
Disks	Typical VCS configurations require that shared disks support the applications that migrate between systems in the cluster. The VCS I/O fencing feature requires that all data and coordinator disks support SCSI-3 Persistent Reservations (PR).
Disk space	See “ Required disk space ” on page 34. Note: VCS may require more temporary disk space during installation than the specified disk space.
Network Interface Cards (NICs)	In addition to the built-in public NIC, VCS requires at least one more NIC per system. Symantec recommends two additional NICs. You can also configure aggregated interfaces. Symantec recommends that you turn off the spanning tree on the LLT switches, and set port-fast on.
Fibre Channel or SCSI host bus adapters	Typical VCS configuration requires at least one SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.
RAM	Each VCS node requires at least 256 megabytes.

Required disk space

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

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

Table 2-2 Disk space requirements and totals

Packages	/	/opt	/usr	/var	Totals
Required	3 MB	241 MB	8 MB	1 MB	250 MB
Optional	1 MB	51 MB	0 MB	7 MB	58 MB

Table 2-2 Disk space requirements and totals (*continued*)

Packages	/	/opt	/usr	/var	Totals
Required and optional total	4 MB	292 MB	8 MB	8 MB	308 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 Linux operating systems

This section lists the supported operating systems for this release of Veritas products.

For important updates regarding this release, review the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/335001>

The Veritas 5.1 release operates on the following operating systems and hardware:

- Red Hat Enterprise Linux 5 (RHEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 10 (SLES 10) with SP2 (2.6.16.60-0.21 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- SUSE Linux Enterprise Server 11 (SLES 11) (2.6.27.19-5 kernel) on AMD Opteron or Intel Xeon EM64T (x86_64)
- Oracle Enterprise Linux (OEL 5) with Update 3 (2.6.18-128.el5 kernel) or later on AMD Opteron or Intel Xeon EM64T (x86_64)

If your system is running an older version of either Red Hat Enterprise Linux, SUSE Linux Enterprise Server, or Oracle Enterprise Linux, you must upgrade it before attempting to install the Veritas Storage Foundation software. Consult the Red Hat, SUSE, or Oracle documentation for more information on upgrading or reinstalling your system.

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

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

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

<http://entsupport.symantec.com/docs/335001>

Required Linux RPMs for VCS

Make sure you installed the following operating system-specific RPMs on the systems where you want to install or upgrade VCS. VCS will support any updates made to the following RPMs, provided the RPMs maintain the ABI compatibility.

[Table 2-3](#) lists the RPMs that VCS requires for a given Linux operating system.

Table 2-3 Required RPMs

Operating system	Required RPMs
RHEL 5	compat-libgcc-296-2.96-138.i386.rpm compat-libstdc++-33-3.2.3-61.i386.rpm compat-libstdc++-296-2.96-138.i386.rpm glibc-2.5-24.i686.rpm libgcc-4.1.2-42.el5.i386.rpm libstdc++-3.4.6-10.i386.rpm compat-libstdc++-33-3.2.3-61.x86_64.rpm glibc-2.5-24.x86_64.rpm glibc-common-2.5-24.x86_64.rpm libgcc-4.1.2-42.el5.x86_64.rpm libstdc++-3.4.6-10.x86_64.rpm java-1.4.2-gcj-compat-1.4.2.0-40jpp.115.noarch.rpm
SLES 10	compat-32bit-2006.1.25-11.2.x86_64.rpm glibc-32bit-2.4-31.54.x86_64.rpm compat-2006.1.25-11.2.x86_64.rpm compat-libstdc++-5.0.7-22.2.x86_64.rpm glibc-2.4-31.54.x86_64.rpm libgcc-4.1.2_20070115-0.21.x86_64.rpm libstdc++-4.1.2_20070115-0.21.x86_64.rpm

Table 2-3 Required RPMs (*continued*)

Operating system	Required RPMs
SLES 11	glibc-2.9-13.2 glibc-32bit-2.9-13.2 libgcc43-4.3.3_20081022-11.18 libgcc43-32bit-4.3.3_20081022-11.18 libstdc++43-4.3.3_20081022-11.18 libstdc++43-32bit-4.3.3_20081022-11.18

Required Solaris patches for VCS

Before you install your Symantec products on Solaris, read the following TechNote and perform the instructions in it:

<http://entsupport.symantec.com/docs/334829>

Supported software

VCS supports the following volume managers and file systems:

- ext2, ext3, reiserfs, NFS, and bind on LVM2, raw disks, and VxVM.
- Veritas Storage Foundation (SF): Veritas Volume Manager (VxVM) with Veritas File System (VxFS)

VCS 5.1 supports the following versions of SF:

- SF 5.0 MP3
 - VxVM 5.0 MP3 with VxFS 5.0 MP3

SF 5.0 MP3 supports only 64-bit architecture on Linux. See the *Veritas Storage Foundation Release Notes* for more details.
- SF 5.1
 - VxVM 5.1 with VxFS 5.1

Note: VCS supports the previous version of SF and the next version of SF to facilitate product upgrades.

I/O fencing requirements

Depending on whether you plan to configure disk-based fencing or server-based fencing, make sure that you meet the requirements for coordination points:

- Coordinator disks
See “[Coordinator disk requirements for I/O fencing](#)” on page 38.
- CP servers
See “[CP server requirements](#)” on page 38.

To configure disk-based fencing or to configure server-based fencing with at least one coordinator disk, make sure a version of Veritas Volume Manager (VxVM) that supports SCSI-3 persistent reservations (SCSI-3 PR) is installed on the VCS cluster.

See the *Veritas Storage Foundation and High Availability Installation Guide*.

Coordinator disk requirements for I/O fencing

Make sure that the I/O fencing coordinator disks meet the following requirements:

- For disk-based I/O fencing, you must have three coordinator disks.
- The coordinator disks can be raw devices, DMP devices, or iSCSI devices. You must use DMP disk policy for iSCSI-based coordinator disks. For the latest information on supported hardware visit the following URL: <http://entsupport.symantec.com/docs/283161>
- Each of the coordinator disks must use a physically separate disk or LUN. Symantec recommends using the smallest possible LUNs for coordinator disks.
- Each of the coordinator disks should exist on a different disk array, if possible.
- The coordinator disks must support SCSI-3 persistent reservations.
- Symantec recommends using hardware-based mirroring for coordinator disks.
- Coordinator disks must not be used to store data or must not be included in disk groups that store user data.
- Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.

CP server requirements

The following requirements must be met for a CP server installation:

- CP server hardware-specific requirements
- OS requirements

- Networking requirements (and recommendations)
- Security requirements

For the basic hardware requirements for the VCS/SFHA cluster to host the CP server, refer to the appropriate VCS or SFHA installation and configuration guide.

[Table 2-4](#) lists additional requirements for hosting the CP server.

Table 2-4 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: <ul style="list-style-type: none"> ■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB) ■ 300 MB in /usr ■ 20 MB in /var
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the CP servers.
RAM	Each CP server requires at least 512 MB.
CP server to client node physical link	A secure TCP/IP connection is required to connect the CP server(s) to the VCS cluster.

[Table 2-5](#) displays the CP server supported operating systems and versions.

Table 2-5 CP server supported operating systems and versions

CP server	Operating system and version
CP server hosted on a VCS single node cluster or	<ul style="list-style-type: none"> ■ Solaris 9 (SPARC) ■ Solaris 10 (SPARC or x86)
CP server hosted on an SFHA cluster	<ul style="list-style-type: none"> ■ Linux (RHEL5, SLES10, SLES11)

For networking requirements, Symantec recommends that network access from the VCS clusters to the CP servers should be made highly-available and redundant. The network connections require either a secure LAN or VPN.

The CP server uses the TCP/IP protocol to connect to and communicate with the VCS cluster(s) by these network paths. The CP server listens for messages from the VCS cluster(s) using TCP port 14250. This is the default port that can be changed during a CP server configuration.

Note: The CP server supports either Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with the VCS clusters. If the CP server is configured to use an IPv6 virtual IP address, then the VCS clusters should also be on the IPv6 network where the CP server is being hosted.

When placing the CP server (s) within a specific network configuration, the number of hops from the different VCS cluster nodes to the CP server (s) should be taken into consideration. As a best practices procedure, Symantec recommends that the number of hops from the different VCS cluster nodes to the CP server(s) should be equal. This ensures that if an event occurs that results in an I/O fencing scenario, there is no bias in the race due to the number of hops between the nodes.

For secure communications between the VCS cluster and CP server, be sure to consider the following requirements and suggestions:

- If security is configured, both VCS and the customized fencing framework can use secure channels for communication. Configuring VCS in secure mode and CP server or VCS cluster in non-secure mode is supported, but configuring VCS in non-secure mode and CP server in secure mode is not supported.
- In a secure communication environment, all CP servers that are used by the VCS cluster must be configured with security enabled. A configuration where the VCS cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.
- The CP server and VCS clusters should also use the same root broker. If the same root broker is not being used, then trust can be established between the cluster nodes and CP server for the secure communication. Trust can be established by the installer when configuring fencing.
- For non-secure communication between CP server and VCS clusters, there is no need to configure Symantec Product Authentication Service. In non-secure mode, authorization is still provided by CP server for the VCS cluster users. The authorization that is performed only ensures that authorized users can perform appropriate actions as per their user privileges on the CP server.

For additional information, see *Veritas Cluster Server Administrator's Guide*.

VCS installation methods

[Table 2-6](#) lists the different methods you can choose to install and configure VCS:

Table 2-6 VCS installation methods

Method	Description
Interactive installation using the script-based installer	<p>You can use one of the following script-based installers:</p> <ul style="list-style-type: none"> ■ Veritas product installer Use to install and configure multiple Veritas products. ■ installvcs program Use to install and configure just VCS. <p>The script-based installer asks you a series of questions and installs and configures VCS based on the information you provide.</p>
Interactive installation using the web-based installer	<p>You can use a web-interface to install and configure VCS.</p>
Automated installation using the VCS response files	<p>At the end of each successful simulated or actual installation and configuration, the installer creates response files. You can use these response files to perform multiple installations to set up a large VCS cluster.</p>
Manual installation using the Linux commands and utilities	<p>You can install VCS using the operating system <code>rpm -i</code> command and then manually configure VCS.</p> <p>You can also install VCS using the Kickstart utility for RHEL.</p>

About the VCS installation program

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

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

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

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

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

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

See [“About preparing to install VCS”](#) on page 51.

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

Features of the script-based installer

The script-based installer supports installing, configuring, upgrading, and uninstalling VCS. In addition, the script-based installer also provides command options to perform the following tasks:

- Check the systems for VCS installation requirements.
See [“Performing automated preinstallation check”](#) on page 62.
- Upgrade VCS if a previous version of VCS currently runs on a cluster.
See [“Upgrading VCS using the script-based installer”](#) on page 234.
- Start or stop VCS processes
See [“Starting and stopping processes for the Veritas products ”](#) on page 398.
- Enable or disable a cluster to run in secure mode using Symantec Product Authentication Service (VxAT)
See the *Veritas Cluster Server Administrator’s Guide*.
- Configure I/O fencing for the clusters to prevent data corruption
See [“Setting up disk-based I/O fencing using installvcs program”](#) on page 129.
See [“Setting up server-based I/O fencing using installvcs program”](#) on page 136.
- Create a single-node cluster
See [“Creating a single-node cluster using the installer program”](#) on page 378.
- Add a node to an existing cluster
See [“Adding nodes using the VCS installer”](#) on page 311.
- Create a kickstart configuration file to install VCS using the Kickstart utility for RHEL.
See [“Installing VCS using Kickstart”](#) on page 195.
- Perform automated installations using the values that are stored in a configuration file.
See [“Installing VCS using response files”](#) on page 161.

See “Configuring VCS using response files” on page 167.

See “Upgrading VCS using response files” on page 253.

Interacting with the `installvcs` program

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

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

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

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

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

About the Web-based installer

The Web-based installer is a convenient GUI method to install the Veritas products. The Web-based installer also enables you to configure the product and verify preinstallation requirements.

The `webinstaller` script is used to start and stop the Veritas XPortal Server `xprtld` process. The `webinstaller` script can also be used to check the status of the XPortal Server.

When the `webinstaller` script starts the `xprtld` process, the script displays a URL. Use this URL to access the Web-based installer from a Web browser such as Internet Explorer or FireFox.

The Web installer creates log files whenever the Web installer is operating. While the installation processes are operating, the log files are located in a session-based directory under the `/var/tmp` directory. After the install process completes, the

log files are located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

The location of the Veritas XPortal Server configuration file is `/var/opt/webinstaller/xprt1id.conf`.

Features supported with Web-based installer

The Web-based installer works similarly to the script installer. For the initial release, certain new or advanced features available in the script installer are not available in the Web-based installer.

The following features are supported in the Web-based installer:

- Installing a product
- Uninstalling a product
- Upgrading a product
- Configuring a clustered product including:
 - Required VCS configuration - Cluster name, Cluster ID, Heartbeat NICs
 - Optional VCS configuration - Users, SMTP Notification, SNMP Notification, GCO required, Virtual IP
 - SFCFS configuration - fencing enabled question
 - Configuring Veritas Volume Manager and Veritas Volume Replicator with the installer is not required for this release.
- Starting a product
- Stopping a product
- Licensing a product
- Performing an installation precheck

About response files

The installer generates a "response file" after performing an installer task such as installation, configuration, uninstallation, or upgrade. These response files contain the details that you provided to the installer questions in the form of values for the response file variables. The response file also contains descriptions and explanations of the variables and their values.

You can also create a response file using the `-makeresponsefile` option of the installer.

The installer displays the location of the response file at the end of each successful installer task. The installer saves the response file in the default location for the

install-related log files: `/opt/VRTS/install/logs`. If you provided a different log path using the `-logpath` option, the installer saves the response file in the path that you specified.

The format of the response file name is:

`/opt/VRTS/install/logs/installscript-YYYYMMDDHHSSxxx`
`/installscript-YYYYMMDDHHSSxxx.response`, where:

- *installscript* may be, for example: `installer`, `webinstaller`, `installvcs` program, or `uninstallvcs` program
- *YYYYMMDDHHSS* is the current date when the *installscript* is run and *xxx* are three random letters that the script generates for an installation instance

For example:

`/opt/VRTS/install/logs/installer-200910101010ldS/installer-200910101010ldS.response`

You can customize the response file as required to perform unattended installations using the `-responsefile` option of the installer. This method of automated installations is useful in the following cases:

- To perform multiple installations to set up a large VCS cluster.
See [“Installing VCS using response files”](#) on page 161.
- To upgrade VCS on multiple systems in a large VCS cluster.
See [“Upgrading VCS using response files”](#) on page 253.
- To uninstall VCS from multiple systems in a large VCS cluster.
See [“Uninstalling VCS using response files”](#) on page 305.

Syntax in the response file

The syntax of the Perl statements that are included in the response file variables varies. It can depend on whether the variables require scalar or list values.

For example, in the case of a string value:

```
$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"];
```

Typical VCS cluster setup models

VCS clusters support different failover configurations, storage configurations, and cluster topologies.

See the *Veritas Cluster Server Administrator's Guide* for more details.

Some of the typical VCS setup models are as follows:

- Basic VCS cluster with two nodes
See [Figure 2-1](#) on page 46.
- VCS clusters in secure mode using Symantec Product Authentication Service (AT)
See [Figure 2-2](#) on page 47.
- VCS clusters centrally managed using Veritas Cluster Server Management Console (VCS MC)
See [Figure 2-3](#) on page 48.
- VCS clusters with I/O fencing for data protection
See [“Typical VCS cluster configuration with disk-based I/O fencing”](#) on page 93.
See [“Typical VCS cluster configuration with server-based I/O fencing”](#) on page 93.
- VCS clusters such as global clusters, replicated data clusters, or campus clusters for disaster recovery
See the *Veritas Cluster Server Administrator's Guide* for disaster recovery cluster configuration models.

[Figure 2-1](#) illustrates a simple VCS cluster setup with two nodes.

Figure 2-1 Typical two-node VCS cluster

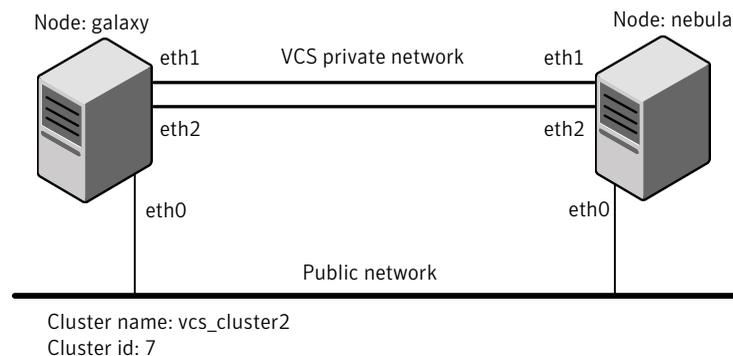
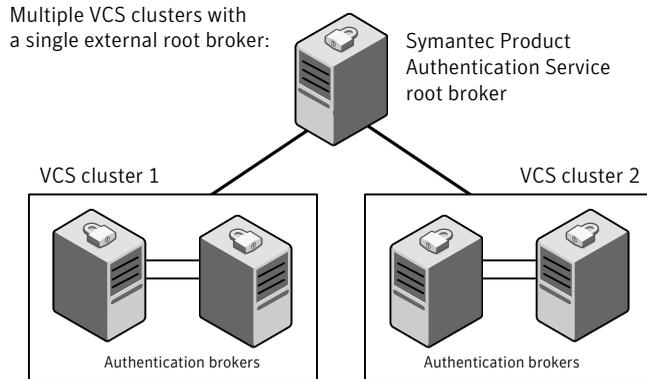


Figure 2-2 illustrates typical configuration of VCS clusters in secure mode. You can use one of the cluster nodes as AT root broker or you can use a stable system outside the cluster as AT root broker.

Figure 2-2 Typical configuration of VCS clusters in secure mode



Single VCS cluster with one of the nodes as root broker:

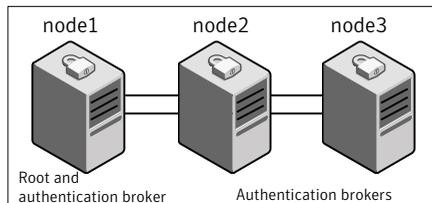
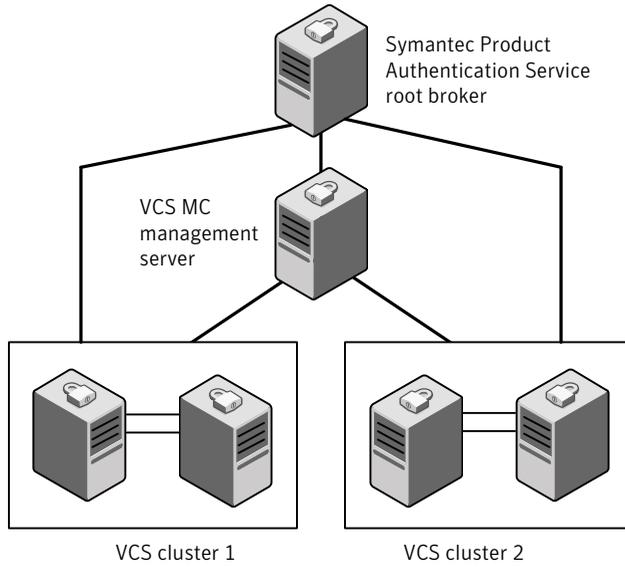


Figure 2-3 illustrates a typical setup of VCS clusters that are centrally managed using VCS Management Console. In this example, the setup uses a single external AT root broker which serves both the management server and the VCS clusters.

Figure 2-3 Typical configuration of VCS MC-managed clusters



Preinstallation tasks

- [Chapter 3. Preparing to install VCS](#)

Preparing to install VCS

This chapter includes the following topics:

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

About preparing to install VCS

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

See [“VCS installation requirements”](#) on page 33.

Performing preinstallation tasks

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

Table 3-1 Preinstallation tasks

Task	Reference
Obtain license keys.	See “Obtaining VCS license keys” on page 52.
Set up the private network.	See “Setting up the private network” on page 53.
Configure persistent network interfaces	See “Configuring persistent network interfaces” on page 55.
Enable communication between systems.	See “Setting up inter-system communication” on page 55.

Table 3-1 Preinstallation tasks (*continued*)

Task	Reference
Set up ssh on cluster systems.	See “Setting up ssh on cluster systems” on page 55.
Set up shared storage for I/O fencing (optional)	See “Setting up shared storage” on page 57.
Set the PATH and the MANPATH variables.	See “Setting the PATH variable” on page 60. See “Setting the MANPATH variable” on page 60.
Set the kernel.panic tunable	See “Setting the kernel.panic tunable” on page 61.
Review basic instructions to optimize LLT media speeds.	See “Optimizing LLT media speed settings on private NICs” on page 61.
Review guidelines to help you set the LLT interconnects.	See “Guidelines for setting the media speed of the LLT interconnects” on page 61.
Mount the product disc	See “Mounting the product disc” on page 62.
Verify the systems before installation	See “Performing automated preinstallation check” on page 62.

Obtaining VCS license keys

If you decide to not use the keyless licensing, you must obtain and install a license key for VCS.

See [“About Veritas product licensing”](#) on page 25.

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

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

<https://licensing.symantec.com>

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

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

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

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

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

Setting up the private network

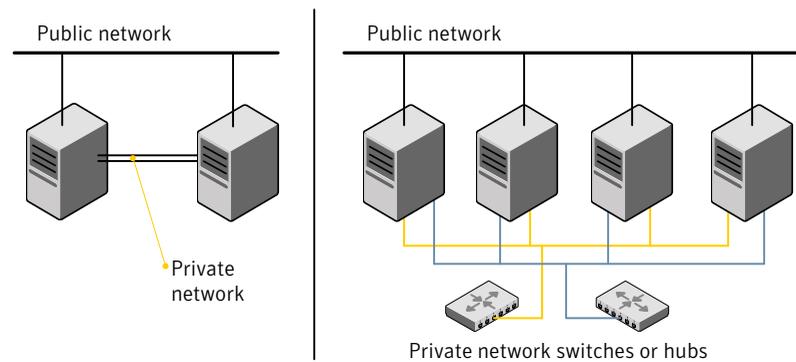
VCS requires you to set up a private network between the systems that form a cluster. You can use either NICs or aggregated interfaces to set up private network.

You can use network switches instead of hubs.

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

Figure 3-1 shows two private networks for use with VCS.

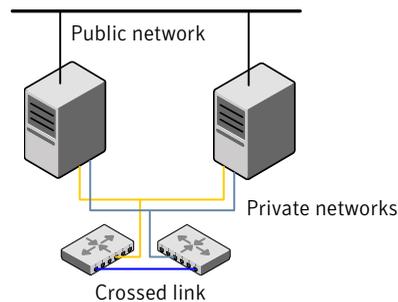
Figure 3-1 Private network setups: two-node and four-node clusters



Symantec recommends configuring two independent networks between the cluster nodes with a network switch for each network. You can also interconnect multiple layer 2 switches for advanced failure protection. Such connections for LLT are called cross-links.

Figure 3-2 shows a private network configuration with crossed links between the network switches.

Figure 3-2 Private network setup with crossed links



To set up the private network

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

Ensure that you meet the following requirements:

- The power to the switches or hubs must come from separate sources.
- On each system, you must use two independent network cards to provide redundancy.
- If a network interface is part of an aggregated interface, you must not configure the network interface under LLT. However, you can configure the aggregated interface under LLT.
- When you configure Ethernet switches for LLT private interconnect, disable the spanning tree algorithm on the ports used for the interconnect.

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

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

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

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

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

Configuring persistent network interfaces

By default, SLES 10 and later uses `udev` to achieve persistent interface names.

For RHEL, SLES, and OEL, refer to the OS documentation for information on configuring persistent interfaces.

Setting up inter-system communication

When you install VCS using the `installvcs` program, to install and configure the entire cluster at one time, make sure that communication between systems exists. By default the installer uses `ssh`. You must grant root privileges for the system where you run `installvcs` program. This privilege facilitates to issue `ssh` or `rsh` commands on all systems in the cluster. If `ssh` is used to communicate between systems, it must be configured in a way such that it operates without requests for passwords or passphrases. Similarly, `rsh` must be configured in such a way to not prompt for passwords.

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

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

Setting up ssh on cluster systems

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

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

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

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

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

Configuring ssh

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

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

To configure ssh

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

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

- 3 Accept the default location of `~/.ssh/id_dsa`.
- 4 When the command prompts, enter a passphrase and confirm it.
- 5 Change the permissions of the `.ssh` directory by typing:

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

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

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

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

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

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

- 8 When the command prompts, enter your DSA passphrase.

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

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

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

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

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

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

Setting up shared storage

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

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

Setting up shared storage: SCSI

Perform the following steps to set up shared storage.

To set up shared storage

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

Set the following:

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

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

Perform the following:

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

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

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

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

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

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

For example, enter the following command:

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

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

7 Power off the disk.

8 Remove the terminator from the disk and connect the disk to the other cluster system.

9 Power on the disk.

10 Boot the second system. The system can now detect the disk.

11 Press `Ctrl+A` to bring up the SCSI BIOS settings for the disk.

Set the following:

- Set Host adapter SCSI ID = 6, or to an appropriate value for your configuration. Note that the SCSI ID should be different from the one configured on the first cluster system.
- Set Host Adapter BIOS in Advanced Configuration Options to Disabled.

12 Verify that you can view the shared disk using the `fdisk` command.

Setting up shared storage: Fiber Channel

Perform the following steps to set up fiber channel.

To set up shared storage for fiber channel

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

- Type the following command:

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

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

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

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

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

For example, enter the following command:

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

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

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

See “[Verifying that the nodes have access to the same disk](#)” on page 133.

Setting the PATH variable

Installation commands as well as other commands reside in the `/opt/VRTS/bin` directory. Add this directory to your PATH environment variable.

If you have any custom scripts located in `/opt/VRTSvcs/bin` directory, make sure to add the `/opt/VRTSvcs/bin` directory to your PATH environment variable.

To set the PATH variable

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

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

```
$ setenv PATH :/opt/VRTS/bin:$PATH
```

Setting the MANPATH variable

Set the MANPATH variable to view the manual pages.

To set the MANPATH variable

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

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

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

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

```
# export LC_ALL=C
```

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

Setting the kernel.panic tunable

By default, the `kernel.panic` tunable is set to zero. Therefore the kernel does not reboot automatically if a node panics. To ensure that the node reboots automatically after it panics, this tunable must be set to a non zero value.

To set the kernel.panic tunable

- 1 Set the `kernel.panic` tunable to a desired value in the `/etc/sysctl.conf` file.
For example, `kernel.panic = 10`, will assign a value 10 seconds to the `kernel.panic` tunable. This step makes the change persistent across reboots.
- 2 Run the command:

```
sysctl -w kernel.panic=10
```

In case of a panic, the node will reboot after 10 seconds.

Optimizing LLT media speed settings on private NICs

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

Guidelines for setting the media speed of the LLT interconnects

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

- Symantec recommends that you manually set the same media speed setting on each Ethernet card on each node.
- If you have hubs or switches for LLT interconnects, then set the hub or switch port to the same setting as used on the cards on each node.
- If you use directly connected Ethernet links (using crossover cables), set the media speed to the highest value common to both cards, typically `1000_Full_Duplex`.

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

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

Mounting the product disc

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

To mount the product disc

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

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

- 4 Navigate to the location of the RPMs.

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

Where *dist* is rhel5, sles10, or sles11, and *arch* is x86_64 for RHEL and SLES.

Performing automated preinstallation check

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

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

You can use the Veritas Operation Services to assess your setup for VCS installation.

See [“About Veritas Operations Services”](#) on page 27.

To check the systems

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

See “[Mounting the product disc](#)” on page 62.

- 2 Start the preinstallation check:

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

The program proceeds in a noninteractive mode to examine the systems for licenses, RPMs, disk space, and system-to-system communications.

- 3 Review the output as the program displays the results of the check and saves the results of the check in a log file.

See “[Command options for installvcs program](#)” on page 349.

Reformatting VCS configuration files on a stopped cluster

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

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

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

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

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

- Upgrade VCS
- Uninstall VCS

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

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

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

```
# haconf -dump
```

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

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

```
# hacf -cftocmd config
```

```
# hacf -cmdtoctf config
```

Getting your VCS installation and configuration information ready

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

When you perform the installation, prepare the following information:

- To install VCS RPMs you need:

The system names where you plan to install VCS Example: **galaxy, nebula**

The required license keys If you decide to use keyless licensing, you do not need to obtain license keys. However, you require to set up management server within 60 days to manage the cluster.

See [“About Veritas product licensing”](#) on page 25.

Depending on the type of installation, keys can 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 52.

- To decide which packages to install
- Minimum RPMs—provides basic VCS functionality.
 - Recommended RPMs—provides full functionality of VCS without advanced features.
 - All RPMs—provides advanced feature functionality of VCS.

The default option is to install the recommended RPMs.

See “[Viewing the list of VCS RPMs](#)” on page 190.

■ To configure VCS you need:

A name for the cluster The cluster name must begin with a letter of the alphabet. The cluster name can contain only the characters "a" through "z", "A" through "Z", the numbers "0" through "9", the hyphen "-", and the underscore "_".

Example: `vcs_cluster27`

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

Example: 7

The device names of the NICs that the private networks use among systems A network interface card that is not part of any aggregated interface, or an aggregated interface.
Do not use the network interface card that is used for the public network, which is typically eth0.

Example: `eth1`, `eth2`

■ To configure virtual IP address of the 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 You can enter either an IPv4 or an IPv6 address. This virtual IP address becomes a resource for use by the ClusterService group. The "Cluster Virtual IP address" can fail over to another cluster system.

Example IPv4 address: 192.168.1.16

Example IPv6 address:

2001:454e:205a:110:203:baff:feee:10

The netmask for the virtual IPv4 address The subnet that you use with the virtual IPv4 address.
Example: 255.255.240.0

The prefix for the virtual IPv6 address The prefix length for the virtual IPv6 address.
Example: 64

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

To decide the root broker system you want to use You can use an external system or one of the nodes in the cluster to serve as root broker.

To decide which configuration mode you want to choose Configuration modes are automatic, semiautomatic, and manual.

If you want one of the nodes in the cluster to serve as root broker, you must choose automatic configuration mode.

See [“Preparing to configure the clusters in secure mode”](#) on page 77.

For automatic mode (default) If you use an external root broker system:

- The name of the root broker system

Example: east

See [“About Symantec Product Authentication Service \(AT\)”](#) on page 29.

- Access to the root broker system without use of a password.

If you use one of the nodes in the cluster as root broker system:

- Decide which node in the cluster you want the installer to configure as root and authentication broker.
- The installer configures all other nodes in the cluster as authentication brokers.

For semiautomatic mode using encrypted files

The path for the encrypted files that you get from the root broker administrator.

See [“Creating encrypted files for the security infrastructure”](#) on page 83.

- For manual mode without using encrypted files
 - The fully-qualified hostname (FQDN) of the root broker . (e.g. east.symantecexample.com)
The given example puts 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 identity and password on each cluster node (e.g. galaxy.symantecexample.com and nebula.symantecexample.com)

■ To add VCS users you need:

- | | |
|---------------------------|--|
| User names | VCS usernames are restricted to 1024 characters.
Example: smith |
| User passwords | VCS passwords are restricted to 255 characters.
Enter the password at the prompt. |
| To decide user privileges | Users have three levels of privileges: A=Administrator, O=Operator, or G=Guest.
Example: A |

■ To configure SMTP email notification (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 |
| 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 name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. Example: eth0
---	--

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 used to configure the virtual IP of the cluster. 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 earlier for the cluster. Otherwise, specify appropriate values for the virtual IP address. Example: 192 . 168 . 1 . 16
-----------------------------------	--

The netmask for the virtual IP address	You can use the same netmask that you used to configure the virtual IP of the cluster. Otherwise, specify appropriate values for the netmask. Example: 255 . 255 . 240 . 0
--	--

■ To configure I/O fencing:

See [“About planning to configure I/O fencing”](#) on page 91.

Installation using the script-based installer

- [Chapter 4. Installing VCS](#)
- [Chapter 5. Preparing to configure VCS](#)
- [Chapter 6. Configuring VCS](#)
- [Chapter 7. Configuring VCS clusters for data integrity](#)

Installing VCS

This chapter includes the following topics:

- [Installing VCS using the installer](#)

Installing VCS using the installer

Perform the following steps to install VCS.

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

To install VCS

- 1 Confirm that you are logged in as the superuser and you mounted the product disc.
See [“Mounting the product disc”](#) on page 62.
- 2 Start the installation program. If you obtained VCS from an electronic download site, which does not include the Veritas product installer, use the `installvcs` program.

Veritas product installer Perform the following steps to start the product installer:

- 1 Start the installer.

```
# ./installer
```

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

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

installvcs program Perform the following steps to start the product installer:

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

```
# cd cdrom_root/cluster_server
```

- 2 Start the installvcs program.

```
# ./installvcs
```

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

- 3 Enter **y** to agree to the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement  
as specified in the EULA.pdf file present on media? [y,n,q,?] y
```

- 4 Choose the VCS RPMs that you want to install.

See "[Veritas Cluster Server installation packages](#)" on page 345.

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

- 1 Installs only the minimal required VCS RPMs that provides basic functionality of the product.
- 2 Installs the recommended VCS RPMs that provides complete functionality of the product. This option does not install the optional VCS RPMs.
Note that this option is the default.
- 3 Installs all the VCS RPMs.
You must choose this option to configure any optional VCS feature.

4 Displays the VCS RPMs for each option.

```
Select the RPMs to be installed on all systems? [1-4,q,?]
(2) 3
```

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

```
Enter the 64 bit operating_system system names separated by spaces:
galaxy nebula
```

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

See [“Creating a single-node cluster using the installer program”](#) on page 378.

The installer does the following for the systems:

- Checks that the local system that runs the installer can communicate with remote systems.
If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases.
If the default communication method ssh fails, the installer attempts to use rsh.
- Makes sure the systems use one of the supported operating systems.
- Makes sure that the systems have the required operating system patches.
If the installer reports that any of the patches are not available, install the patches on the system before proceeding with the VCS installation.
See [“Required Linux RPMs for VCS”](#) on page 36.
See [“Required Solaris patches for VCS”](#) on page 37.
- Checks for product licenses.
- 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 5.1.
See [“About upgrading to VCS 5.1”](#) on page 227.
- 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 installer stops and indicates the actions that you must perform to proceed with the process.
- Checks whether any of the RPMs already exists on a system.

If the current version of any RPM exists, the installer removes the RPM from the installation list for the system. If a previous version of any RPM exists, the installer replaces the RPM with the current version.

6 Review the list of RPMs that the installer would install on each node.

The installer installs the VCS RPMs on the systems galaxy and nebula.

7 Select the license type.

- 1) Enter a valid license key
- 2) Enable keyless licensing and complete system licensing later

How would you like to license the systems? [1-2,q] (2)

Based on what license type you want to use, enter one of the following:

1 You must have a valid license key. Enter the license key at the prompt:

```
Enter a VCS license key: [b,q,?]  
XXXX-XXXX-XXXX-XXXX-XXXX
```

If you plan to configure global clusters, enter the corresponding license keys when the installer prompts for additional licenses.

```
Do you wish to enter additional licenses? [y,n,q,b] (n) y
```

2 The keyless license option enables you to install VCS without entering a key. However, to ensure compliance, keyless licensing requires that you manage the systems with a management server.

For more information, go to the following website:

<http://go.symantec.com/sfhakeyless>

Note that this option is the default.

If you plan to set up global clusters, enter **y** at the prompt.

```
Would you like to enable the Global Cluster Option? [y,n,q] (n) y
```

The installer registers the license and completes the installation process.

8 To configure VCS, enter **y** at the prompt. You can also configure VCS later.

```
Would you like to configure VCS on galaxy nebula [y,n,q] (n) n
```

See “[Overview of tasks for VCS configuration using installvcs program](#)” on page 109.

9 Enter **y** at the prompt to send the installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] (y) **y**

The installer provides an option to collect data about the installation process each time you complete an installation, upgrade, configuration, or uninstall of the product. The installer transfers the contents of the install log files to an internal Symantec site. The information is used only to gather metrics about how you use the installer. No personal customer data is collected, and no information will be shared by any other parties. Information gathered may include the product and the version installed or upgraded, how many systems were installed, and the time spent in any section of the install process.

- 10** After the installation, note the location of the installation log files, the summary file, and the response file for future reference.

The files provide the useful information that can assist you with the configuration and can also assist future configurations.

summary file	Lists the RPMs that are installed on each system.
log file	Details the entire installation.
response file	Contains the installation information that can be used to perform unattended or automated installations on other systems. See “Installing VCS using response files” on page 161.

Preparing to configure VCS

This chapter includes the following topics:

- [Preparing to configure the clusters in secure mode](#)
- [About configuring VCS clusters for data integrity](#)
- [Setting up the CP server](#)

Preparing to configure the clusters in secure mode

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

If you want to enable or disable AT in a cluster that is online, run the following command:

```
# /opt/VRTS/install/installvcs -security
```

See the *Veritas Cluster Server Administrator's Guide* for instructions.

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

- A system in your enterprise that serves as root broker (RB).
You can either use an external system as root broker, or use one of the cluster nodes as root broker.
 - To use an external root broker, identify an existing root broker system in your enterprise or install and configure root broker on a stable system. See [“Installing the root broker for the security infrastructure”](#) on page 81.
 - To use one of the cluster nodes as root broker, the installer does not require you to do any preparatory tasks.
When you configure the cluster in secure mode using the `installvcs` program, choose the automatic mode and choose one of the nodes for the installer to configure as root broker.

Symantec recommends that you configure a single root broker system for your entire enterprise. If you use different root broker systems, then you must establish trust between the root brokers. For example, if the management server and the cluster use different root brokers, then you must establish trust.

- For external root broker, an authentication broker (AB) account for each node in the cluster is set up on the root broker system.
See [“Creating authentication broker accounts on root broker system”](#) on page 82.
- The system clocks of the external root broker and authentication brokers must be in sync.

The `installvcs` program provides the following configuration modes:

Automatic mode	The external root broker system must allow rsh or ssh passwordless login to use this mode.
Semi-automatic mode	This mode requires encrypted files (BLOB files) from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See “Setting up inter-system communication” on page 55.
Manual mode	This mode requires <code>root_hash</code> file and the root broker information from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See “Setting up inter-system communication” on page 55.

[Figure 5-1](#) depicts the flow of configuring VCS cluster in secure mode.

Figure 5-1 Workflow to configure VCS cluster in secure mode

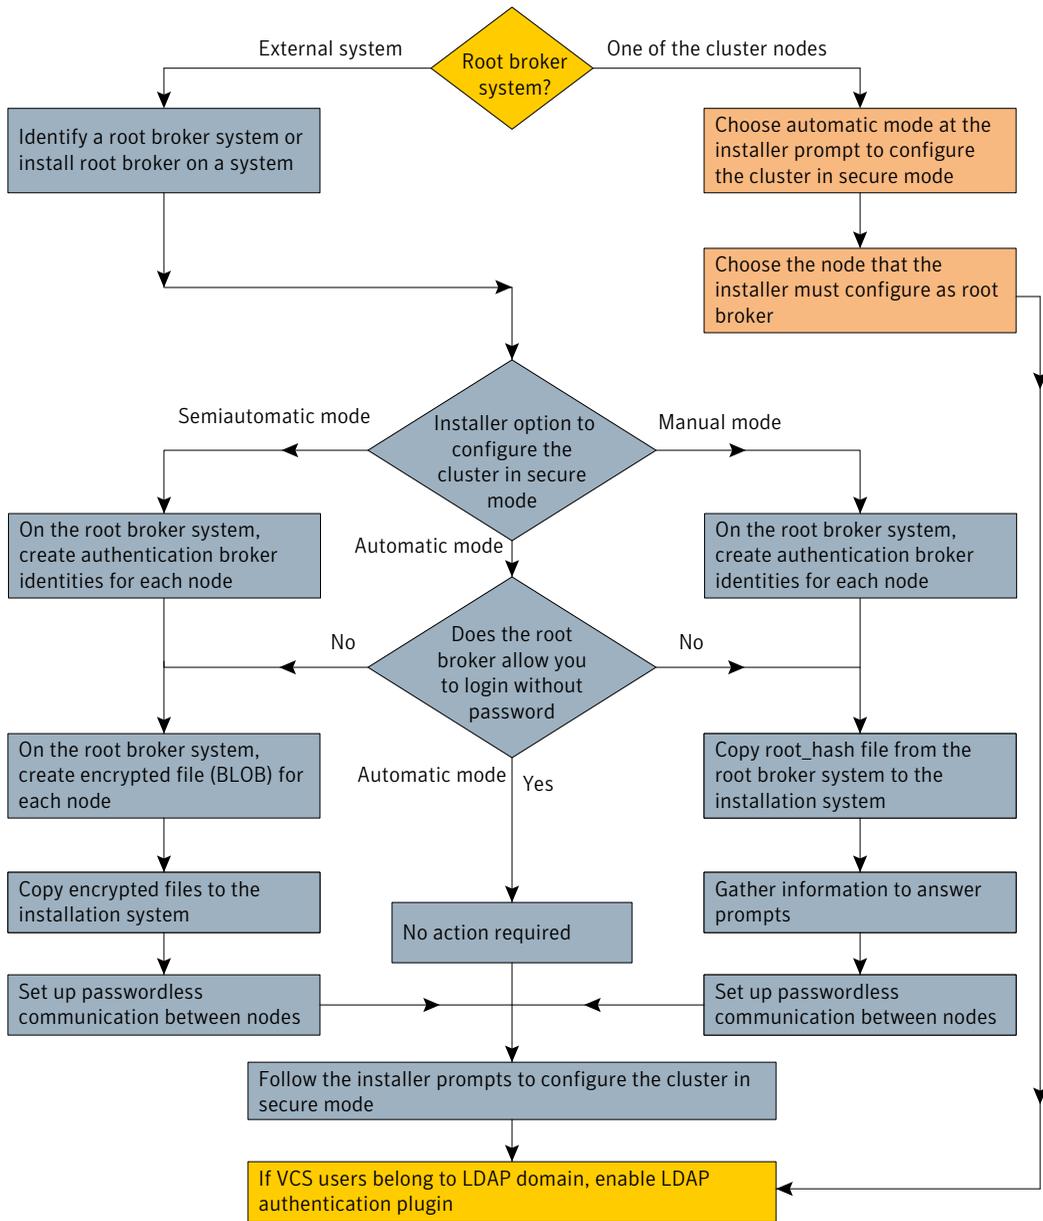


Table 5-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform. These preparatory tasks apply only when you use an external root broker system for the cluster.

Table 5-1 Preparatory tasks to configure a cluster in secure mode (with an external root broker)

Tasks	Who performs this task
<p>Decide one of the following configuration modes to set up a cluster in secure mode:</p> <ul style="list-style-type: none"> ■ Automatic mode ■ Semi-automatic mode ■ Manual mode 	VCS administrator
<p>Install the root broker on a stable system in the enterprise.</p> <p>See “Installing the root broker for the security infrastructure” on page 81.</p>	AT administrator
<p>To use the semi-automatic mode or the manual mode, on the root broker system, create authentication broker accounts for each node in the cluster.</p> <p>See “Creating authentication broker accounts on root broker system” on page 82.</p> <p>AT administrator requires the following information from the VCS administrator:</p> <ul style="list-style-type: none"> ■ Node names that are designated to serve as authentication brokers ■ Password for each authentication broker 	AT administrator
<p>To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator.</p> <p>See “Creating encrypted files for the security infrastructure” on page 83.</p> <p>AT administrator requires the following additional information from the VCS administrator:</p> <ul style="list-style-type: none"> ■ Administrator password for each authentication broker Typically, the password is the same for all nodes. 	AT administrator
<p>To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.</p>	AT administrator
<p>Copy the files that are required to configure a cluster in secure mode to the system from where you plan to install and configure VCS.</p> <p>See “Preparing the installation system for the security infrastructure” on page 85.</p>	VCS administrator

Installing the root broker for the security infrastructure

Install the root broker only if you plan to use AT to configure the cluster in secure mode. You can use a system outside the cluster or one of the systems within the cluster as root broker. If you plan to use an external broker, 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 also identify an existing root broker system in the data center to configure the cluster in secure mode. The root broker system can run AIX, HP-UX, Linux, or Solaris operating system.

See Symantec Product Authentication Service documentation for more information.

See [“About Symantec Product Authentication Service \(AT\)”](#) on page 29.

To install the root broker

- 1 Mount the product disc and start the installer.

```
# ./installer
```

- 2 From the Task Menu, choose I for "Install a Product."
- 3 From the displayed list of products to install, choose: Symantec Product Authentication Service (AT).
- 4 Enter **y** to agree to the End User License Agreement (EULA).
- 5 Enter 2 to install the recommended packages.
- 6 Enter the name of the system where you want to install the Root Broker.

```
Enter the operating system system names separated by spaces: venus
```

- 7 Review the output as the installer does the following:
 - Checks to make sure that VCS supports the operating system
 - Checks if the RPMs are already on the system.

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

- 8 Review the output as the installer installs the root broker on the system.
- 9 After the installation, configure the root broker.

- 10 Select to configure the root broker from the three choices that the installer presents:

```
1) Root+AB Mode
2) Root Mode
3) AB Mode
```

```
Enter the mode in which you would like AT to be configured? [1-3,q] 2
```

```
Do you want the installer to do cluster configuration? [y,n,q] (n) n
```

- 11 Press Enter to continue and review the output as the installer starts the Authentication Service.

Creating authentication broker accounts on root broker system

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

To create authentication broker accounts on root broker system

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

```
venus> # vssat showalltrustedcreds
```

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

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

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

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

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

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

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

"Failed To Get Attributes For Principal"

Proceed to step 3.

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

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

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

Creating encrypted files for the security infrastructure

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

To create encrypted files

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

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

```
venus> # vssat showbrokerhash
```

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

```
venus> # vssat showalltrustedcreds
```

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

identity	<p>The value for the authentication broker identity, which you provided to create authentication broker principal on the root broker system.</p> <p>This is the value for the <code>--prplname</code> option of the <code>addprpl</code> command.</p> <p>See “Creating authentication broker accounts on root broker system” on page 82.</p>
password	<p>The value for the authentication broker password, which you provided to create authentication broker principal on the root broker system.</p> <p>This is the value for the <code>--password</code> option of the <code>addprpl</code> command.</p> <p>See “Creating authentication broker accounts on root broker system” on page 82.</p>
broker_admin_password	<p>The value for the authentication broker password for Administrator account on the node. This password must be at least five characters.</p>

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

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

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

[configab]
identity=galaxy
password=password
root_domain=vx:root@venus.symantecexample.com
root_broker=venus.symantecexample.com:2821

start_broker=false
enable_pbx=false
```

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

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

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

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

For example:

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

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

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

Preparing the installation system for the security infrastructure

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

To prepare the installation system for the security infrastructure

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

Automatic mode Do the following:

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

Semi-automatic mode Do the following:

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

Manual mode

Do the following:

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

About configuring VCS clusters for data integrity

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

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

- **Broken set of private networks**
If a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects. The remaining node then takes corrective action. The failure of the private interconnects, instead of the actual nodes, presents identical symptoms and causes each node to determine its peer has departed. This situation typically results in data corruption because both nodes try to take control of data storage in an uncoordinated manner
- **System that appears to have a system-hang**
If a system is so busy that it appears to stop responding, the other nodes could declare it as dead. This declaration may also occur for the nodes that use the hardware that supports a "break" and "resume" function. When a node drops to PROM level with a break and subsequently resumes operations, the other nodes may declare the system dead. They can declare it dead even if the system later returns and begins write operations.

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

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

You can configure disk-based I/O fencing or server-based I/O fencing either manually or using the `installvcs` program.

About I/O fencing components

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

- **Data disks**—Store shared data
 See [“About data disks”](#) on page 87.
- **Coordination points**—Act as a global lock during membership changes
 See [“About coordination points”](#) on page 87.

About data disks

Data disks are standard disk devices for data storage and are either physical disks or RAID Logical Units (LUNs).

These disks must support SCSI-3 PR and must be part of standard VxVM disk groups. VxVM is responsible for fencing data disks on a disk group basis. Disks that are added to a disk group and new paths that are discovered for a device are automatically fenced.

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. Racing for control of the coordination points to fence data disks is the key to understand how fencing prevents split-brain.

The coordination points can either be disks or servers or both. Typically, a cluster must have three coordination points.

- **Coordinator disks**
 Disks that act as coordination points are called coordinator disks. Coordinator disks are three standard disks or LUNs set aside for I/O fencing during cluster reconfiguration. Coordinator disks do not serve any other storage purpose in the VCS configuration.
 You can configure coordinator disks to use Veritas Volume Manager Dynamic Multipathing (DMP) feature. Dynamic Multipathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. So, you can configure I/O fencing to use either

DMP devices or the underlying raw character devices. I/O fencing uses SCSI-3 disk policy that is either raw or dmp based on the disk device that you use. The disk policy is dmp by default.

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

■ Coordination point servers

The coordination point server (CP server) is a software solution which runs on a remote system or cluster. CP server provides arbitration functionality by allowing the VCS cluster nodes to perform the following tasks:

- Self-register to become a member of an active VCS cluster (registered with CP server) with access to the data drives
- Check which other nodes are registered as members of this activeVCS cluster
- Self-unregister from this activeVCS cluster
- Forcefully unregister other nodes (preempt) as members of this active VCS cluster

In short, the CP server functions as another arbitration mechanism that integrates within the existing I/O fencing module.

Note: With the CP server, the fencing arbitration logic still remains on the VCS cluster.

Multiple VCS clusters running different operating systems can simultaneously access the CP server. TCP/IP based communication is used between the CP server and the VCS clusters.

About I/O fencing configuration files

[Table 5-2](#) lists the I/O fencing configuration files.

Table 5-2 I/O fencing configuration files

File	Description
/etc/sysconfig/vxfen	<p>This file stores the start and stop environment variables for I/O fencing:</p> <ul style="list-style-type: none"> ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: <ul style="list-style-type: none"> 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: <ul style="list-style-type: none"> 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of VCS configuration.</p> <p>If you manually configured VCS, you must make sure to set the values of these environment variables to 1.</p>
/etc/vxfendg	<p>This file includes the coordinator disk group information.</p> <p>This file is not applicable for server-based fencing.</p>

Table 5-2 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfenmode	<p>This file contains the following parameters:</p> <ul style="list-style-type: none"> ■ vxfen_mode <ul style="list-style-type: none"> ■ scsi3—For disk-based fencing ■ customized—For server-based fencing ■ disabled—To run the I/O fencing driver but not do any fencing operations. ■ vxfen_mechanism <p>This parameter is applicable only for server-based fencing. Set the value as cps.</p> ■ scsi3_disk_policy <ul style="list-style-type: none"> ■ dmp—Configure the vxfen module to use DMP devices The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp. ■ raw—Configure the vxfen module to use the underlying raw character devices <p>Note: You must use the same SCSI-3 disk policy on all the nodes.</p> ■ security <p>This parameter is applicable only for server-based fencing.</p> <p>1—Indicates that Symantec Product Authentication Service is used for CP server communications. This setting is the default.</p> <p>0—Indicates that communication with the CP server is in non-secure mode.</p> <p>Note: The CP server and the VCS clusters must have the same security setting.</p> ■ List of coordination points <p>This list is required only for server-based fencing configuration.</p> <p>Coordination points in a server-based fencing can include coordinator disks, CP servers, or a mix of both. If you use coordinator disks, you must create a coordinator disk group with the coordinator disk names.</p> <p>Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points.</p>

Table 5-2 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	<p>When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfermode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.</p> <p>Note: The /etc/vxfentab file is a generated file; do not modify this file.</p> <p>For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:</p> <ul style="list-style-type: none"> ■ Raw disk: <ul style="list-style-type: none"> /dev/sdx /dev/sdy /dev/sdz ■ DMP disk: <ul style="list-style-type: none"> /dev/vx/rdump/sdx /dev/vx/rdump/sdy /dev/vx/rdump/sdz <p>For server-based fencing, the /etc/vxfentab file also includes the security settings information.</p>

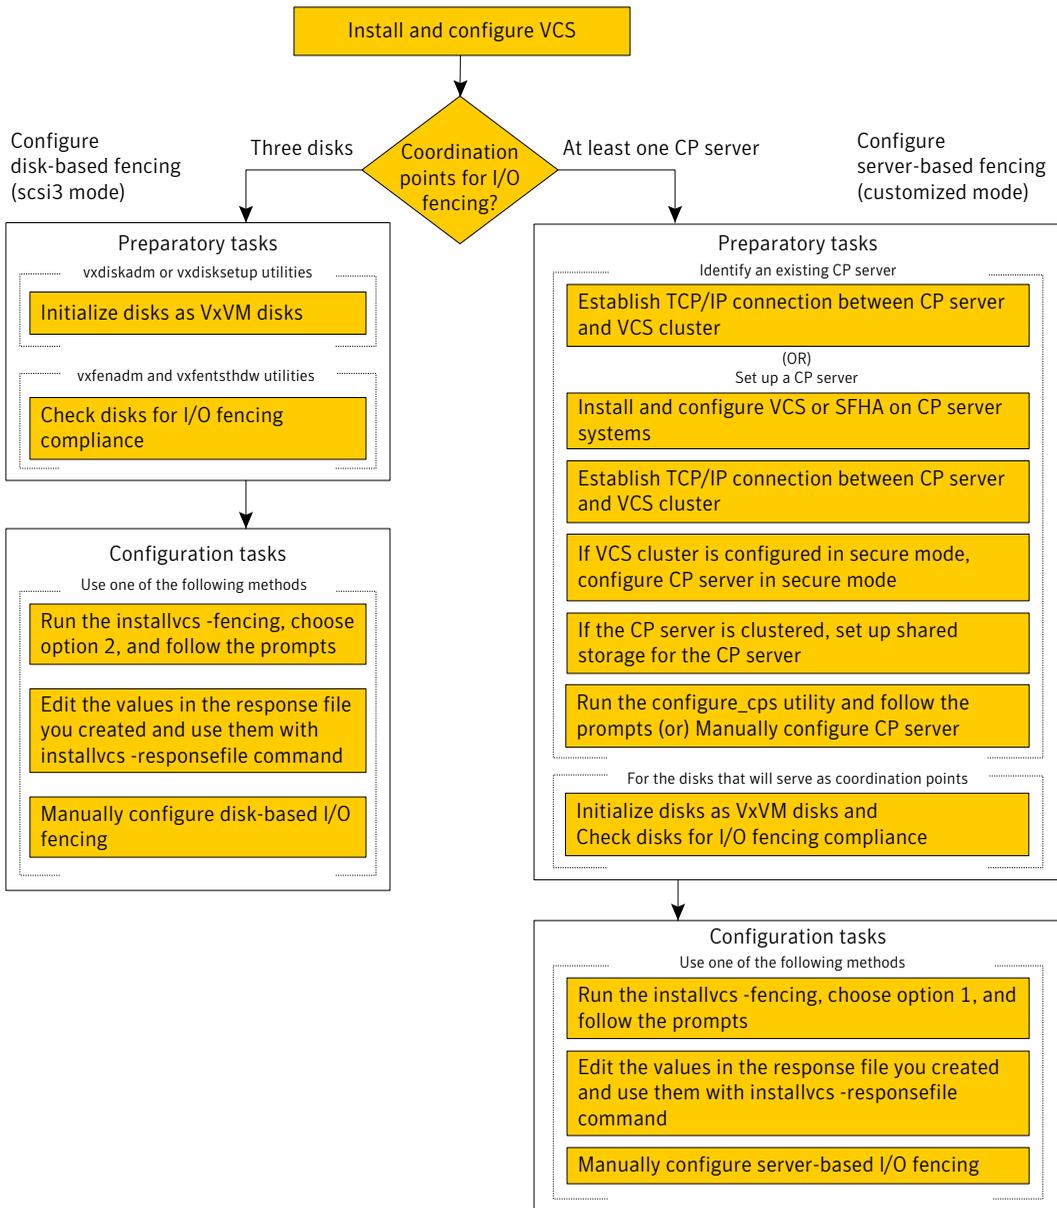
About planning to configure I/O fencing

After you configure VCS with the installer, the installer starts VCS with I/O fencing in disabled mode. To use I/O fencing in the cluster for data integrity, you must configure I/O fencing.

You can configure either disk-based I/O fencing or server-based I/O fencing. If your enterprise setup has multiple clusters that use VCS for clustering, Symantec recommends you to configure server-based I/O fencing. After you perform the preparatory tasks, you can use the installvcs program to configure I/O fencing. You can also use response files or manually configure I/O fencing.

Figure 5-2 illustrates a high-level flowchart to configure I/O fencing for the VCS cluster.

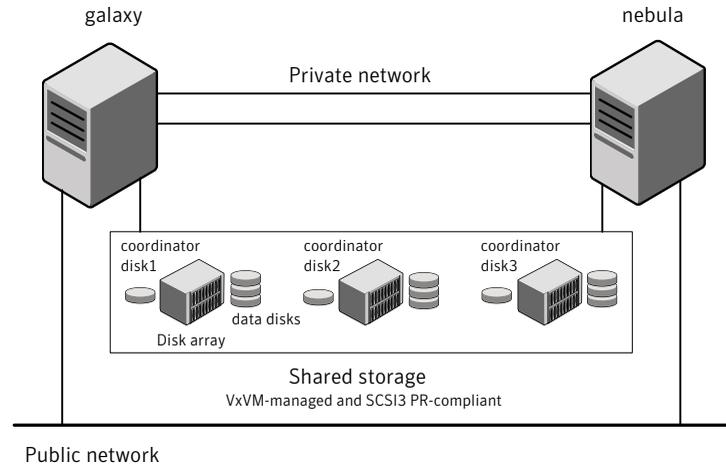
Figure 5-2 Workflow to configure I/O fencing



Typical VCS cluster configuration with disk-based I/O fencing

Figure 5-3 displays a typical VCS configuration with two nodes and shared storage. The configuration uses three coordinator disks for I/O fencing.

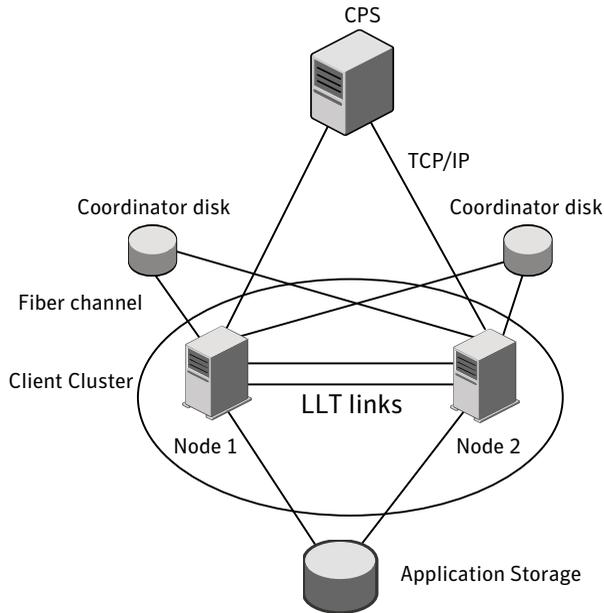
Figure 5-3 Typical VCS cluster configuration with disk-based I/O fencing



Typical VCS cluster configuration with server-based I/O fencing

Figure 5-4 displays a configuration using a VCS cluster (with two nodes), a single CP server, and two coordinator disks. The nodes within the VCS cluster are connected to and communicate with each other using LLT links.

Figure 5-4 CP server, VCS cluster, and coordinator disks



Recommended CP server configurations

This section discusses the following recommended CP server configurations:

- A CP server configuration where multiple VCS clusters use 3 CP servers as their coordination points
- A CP server configuration where multiple VCS clusters use a single CP server and multiple pairs of coordinator disks (2) as their coordination points

Note: Although the recommended CP server configurations use three coordination points, three or more odd number of coordination points may be used for I/O fencing. In a configuration where multiple VCS clusters share a common set of CP server coordination points, the VCS VCS cluster as well as the CP server use a Universally Unique Identifier(UUID) to uniquely identify a VCS cluster.

Figure 5-5 displays a configuration using a single CP server that is connected to multiple VCS clusters with each VCS cluster also using two coordinator disks.

Figure 5-5 Single CP server connecting to multiple VCS clusters

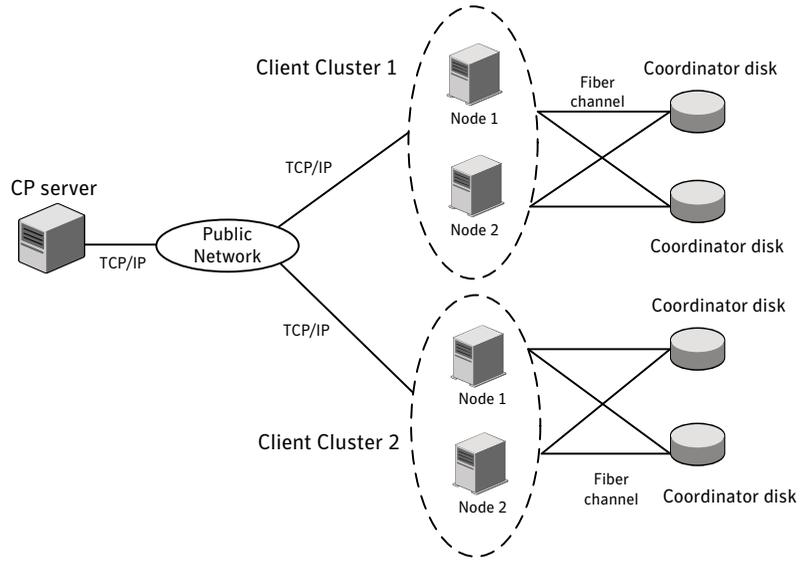
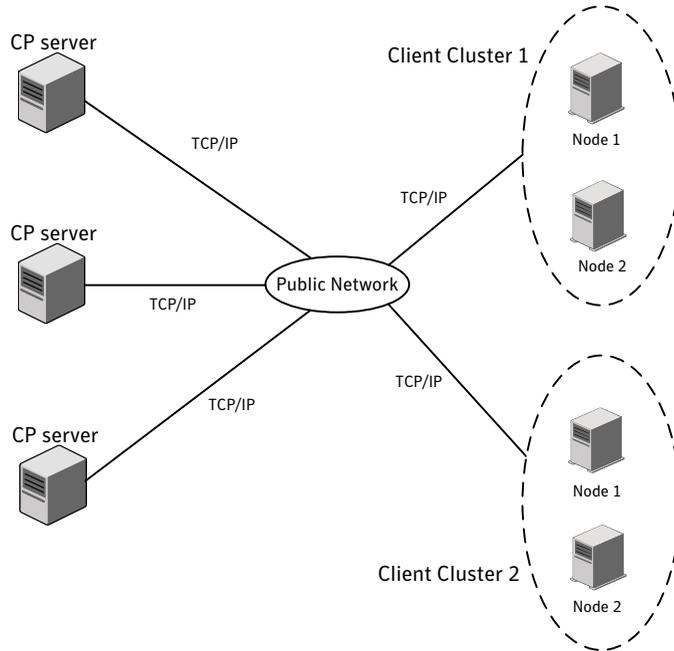


Figure 5-6 displays a configuration using 3 CP servers that are connected to multiple VCS clusters.

Figure 5-6 Three CP servers connecting to multiple VCS clusters



For detailed deployment diagrams for server-based fencing:

See [“Configuration diagrams for setting up server-based I/O fencing”](#) on page 405.

Setting up the CP server

The following preparations must be taken before running the configuration utility.

To prepare to configure the CP server

- 1 Ensure that VCS is installed and configured for hosting CP server on a single node VCS cluster, or that SFHA is installed and configured for hosting CP server on an SFHA cluster.

Refer to the appropriate VCS or SFHA installation and configuration guide to configure the VCS or SFHA cluster using the installer.

- 2 If the CP server is hosted on an SFHA cluster, configure fencing in enabled mode during the SFHA configuration using either the installer or manually.

- 3 Decide if you want to secure the communication between the CP server and VCS clusters using the Symantec Product Authentication Service (AT).

Symantec recommends setting up security for the CP server and VCS cluster communications.

For information about configuring security on the CP server:
See [“Configuring security on the CP server”](#) on page 98.
- 4 Choose a name for the CP server.

The CP server name should not contain any special characters.
- 5 Choose a port number for the CP server.

Allocate a TCP/IP port for use by the CP server.

The default port number is 14250. Alternatively, the user can specify any other valid port from the following valid port range: 49152-65535.
- 6 If CP server is hosted on an SFHA cluster, then set up shared storage for the CP server database.

For information about setting up shared storage for the CP server database:
See [“Setting up shared storage for the CP server database”](#) on page 99.
- 7 Choose a valid virtual IP address, network interface, and netmask for the CP server.

Installing the CP server using the installer

This section describes how to use the installer to install all CP server-related packages on a single node or SFHA cluster hosting the CP server. This installation procedure also installs the packages that are required to provide secure communication between the VCS cluster and CP server.

The installation is performed from the common VCS or SFHA DVD, so that the user can proceed to configure CP server on that node or cluster.

The following procedure describes how to install CP server on a single node or cluster.

To install CP server using the VCS installer on a single node or the SFHA installer on an SFHA cluster

- 1 Review the CP server hardware and networking requirements, and set up the CP server hardware and network.
- 2 Establish network connections between the CP server(s) and the VCS clusters through the TCP/IP network. This step requires that you have valid IP addresses, hostnames, and netmasks set up for the CP servers.

- 3 For installing CP server on a single node:
 - Install VCS 5.1 onto the system where you are installing the CP server. Installing VCS 5.1 also installs CP server on the system.

When installing VCS 5.1, be sure to select the complete installation option and not the minimum package installation option. The VRTScps package is only part of the complete installation.
- 4 For installing CP server to be hosted on an SFHA cluster:
 - Install SFHA 5.1 onto each system where you are installing CP server to be hosted on a cluster. Installing SFHA 5.1 also installs CP server on the system. Refer to the *Veritas Storage Foundation™ and High Availability Installation Guide* for instructions on installing SFHA 5.1.

When installing SFHA 5.1, be sure to select the complete installation option and not the minimum package installation option. The VRTScps package is only part of the complete installation.
- 5 Proceed to configure the single node or SFHA cluster for CP server.

Configuring security on the CP server

This section describes configuring security on the CP server. You must configure security on the CP server only if you want to secure the communication between the CP server and the VCS cluster.

Note: If Symantec™ Product Authentication Service has already been configured during VCS configuration, skip this section.

The CP server cluster needs to be configured for security with Symantec™ Product Authentication Service using the installer (`installvcs -security` command). This step secures the HAD communication, besides ensuring that the service group configuration for making the authentication broker (essentially VxSS service group) is highly available.

For additional information:

See [“Preparing to configure the clusters in secure mode”](#) on page 77.

Setting up shared storage for the CP server database

To set up shared storage for the CP server database

- 1 Create a disk group containing the disk(s). Two disks are required for creating a mirrored volume.

For a command example:

```
# vxvg init cps_dg disk1 disk2
```

- 2 Import the disk group if it's not already imported.

For a command example:

```
# vxvg import cps_dg
```

- 3 Create a mirrored volume over the disk group.

Symantec recommends a mirrored volume for hosting the CP server database.

For a command example:

```
# vxassist -g cps_dg make cps_vol volume size layout=mirror
```

- 4 Create a file system over the volume.

The CP server configuration utility only supports vxfs file system type. If you use an alternate file system, then configure CP server manually.

Symantec recommends the vxfs file system type.

If your CP server runs on a Solaris system, enter the following command:

```
# mkfs -F vxfs /dev/vx/rdmp/cps_dg/cps_volume
```

If your CP server runs on a Linux system, enter the following command::

```
# mkfs -t vxfs /dev/vx/rdmp/cps_dg/cps_volume
```

Configuring the CP server using the configuration utility

Ensure that the preparatory steps for configuring a CP server have been performed.

The configuration utility can be used to configure the CP server. The configuration utility is part of the VRTScps package. The following procedure describes how to configure CP server on a single node VCS cluster or on an SFHA cluster.

If the CP server is being hosted on SFHA cluster, ensure that passwordless ssh/rsh is configured on the cluster nodes.

Note: CP server is supported on Linux and Solaris operating systems only.

To configure hosting for the CP server on a single node VCS cluster or on an SFHA cluster

- 1 Ensure that the tasks required to prepare the CP server for configuration are completed:

See “[Setting up the CP server](#)” on page 96.

- 2 To run the configuration script, enter the following command on the node where you want to configure the CP server:

```
# /opt/VRTScps/bin/configure_cps.pl
```

If the CP server is being configured on SFHA cluster, the utility uses ssh by default for communication with the other nodes.

Use the -n option for using rsh communication.

- 3 The Veritas Coordination Point Server Configuration utility appears with an option menu and note.

```
VERITAS COORDINATION POINT SERVER CONFIGURATION UTILITY
```

```
=====
```

```
Select one of the following:
```

```
[1] Configure Coordination Point Server on single node VCS system
```

```
[2] Configure Coordination Point Server on SFHA cluster
```

```
[3] Unconfigure Coordination Point Server
```

```
Enter the option:
```

```
NOTE: For configuring CP server on SFHA cluster, the CP server database should reside on shared storage. Please refer to documentation for information on setting up of shared storage for CP server database.
```

- 4 Depending upon your configuration, select either option 1 or option 2. The configuration utility then runs the following preconfiguration checks:

- Checks to see if a single node VCS cluster or an SFHA cluster is running with the supported platform. (only Solaris and Linux platforms are supported)
- Checks to see if the CP server is already configured on the system. If the CP server is already configured, then the configuration utility informs the user and requests that the user unconfigure the server before trying to configure it.
- Checks to see if VCS is installed and configured on the system. The CP server requires VCS to be installed and configured before its configuration.

5 Enter the name of the CP server.

For example:

```
Enter the name of the CP Server: mycps1.symantecexample.com
```

6 Enter a valid Virtual IP address on which the CP server process should depend on.

For example:

```
Enter a valid Virtual IP address on which  
the CP Server process should depend on:  
10.209.83.85
```

7 Enter the CP server port number or press Enter to accept the default value (14250).

For example:

```
Enter a port number in range [49152 - 65535], or  
press <enter> for default port (14250)
```

- 8 Choose if the communication between the VCS clusters and the CP server has to be made secure.

This requires Symantec Product Authentication Service to be configured on the CP server.

For example:

```
Veritas recommends secure communication between the CP server and application clusters. Enabling security requires Symantec Product Authentication Service to be installed and configured on the cluster.
```

```
Do you want to enable Security for the communications? (y/n)  
(Default:y) :
```

The above note indicates that Symantec Product Authentication Service (AT) must be configured on the CP server cluster, if you want to enable security for communication between the VCS clusters and CP server.

If security is chosen but not already configured on the system, then the script immediately exits. You can configure security with VCS and later rerun the configuration script.

Symantec recommends enabling security for communication between CP server and the VCS clusters.

For information about configuring security on the CP server:

See [“Configuring security on the CP server”](#) on page 98.

- 9** Enter the absolute path of the CP server database or press Enter to accept the default value (/etc/VRTScps/db).

Depending upon your configuration, you are presented with one of the following examples.

For a single node VCS configuration for CP server example:

CP Server uses an internal database to store the client information.

Note: As the CP Server is being configured on a single node VCS, the database can reside on local file system.

Enter absolute path of the database (Default:/etc/VRTScps/db):

For configuring CP server on an SFHA cluster example:

CP Server uses an internal database to store the client information.

Note: As the CP Server is being configured on SFHA cluster, the database should reside on shared storage with vxfs file system.

Please refer to documentation for information on setting up of shared storage for CP server database.

Enter absolute path of the database (Default:/etc/VRTScps/db):

- 10** Review the displayed CP server configuration information.

If you want to change the current configuration, press b. If you want to continue, press Enter.

For example:

Following is the CP Server configuration information:

```

-----
(a) CP Server Name: mycps1.symantecexample.com
(b) CP Server Virtual IP: 10.209.83.85
(c) CP Server Port: 14250
(d) CP Server Security : 1
(e) CP Server Database Dir: /etc/VRTScps/db
-----

```

Press b if you want to change the configuration, <enter> to continue :

- 11 The configuration utility proceeds with the configuration process. A `vxcps.conf` configuration file is created. Depending upon your configuration, one of the following messages appear.

For a single node VCS configuration for CP server example:

```
Successfully generated the /etc/vxcps.conf configuration file.  
Successfully created directory /etc/VRTScps/db.
```

```
Configuring CP Server Service Group (CPSSG) for this cluster  
-----
```

```
NOTE: Please ensure that the supplied network interface is a  
public NIC
```

For configuring CP server on an SFHA cluster example:

```
Successfully generated the /etc/vxcps.conf  
configuration file.  
Successfully created directory /etc/VRTScps/db.  
Creating mount point /etc/VRTScps/db on  
mycps1.symantecexample.com.  
Copying configuration file /etc/vxcps.conf to  
mycps1.symantecexample.com
```

```
Configuring CP Server Service Group (CPSSG) for this cluster  
-----
```

- 12 For configuring CP server on an SFHA cluster, you are prompted to use the same NIC name for the virtual IP on all the systems in the cluster. For example:

```
Is the name of NIC for virtual IP 10.209.83.85 same on all the systems?  
[y/n] : y
```

```
NOTE: Please ensure that the supplied network interface is a  
public NIC
```

13 Enter a valid interface for virtual IP address for the CP server process.

For a single node VCS configuration for CP server example:

```
Enter a valid network interface for virtual IP 10.209.83.85
on mycps1.symantecexample.com: bge0
```

For configuring CP server on an SFHA cluster example:

```
Enter a valid interface for virtual IP 10.209.83.85
on all the systems : bge0
```

14 Enter the netmask for the virtual IP address. For example:

```
Enter the netmask for virtual IP 10.209.83.85 :
255.255.252.0
```

15 For configuring CP server on an SFHA cluster, enter the name of the disk group for the CP server database. For example:

```
Enter the name of diskgroup for cps database :
cps_dg
```

16 For configuring CP server on an SFHA cluster, enter the name of the volume that is created on the above disk group. For example:

```
Enter the name of volume created on diskgroup cps_dg :
cps_volume
```

- 17** After the configuration process has completed, a success message appears. For example:

```
Successfully added the CPSSG service group to  
VCS configuration. Bringing the CPSSG service  
group online. Please wait...
```

```
The Veritas Coordination Point Server has been  
configured on your system.
```

- 18** Run the `hagrp -state` command to ensure that the CPSSG service group has been added.

For example:

```
# hagrp -state CPSSG
```

#Group	Attribute	System	Value
CPSSG	State	mycps1.symantecexample.com	ONLINE

It also generates the configuration file for CP server (`/etc/vxcps.conf`).

The configuration utility adds the `vxcpserv` process and other resources to the VCS configuration in the CP server service group (CPSSG).

For information about the CPSSG, refer to the *Veritas Cluster Server Administrator's Guide*.

In addition, the `main.cf` samples contain details about the `vxcpserv` resource and its dependencies:

See [“Sample configuration files for CP server”](#) on page 369.

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

To manually configure the CP server

- 1 Ensure that the CP server preparation procedures have been performed:
- 2 Stop VCS on each node by using the following command:

```
# hastop -local
```

- 3 Edit the main.cf to add the CPSSG service group on any node. Use the CPSSG service group in the main.cf as an example:

See [“Sample configuration files for CP server”](#) on page 369.

Customize the resources under the CPSSG service group as per your configuration.

- 4 Verify the main.cf using the following command:

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

If successfully verified, proceed to copy this main.cf to all other cluster nodes.

- 5 Create the vxcps.conf file using the sample configuration file provided at /etc/vxcps/vxcps.conf.sample.

Confirm that security for communication has been established between the application clusters and the CP server. If security is to be disabled, set the security parameter to 0 in /etc/vxcps.conf file. If security parameter is set to 1 and security is not already configured, then CP server start-up fails. You can configure security and set security parameter to 1 in /etc/vxcps.conf file.

For more information about configuring security on the CP server:

See [“Configuring security on the CP server”](#) on page 98.

Symantec recommends enabling security for communication between CP server and the application clusters.

- 6 Start VCS on all the cluster nodes.

Enter the following command:

```
# hstart
```

- 7 Verify that the CP server service group (CPSSG) is online.

Enter the following command:

```
# hagrps -state CPSSG
```

Output similar to the following should appear:

```
# Group Attribute System Value
CPSSG State mycps1.symantecexample.com |ONLINE|
```

Verifying the CP server configuration

During the CP server configuration process, individual files are updated on the node or nodes hosting the CP server. After your configuration, you should check for the following files on your CP server node or nodes:

- `/etc/vxcps.conf` (CP server configuration file)
- `/etc/VRTSvcs/conf/config/main.cf`
- `/etc/VRTSvcs/db` (default location for CP server database)

Additionally, use the `cpsadm` command to check if the `vxcpserv` process is listening on the configured Virtual IP. For example, run the following command:

```
# cpsadm -s cp_server -a ping_cps
```

where *cp_server* is the virtual IP/ virtual hostname of the CP server.

Configuring VCS

This chapter includes the following topics:

- [Overview of tasks for VCS configuration using installvcs program](#)
- [Starting the software configuration](#)
- [Specifying systems for configuration](#)
- [Configuring the basic cluster](#)
- [Configuring the virtual IP of the cluster](#)
- [Configuring the cluster in secure mode](#)
- [Adding VCS users](#)
- [Configuring SMTP email notification](#)
- [Configuring SNMP trap notification](#)
- [Configuring global clusters](#)
- [Completing the VCS configuration](#)
- [Verifying and updating licenses on the system](#)

Overview of tasks for VCS configuration using installvcs program

Tasks involved in configuring VCS are as follows:

- Start the software configuration
See [“Starting the software configuration”](#) on page 110.
- Specify the systems where you want to configure VCS

See [“Specifying systems for configuration”](#) on page 111.

- Configure the basic cluster
See [“Configuring the basic cluster”](#) on page 111.
- Configure virtual IP address of the cluster (optional)
See [“Configuring the virtual IP of the cluster”](#) on page 113.
- Configure the cluster in secure mode (optional)
See [“Configuring the cluster in secure mode”](#) on page 115.
- Add VCS users (required if you did not configure the cluster in secure mode)
See [“Adding VCS users”](#) on page 118.
- Configure SMTP email notification (optional)
See [“Configuring SMTP email notification”](#) on page 118.
- Configure SNMP email notification (optional)
See [“Configuring SNMP trap notification”](#) on page 120.
- Configure global clusters (optional)
You must have enabled Global Cluster Option when you installed VCS.
See [“Configuring global clusters”](#) on page 122.
- Complete the software configuration
See [“Completing the VCS configuration”](#) on page 123.

Starting the software configuration

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

To configure VCS using the product installer

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

```
# ./installer
```

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

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

To configure VCS using the `installvcs` program

- 1 Confirm that you are logged in as the superuser.
- 2 Start the `installvcs` program.

```
# /opt/VRTS/install/installvcs -configure
```

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

Specifying systems for configuration

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

To specify system names for installation

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

```
Enter the system names separated by spaces: [q,?]
(galaxy) galaxy nebula
```

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

The installer does the following tasks:

- Checks that the local node running the installer can communicate with remote nodes
If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases.
- Makes sure the systems use the proper operating system
- Checks whether VCS is installed
- Exits if VCS 5.1 is not installed

Configuring the basic cluster

Enter the cluster information when the installer prompts you.

To configure the cluster

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

```
Enter the unique cluster name: [q,?] clus1
Enter a unique Cluster ID number between 0-65535: [b,q,?] (0) 7
```

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

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

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

```
Is eth1 a bonded NIC? [y,n,q] (n)
```

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 link on galaxy:
[b,q,?] eth1
eth1 has an IP address configured on it. It could be a
public NIC on galaxy.
Are you sure you want to use eth1 for the first private
heartbeat link? [y,n,q,b,?] (n) y
Is eth1 a bonded NIC? [y,n,q] (n)
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on galaxy:
[b,q,?] eth2
eth2 has an IP address configured on it. It could be a
public NIC on galaxy.
Are you sure you want to use eth2 for the first private
heartbeat link? [y,n,q,b,?] (n) y
Is eth2 a bonded NIC? [y,n,q] (n)
Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n)
Do you want to configure an additional low priority heartbeat
link? [y,n,q,b,?] (n)
```

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

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

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

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

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

Configuring the virtual IP of the cluster

You can configure the virtual IP of the cluster to use to connect to the Cluster Manager (Java Console) or to specify in the RemoteGroup resource.

See the *Veritas Cluster Server Administrator's Guide* for information on the Cluster Manager.

See the *Veritas Cluster Server Bundled Agents Reference Guide* for information on the RemoteGroup agent.

To configure the virtual IP of the cluster

- 1 Review the required information to configure the virtual IP of the cluster.
- 2 To configure virtual IP, enter **y** at the prompt.
- 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 galaxy: eth0
Enter the NIC for Virtual IP of the Cluster to use on galaxy:
[b,q,?] (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.

You can enter either an IPv4 address or an IPv6 address.

- For IPv4:
- Enter the virtual IP address.

```
Enter the Virtual IP address for the Cluster:  
[b,q,?] 192.168.1.16
```

- Confirm the default netmask or enter another one:

```
Enter the netmask for IP 192.168.1.16: [b,q,?]  
(255.255.240.0)
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

```
NIC: eth0  
IP: 192.168.1.16  
Netmask: 255.255.240.0
```

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

- For IPv6
- Enter the virtual IP address.

```
Enter the Virtual IP address for the Cluster:  
[b,q,?] 2001:454e:205a:110:203:baff:feee:10
```

- Enter the prefix for the virtual IPv6 address you provided. For example:

```
Enter the Prefix for IP  
2001:454e:205a:110:203:baff:feee:10: [b,q,?] 64
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

```
NIC: eth0  
IP: 2001:454e:205a:110:203:baff:feee:10  
Prefix: 64
```

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

Configuring the cluster in secure mode

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

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

See “[Preparing to configure the clusters in secure mode](#)” on page 77.

To configure the cluster in secure mode

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

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

- If you want to configure the cluster in secure mode, make sure you meet the prerequisites and enter **y**.
- If you do not want to configure the cluster in secure mode, enter **n**. You must add VCS users when the configuration program prompts. See “[Adding VCS users](#)” on page 118.

- 2 Select one of the options to enable security.

Before you choose any of the options, make sure that all the nodes in the cluster can successfully ping the root broker system.

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

```
Security Menu
```

- 1) Configure security completely automatically
- 2) Provide AB credentials using BLOBs
- 3) Provide AB credentials without using BLOBs
- b) Back to previous menu

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

Option 1.
Automatic
configuration

Based on the root broker you want to use, do one of the following:

- To use an external root broker:

Enter the name of the root broker system when prompted.

Requires remote access to the root broker. Make sure that all the nodes in the cluster can successfully ping the root broker system.

Review the output as the installer verifies communication with the root broker system, checks vxatd process and version, and checks security domain.

- To configure one of the nodes as root broker:

- Press Enter at the following installer prompt:

```
If you already have an external
RB(Root Broker) installed and configured, enter
the RB name, or press Enter to skip: [b]
```

- Choose the node that the installer must configure as root and authentication broker. The installer configures the other nodes as authentication brokers.

At the installer prompt, you can choose the first node in the cluster to configure as RAB, or you can enter n to configure another node as RAB. For example:

```
Do you want to configure <galaxy> as RAB,
and other nodes as AB? [y,n,q,b] (y) n
Enter the node name which you want to
configure as RAB: nebula
```

Option 2.
Semiautomatic
configuration

Enter the path of the encrypted file (BLOB file) for each node when prompted.

Option 3.
Manual
configuration

Enter the following Root Broker information as the installer prompts you:

```
Enter root broker name: [b]
east.symantecexample.com
Enter root broker FQDN: [b]
(symantecexample.com)
symantecexample.com
Enter the root broker domain name for the
Authentication Broker's identity: [b]
root@east.symantecexample.com
Enter root broker port: [b] 2821
Enter path to the locally accessible root hash [b]
(/var/tmp/installvcs-200910221810ROA/root_hash)
/var/tmp/installvcs-200910221810ROA/root_hash
```

Enter the following Authentication Broker information as the installer prompts you for each node:

```
Enter Authentication broker's identity on
galaxy [b]
(galaxy.symantecexample.com)
galaxy.symantecexample.com
Enter the password for the Authentication broker's
identity on galaxy:
Enter Authentication broker's identity on
nebula [b]
(nebula.symantecexample.com)
nebula.symantecexample.com
Enter the password for the Authentication broker's
identity on nebula:
```

- 3 After you provide the required information to configure the cluster in secure mode, the program prompts you to configure SMTP email notification.

Note that the installer does not prompt you to add VCS users if you configured the cluster in secure mode. However, you must add VCS users later.

See the *Veritas Cluster Server Administrator's Guide* for more information.

Adding VCS users

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

To add VCS users

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

```
Do you want to set the username and/or password for the Admin user
(default username = 'admin', password='password')? [y,n,q] (n) y
Enter the user name: [b,q,?] (admin)
Enter the 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: [b,q,?] smith
Enter New Password:*****
Enter Again:*****
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [?] a
```

- 5 Enter **n** at the prompt if you have finished adding users.

```
Would you like to add another user? [y,n,q] (n)
```

- 6 Review the summary of the newly added users and confirm the information.

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

To configure SMTP email notification

- 1 Review the required information to configure the SMTP email notification.
- 2 Specify whether you want to configure the SMTP notification.

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

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

See [“Configuring SNMP trap notification”](#) on page 120.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (eth0)
Is eth0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SMTP server’s host name.

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

- Enter the email address of each recipient.

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

```
NIC: eth0
```

```
SMTP Address: smtp.example.com
```

```
Recipient: ozzie@example.com receives email for Warning or  
higher events
```

```
Recipient: harriet@example.com receives email for Error or  
higher events
```

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

Configuring SNMP trap notification

You can choose to configure VCS to send event notifications to SNMP management consoles. You need to provide the SNMP management console name to be notified and message severity levels.

Note that you can also configure the notification after installation.

Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

To configure the SNMP trap notification

- 1 Review the required information to configure the SNMP notification feature of VCS.
- 2 Specify whether you want to configure the SNMP notification.

```
Do you want to configure SNMP notification? [y,n,q] (n) y
```

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

See [“Configuring global clusters”](#) on page 122.

3 Provide information to configure SNMP trap notification.

Provide the following information:

■ Enter the NIC information.

```
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (eth0)
Is eth0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

■ Enter the SNMP trap daemon port.

```
Enter the SNMP trap daemon port: [b,q,?] (162)
```

■ Enter the SNMP console system name.

```
Enter the SNMP console system name: [b,q,?] 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,q,?] 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,q,?] 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,q,?] 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.

```
NIC: eth0
```

```
SNMP Port: 162
```

```
Console: saturn receives SNMP traps for Error or  
higher events
```

```
Console: jupiter receives SNMP traps for SevereError or  
higher events
```

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

Configuring global clusters

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster. If not, the installer proceeds to configure VCS based on the configuration details you provided. You can also run the `gcoconfig` utility in each cluster later to update the VCS configuration file for global cluster.

You can configure global clusters to link clusters at separate locations and enable wide-area failover and disaster recovery. The installer adds basic global cluster information to the VCS configuration file. You must perform additional configuration tasks to set up a global cluster.

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

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

To configure the global cluster option

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

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

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

3 Provide information to configure this cluster as global cluster.

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

If you had entered virtual IP address details, the installer discovers the values you entered. You can use the same virtual IP address for global cluster configuration or enter different values.

You can also enter an IPv6 address as a virtual IP address.

4 Verify and confirm the configuration of the global cluster.

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

```
NIC: eth0  
IP: 192.168.1.16  
Netmask: 255.255.240.0
```

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

Completing the VCS configuration

After you enter the VCS configuration information, the installer prompts to stop the VCS processes to complete the configuration process. The installer continues to create configuration files and copies them to each system. The installer also configures a cluster UUID value for the cluster at the end of the configuration. After the installer successfully configures VCS, it restarts VCS.

If you chose to configure the cluster in secure mode, the installer then does the following before it starts VCS in secure mode:

- Depending on the security mode you chose to set up Authentication Service, the installer does one of the following:
 - Creates the security principal
 - Executes the encrypted file to create security principal on each node in the cluster
- Creates the VxSS service group
- Creates the Authentication Server credentials on each node in the cluster
- Creates the Web credentials for VCS users
- Sets up trust with the root broker

To complete the VCS configuration

- 1 Press Enter at the following prompt.

```
Do you want to stop VCS processes now? [y,n,q,?] (y)
```

- 2 Review the output as the installer stops various processes and performs the configuration. The installer then restarts VCS.
- 3 Enter y at the prompt to send the installation information to Symantec.

```
Would you like to send the information about this installation  
to Symantec to help improve installation in the future? [y,n,q,?] (y) y
```

- 4 After the installer configures VCS successfully, note the location of summary, log, and response files that installer creates.

The files provide the useful information that can assist you with the configuration and can also assist future configurations.

summary file	Describes the cluster and its configured resources.
log file	Details the entire configuration.
response file	Contains the configuration information that can be used to perform secure or unattended installations on other systems. See “Configuring VCS using response files” on page 167.

Verifying the NIC configuration

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

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

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

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

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

```
No config file for NIC with MAC address 00:11:43:33:17:29
(present name eth1) found!
CPI WARNING V-9-122-1022
No config file for NIC with MAC address 00:04:23:ac:25:1f
(present name eth3) found!
```

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

Set the PERSISTENT_NAME for all the NICs.

See [“Performing preinstallation tasks”](#) on page 51.

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

Verifying and updating licenses on the system

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

Checking licensing information on the system

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

To check licensing information

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

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

- 2 Review the following output to determine the following information:

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

```
License Key           = xxx-xxx-xxx-xxx-xxx
Product Name          = Veritas Cluster Server
```

```

Serial Number          = 1249
License Type           = PERMANENT
OEM ID                 = 478

Features :=
Platform              = Linux
Version               = 5.1
Tier                  = 0
Reserved              = 0
Mode                   = VCS

```

Updating product licenses using vxlicinst

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

See [“Replacing a VCS demo license with a permanent license”](#) on page 126.

To update product licenses

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

```

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

```

Replacing a VCS demo license with a permanent license

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

To replace a demo key

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

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

This command does not shut down any running applications.

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

```

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

```

- 4 Make sure demo licenses are replaced on all cluster nodes before starting VCS.

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

- 5 Start VCS on each node:

```
# hstart
```


Configuring VCS clusters for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing using installvcs program](#)
- [Setting up server-based I/O fencing using installvcs program](#)

Setting up disk-based I/O fencing using installvcs program

You can configure I/O fencing using the `-fencing` option of the `installvcs` program.

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

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

```
# fdisk -l
```

- 2 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
```

The example specifies the CDS format:

```
# vxdisksetup -i sdr
```

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

Configuring disk-based I/O fencing using installvcs program

Note: The installer stops and starts VCS to complete I/O fencing configuration. Make sure to unfreeze any frozen VCS service groups in the cluster for the installer to successfully stop VCS.

To set up disk-based I/O fencing using the installvcs program

- 1 Start the installvcs program with `-fencing` option.

```
# /opt/VRTS/install/installvcs -fencing
```

The installvcs program starts with a copyright message and verifies the cluster information.

- 2 Confirm that you want to proceed with the I/O fencing configuration at the prompt.

The program checks that the local node running the script can communicate with remote nodes and checks whether VCS 5.1 is configured properly.

- 3 Review the I/O fencing configuration options that the program presents. Type **2** to configure disk-based I/O fencing.

```
Select the fencing mechanism to be configured in this  
Application Cluster  
[1-3,b,q] 2
```

- 4 Review the output as the configuration program checks whether VxVM is already started and is running.
 - If the check fails, configure and enable VxVM before you repeat this procedure.
 - If the check passes, then the program prompts you for the coordinator disk group information.
- 5 Choose whether to use an existing disk group or create a new disk group to configure as the coordinator disk group.

The program lists the available disk group names and provides an option to create a new disk group. Perform one of the following:

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.

The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.

- To create a new disk group, perform the following steps:
 - Enter the number corresponding to the **Create a new disk group** option. The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks. Symantec recommends to use three disks as coordination points for disk-based I/O fencing.
 - Enter the numbers corresponding to the disks that you want to use as coordinator disks.
 - Enter the disk group name.

6 Verify that the coordinator disks you chose meet the I/O fencing requirements.

You must verify that the disks are SCSI-3 PR compatible using the `vxfcntlsthdw` utility and then return to this configuration program.

See [“Checking shared disks for I/O fencing”](#) on page 132.

7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.

8 Enter the I/O fencing disk policy that you chose to use. For example:

```
Enter fencing mechanism name (raw/dmp): [b,q,?] raw
```

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
- Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information

9 Verify and confirm the I/O fencing configuration information that the installer summarizes.

10 Review the output as the configuration program does the following:

- Stops VCS and I/O fencing on each node.
- Configures disk-based I/O fencing and starts the I/O fencing process.
- Updates the VCS configuration file `main.cf` if necessary.

- Copies the `/etc/vxfenmode` file to a date and time suffixed file `/etc/vxfenmode-date-time`. This backup file is useful if any future fencing configuration fails.
 - Starts VCS on each node to make sure that the VCS is cleanly configured to use the I/O fencing feature.
- 11 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.
 - 12 Configure the Coordination Point agent to monitor the coordinator disks.
See [“Configuring Coordination Point agent to monitor coordination points”](#) on page 222.

Checking shared disks for I/O fencing

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

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

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

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

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

- Verifying the Array Support Library (ASL)
See [“Verifying Array Support Library \(ASL\)”](#) on page 133.
- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 133.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfentsthdw utility”](#) on page 134.

Verifying Array Support Library (ASL)

Make sure that the Array Support Library (ASL) for the array that you add is installed.

To verify Array Support Library (ASL)

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

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

- 2 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	PID
libvxhitachi.so	HITACHI	DF350, DF400, DF400F, DF500, DF500F
libvxxpl281024.so	HP	All
libvxxpl2k.so	HP	All
libvxdds2a.so	DDN	S2A 9550, S2A 9900, S2A 9700
libvxpurple.so	SUN	T300
libvxxiotechE5k.so	XIOTECH	ISE1400
libvxcopan.so	COPANSYS	8814, 8818
libvxibm8k.so	IBM	2107

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

```
# vxdisk scandisks
```

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

Verifying that the nodes have access to the same disk

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

To verify that the nodes have access to the same disk

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed VCS.
- 2 Ensure that both nodes are connected to the same disk during the testing. Use the `vxfenadm` command to verify the disk serial number.

```
vxfenadm -i diskpath
```

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

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

From node A, enter:

```
vxfenadm -i /dev/sdx
```

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

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

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

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

```
# vxfenadm -i /dev/sdz
```

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

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

Testing the disks using `vxfcntlsthdw` utility

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

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

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

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

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

To test the disks using `vxfcntlsthdw` utility

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

See “[Setting up inter-system communication](#)” on page 55.

- 2 From one node, start the utility.

Run the utility with the `-n` option if you use `rsh` for communication.

```
# vxfcntlsthdw [-n]
```

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

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

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

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

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

```
Enter the disk name to be checked for SCSI-3 PGR on node
galaxy in the format:
```

```
for dmp: /dev/vx/rdmp/sdx
```

```
for raw: /dev/sdx
```

```
Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/sdx
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
nebula in the format:
```

```
for dmp: /dev/vx/rdmp/sdx
```

```
for raw: /dev/sdx
```

```
Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/sdx
```

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

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

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

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

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

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

Setting up server-based I/O fencing using installvcs program

If VCS cluster is configured to run in secure mode, then verify that the configuration is correct before you configure CP server-based I/O fencing.

See [“Verifying security configuration on VCS cluster to use CP server coordination point”](#) on page 137.

See [“Configuring server-based I/O fencing”](#) on page 139.

Verifying security configuration on VCS cluster to use CP server coordination point

After configuring security using the `installvcs -security` command, follow the procedure below on each VCS cluster node to confirm that security is correctly configured.

To verify the security configuration on VCS cluster to use CP server coordination point

- 1 Run the following command:

```
# /opt/VRTScps/bin/cpsat listpd -t local
```

```
Domain(s) Found 1
```

```
*****
```

```
Domain Name HA_SERVICES@galaxy.symantec.com
```

```
Expiry Interval 0
```

```
*****
```

- 2 There should be a domain name entry with the following format:

```
HA_SERVICES@hostname.domainname
```

or

```
HA_SERVICES@hostname
```

3 There should not be duplicate entries for HA_SERVICES domain.

An example of incorrect configuration is given below.

```
showdomains

Domain(s) Found :          3

*****

Domain Name:      HA_SERVICES@galaxy.symantec.com

Domain Type:     vx

*****

Domain Name:      broker@galaxy.symantec.com

Domain Type:     vx

*****

Domain Name:      HA_SERVICES@galaxy

Domain Type:     vx

*****
```

Proceed to reconfigure security in case duplicate entries appear as shown in the above example.

Configuring server-based I/O fencing

This section describes how to configure server-based I/O fencing for the VCS cluster. With server-based I/O fencing, a combination of CP servers and SCSI-3 compliant coordinator disks can act as coordination points for I/O fencing.

To configure the VCS cluster with server-based I/O fencing

- 1 Ensure that the CP server(s) are configured and reachable from the cluster. If coordinator disks are to be used as coordination points, ensure that they are SCSI-3 compliant.
- 2 Run the `installvcs -fencing` command to configure fencing.

For example:

```
/opt/VRTS/install/installvcs -fencing
```

The installer creates a `vxfenmode` file on each node. The file is located at `/etc/vxfenmode`.

The following procedure can be used as an example to configure server-based I/O fencing. In this procedure example, there is one CP server and two disks acting as the coordination points.

To configure fencing configuration using the installer - CP client-based fencing

- 1 After installing and configuring VCS on the VCS cluster, the user issues the following command for configuring fencing:

```
/opt/VRTS/install/installvcs -fencing
```

- 2 After issuing the command, the installer displays Symantec copyright information and the location of log files for the configuration process.

Access and review these log files if there is any problem with the installation process. The following is an example of the command output:

```
Logs for installvcs are being created in /var/tmp/installvcs-LqwKwB.
```

- 3 Next, the installer displays the current cluster information for verification purposes. The following is an example of the command output:

```
Cluster information verification:
```

```
Cluster Name: clus1  
Cluster ID Number: 4445  
Systems: galaxy nebula
```

The cluster name, systems, and ID number are all displayed.

You are then asked whether you want to configure I/O fencing for the cluster. Enter "y" for yes. The rsh (or ssh) communication with the cluster nodes is then checked by the installer.

- 4** Next, you are prompted to select one of the following options for your fencing configuration:

```
Fencing configuration
```

- 1) Configure CP client based fencing
- 2) Configure disk based fencing
- 3) Configure fencing in disabled mode

```
Select the fencing mechanism to be configured in this  
Application Cluster [1-3,q]
```

Select the first option for CP client-based fencing.

- 5** Enter the total number of coordination points including both servers and disks. This number should be at least 3.

For example:

```
Enter the total number of co-ordination points including both  
CP servers and disks: [b] (3)
```

- 6** Enter the total number of coordinator disks among the coordination points. In this example, there are two coordinator disks.

For example:

```
Enter the total number of disks among these:  
[b] (0) 2
```

- 7** Enter the Virtual IP addresses or host names of the virtual IP address for each of the Coordination Point servers.

Note: The installer assumes these values to be the identical as viewed from all the client cluster nodes.

For example:

```
Enter the Virtual IP address/fully qualified host name  
for the Co-ordination Point Server #1::  
[b] 10.209.80.197
```

- 8** Enter the port that the CP server would be listening on.

For example:

```
Enter the port in the range [49152, 65535] which the
Co-ordination Point Server 10.209.80.197
would be listening on or simply accept the default port suggested:
[b] (14250)
```

- 9** Enter the fencing mechanism for the disk or disks.

For example:

```
Enter fencing mechanism for the disk(s) (raw/dmp):
[b,q,?] raw
```

- 10** The installer then displays a list of available disks to choose from to set up as coordinator points.

```
Select disk number 1 for co-ordination point
```

```
1) c3t0d0s2
2) c3t1d0s3
3) c3t2d0s4
```

```
Please enter a valid disk which is available from all the
cluster nodes for co-ordination point [1-3,q] 1
```

Select a disk from the displayed list.

Ensure that the selected disk is available from all the VCS cluster nodes.

- 11** Read the displayed recommendation from the installer to verify the disks prior to proceeding:

```
It is strongly recommended to run the 'VxFen Test Hardware' utility
located at '/opt/VRTSvcs/vxfen/bin/vxfentsthdw' in another window
before continuing. The utility verifies if the shared storage
you intend to use is configured to support I/O
fencing. Use the disk you just selected for this
verification. Come back here after you have completed
the above step to continue with the configuration.
```

Symantec recommends that you verify that the disks you are using as coordination points have been configured to support I/O fencing. Press Enter to continue.

You are then prompted to confirm your disk selection after performing a 'vxfentsthdw' test.

Press Enter to accept the default (y) and continue.

- 12** The installer then displays a list of available disks to choose from to set up as coordinator points.

Select a disk from the displayed list for the second coordinator point.

Ensure that the selected disk is available from all the VCS cluster nodes.

- 13** Proceed to read the displayed recommendation from the installer to verify the disks prior to proceeding.

Press Enter to continue.

- 14** You are then prompted to confirm your disk selection after performing a 'vxfentsthdw' test.

Press Enter to accept the default (y) and continue.

- 15** Proceed to enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):
[b] (vxfencoorddg)
```

- 16** The installer now begins verification of the coordination points. At the end of the verification process, the following information is displayed:

- Total number of coordination points being used
- CP Server Virtual IP/hostname and port number
- SCSI-3 disks

- Disk Group name for the disks in customized fencing
- Disk mechanism used for customized fencing

For example:

```
Total number of coordination points being used: 3
CP Server (Port):
  1. 10.209.80.197 (14250)
SCSI-3 disks:
  1. c3t0d0s2
  2. c3t1d0s3
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk mechanism used for customized fencing: raw
```

You are then prompted to accept the above information. Press Enter to accept the default (y) and continue.

The disks and disk group are initialized and the disk group deported on the VCS cluster node.

- 17** The installer now automatically determines the security configuration of the CP server's side and takes the appropriate action:
- If the CP server's side is configured for security, then the VCS cluster's side will be configured for security.
 - If the CP server's side is not configured for security, then the VCS cluster's side will not be configured for security.

For example:

```
While it is recommended to have secure communication
configured between CP Servers and CP client cluster, the client
cluster must be in the same mode (secure or non-secure) as the
CP servers are.
```

```
Since the CP servers are configured in secure mode, the installer
will configure the client cluster also as a secure cluster.
```

```
Press [Enter] to continue:
```

```
Trying to configure Security on the cluster:
```

```
All systems already have established trust within the
```

```
Symantec Product Authentication Service domain  
root@galaxy.symantec.com
```

- 18** Enter whether you are using different root brokers for the CP servers and VCS clusters.

If you are using different root brokers, then the installer tries to establish trust between the authentication brokers of the CP servers and the VCS cluster nodes for their communication.

After entering "y" for yes or "n" for no, press Enter to continue.

- 19** If you entered "y" for yes in step 18, then you are also prompted for the following information:
- Hostname for the authentication broker for any one of the CP servers
 - Port number where the authentication broker for the CP server is listening for establishing trust
 - Hostname for the authentication broker for any one of the VCS cluster nodes
 - Port number where the authentication broker for the VCS cluster is listening for establishing trust

Press Enter to continue.

- 20** The installer then displays your I/O fencing configuration and prompts you to indicate whether the displayed I/O fencing configuration information is correct.

If the information is correct, enter "y" for yes.

For example:

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm  
Cluster ID: 2122  
Cluster Name: clus1  
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

21 The installer then updates the VCS cluster information on each of the CP Servers to ensure connectivity between them.

The installer then populates the file `/etc/vxfenmode` with the above details in each of the CP VCS cluster nodes.

For example:

```
Updating client cluster information on CP Server 10.210.80.199

Adding the client cluster to the CP Server 10.210.80.199 ..... Done

Registering client node galaxy with CP Server 10.210.80.199..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Registering client node nebula with CP Server 10.210.80.199 ..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Updating /etc/vxfenmode file on galaxy ..... Done
Updating /etc/vxfenmode file on nebula ..... Done
```

For additional information about the `vxfenmode` file in mixed disk and CP server mode, or pure server-based mode:

See [“About I/O fencing configuration files”](#) on page 88.

22 You are then prompted to configure the CP agent on the client cluster.

```
Do you want to configure CP Agent on the client cluster? [y,n,q]
(y)

Enter a non-existing name for the service group for CP Agent:
[b] (vxfen)

Adding CP Agent via galaxy ..... Done
```

23 The VCS and the fencing process are then stopped and restarted on each VCS cluster node, and the I/O configuration process then finished.

```
Stopping VCS on galaxy ..... Done
Stopping Fencing on galaxy ..... Done
Stopping VCS on nebula ..... Done
Stopping Fencing on nebula ..... Done
```

24 At the end of this process, the installer then displays the location of the configuration log files, summary files, and response files.

4

Section

Installation using the Web-based installer

- [Chapter 8. Installing VCS](#)
- [Chapter 9. Configuring VCS](#)

Installing VCS

This chapter includes the following topics:

- [Before using the Veritas Web-based installer](#)
- [Starting the Veritas Web-based installer](#)
- [Obtaining a security exception on Mozilla Firefox](#)
- [Performing a pre-installation check with the Veritas Web-based installer](#)
- [Installing VCS with the Veritas Web-based installer](#)

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 8-1 Web-based installer requirements

System	Function	Requirements
Target system	The system(s) where the Veritas products will be installed.	Must be a supported platform for VCS 5.1
Installation server	The server from which to initiate the installation. The installation media is mounted and accessible from the installation server.	Must be the same OS as the system(s) on which to install.
Administrative system	The system on which you run the web browser to perform the installation.	Web browser

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprt1d`, on the installation server:

```
# ./webinstaller start
```

The `webinstaller` script displays a URL.

- 2 On the administrative server, start the Web browser.
- 3 Navigate to the URL displayed from step 1.
- 4 The browser may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

- 5 When prompted, enter `root` and `root`'s password of the installation server.

Obtaining a security exception on Mozilla Firefox

You may need to get a security exception on Mozilla Firefox.

To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **Add Exception** button.
- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.
- 5 Click **Confirm Security Exception** button.
- 6 Enter `root` in User Name field and `root` password of the web server in the Password field.

Performing a pre-installation check with the Veritas Web-based installer

This section describes performing a pre-installation check with the Veritas Web-based installer.

To perform a pre-installation check

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 152.
- 2 On the Select a task and a product page, select **Perform a Pre-installation check** from the **Task** drop-down list.
- 3 Select the product from the **Product** drop-down list, and click **Next**.
- 4 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Validate**.
- 5 The installer performs the precheck and displays the results.
- 6 If the validation completes successfully, click **Next**. The installer prompts you to begin the installation. Click **Ok** to install VCS on the selected system. Click **Cancel** to install later.
- 7 Click **Finish**. The installer prompts you for another task.

Installing VCS with the Veritas Web-based installer

This section describes installing VCS with the Veritas Web-based installer.

To install VCS

- 1 Perform preliminary steps.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 152.
- 3 On the Select a task and a product page, select **Install a Product** from the **Task** drop-down list.
- 4 Select VCS from the Product drop-down list, and click **Next**.
- 5 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 6 Choose minimal, recommended, or all packages. Click **Next**.
- 7 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
- 8 After the validation completes successfully, click **Next** to install VCS on the selected system.
- 9 After the installation completes, you must choose your licensing method.
On the license page, select one of the following tabs:

- Keyless licensing

Note: The keyless license option enables you to install without entering a key. However, in order to ensure compliance you must manage the systems with a management server.

For more information, go to the following website:

<http://go.symantec.com/sfhakeyless>

Complete the following information:

Choose whether you want to enable Global Cluster option.

Click Register.

- Enter license key

If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

10 The installer prompts you to configure the cluster.

If you select n, you can exit the installer. You must configure the product before you can use VCS.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

11 Select the checkbox to specify whether you want to send your installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future?

Click **Finish**. The installer prompts you for another task.

Configuring VCS

This chapter includes the following topics:

- [Configuring VCS using the web-based installer](#)

Configuring VCS using the web-based installer

This section describes the procedure to configure VCS using the web-based installer. Before you begin with the procedure, review the requirements for configuring VCS.

See “[Getting your VCS installation and configuration information ready](#)” on page 64.

To configure VCS on a cluster

- 1 Start the web-based installer.
- 2 Select the following on the **Select Product/Task** screen:
 - From the list of tasks, select **Configure a product**.
 - From the list of products, select **Veritas Cluster Server**.
 - By default, the communication between the systems is selected as SSH. If SSH is used for communication between systems, the SSH commands execute without prompting for passwords or confirmations.
 - Click **Next**.

Note: You can click **Quit** to quit the web-installer at any time during the configuration process.

- 3 Select the following on the **Select Systems** screen:

- Enter the system names on which VCS is to be configured, and then click **Validate**.
Example: **galaxy nebula**
The installer performs the initial system verification. It checks that communication between systems has been set up. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.
 - Click **Next** after the installer completes the system verification successfully.
- 4 Select the following on the **Set Cluster Name/ ID** screen.
- Enter the unique cluster name and Cluster ID number.
 - Select the number of heartbeat links.
 - Select **Low priority heartbeat** if you want to configure one heartbeat link as a low priority link.
 - Select **Unique NICs per system** if you do not want to use the same NIC details to configure private heartbeat links on other systems.
 - Click **Next**.
- 5 Select the following on the **Set Cluster Heartbeat** screen.
- If you are using the same NICs to configure private heartbeat links on all the systems, select the NIC for the first private heartbeat NIC on each system.
Select the NIC for the second private heartbeat NIC on each system.
 - If you have selected **Unique NICs per system** in the previous screen, provide the NIC details for each system.
 - Click **Next**.
- 6 In the **VCS Optional Configure** screen, select the VCS options that you want to configure, namely Virtual IP, User, SMTP, SNMP, and GCO. Depending on the options that you select, you can enter the details regarding each option.
- To configure the virtual IP, do the following:
 - Select **Configure Virtual IP**.
 - If each system uses a separate NIC, select **Configure NICs for every system separately**.
 - Select the interface on which you want to configure the virtual IP.
 - Enter a virtual IP address and value for the netmask.
 - To configure the VCS users, enter the following information:

- Reset the password for the Admin user, if necessary.
- Click **Add** to add a new user.
Specify the user name, password, and user privileges for this user.
- To configure SMTP notification, enter the following information:
 - If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
 - Enter the domain-based hostname of the SMTP server. Example: smtp.yourcompany.com
 - Enter the full email address of the SMTP recipient. Example: user@yourcompany.com.
 - Select the minimum security level of messages to be sent to each recipient.
 - Click **Add** to add more SMTP recipients, if necessary.
- To configure SNMP notification, enter the following information.
 - If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
 - Enter the SNMP trap daemon port: (162).
 - Enter the SNMP console system name.
 - Select the minimum security level of messages to be sent to each console.
 - Click **Add** to add more SNMP consoles, if necessary.
- If you installed a valid HA/DR license, you can select the gco option to configure this cluster as a global cluster.
See *Veritas Cluster Server Administrator's Guide* for instructions to set up VCS global clusters.
 - Select a NIC.
 - Enter a virtual IP address and value for the netmask.
 - Click **Next**.
The installer proceeds to configure VCS based on the configuration details you provided.

- 7 In the **Starting Processes** screen, the installer completes the VCS configuration.

The installer starts VCS and its components on each system.

After the startup process is complete, click **Next** to move to the next screen.

- 8 Click **Next** to complete the process of configuring VCS.

View the summary file, log file, or response file, if needed, to confirm the configuration.

- 9 Select the checkbox to specify whether you want to send your installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future?

Click **Finish**. The installer prompts you for another task.

Installation using response files

- [Chapter 10. Performing automated VCS installation](#)
- [Chapter 11. Performing automated VCS configuration](#)
- [Chapter 12. Performing automated I/O fencing configuration for VCS](#)

Performing automated VCS installation

This chapter includes the following topics:

- [Installing VCS using response files](#)
- [Response file variables to install VCS](#)
- [Sample response file for installing VCS](#)

Installing VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS installation on one cluster to install VCS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To install VCS using response files

- 1 Make sure the systems where you want to install VCS meet the installation requirements.
See [“VCS installation requirements”](#) on page 33.
- 2 Make sure the preinstallation tasks are completed.
See [“Performing preinstallation tasks”](#) on page 51.
- 3 Copy the response file to one of the cluster systems where you want to install VCS.
See [“Sample response file for installing VCS”](#) on page 164.
- 4 Edit the values of the response file variables as necessary.
See [“Response file variables to install VCS”](#) on page 162.

- 5 Mount the product disc and navigate to the folder that contains the installation program.
- 6 Start the installation from the system to which you copied the response file. For example:

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

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

Response file variables to install VCS

[Table 10-1](#) lists the response file variables that you can define to install VCS.

Table 10-1 Response file variables specific to installing VCS

Variable	List or Scalar	Description
CFG{opt}{install}	Scalar	Installs VCS RPMs. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{opt}{stopfail_allow}	Scalar	Decides whether or not to proceed if the installer fails while stopping the processes or while unloading the drivers. (Optional)
CFG{systems}	List	List of systems on which the product is to be installed. Required
CFG{prod}	Scalar	Defines the product to be installed. The value is VCS51 for VCS. (Required)

Table 10-1 Response file variables specific to installing VCS (*continued*)

Variable	List or Scalar	Description
CFG{opt}installallpkgs or CFG{opt}installrecpkgs or CFG{opt}installminpkgs	Scalar	Instructs the installer to install VCS RPMs based on the variable that has the value set to 1: <ul style="list-style-type: none"> ■ installallpkgs: Installs all RPMs ■ installrecpkgs: Installs recommended RPMs ■ installminpkgs: Installs minimum RPMs <p>Note: The installer requires only one of these variable values to be set to 1.</p> (Required)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. (Optional)
CFG{opt}{gco}	Scalar	Defines that the installer must enable the global cluster option. You must set this variable value to 1 if you want to configure global clusters. (Optional)
CFG{opt}{keyfile}	Scalar	Defines the location of an <i>ssh</i> keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{patchpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems. (Optional)

Table 10-1 Response file variables specific to installing VCS (*continued*)

Variable	List or Scalar	Description
CFG{opt}{pkgpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product RPMs. The location must be accessible from all target systems. (Optional)
CFG{opt}{tmppath}	Scalar	Defines the location where a working directory is created to store temporary files and the depots that are needed during the install. The default location is /var/tmp. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
\$CFG{opt}{vxkeyless}	Scalar	Installs the product with keyless license if the value is set to 1. If the value is set to 0, you must define the CFG{keys}{system} variable with the license keys. (Optional)
CFG{keys}{system}	Scalar	List of keys to be registered on the system if the variable \$CFG{opt}{vxkeyless} is set to 0. (Optional)

Sample response file for installing VCS

Review the response file variables and their definitions.

See [“Response file variables to install VCS”](#) on page 162.

```
#  
# Configuration Values:  
#  
our %CFG;  
  
$CFG{accepteula}=1;  
$CFG{opt}{install}=1;  
$CFG{opt}{installrecpkgs}=1;  
$CFG{prod}="VCS51";  
$CFG{systems}=[ qw(galaxy nebula) ];
```


Performing automated VCS configuration

This chapter includes the following topics:

- [Configuring VCS using response files](#)
- [Response file variables to configure VCS](#)
- [Sample response file for configuring VCS](#)

Configuring VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS configuration on one cluster to configure VCS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure VCS using response files

- 1 Make sure the VCS RPMs are installed on the systems where you want to configure VCS.
- 2 Copy the response file to one of the cluster systems where you want to configure VCS.

See [“Sample response file for configuring VCS”](#) on page 174.

- 3 Edit the values of the response file variables as necessary.
 To configure optional features, you must define appropriate values for all the response file variables that are related to the optional feature.
 See “[Response file variables to configure VCS](#)” on page 168.
- 4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installvcs -responsefile /tmp/response_file
```

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

Response file variables to configure VCS

[Table 11-1](#) lists the response file variables that you can define to configure VCS.

Table 11-1 Response file variables specific to configuring VCS

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the RPMs are already installed. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{opt}{stopfail_allow}	Scalar	Decides whether or not to proceed if the installer fails while stopping the processes or while unloading the drivers. (Optional)
CFG{systems}	List	List of systems on which the product is to be configured. (Required)
CFG{prod}	Scalar	Defines the product to be configured. The value is VCS51 for VCS. (Required)

Table 11-1 Response file variables specific to configuring VCS (*continued*)

Variable	List or Scalar	Description
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of ssh as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are 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 smtpsev), the SNMP trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

[Table 11-2](#) lists the response file variables that specify the required information to configure a basic VCS cluster.

Table 11-2 Response file variables specific to configuring a basic VCS cluster

Variable	List or Scalar	Description
CFG{vcs_clusterid}	Scalar	An integer between 0 and 65535 that uniquely identifies the cluster. (Required)
CFG{vcs_clustername}	Scalar	Defines the name of the cluster. (Required)

Table 11-2 Response file variables specific to configuring a basic VCS cluster
(continued)

Variable	List or Scalar	Description
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)
CFG{vcs_lltlink#} {system}	Scalar	Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. (Required)
CFG{vcs_lltlinklowpri} {system}	Scalar	Defines a low priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication. (Optional)

Table 11-3 lists the response file variables that specify the required information to configure virtual IP for VCS cluster.

Table 11-3 Response file variables specific to configuring virtual IP for VCS cluster

Variable	List or Scalar	Description
CFG{vcs_csgnic} {system}	Scalar	Defines the NIC device to use on a system. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_csgvip}	Scalar	Defines the virtual IP address for the cluster. (Optional)

Table 11-3 Response file variables specific to configuring virtual IP for VCS cluster (*continued*)

Variable	List or Scalar	Description
CFG{vcs_csgnetmask}	Scalar	Defines the Netmask of the virtual IP address for the cluster. (Optional)

Table 11-4 lists the response file variables that specify the required information to configure the VCS cluster in secure mode.

Table 11-4 Response file variables specific to configuring VCS cluster in secure mode

Variable	List or Scalar	Description
CFG{at_rootdomain}	Scalar	Defines the name of the system where the root broker is installed. (Optional)
CFG{vcs_securitymenuopt}	Scalar	Specifies the menu option to choose to configure the cluster in secure mode. <ul style="list-style-type: none"> ■ 1—Automatic ■ 2—Semi-automatic ■ 3—Manual (Optional)
CFG{vcs_vssdefport}	Scalar	Specifies the default port address of the root broker. (Optional)
CFG{vcs_roothashpath}	Scalar	Specifies the path of the root hash file. (Optional)
CFG{vcs_ab_prplname} {system}	Scalar	Specifies the authentication broker's principal name on system. (Optional)
CFG{vcs_ab_password} {system}	Scalar	Specifies the authentication broker's password on system. (Optional)

Table 11-4 Response file variables specific to configuring VCS cluster in secure mode (*continued*)

Variable	List or Scalar	Description
CFG{vcs_blobpath} {system}	Scalar	Specifies the path of the encrypted BLOB file for system. (Optional)

Table 11-5 lists the response file variables that specify the required information to configure VCS users.

Table 11-5 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	List of encoded passwords for VCS users The value in the list can be "Administrators Operators Guests" Note: The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list. (Optional)
CFG{vcs_username}	List	List of names of VCS users (Optional)
CFG{vcs_userpriv}	List	List of privileges for VCS users Note: The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list. (Optional)

Table 11-6 lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 11-6 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)
CFG{vcs_smtpsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

[Table 11-7](#) lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 11-7 Response file variables specific to configuring VCS notifications using SNMP

Variable	List or Scalar	Description
CFG{vcs_snmpport}	Scalar	Defines the SNMP trap daemon port (default=162). (Optional)
CFG{vcs_snmpcons}	List	List of SNMP console system names (Optional)

Table 11-7 Response file variables specific to configuring VCS notifications using SNMP (*continued*)

Variable	List or Scalar	Description
CFG{vcs_snmpcsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

Table 11-8 lists the response file variables that specify the required information to configure VCS global clusters.

Table 11-8 Response file variables specific to configuring VCS global clusters

Variable	List or Scalar	Description
CFG{vcs_gconic} {system}	Scalar	Defines the NIC for the Virtual IP that the Global Cluster Option uses. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_gcovip}	Scalar	Defines the virtual IP address to that the Global Cluster Option uses. (Optional)
CFG{vcs_gconetmask}	Scalar	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. (Optional)

Sample response file for configuring VCS

Review the response file variables and their definitions.

See [“Response file variables to configure VCS”](#) on page 168.

```
#
# Configuration Values:
```

```
#
our %CFG;

$CFG{at_rootdomain}="root\@east.symantecexample.com";
$CFG{rootbroker}="east.symantecexample.com";
$CFG{opt}{configure}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{ha}=1;
$CFG{prod}="VCS51";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_csgnetmask}="255.255.255.0";
$CFG{vcs_csgnic}{all}="eth0";
$CFG{vcs_csgvip}="10.10.12.1";
$CFG{vcs_gconetmask}="255.255.255.0";
$CFG{vcs_gcovip}="10.10.12.1";
$CFG{vcs_lltlink1}{galaxy}="eth1";
$CFG{vcs_lltlink1}{nebula}="eth1";
$CFG{vcs_lltlink2}{galaxy}="eth2";
$CFG{vcs_lltlink2}{nebula}="eth2";

$CFG{vcs_securitymenuopt}=1;
$CFG{vcs_smtprecp}=[ qw(earnie@symantecexample.com) ];
$CFG{vcs_smtprsev}=[ qw(SevereError) ];
$CFG{vcs_smtprserver}="smtp.symantecexample.com";
$CFG{vcs_snmpcons}=[ qw(neptune) ];
$CFG{vcs_snmpcsev}=[ qw(SevereError) ];
$CFG{vcs_snmpport}=162;
```


Performing automated I/O fencing configuration for VCS

This chapter includes the following topics:

- [Configuring I/O fencing using response files](#)
- [Response file variables to configure disk-based I/O fencing](#)
- [Sample response file for configuring disk-based I/O fencing](#)
- [Response file variables to configure server-based I/O fencing](#)
- [Sample response file for configuring server-based I/O fencing](#)

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for VCS. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure I/O fencing using response files

- 1 Make sure that VCS is configured.
- 2 Based on whether you want to configure disk-based or server-based I/O fencing, make sure you have completed the preparatory tasks.

See “[About planning to configure I/O fencing](#)” on page 91.

- 3 Copy the response file to one of the cluster systems where you want to configure I/O fencing.

See “[Sample response file for configuring disk-based I/O fencing](#)” on page 179.

See “[Sample response file for configuring server-based I/O fencing](#)” on page 182.

- 4 Edit the values of the response file variables as necessary.

See “[Response file variables to configure disk-based I/O fencing](#)” on page 178.

See “[Response file variables to configure server-based I/O fencing](#)” on page 180.

- 5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installvcs -responsefile /tmp/response_file
```

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

Response file variables to configure disk-based I/O fencing

[Table 12-1](#) lists the response file variables that specify the required information to configure disk-based I/O fencing for VCS.

Table 12-1 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{vxfen_config_fencing_option}	Scalar	Specifies the I/O fencing configuration mode. <ul style="list-style-type: none"> ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode (Required)
CFG {vxfen_config_fencing_mechanism}	Scalar	Specifies the I/O fencing mechanism. (Optional)

Table 12-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG{vxfen_config_fencing_dg}	Scalar	Specifies the disk group for I/O fencing. (Optional) Note: You must define either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.
CFG{vxfen_config_fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) Note: You must define either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.

Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions. See [“Response file variables to configure disk-based I/O fencing”](#) on page 178.

```
#
# Configuration Values:
#
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;

$CFG{prod}="VCS51";

$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{vxfen_config_fencing_dg}="fendg";
$CFG{vxfen_config_fencing_mechanism}="dmp";
$CFG{vxfen_config_fencing_newdg_disks}=
```

```
[ qw(sdx sdy sdz) ];  
$CFG{vxfen_config_fencing_option}=2;
```

Response file variables to configure server-based I/O fencing

You can use a CP server response file to configure server-based customized I/O fencing. The installer uses the CP server response file for the following types of I/O fencing configurations:

- **Client cluster fencing (server-based I/O fencing configuration itself)**
The installer configures server-based customized I/O fencing on the VCS cluster without prompting for user input.
- **Disk-based fencing with the disk group already created**
The installer configures fencing in disk-based mode on the VCS cluster without prompting for user input.
Disk-based fencing configuration is one in which SCSI-3 disks are used as the only coordination points.
Disk-based fencing with the disk group already created means that the disk group consisting of the coordinating disks already exists on the VCS cluster nodes.
- **Disk-based fencing with the disk group to be created**
The installer creates the disk group and configures fencing properly on all the nodes in the VCS cluster without user intervention.
Disk-based fencing with the disk group to be created means that the disk group does not exist yet, but will be created with the disks mentioned as coordination point.

[Table 12-2](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table 12-2 CP server response file definitions

Response file field	Definition
fencing_cpc_config_cpagent	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.

Table 12-2 CP server response file definitions (*continued*)

Response file field	Definition
fencing_cpc_cpagentgrp	<p>Name of the service group which will have the Coordination Point agent resource as part of it.</p> <p>Note: This field is obsolete if the <code>fencing_cpc_config_cpagent</code> field is given a value of '0'.</p>
fencing_cpc_cps	<p>Virtual IP address or Virtual hostname of the CP servers.</p>
fencing_cpc_reusedg	<p>This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks).</p> <p>Enter either a "1" or "0".</p> <p>Entering a "1" indicates reuse, and entering a "0" indicates do not reuse.</p> <p>When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text, such as <code>"\$CFG{fencing_cpc_reusedg}=0"</code> or <code>"\$CFG{fencing_cpc_reusedg}=1"</code> before proceeding with a silent installation.</p>
fencing_cpc_dgname	<p>The name of the disk group to be used in the customized fencing, where at least one disk is being used.</p>
fencing_cpc_diffab	<p>This response field indicates whether the CP servers and VCS clusters use different root brokers.</p> <p>Entering a "1" indicates that they are using different root brokers.</p> <p>Entering a "0" indicates that they are not using different root brokers.</p>
fencing_cpc_disks	<p>The disks being used as coordination points if any.</p>
fencing_cpc_ncps	<p>Total number of coordination points being used, including both CP servers and disks.</p>
fencing_cpc_ndisks	<p>The number of disks being used.</p>
fencing_cpc_ports	<p>The port of the CP server that is denoted by <i>cps</i>.</p>

Table 12-2 CP server response file definitions (*continued*)

Response file field	Definition
fencing_cpc_ccab	The name of the authentication broker (AB) for any one of the VCS cluster nodes.
fencing_cpc_cpsabport	The port at which the authentication broker (AB) mentioned above listens for authentication..
fencing_cpc_ccabport	The port at which the authentication broker (AB) mentioned above listens for authentication.
fencing_cpc_mechanism	The disk mechanism that is used by customized fencing. The value for this field is either "raw" or "dmp"
fencing_cpc_cpsab	The name of the authentication broker (AB) for any one of the CP servers.
fencing_cpc_security	This field indicates whether security is enabled or not Entering a "1" indicates that security is enabled. Entering a "0" indicates that security has not been enabled.

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing :

```

$CFG{fencing_cpc_config_cpagent}=0;
$CFG{fencing_cpc_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cpc_dgname}="vxfencoorddg";
$CFG{fencing_cpc_diffab}=0;
$CFG{fencing_cpc_disks}=[ qw(emc_clariion0_37 emc_clariion0_13) ];
$CFG{fencing_cpc_mechanism}="raw";
$CFG{fencing_cpc_ncps}=3;
$CFG{fencing_cpc_ndisks}=2;
$CFG{fencing_cpc_ports}{"10.200.117.145"}=14250;
$CFG{fencing_cpc_reusedg}=1;
$CFG{fencing_cpc_security}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="VCS51";

```

```
$CFG{systems}=[ qw(galaxy nebula) ];  
$CFG{vcs_clusterid}=1256;  
$CFG{vcs_clustername}="clus1";  
$CFG{vxfen_config_fencing_option}=1;
```


Manual installation

- [Chapter 13. Performing preinstallation tasks](#)
- [Chapter 14. Manually installing VCS](#)
- [Chapter 15. Manually configuring VCS](#)
- [Chapter 16. Manually configuring the clusters for data integrity](#)

Performing preinstallation tasks

This chapter includes the following topics:

- [Preparing for a manual installation](#)
- [Requirements for installing VCS](#)

Preparing for a manual installation

Before you start installation, log in as the superuser. Mount the disc, copy the files to a temporary location locally for your convenience. Each operating system occupies an entire disc. Each disc has an identical directory structure.

To prepare for installation

- 1 Log in as the superuser.
- 2 Mount the appropriate disc.
See [“Mounting the product disc”](#) on page 62.
- 3 Copy the files to a temporary location on the system.

```
# cp -r pkgs/* /tmp/install
```

Requirements for installing VCS

Review requirements before you install.

See [“VCS installation requirements”](#) on page 33.

Manually installing VCS

This chapter includes the following topics:

- [About VCS manual installation](#)
- [Installing VCS software manually](#)
- [Installing VCS using Kickstart](#)

About VCS manual installation

You can manually install and configure VCS instead of using the `installvcs` program.

A manual installation takes a lot of time, patience, and care. Symantec recommends that you use the `installvcs` program instead of the manual installation when possible.

Installing VCS software manually

[Table 14-1](#) lists the tasks that you must perform when you manually install and configure VCS 5.1.

Table 14-1 Manual installation tasks for VCS 5.1

Task	Reference
Install VCS software manually on each node in the cluster.	See “Installing VCS RPMs for a manual installation” on page 191.
Add a license key.	See “Adding a license key for a manual installation” on page 193.

Table 14-1 Manual installation tasks for VCS 5.1 (*continued*)

Task	Reference
Copy the installation guide to each node.	See “Copying the installation guide to each node” on page 195.
Configure LLT and GAB.	<ul style="list-style-type: none"> ■ See “Configuring LLT for a manual installation” on page 201. ■ See “Configuring GAB for a manual installation” on page 204.
Configure VCS.	See “Configuring VCS” on page 204.
Start LLT, GAB, and VCS services.	See “Starting LLT, GAB, and VCS for a manual installation” on page 206.
Modify the VCS configuration.	See “Modifying the VCS configuration” on page 207.
Replace demo license with a permanent license.	See “Replacing a VCS demo license with a permanent license for manual installations” on page 195.

Viewing the list of VCS RPMs

During the VCS installation, the installer prompts you with an option to choose the VCS RPMs to install. You can view the list of RPMs that each of these options would install using the installer command-line option.

Manual installation or upgrade of the product requires you to install the RPMs in a specified order. For example, you must install some RPMs before other RPMs because of various product dependencies. The following installer command options list the RPMs in the order in which you must install these RPMs.

[Table 14-2](#) describes the VCS RPM installation options and the corresponding command to view the list of RPMs.

Table 14-2 Installer command options to view VCS RPMs

Option	Description	Command option to view the list of RPMs
1	Installs only the minimal required VCS RPMs that provide basic functionality of the product.	<code>installvcs -minpkgs</code>
2	Installs the recommended VCS RPMs that provide complete functionality of the product. This option does not install the optional VCS RPMs.	<code>installvcs -recpkgs</code>

Table 14-2 Installer command options to view VCS RPMs (*continued*)

Option	Description	Command option to view the list of RPMs
3	Installs all the VCS RPMs. You must choose this option to configure any optional VCS feature.	<code>installvcs -allpkgs</code>

To view the list of VCS RPMs

- 1 Navigate to the directory where you can start the `installvcs` program.

```
# cd cluster_server
```

- 2 Run the following command to view the list of RPMs. Based on what RPMs you want to install, enter the appropriate command option:

```
# ./installvcs -minpkgs
```

Or

```
# ./installvcs -recpkgs
```

Or

```
# ./installvcs -allpkgs
```

Installing VCS RPMs for a manual installation

All RPMs are installed into the `/opt` directory.

You can create lists of the RPMs to install.

See [“Viewing the list of VCS RPMs”](#) on page 190.

If you copied these files to `/tmp/install`, navigate to the directory and perform the following on each system:

To install VCS RPMs on a node

- ◆ Install the RPMs in the order shown. Do not install any RPMs already installed on the system. Pay special attention to operating system distribution and architecture.

- RHEL5:

```
# rpm -i VRTSvlic-3.02.51.0000-0.x86_64.rpm
# rpm -i VRTSperl-5.10.0.7-RHEL5.3.x86_64.rpm
```

```
# rpm -i VRTSspt-5.5.00.0-GA.noarch.rpm
# rpm -i VRTSllt-5.1.00.00-GA_RHEL5.x86_64.rpm
# rpm -i VRTSgab-5.1.00.00-GA_RHEL5.x86_64.rpm
# rpm -i VRTSvxfen-5.1.00.00-GA_RHEL5.x86_64.rpm
# rpm -i VRTSvcs-5.1.00.00-GA_RHEL5.i686.rpm
# rpm -i VRTScps-5.1.00.00-GA_RHEL5.i686.rpm
# rpm -i VRTSvcsag-5.1.00.00-GA_RHEL5.i686.rpm
# rpm -i VRTSvcsdr-5.1.00.00-GA_RHEL5.x86_64.rpm
# rpm -i VRTScutil-5.1.00.00-GA_GENERIC.noarch.rpm
# rpm -i VRTSatClient-5.0.31.0-0.x86_64.rpm
# rpm -i VRTSatServer-5.0.31.0-0.x86_64.rpm
# rpm -i VRTSvcsea-5.1.00.00-GA_RHEL5.i686.rpm
```

■ SLES10:

```
# rpm -i VRTSvlic-3.02.51.0000-0.x86_64.rpm
# rpm -i VRTSperl-5.10.0.7-SLES10.x86_64.rpm
# rpm -i VRTSspt-5.5.00.0-GA.noarch.rpm
# rpm -i VRTSllt-5.1.00.00-GA_SLES10.x86_64.rpm
# rpm -i VRTSgab-5.1.00.00-GA_SLES10.x86_64.rpm
# rpm -i VRTSvxfen-5.1.00.00-GA_SLES10.x86_64.rpm
# rpm -i VRTSvcs-5.1.00.00-GA_SLES10.i586.rpm
# rpm -i VRTScps-5.1.00.00-GA_SLES10.i686.rpm
# rpm -i VRTSvcsag-5.1.00.00-GA_SLES10.i586.rpm
# rpm -i VRTSvcsdr-5.1.00.00-GA_SLES10.x86_64.rpm
# rpm -i VRTScutil-5.1.00.00-GA_GENERIC.noarch.rpm
# rpm -i VRTSatClient-5.0.31.0-0.x86_64.rpm
# rpm -i VRTSatServer-5.0.31.0-0.x86_64.rpm
# rpm -i VRTSvcsea-5.1.00.00-GA_SLES10.i586.rpm
```

■ SLES11:

```
# rpm -i VRTSvlic-3.02.51.0000-0.x86_64.rpm
# rpm -i VRTSperl-5.10.0.5-SLES11.x86_64.rpm
# rpm -i VRTSspt-5.5.00.0-GA.noarch.rpm
# rpm -i VRTSllt-5.1.00.00-GA_SLES11.x86_64.rpm
# rpm -i VRTSgab-5.1.00.00-GA_SLES11.x86_64.rpm
# rpm -i VRTSvxfen-5.1.00.00-GA_SLES11.x86_64.rpm
# rpm -i VRTSvcs-5.1.00.00-GA_SLES11.i686.rpm
# rpm -i VRTScps-5.1.00.00-GA_SLES11.i686.rpm
# rpm -i VRTSvcsag-5.1.00.00-GA_SLES11.i686.rpm
# rpm -i VRTSvcsdr-5.1.00.00-GA_SLES11.x86_64.rpm
```

```
# rpm -i VRTScutil-5.1.00.00-GA_GENERIC.noarch.rpm
# rpm -i VRTSatClient-5.0.31.0-0.x86_64.rpm
# rpm -i VRTSatServer-5.0.31.0-0.x86_64.rpm
# rpm -i VRTSvcsea-5.1.00.00-GA_SLES11.i686.rpm
```

See “[Veritas Cluster Server installation packages](#)” on page 345.

Adding a license key for a manual installation

You can either add the VCS license keys or use keyless licensing for VCS.

See “[Setting or changing the product level for keyless licensing](#)” on page 193.

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
```

Setting or changing the product level for keyless licensing

The keyless licensing method uses product levels to determine the Veritas products and functionality that are licensed. In order to use keyless licensing, you must set up a Management Server to manage your systems.

For more information and to download the management server, see the following URL:

<http://go.symantec.com/vom>

When you set the product license level for the first time, you enable keyless licensing for that system. If you install with the product installer and select the keyless option, you are prompted to select the product and feature level that you want to license.

After you install, you can change product license levels at any time to reflect the products and functionality that you want to license. When you set a product level, you agree that you have the license for that functionality.

To set or change the product level

- 1 View the current setting for the product level.

```
# vxkeyless [-v] display
```

- 2 View the possible settings for the product level.

```
# vxkeyless displayall
```

- 3 Set the desired product level.

```
# vxkeyless [-q] set prod_levels
```

where *prod_levels* is a comma-separated list of keywords, as shown in step 2

If you want to remove keyless licensing and enter a key, you must clear the keyless licenses. Use the NONE keyword to clear all keys from the system.

Warning: Clearing the keys disables the Veritas products until you install a new key or set a new product level.

To clear the product license level

- 1 View the current setting for the product license level.

```
# vxkeyless [-v] display
```

- 2 If there are keyless licenses installed, remove all keyless licenses:

```
# vxkeyless [-q] set NONE
```

For more details on using the `vxkeyless` utility, see the `vxkeyless(1m)` manual page.

Checking licensing information on the system for a manual installation

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 following:

- The license key

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

Replacing a VCS demo license with a permanent license for manual installations

When a VCS demo key license expires, you can replace it with a permanent license using the `vxlicinst` program.

See “[Checking licensing information on the system](#)” on page 125.

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.

Installing VCS using Kickstart

These installation instructions using Kickstart assume a working knowledge of the Kickstart procedure. See the Kickstart documentation that came with your operating system for details. Kickstart is supported only for Red Hat Enterprise Linux (RHEL); therefore, in Release 5.1, installing Veritas product RPMs using the Kickstart installation method only applies to RHEL.

Only fresh installs of Veritas products are supported using Kickstart. Upgrading is not supported.

The following procedure assumes a stand-alone configuration.

About Kickstart

The Kickstart feature is an automatic operating system installation method available for the Red Hat Linux operating system. All Red Hat Linux base installations require some basic configuration. The Kickstart feature enables the system administrator to automatically install systems based on predefined customized configurations.

Benefits of Kickstart automated installation:

- Simplifies installations
- Installs faster than manual CD/DVD ROM or `askmethod` type installations.

- Allows unattended installation
- Enables administrators to configure systems across the enterprise the same way

Summary of Kickstart installation tasks

For detailed instructions, follow the Kickstart documentation that came with your operating system. These steps are provided as a summary only.

To install the Veritas packages on a Kickstart server

- 1 Create a Red Hat Linux Kickstart Install Server.
- 2 Copy the Veritas rpm files to the Kickstart install server.
Note the subdirectory with the relevant packages for your installation.
- 3 Create the Kickstart installation scripts for the Veritas product.
The Veritas product installer provides an option to generate a Kickstart script. Modify the generated script with the necessary install server information.
- 4 Create a Kickstart configuration file that includes the Veritas Kickstart scripts.
- 5 Run Kickstart installation using the customized Kickstart configuration file to install Red Hat Enterprise Linux and the Veritas products on the target system.
- 6 Configure the Veritas software.

Setting up an NFS exported location for Veritas product RPMs

To set up an NFS exported location for Veritas product RPMs

- 1 Set up an NFS exported location that the Kickstart client can access.
- 2 Copy the VRTS RPMs from the `rpms` directory of the installation media to an NFS location.
- 3 Copy the generated product stack `installprod` and `uninstallprod` scripts to the same NFS location.

Copying the rpms to the install/NFS server

Before you can install the rpms, you must copy the Veritas rpm files to the install server. This procedure is only required when you install using Kickstart.

To copy the rpms

- 1 Load and mount the Veritas software disc.

```
# mount /mnt/cdrom
```

- 2 Create a temporary directory for installation.

```
# mkdir /path/to/rpms_and_scripts
```

- 3 Copy the rpms from the software disc to the temporary directory.

```
# cp -rp /mnt/cdrom/rhel5_x86_64/rpms/* /path/to/rpms_and_scripts
```

- 4 Use this directory path for the packages and Veritas script files for the manual installation procedure.

Make sure the `/path/to/rpms_and_scripts` directory is NFS exported and is accessible from the NFS client.

Creating the Veritas Kickstart script file

The operations to install the rpms must be coded in a script that can be used with the Kickstart server. Use the following instructions to create the Veritas Kickstart script.

To create the Veritas Kickstart script file

- 1 Make a temporary directory for the Veritas Kickstart script files:

```
# mkdir -p /tmp/kickstart
```

- 2 Move to the Veritas software disc.

```
# cd /mnt/cdrom/rhel5_x86_64
```

- 3 Generate the Kickstart configuration files, which have a .ks extension. Run the installer with the `-kickstart` option to create a sample Kickstart script for the Veritas product that you plan to install. The command has the following format:

```
# ./installprod -kickstart dir_path
```

For example, to create a sample file for Veritas Cluster Server:

```
# cd cluster_server
# ./installvcs -kickstart dir_path
```

The *dir_path* indicates the path to the directory in which to create the finish file. In this example, the *dir_path* is `/tmp/kickstart` and the output of above command generates three Veritas script files: `kickstart_vcs51.ks`, `installvcs`, and `uninstallvcs` to the `/tmp/kickstart` location.

- 4 Modify the Veritas Kickstart script file (`kickstart_vcs51.ks`) to include the name of the install server and the location of the files.

Edit the line beginning with `BUILDSRC=` as follows:

```
BUILDSRC="hostname_or_ip:/path/to/rpms_and_scripts"
```

Where *hostname_or_ip* is the hostname or IP address of the install/NFS Server.

Creating and modifying the Kickstart configuration file

Refer to the operating system documentation for details about creating and modifying the Kickstart configuration file. The following procedure discusses the requirements for the Veritas software.

To modify the Kickstart configuration file

- 1 Modify the `BUILDSRC` variable in the Kickstart configuration file so that the variable points to your actual NFS location.
- 2 Append the modified contents of the Kickstart configuration file (for example `kickstart_vcs51.ks`) to the main RHEL operating system `ks.cfg` file.

Completing the Kickstart process

Complete the Kickstart installation on the target system. The following procedure provides the high-level process. Detailed steps for the Kickstart installation process are beyond the scope of this document. Refer to the Red Hat Linux documentation for details.

The Kickstart process installs the Red Hat Linux and installs the Veritas rpms.

To complete the Kickstart process

- 1** Launch the Kickstart installation for the RHEL operating system.
Refer to the operating system's documentation for a Kickstart-based installation.
- 2** After the operating system installation is complete, check the file `/var/tmp/kickstart.log` for any errors related to the installation of VRTS RPMs and CPI scripts. Also determine whether all the RPMs have been installed. Enter the following command:

```
# rpm -qa | grep -i vrts
```
- 3** If you do not find any issues or errors, configure the product stack using the installer script on the product media.
- 4** Verify that all the configured llt links and gab ports have successfully started.
- 5** If the node is configured in a secured mode, verify the VxSS service group status.

Manually configuring VCS

This chapter includes the following topics:

- [Configuring LLT for a manual installation](#)
- [Configuring GAB for a manual installation](#)
- [Configuring VCS](#)
- [Starting LLT, GAB, and VCS for a manual installation](#)
- [Modifying the VCS configuration](#)

Configuring LLT for a manual installation

VCS uses the Low Latency Transport (LLT) protocol for all cluster communications as a high-performance, low-latency replacement for the IP stack. LLT has two major functions.

It handles the following tasks:

- Traffic distribution
- Heartbeat traffic

To configure LLT, perform the following steps on each node in the cluster:

- Set up the file `/etc/llthosts`.
See [“Setting up /etc/llthosts for a manual installation”](#) on page 202.
- Set up the file `/etc/llttab`.
See [“Setting up /etc/llttab for a manual installation”](#) on page 202.
- Edit the following file on each node in the cluster to change the values of the `LLT_START` and the `LLT_STOP` environment variables to 1:
`/etc/sysconfig/llt`

Setting up /etc/llthosts for a manual installation

The file `llthosts(4)` is a database. It contains one entry per system that links the LLT system ID (in the first column) with the LLT host name. You must ensure that contents of this file are identical on all the nodes in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.

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

```
0 galaxy
1 nebula
```

Setting up /etc/llttab for a manual installation

The `/etc/llttab` file must specify the system's ID number (or its node name), its cluster ID, 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 for a manual installation](#)” on page 202.

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

```
set-node node name
set-cluster cluster_id
link eth1 eth-MAC_address - ether - -
link eth2 eth-MAC_address - ether - -
```

The first line must identify the system where the file exists. In the example, the value for `set-node` can be: `galaxy`, `0`, or the file name `/etc/nodename`. The file needs to contain the name of the system (`galaxy` 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 `llttab` file in `/opt/VRTSllt`.

LLT directives for a manual installation

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

[Table 15-1](#) contains the LLT directives for a manual installation.

Table 15-1 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, which is 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> specifies the network interface to use. For bonds or <code>vlan</code> interfaces, use the interface name. For standard network interfaces, Symantec recommends the usage of <code>eth-mac</code> to specify the corresponding network interface.</p> <p>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. In addition to enabling VCS communication, it 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 for a manual installation

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

Configuring GAB for a manual installation

VCS uses the Group Membership Services/Atomic Broadcast (GAB) protocol for cluster membership and reliable cluster communications. GAB has two major functions.

It handles the following tasks:

- Cluster membership
- Cluster communications

To configure GAB

- 1 Set up an `/etc/gabtab` configuration file on each node in the cluster using `vi` or another editor. The following example shows an `/etc/gabtab` file:

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

Where the `-c` option configures the driver for use. The `-nN` option specifies that the cluster is not formed until at least `N` systems are ready to form the cluster. Symantec recommends that you set `N` to be the total number of systems in the cluster.

Warning: Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. Using `-c -x` can lead to a split-brain condition.

- 2 Edit the following file on each node in the cluster to change the values of the `GAB_START` and the `GAB_STOP` environment variables to 1:

```
/etc/sysconfig/gab
```

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.

types.cf file Note that the "include" statement in main.cf refers to the types.cf file. 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.

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

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

For a full description of the main.cf file, and how to edit and verify it, refer to the *Veritas Cluster Server Administrator's Guide*.

To configure VCS manually

- 1 Log on as superuser, and move to the directory that contains the configuration file:

```
# cd /etc/VRTSvcs/conf/config
```

- 2 Use vi or another text editor to edit the main.cf file, defining your cluster name and system names. Refer to the following example.

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

```
include "types.cf"
cluster VCSCluster2 ( )
system galaxy ( )
system nebula ( )
```

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

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

- 3 Save and close the main.cf file.
- 4 Edit the following file on each node in the cluster to change the values of the VCS_START and the VCS_STOP environment variables to 1:

```
/etc/sysconfig/vcs
```

Configuring the cluster UUID when creating a cluster manually

You need to configure the cluster UUID when you manually create a cluster.

To configure the cluster UUID when you create a cluster manually

- ◆ On one node in the cluster, perform the following command to populate the cluster UUID on each node in the cluster.

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -configure nodeA  
nodeB ... nodeN
```

Where *nodeA*, *nodeB*, through *nodeN* are the names of the cluster nodes.

Starting LLT, GAB, and VCS for a manual installation

Start LLT, GAB, and VCS.

To start LLT

- ◆ On each node, type:

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

If LLT is configured correctly on each node, the console output resembles:

```
Loading LLT Driver...  
Starting LLT...  
Starting LLT done.
```

See [“Verifying LLT”](#) on page 285.

To start GAB

- ◆ On each node, type:

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

If GAB is configured correctly on each node, the console output resembles:

```
GAB: Starting  
GAB: Starting Done
```

See [“Verifying GAB”](#) on page 288.

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: galaxy
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 (galaxy) 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 (nebula) is in Regular Membership
- Membership: 0x3
```

See “[Verifying the cluster](#)” on page 290.

Modifying the VCS configuration

After the successful installation of VCS, you can modify the configuration of VCS using several methods. You can dynamically modify the configuration from the command line, Veritas Cluster Server Management Console, or the Cluster Manager (Java Console). For information on management tools, refer to the *Veritas Cluster Server Administrator's Guide*.

You can also edit the main.cf file directly. For information on the structure of the main.cf file, refer to the *Veritas Cluster Server Administrator's Guide*.

Configuring the ClusterService group

When you have installed VCS, and verified that LLT, GAB, and VCS work, you can create a service group to include the optional features. These features include the VCS notification components, and the Global Cluster option. If you manually

added VCS to your cluster systems, you must manually create the ClusterService group. Presented in this guide is a reference configuration example of a system with a ClusterService group.

See [“Sample main.cf file for VCS clusters”](#) on page 279.

Manually configuring the clusters for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing manually](#)
- [Setting up server-based I/O fencing manually](#)

Setting up disk-based I/O fencing manually

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

Table 16-1 Tasks to set up I/O fencing manually

Action	Description
Initializing disks as VxVM disks	See “Initializing disks as VxVM disks” on page 129.
Identifying disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 210.
Checking shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 132.
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 210.
Creating I/O fencing configuration files	See “Creating I/O fencing configuration files” on page 211.
Modifying VCS configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 212.

Table 16-1 Tasks to set up I/O fencing manually (*continued*)

Action	Description
Configuring Coordination Point agent to monitor coordination points	See “Configuring Coordination Point agent to monitor coordination points” on page 222.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 213.

Identifying disks to use as coordinator disks

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

See [“Initializing disks as VxVM disks”](#) on page 129.

To identify the coordinator disks

- 1 List the disks on each node.

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

```
# vxdisk -o alldgs list
```

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

See [“Checking shared disks for I/O fencing”](#) on page 132.

Setting up coordinator disk groups

From one node, create a disk group named `vxencoorddg`. 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.

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

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

To create the vxfencoorddg disk group

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

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

- 2 Set the coordinator attribute value as "on" for the coordinator disk group.

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

- 1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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

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

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

- 4 Edit the following file on each node in the cluster to change the values of the VXFEN_START and the VXFEN_STOP environment variables to 1:

```
/etc/sysconfig/vxfen
```

Modifying VCS configuration to use I/O fencing

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

To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

- 3 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

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

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

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

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

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

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

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

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

For example, on each remaining node, enter:

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

- 9 Start the I/O fencing driver and VCS. Perform the following steps on each node:

- Start the I/O fencing driver.

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

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

- Start VCS.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

- ◆ On one of the nodes, type:

```
# vxfenadm -d

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

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

    * 0 (galaxy)
      1 (nebula)

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

Setting up server-based I/O fencing manually

Tasks that are involved in setting up server-based I/O fencing manually include:

Table 16-2 Tasks to set up server-based I/O fencing manually

Action	Description
Preparing the CP servers for use by the VCS cluster	See “Preparing the CP servers manually for use by the VCS cluster” on page 215.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See “Configuring server-based fencing on the VCS cluster manually” on page 218.
Configuring Coordination Point agent to monitor coordination points	See “Configuring Coordination Point agent to monitor coordination points” on page 222.
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 224.

Preparing the CP servers manually for use by the VCS cluster

Use this procedure to manually prepare the CP server for use by the VCS cluster or clusters.

[Table 16-3](#) displays the sample values used in this procedure.

Table 16-3 Sample values in procedure

CP server configuration component	Sample name
CP server	mycps1.symantecexample.com
Node #1 - VCS cluster	galaxy
Node #2 - VCS cluster	nebula
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the VCS cluster

- 1 Determine the cluster name and uuid on the VCS cluster.

For example, issue the following commands on one of the VCS cluster nodes (galaxy):

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster clus1

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2}
```

- 2 Check whether the VCS cluster and nodes are present in the CP server.

```
# cpsadm -s mycps1.symantecexample.com -a list_nodes

ClusName  UUID                               Hostname(Node ID)  Registered
clus1     {f0735332-1dd1-11b2} galaxy(0)           0
clus1     {f0735332-1dd1-11b2} nebula(1)          0
```

If the output does not show the cluster and nodes, then add them as described in the next step.

3 Add the VCS cluster and nodes to each CP server.

For example, issue the following command on the CP server (mycps1.symantecexample.com) to add the cluster:

```
# cpsadm -s mycps1.symantecexample.com -a add_clus\  
-c clus1 -u {f0735332-1dd1-11b2}
```

```
Cluster clus1 added successfully
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the first node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c clus1 -u {f0735332-1dd1-11b2} -h galaxy -n0
```

```
Node 0 (galaxy) successfully added
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the second node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c clus1 -u {f0735332-1dd1-11b2} -h nebula -n1
```

```
Node 1 (nebula) successfully added
```

4 If security is to be enabled, check whether the `_HA_VCS_users` are created in the CP server.

If the output below does not show the users, then add them as described in the next step.

```
# cpsadm -s mycps1.symantecexample.com -a list_users
```

Username/Domain Type	Cluster Name / UUID	Role
<code>_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com/vx</code>	<code>clus1/{f0735332-1dd1-11b2}</code>	Operator
<code>_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com/vx</code>	<code>clus1/{f0735332-1dd1-11b2}</code>	Operator

If security is to be disabled, then add the user name "cpsclient@hostname" to the server instead of the `_HA_VCS_users` (for example, cpsclient@galaxy).

The CP server can only run in either secure mode or non-secure mode, both connections are not accepted at the same time.

5 Add the users to the CP server.

First, determine the `user@domain` to be added.

The user for fencing should be of the form `_HA_VCS_`*short-hostname* and domain name is that of HA_SERVICES user in the output of command:

```
# /opt/VRTScps/bin/cpsat listpd -t local
```

Next, issue the following commands on the CP server (`mycps1.symantecexample.com`):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com  
successfully added
```

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com  
successfully added
```

- 6 Authorize the CP server user to administer the VCS cluster. You must perform this task for the CP server users corresponding to each node in the VCS cluster.

For example, issue the following command on the CP server (mycps1.symantecexample.com) for VCS cluster clus1 with two nodes galaxy and nebula:

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c clus1\  
-u {f0735332-1dd1-11b2}\  
-e _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-f cps_operator -g vx
```

```
Cluster successfully added to user  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com privileges.
```

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c clus1\  
-u {f0735332-1dd1-11b2}\  
-e _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com\  
-f cps_operator -g vx
```

```
Cluster successfully added to user  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com privileges.
```

Configuring server-based fencing on the VCS cluster manually

The configuration process for the client or VCS cluster to use CP server as a coordination point requires editing the /etc/vxfenmode file. You need to edit this file to specify the following information for your configuration:

- Fencing mode
- Fencing mechanism
- Fencing disk policy (if applicable to your I/O fencing configuration)
- Appropriate value for the security configuration
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)

Whenever coordinator disks are used as coordination points in your I/O fencing configuration, a disk group (vxfsd) has to be created. This disk group has to be

specified in the `/etc/vxfenmode` file. For information about creating the disk group, see the *Veritas™ Cluster Server Installation Guide*.

The customized fencing framework also generates the `/etc/vxfentab` file which has security setting and the coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

Edit the following file on each node in the cluster to change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1:

`/etc/sysconfig/vxfen`

Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

The following file output provides an example of what the `/etc/vxfenmode` file contains:

```
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps        - use a coordination point server with optional script
#              controlled scsi3 disks
#
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
```

```
scsi3_disk_policy=dmp

# security when enabled uses secure communication to the cp server
# using VxAT (Veritas Authentication Service)
# available options:
# 0 - don't use Veritas Authentication Service for cp server
#   communication
# 1 - use Veritas Authentication Service for cp server
#   communication
security=1

#
# Specify 3 or more odd number of coordination points in this file,
# one in each row. They can be all-CP servers, all-SCSI-3 compliant
# coordinator disks, or a combination of CP servers and SCSI-3
# compliant coordinator disks. Please ensure that the CP server
# coordination points are numbered sequentially and in the same
# order on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as:
#
# cps<number>=<Virtual IP/ Virtual hostname of cp server> in
# square brackets ([]), followed by ":" and CPS port number.
#
# Examples:
# cps1=[192.168.0.23]:14250
# cps2=[mycps.company.com]:14250
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoorddg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
```

```
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
```

[Table 16-4](#) defines the vxfenmode parameters that must be edited.

Table 16-4 vxfenmode file parameters

vxfenmode File Parameter	Description
vxfen_mode	Fencing mode of operation. This parameter must be set to "customized".
vxfen_mechanism	Fencing mechanism. This parameter defines the mechanism that is used for fencing. If one of the three coordination points is a CP server, then this parameter must be set to "cps".
scsi3_disk_policy	Configure the vxfen module to use either DMP devices, "dmp" or the underlying raw character devices, "raw". Note: The configured disk policy is applied on all the nodes.
security	Security parameter 1 indicates that Symantec Product Authentication Service is used for CP server communications. Security parameter 0 indicates that communication with the CP server is made in non-secure mode. The default security value is 1. Note: Symantec only supports a configuration where both the CP server and client sides have the same security setting. The security setting on both sides must be either enabled or disabled.

Table 16-4 vxfenmode file parameters (*continued*)

vxfenmode File Parameter	Description
cps1, cps2, cps3, or vxfendg	Coordination point parameters. Enter either the Virtual IP address or FQHN (whichever is accessible) of the CP server. Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfendg) and specified in the /etc/vxfenmode file. Additionally, the customized fencing framework also generates the /etc/vxfentab file which specifies the security setting and the coordination points (all the CP servers and the disks from disk group specified in /etc/vxfenmode file).

After editing the /etc/vxfenmode file, run the vxfen init script to start fencing.

For example:

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

Configuring Coordination Point agent to monitor coordination points

The following procedure describes how to manually configure the Coordination Point agent to monitor coordination points (CP server or SCSI-3 disks).

See the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on the agent.

To configure Configuration Point agent to monitor coordination points

- 1 Ensure that your VCS cluster has been properly installed and configured with fencing enabled.
- 2 Create a parallel service group `vxfen` and add a `coordpoint` resource to the `vxfen` service group as follows:

```
# haconf -makerw
# hagr -add vxfen
# hagr -modify vxfen SystemList system1 0 system2 1
# hagr -modify vxfen AutoFailOver 0
# hagr -modify vxfen Parallel 1
# hagr -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 1
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

- 3 Verify the status of the agent on the VCS cluster using the `hares` commands.

For example:

```
# hares -state coordpoint
```

The following is an example of the command and output::

```
# hares -state

# Resource      Attribute      System      Value
coordpoint     State          galaxy      ONLINE
```

- 4 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed Coordination Point agent monitoring information; including information about whether the Coordination Point agent is able to access all the coordination points, information to check on which coordination points the Coordination Point agent is reporting missing keys, etc.

To view all such information in the engine log, change the `dbg` level for that node using the following command:

```
# hatype -modify coordpoint LogDbg 10
```

The agent log can now be viewed at the following location:

```
/var/VRTSvcS/log/engine_A.log
```

Verifying server-based I/O fencing configuration

During the VCS cluster installation, the installer populates the following files based on inputs that are received during the configuration phase:

- `/etc/vxfenmode` (edited for CP server)
- `/etc/vxfentab` (edited for CP server)

Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

For troubleshooting server-based I/O fencing configuration issues, refer to the *Veritas Cluster Server Administrator's Guide*.

Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```

Upgrading VCS

- [Chapter 17. Planning to upgrade VCS](#)
- [Chapter 18. Performing a typical VCS upgrade using the installer](#)
- [Chapter 19. Performing a phased upgrade](#)
- [Chapter 20. Performing an automated VCS upgrade using response files](#)

Planning to upgrade VCS

This chapter includes the following topics:

- [About upgrading to VCS 5.1](#)
- [VCS supported upgrade paths](#)
- [Upgrading VCS in secure enterprise environments](#)
- [About phased upgrade](#)

About upgrading to VCS 5.1

You can upgrade VCS using one of the following methods:

- Typical upgrade using Veritas product installer or the `installvcs` program
See [“VCS supported upgrade paths”](#) on page 228.
See [“Upgrading VCS using the script-based installer”](#) on page 234.
- Typical upgrade Veritas Web installer
See [“VCS supported upgrade paths”](#) on page 228.
See [“Upgrading VCS with the Veritas Web-based installer”](#) on page 235.
- Phased upgrade to reduce downtime
See [“Performing a phased upgrade from VCS 5.0 MP3”](#) on page 237.
- Automated upgrade using response files
See [“VCS supported upgrade paths”](#) on page 228.
See [“Upgrading VCS using response files”](#) on page 253.

You can upgrade VCS 5.1 to Storage Foundation High Availability 5.1 using Veritas product installer or response files.

See the *Veritas Storage Foundation and High Availability Installation Guide*.

VCS supported upgrade paths

If you are currently running a cluster with any earlier VCS versions that is supported for upgrade, you can run the installer to upgrade to VCS 5.1.

Review the supported upgrade path tables for VCS clusters on RHEL and SLES operating systems.

The following variations apply to the upgrade paths:

- To upgrade VCS 4.1 MP4 on RHEL 5:
 - Upgrade to RHEL 5 U3.
 - Upgrade VCS to 4.1 MP4 RP3.
 - Upgrade to VCS 5.1
- To upgrade to VCS 5.1 on SLES10 SP2
 - Upgrade to SLES 10 SP2
 - Upgrade VCS to 4.1 MP4 RP3
 - Upgrade to VCS 5.1

[Table 17-1](#) lists the supported upgrade paths for Red Hat Enterprise Linux and Oracle Enterprise Linux.

Table 17-1 Supported upgrade paths for RHEL and OEL

Upgrade scenarios	From VCS-RHEL	To VCS-RHEL
VCS upgrade and RHEL upgrade	VCS 4.1 MP4 on RHEL 5 VCS 4.1 MP4 RP3 on RHEL 5 U2 VCS 5.0 MP3 on RHEL 5 U1 VCS 5.0 MP3 on RHEL 5 U2	VCS 5.1 on RHEL 5 U3
VCS upgrade and OEL upgrade	VCS 5.0 MP3 on OEL 5 U1	VCS 5.1 on OEL 5 U3

[Table 17-2](#) lists the supported upgrade paths for SUSE Linux Enterprise Server

Table 17-2 Supported upgrade paths for SUSE Linux Enterprise Server

Upgrade scenarios	From VCS-SLES	To VCS-SLES
VCS upgrade and SLES upgrade	VCS 4.1MP3 on SLES 10 VCS 4.1 MP4 on SLES 10 SP1	VCS 5.1 on SLES 10 SP2

Table 17-2 Supported upgrade paths for SUSE Linux Enterprise Server
(continued)

Upgrade scenarios	From VCS-SLES	To VCS-SLES
VCS upgrade	VCS 5.0 MP3 on SLES 10 SP2	VCS 5.1 on SLES 10 SP2
	VCS 5.0 RU1 on SLES 11	VCS 5.1 on SLES 11

Upgrading VCS in secure enterprise environments

In secure enterprise environments, ssh or rsh communication is not allowed between systems. In such cases, the `installvcs` program can upgrade VCS only on systems with which it can communicate (most often the local system only).

Run the `installvcs` program on each node to upgrade the cluster to VCS 5.1. On each node, the `installvcs` program updates the configuration, stops the cluster, and then upgrades VCS on the node. After the last node is upgraded and started, the upgrade is complete.

Warning: If you are running the cluster in secure mode, make sure to remove the file `/tmp/disable_selfcont` from the cluster before upgrading to VCS 5.1.

About phased upgrade

Perform a phased upgrade to minimize the downtime for the cluster. Depending on the situation, you can calculate the approximate downtime as follows:

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

Prerequisites for a phased upgrade

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

Planning for a phased upgrade

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

Some rough guidelines follow:

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

Phased upgrade limitations

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

- 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.

Phased 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. For the purposes of this example, the cluster is split into two subclusters. The nodes node01 and node02 are in the first subcluster, which you first upgrade. The nodes node03 and node04 are in the second subcluster, which you upgrade last.

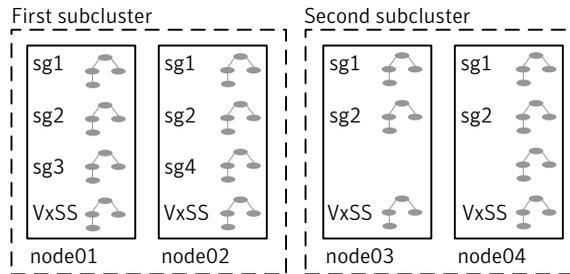
Each service group is running on the nodes as follows:

- sg1 and sg2 are parallel service groups and run on all the nodes.
- sg3 and sg4 are failover service groups. sg3 runs on node01 and sg4 runs on node02.
- VxSS service group runs on all nodes (secure mode is enabled)

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

- sg1 and sg2 are running on all the nodes.
- sg3 and sg4 can fail over to any of the nodes in the cluster.
- VxSS service group runs on all nodes

Figure 17-1 Example of phased upgrade set up



Phased upgrade example overview

This example's upgrade path follows:

- Move all the service groups from the first subcluster to the second subcluster.
- Upgrade the operating system on the first subcluster's nodes, if required.
- On the first subcluster, start the upgrade using the installation program.
- Get the second subcluster ready.
- Activate the first subcluster.
- Upgrade the operating system on the second subcluster's nodes, if required.
- On the second subcluster, start the upgrade using the installation program.
- Activate the second subcluster.

See [“Performing a phased upgrade from VCS 5.0 MP3”](#) on page 237.

Performing a typical VCS upgrade using the installer

This chapter includes the following topics:

- [Before upgrading from 4.x using the script-based or Web-based installer](#)
- [Upgrading VCS using the script-based installer](#)
- [Upgrading VCS with the Veritas Web-based installer](#)

Before upgrading from 4.x using the script-based or Web-based installer

Before you upgrade VCS, perform the following steps if you are upgrading from VCS 4.x. You first need to remove deprecated resource types and modify changed values.

To prepare to upgrade to VCS 5.1 from VCS 4.x

- 1 Remove deprecated resources and modify attributes. The installer program can erase obsolete types and resources can be erased from the system or you can manually remove them.

See [“Manually removing deprecated resource types and modifying attributes”](#) on page 365.

- 2 Stop the application agents that are installed on the VxVM disk (for example the NBU agent).

Perform the following steps to stop the application agents:

- Take the resources offline on all systems that you want to upgrade.

```
# hares -offline resname -sys sysname
```

- Stop the application agents that are installed on VxVM disk on all the systems.

```
# haagent -stop agentname -sys sysname
```

- Ensure that the agent processes are not running.

```
# ps -ef | grep Agent
```

This command does not list any processes in the VxVM installation directory.

- 3 Make sure that LLT, GAB, and VCS are running on all of the nodes in the cluster. The installer program cannot proceed unless these processes are running.

```
# lltconfig
```

LLT is running

```
# gabconfig -a
```

```
=====  
Port a gen cc701 membership 01  
Port h gen cc704 membership 01
```

Upgrading VCS using the script-based installer

You can use the product installer to upgrade VCS.

To upgrade VCS using the product installer

- 1 Log in as superuser and mount the product disc.
- 2 Start the installer.

```
# ./installer
```

The installer starts the product installation program with a copyright message. It then specifies where it creates the logs. Note the log's directory and name.

- 3 From the opening Selection Menu, choose: **G** for "Upgrade a Product."
- 4 Enter the names of the nodes that you want to upgrade. Use spaces to separate node names. Press the Enter key to proceed.

The installer runs some verification checks on the nodes.

- 5 When the verification checks are complete, press the Enter key to continue.
The installer lists the RPMs to upgrade.
- 6 The installer asks if you want to stop VCS processes. Press the Enter key to continue.
The installer stops VCS processes, uninstalls RPMs, installs, upgrades, and configures VCS.
- 7 The installer lists the nodes that Symantec recommends you restart.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native OS accounts.

See [“Creating new VCS accounts if you used native operating system accounts”](#) on page 366.

Upgrading VCS with the Veritas Web-based installer

This section describes upgrading VCS with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

To upgrade VCS

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 152.
- 3 Select **Upgrade**.
The installer detects the product that is installed on the specified system.
- 4 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.
- 5 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Validate**.
- 6 Click **Next** to complete the upgrade.
After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.
- 7 Click **Finish**. The installer prompts you for another task.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native operating system accounts.

See [“Creating new VCS accounts if you used native operating system accounts”](#) on page 366.

Performing a phased upgrade

This chapter includes the following topics:

- [Performing a phased upgrade from VCS 5.0 MP3](#)

Performing a phased upgrade from VCS 5.0 MP3

This section explains how to perform a phased upgrade of VCS on four nodes with four service groups. Note that in this scenario, the service groups cannot stay online during the upgrade of the second subcluster. Do not add, remove, or change resources or service groups on any nodes during the upgrade. These changes are likely to get lost after the upgrade. The following example illustrates the steps to perform a phased upgrade. The phased upgrade is from VCS 5.0 MP3 in a secure cluster to VCS 5.1 in a secure cluster.

See [“About phased upgrade”](#) on page 229.

Moving the service groups to the second subcluster

Perform the following steps to establish the service group's status and to switch the service groups.

To move service groups to the second subcluster

- 1 On the first subcluster, determine where the service groups are online.

```
# hagr -state
```

The output resembles:

```
#Group  Attribute System Value
sg1     State     node01  |ONLINE|
sg1     State     node02  |ONLINE|
sg1     State     node03  |ONLINE|
sg1     State     node04  |ONLINE|
sg2     State     node01  |ONLINE|
sg2     State     node02  |ONLINE|
sg2     State     node03  |ONLINE|
sg2     State     node04  |ONLINE|
sg3     State     node01  |ONLINE|
sg3     State     node02  |OFFLINE|
sg3     State     node03  |OFFLINE|
sg3     State     node04  |OFFLINE|
sg4     State     node01  |OFFLINE|
sg4     State     node02  |ONLINE|
sg4     State     node03  |OFFLINE|
sg4     State     node04  |OFFLINE|
VxSS    State     node01  |ONLINE|
VxSS    State     node02  |ONLINE|
VxSS    State     node03  |ONLINE|
VxSS    State     node04  |ONLINE|
```

- 2 Offline the parallel service groups (sg1 and sg2) and the VXSS group from the first subcluster. Switch the failover service groups (sg3 and sg4) from the first subcluster (node01 and node02) to the nodes on the second subcluster (node03 and node04).

```
# hagr -offline sg1 -sys node01
# hagr -offline sg2 -sys node01
# hagr -offline sg1 -sys node02
# hagr -offline sg2 -sys node02
# hagr -offline VxSS -sys node01
# hagr -offline VxSS -sys node02
# hagr -switch sg3 -to node03
# hagr -switch sg4 -to node04
```

- 3 On the nodes in the first subcluster, unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -h
Filesystem              Size  Used Avail Use% Mounted on
/dev/sda1                26G   3.3G   22G   14% /
udev                    1007M   352K 1006M    1% /dev
tmpfs                    4.0K      0   4.0K    0% /dev/vx
/dev/vx/dsk/dg2/dg2vol1 3.0G    18M   2.8G    1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2 1.0G    18M   944M    2% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3 10G    20M   9.4G    1% /mnt/dg2/dg2vol3

# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

- 4 On the nodes in the first subcluster, stop all VxVM volumes (for each disk group) that VCS does not manage.
- 5 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

- 6 Freeze the nodes in the first subcluster.

```
# hasys -freeze -persistent node01
# hasys -freeze -persistent node02
```

- 7 Dump the configuration and make it read-only.

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

- 8 Verify that the service groups are offline on the first subcluster that you want to upgrade.

```
# hagrps -state
```

Output resembles:

```
#Group Attribute System Value
sg1 State node01 |OFFLINE|
sg1 State node02 |OFFLINE|
sg1 State node03 |ONLINE|
sg1 State node04 |ONLINE|
sg2 State node01 |OFFLINE|
sg2 State node02 |OFFLINE|
sg2 State node03 |ONLINE|
sg2 State node04 |ONLINE|
sg3 State node01 |OFFLINE|
sg3 State node02 |OFFLINE|
sg3 State node03 |ONLINE|
sg3 State node04 |OFFLINE|
sg4 State node01 |OFFLINE|
sg4 State node02 |OFFLINE|
sg4 State node03 |OFFLINE|
sg4 State node04 |ONLINE|
VxSS State node01 |OFFLINE|
VxSS State node02 |OFFLINE|
VxSS State node03 |ONLINE|
VxSS State node04 |ONLINE|
```

- 9 Perform this step on the nodes (node01 and node02) in the first subcluster if the cluster uses I/O Fencing. Use an editor of your choice and change the following:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `scsi3` to `disabled`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. You want the line in the `main.cf` file to resemble:

```
UseFence = NONE
```

- 10 Back up the llttab, llthosts, gabtab, types.cf, main.cf and AT configuration files on the first subcluster.

```
# cp /etc/llttab /etc/llttab.bkp
# cp /etc/llthosts /etc/llthosts.bkp
# cp /etc/gabtab /etc/gabtab.bkp
# cp /etc/VRTSvcs/conf/config/main.cf \
    /etc/VRTSvcs/conf/config/main.cf.bkp
# cp /etc/VRTSvcs/conf/config/types.cf \
    /etc/VRTSvcs/conf/config/types.cf.bkp
# /opt/VRTSat/bin/vssat showbackuplist
B|/var/VRTSat/.VRTSat/profile/VRTSatlocal.conf
B|/var/VRTSat/.VRTSat/profile/certstore
B|/var/VRTSat/ABAuthSource
B|/etc/vx/vss/VRTSat.conf
Quiescing ...
Snapshot Directory :/var/VRTSatSnapShot
```

Upgrading the operating system on the first subcluster

You can perform the operating system upgrade on the first subcluster, if required. Refer to the operating system's documentation for more information.

Upgrading the first subcluster

You now navigate to the installer program and start it.

To start the installer for the phased upgrade

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

```
# cd /cluster_server
```

- 3 Make sure that VCS is running. Start the installvcs program, specify the nodes in the first subcluster (node1 and node2).

```
# ./installvcs node1 node2
```

The program starts with a copyright message and specifies the directory where it creates the logs.

4 Enter **y** to agree to the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement
as specified in the EULA.pdf file present on media? [y,n,q,?] y
```

5 Review the available installation options.

See “[Veritas Cluster Server installation packages](#)” on page 345.

- 1 Installs only the minimal required VCS RPMs that provides basic functionality of the product.
- 2 Installs the recommended VCS RPMs that provides complete functionality of the product. This option does not install the optional VCS RPMs.
Note that this option is the default.
- 3 Installs all the VCS RPMs.
You must choose this option to configure any optional VCS feature.
- 4 Displays the VCS RPMs for each option.

For this example, select **3** for all RPMs.

```
Select the RPMs to be installed on all systems? [1-4,q,?]
(2) 3
```

6 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility.

7 When you are prompted, reply **y** to continue with the upgrade.

```
Do you want to continue? [y,n,q] (y)
```

- 8** When you are prompted, reply **y** to stop appropriate processes.

```
Do you want to stop VCS processes? [y,n,q] (y)
```

- 9** The installer ends for the first subcluster with the following output:

```
Configuring VCS: 100%

Estimated time remaining: 0:00

Performing VCS upgrade configuration ..... Done

Veritas Cluster Server Configure completed successfully

You are performing phased upgrade (Phased 1) on the systems.
Follow the steps in install guide to upgrade the remaining
systems.

Would you like to send the information about this installation to
Symantec to help improve installation in the future? [y,n,q,?] (y)

The upgrade is finished on the first subcluster. Do not reboot the nodes in
the first subcluster until you complete the Preparing the second subcluster
procedure.
```

Preparing the second subcluster

Perform the following steps on the second subcluster before rebooting nodes in the first subcluster.

To prepare to upgrade the second subcluster

1 Get the summary of the status of your resources.

```
# hastatus -summ
-- SYSTEM STATE
-- System                State                Frozen

A  node01                EXITED                1
A  node02                EXITED                1
A  node03                RUNNING              0
A  node04                RUNNING              0

-- GROUP STATE
-- Group                System  Probed    AutoDisabled  State

B  SG1                  node01  Y         N              OFFLINE
B  SG1                  node02  Y         N              OFFLINE
B  SG1                  node03  Y         N              ONLINE
B  SG1                  node04  Y         N              ONLINE
B  SG2                  node01  Y         N              OFFLINE
B  SG2                  node02  Y         N              OFFLINE
B  SG2                  node03  Y         N              ONLINE
B  SG2                  node04  Y         N              ONLINE
B  SG3                  node01  Y         N              OFFLINE
B  SG3                  node02  Y         N              OFFLINE
B  SG3                  node03  Y         N              ONLINE
B  SG3                  node04  Y         N              OFFLINE
B  SG4                  node01  Y         N              OFFLINE
B  SG4                  node02  Y         N              OFFLINE
B  SG4                  node03  Y         N              OFFLINE
B  SG4                  node04  Y         N              ONLINE
B  VxSS                 node01  Y         N              OFFLINE
B  VxSS                 node02  Y         N              OFFLINE
B  VxSS                 node03  Y         N              ONLINE
B  VxSS                 node04  Y         N              ONLINE
```

- 2 Unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -h
Filesystem              Size  Used Avail Use% Mounted on
/dev/sda1                26G   3.3G   22G  14% /
udev                    1007M   352K 1006M   1% /dev
tmpfs                    4.0K     0   4.0K   0% /dev/vx
/dev/vx/dsk/dg2/dg2vol1 3.0G   18M   2.8G   1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2 1.0G   18M   944M   2% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3 10G   20M   9.4G   1% /mnt/dg2/dg2vol3

# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

- 3 Stop all VxVM volumes (for each disk group) that VCS does not manage.

- 4 Make the configuration writable on the second subcluster.

```
# haconf -makerw
```

- 5 Unfreeze the service groups.

```
# hagrps -unfreeze sg1 -persistent
# hagrps -unfreeze sg2 -persistent
# hagrps -unfreeze sg3 -persistent
# hagrps -unfreeze sg4 -persistent
# hagrps -unfreeze VxSS -persistent
```

- 6 Dump the configuration and make it read-only.

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

7 Take the service groups offline on node03 and node04.

```
# hagr -offline sg1 -sys node03
# hagr -offline sg1 -sys node04
# hagr -offline sg2 -sys node03
# hagr -offline sg2 -sys node04
# hagr -offline sg3 -sys node03
# hagr -offline sg4 -sys node04
# hagr -offline VxSS -sys node03
# hagr -offline VxSS -sys node04
```

8 Verify the state of the service groups.

```
# hagr -state
#Group      Attribute  System  Value
SG1         State     node01  |OFFLINE|
SG1         State     node02  |OFFLINE|
SG1         State     node03  |OFFLINE|
SG1         State     node04  |OFFLINE|
SG2         State     node01  |OFFLINE|
SG2         State     node02  |OFFLINE|
SG2         State     node03  |OFFLINE|
SG2         State     node04  |OFFLINE|
SG3         State     node01  |OFFLINE|
SG3         State     node02  |OFFLINE|
SG3         State     node03  |OFFLINE|
SG3         State     node04  |OFFLINE|
VxSS       State     node01  |OFFLINE|
VxSS       State     node02  |OFFLINE|
VxSS       State     node03  |OFFLINE|
VxSS       State     node04  |OFFLINE|
```

9 Perform this step on node03 and node04 if the cluster uses I/O Fencing. Use an editor of your choice and change the following:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `scsi3` to `disabled`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. You want the line in the `main.cf` file to resemble:

```
UseFence = NONE
```

10 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

```
# hstop -local
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

11 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 not loaded.

```
# /etc/init.d/vxfen status
VXFEN module is not loaded

# /etc/init.d/gab status
GAB module is not loaded

# /etc/init.d/llt status
LLT module is not loaded
```

Activating the first subcluster

Get the first subcluster ready for the service groups.

Note: These steps fulfill part of the installer's output instructions, see [Upgrading the first subcluster](#) step 9.

To activate the first subcluster

1 Perform this step on node01 and node02 if the cluster uses I/O Fencing. Use an editor of your choice and revert the following to an enabled state before you reboot the first subcluster's nodes:

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `NONE` to `SCSI3`. You want the line in the `main.cf` file to resemble:

```
UseFence = SCSI3
```

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `disabled` to `scsi3`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=scsi3
```

- 2 Reboot the node01 and node02 in the first subcluster.

```
# /sbin/shutdown -r now
```

- 3 Seed node01 and node02 in the first subcluster.

```
# gabconfig -xc
```

- 4 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

- 5 Unfreeze the nodes in the first subcluster.

```
# hasys -unfreeze -persistent node01  
# hasys -unfreeze -persistent node02
```

- 6 Dump the configuration and make it read-only.

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

- 7 Bring the service groups online on node01 and node02.

```
# hagrps -online sg1 -sys node01  
# hagrps -online sg1 -sys node02  
# hagrps -online sg2 -sys node01  
# hagrps -online sg2 -sys node02  
# hagrps -online sg3 -sys node01  
# hagrps -online sg4 -sys node02  
# hagrps -online VxSS -sys node01  
# hagrps -online VxSS -sys node02
```

Upgrading the operating system on the second subcluster

You can perform the operating system upgrade on the second subcluster, if required. Refer to the operating system's documentation for more information.

Before you perform the operating system upgrade, make sure to disable VCS, VXFEN, GAB, and LLT.

To disable VCS, VXFEN, GAB, and LLT

- ◆ On the second subcluster, perform the following commands:

```
# chkconfig vcs off
# chkconfig vxfen off
# chkconfig gab off
# chkconfig llt off
```

Perform the operating system upgrade. After you finish the operating system, enable VCS, VXFEN, GAB and LLT.

To enable VCS, VXFEN, GAB and LLT

- ◆ On second subcluster, perform following commands:

```
# chkconfig llt on
# chkconfig gab on
# chkconfig vxfen on
# chkconfig vcs on
```

Upgrading the second subcluster

Perform the following procedure to upgrade the second subcluster (node03 and node04).

To start the installer to upgrade the second subcluster

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

```
# cd /cluster_server
```

- 3 Confirm that VCS is stopped on node03 and node04. Start the `installvcs` program, specify the nodes in the second subcluster (node3 and node4).

```
# ./installvcs node3 node4
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 4 Enter **y** to agree to the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement
as specified in the EULA.pdf file present on media? [y,n,q,?] y
```

5 Review the available installation options.

See [“Veritas Cluster Server installation packages”](#) on page 345.

- 1 Installs only the minimal required VCS RPMs that provides basic functionality of the product.
- 2 Installs the recommended VCS RPMs that provides complete functionality of the product. This option does not install the optional VCS RPMs.
Note that this option is the default.
- 3 Installs all the VCS RPMs.
You must choose this option to configure any optional VCS feature.
- 4 Displays the VCS RPMs for each option.

For this example, select **3** for all RPMs.

```
Select the RPMs to be installed on all systems? [1-4,q,?]  
(2) 3
```

6 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility.

7 When you are prompted, reply **y** to continue with the upgrade.

```
Do you want to continue? [y,n,q] (y)
```

8 When you are prompted, reply **y** to stop VCS processes.

```
Do you want to stop VCS processes? [y,n,q] (y)
```

9 Monitor the installer program answering questions as appropriate until the upgrade completes.

Finishing the phased upgrade

You now have to reboot the nodes in the second subcluster.

To finish the upgrade

- 1 Verify that the cluster UUID is the same on the nodes in the second subcluster and the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh]
-clus -display node1 [node2 ...]
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first subcluster to the nodes in the second subcluster. For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus
-copy -from_sys node01 -to_sys node03 node04
```

- 2 Perform this step on node03 and node04 if the cluster uses I/O Fencing. Use an editor of your choice and revert the following to an enabled state before you reboot the second subcluster's nodes:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from disabled to `scsi3`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=scsi3
```

- 3 Reboot the node03 and node04 in the second subcluster.

```
# /sbin/shutdown -r now
```

The nodes in the second subcluster join the nodes in the first subcluster.

- 4 Check to see if VCS and its components are up.

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen  nxxxxnn membership 0123
Port b gen  nxxxxnn membership 0123
Port h gen  nxxxxnn membership 0123
```

5 Run an `hastatus -sum` command to determine the status of the nodes, service groups, and cluster.

```
# hastatus -sum

-- SYSTEM STATE
-- System          State          Frozen

A node01          RUNNING        0
A node02          RUNNING        0
A node03          RUNNING        0
A node04          RUNNING        0

-- GROUP STATE
-- Group           System         Probed   AutoDisabled  State

B VxSS            node01        Y        N              ONLINE
B VxSS            node02        Y        N              ONLINE
B VxSS            node03        Y        N              ONLINE
B VxSS            node04        Y        N              ONLINE
B sg1             node01        Y        N              ONLINE
B sg1             node02        Y        N              ONLINE
B sg1             node03        Y        N              ONLINE
B sg1             node04        Y        N              ONLINE
B sg2             node01        Y        N              ONLINE
B sg2             node02        Y        N              ONLINE
B sg2             node03        Y        N              ONLINE
B sg2             node04        Y        N              ONLINE
B sg3             node01        Y        N              OFFLINE
B sg3             node02        Y        N              OFFLINE
B sg3             node03        Y        N              OFFLINE
B sg3             node04        Y        N              OFFLINE
B sg4             node01        Y        N              OFFLINE
B sg4             node02        Y        N              ONLINE
B sg4             node03        Y        N              OFFLINE
B sg4             node04        Y        N              OFFLINE
```

6 After the upgrade is complete, mount the VxFS file systems and start the VxVM volumes (for each disk group) that VCS does not manage.

In this example, you have performed a phased upgrade of VCS. The service groups were down when you took them offline on node03 and node04, to the time VCS brought them online on node01 or node02.

Performing an automated VCS upgrade using response files

This chapter includes the following topics:

- [Upgrading VCS using response files](#)
- [Response file variables to upgrade VCS](#)
- [Sample response file for upgrading VCS](#)

Upgrading VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS upgrade on one cluster to upgrade VCS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To perform automated VCS upgrade

- 1 Make sure the systems where you want to upgrade VCS meet the upgrade requirements.
- 2 Make sure the pre-upgrade tasks are completed.
- 3 Copy the response file to one of the cluster systems where you want to upgrade VCS.
See [“Sample response file for upgrading VCS”](#) on page 256.
- 4 Edit the values of the response file variables as necessary.
See [“Response file variables to upgrade VCS”](#) on page 254.

- 5 Mount the product disc and navigate to the folder that contains the installation program.
- 6 Start the upgrade from the system to which you copied the response file. For example:

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

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

Response file variables to upgrade VCS

[Table 20-1](#) lists the response file variables that you can define to upgrade VCS.

Table 20-1 Response file variables specific to upgrading VCS

Variable	List or Scalar	Description
CFG{opt}{upgrade}	Scalar	Upgrades VCS RPMs. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{opt}{stopfail_allow}	Scalar	Decides whether or not to proceed if the installer fails while stopping the processes or while unloading the drivers. (Optional)
CFG{opt}{systems}	List	List of systems on which the product is to be upgraded. (Optional)
CFG{prod}	Scalar	Defines the product to be upgraded. The value is VCS51 for VCS. (Required)

Table 20-1 Response file variables specific to upgrading VCS (*continued*)

Variable	List or Scalar	Description
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{patchpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems. (Optional)
CFG{opt}{pkgpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product RPMs. The location must be accessible from all target systems. (Optional)
CFG{opt}{tmppath}	Scalar	Defines the location where a working directory is created to store temporary files and the depots that are needed during the install. The default location is /var/tmp. (Optional)

Table 20-1 Response file variables specific to upgrading VCS (*continued*)

Variable	List or Scalar	Description
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. (Optional)

Sample response file for upgrading VCS

Review the response file variables and their definitions.

See [“Response file variables to upgrade VCS”](#) on page 254.

```
#
# Configuration Values:
#
our %CFG;

$CFG{accepteula}=1;
$CFG{vcs_allowcomms}=1;
$CFG{opt}{upgrade}=1;
$CFG{prod}="VCS51";
$CFG{systems}=[ qw( galaxy nebula ) ];
```

Post-installation tasks

- [Chapter 21. Performing post-installation tasks](#)
- [Chapter 22. Installing or upgrading VCS components](#)
- [Chapter 23. Verifying the VCS installation](#)

Performing post-installation tasks

This chapter includes the following topics:

- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [Accessing the VCS documentation](#)
- [Removing permissions for communication](#)

About enabling LDAP authentication for clusters that run in secure mode

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

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

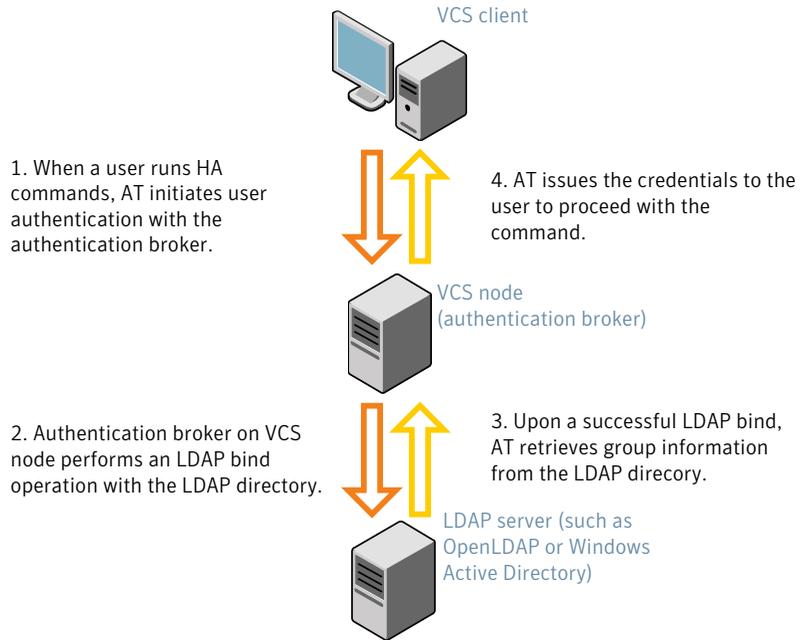
See [“Enabling LDAP authentication for clusters that run in secure mode”](#) on page 261.

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

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

[Figure 21-1](#) depicts the VCS cluster communication with the LDAP servers when clusters run in secure mode.

Figure 21-1 Client communication with LDAP servers



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

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

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

- The type of LDAP schema used (the default is RFC 2307)
 - UserObjectClass (the default is `posixAccount`)
 - UserObject Attribute (the default is `uid`)
 - User Group Attribute (the default is `gidNumber`)
 - Group Object Class (the default is `posixGroup`)
 - GroupObject Attribute (the default is `cn`)
 - Group GID Attribute (the default is `gidNumber`)
 - Group Membership Attribute (the default is `memberUid`)
- URL to the LDAP Directory

- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)
- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 5.0.32.0 or later.

```
# /opt/VRTSAt/bin/vssat showversion
vssat version: 5.0.32.0
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

To enable OpenLDAP authentication for clusters that run in secure mode

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
# /opt/VRTSsat/bin/vssat addldapdomain \  
--domainname "MYENTERPRISE.symantecdomain.com"\  
--server_url "ldap://my_openldap_host.symantecexample.com"\  
--user_base_dn "ou=people,dc=symantecdomain,dc=myenterprise,dc=com"\  
--user_attribute "cn" --user_object_class "account"\  
--user_gid_attribute "gidNumber"\  
--group_base_dn "ou=group,dc=symantecdomain,dc=myenterprise,dc=com"\  
--group_attribute "cn" --group_object_class "posixGroup"\  
--group_gid_attribute "member"\  
--admin_user "cn=manager,dc=symantecdomain,dc=myenterprise,dc=com"\  
--admin_user_password "password" --auth_type "FLAT"
```

- 2 Verify that you can successfully authenticate an LDAP user on the VCS nodes.

You must have a valid LDAP user ID and password to run the command. In the following example, authentication is verified for the MYENTERPRISE domain for the LDAP user, `vcsadmin1`.

```
galaxy# /opt/VRTSsat/bin/vssat authenticate  
--domain ldap:MYENTERPRISE.symantecdomain.com  
--prplname vcsadmin1 --broker galaxy:2821
```

```
Enter password for vcsadmin1: #####
```

```
authenticate  
-----  
-----
```

```
Authenticated User vcsadmin1  
-----
```

3 Add the LDAP user to the main.cf file.

```
# haconf makerw
# hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
# haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
# hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

4 Verify that the main.cf file has the following lines:

```
# cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster clus1 (
  SecureClus = 1
  Administrators = {
    "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/
    DC=com@myenterprise.symantecdomain.com" }
  AdministratorGroups = {
    "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/
    DC=com@myenterprise.symantecdomain.com " }
  )
...
...
```

5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=myenterprise.symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=myenterprise.symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute    Value
galaxy       Attribute    RUNNING
nebula       Attribute    RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the VCS node using the VCS Cluster Manager (Java Console).

7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

To enable Windows Active Directory authentication for clusters that run in secure mode

- 1 Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSat/bin/atldapconf -d
-s domain_controller_name_or_ipaddress
-u domain_user -g domain_group
```

For example:

```
# /opt/VRTSat/bin/atldapconf -d -s 192.168.20.32 \
-u Administrator -g "Domain Admins"
Search User provided is invalid or Authentication is required to
proceed further.
Please provide authentication information for LDAP server.
```

```
Username/Common Name: symantecdomain\administrator
Password:
```

Attribute file created.

- 2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSat/bin/atldapconf -c -d windows_domain_name
```

For example:

```
# /opt/VRTSat/bin/atldapconf -c -d symantecdomain.com
Attribute list file not provided, using default AttributeList.txt.
CLI file name not provided, using default CLI.txt.
```

CLI for addldapdomain generated.

- 3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSat/bin/atldapconf -x
```

- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSat/bin/vssat listldapdomains
```

```
Domain Name :          symantecdomain.com
Server URL  :          ldap://192.168.20.32:389
SSL Enabled :          No
User Base DN :         CN=people,DC=symantecdomain,DC=com
User Object Class :    account
User Attribute :      cn
User GID Attribute :  gidNumber
Group Base DN :       CN=group,DC=symantecdomain,DC=com
Group Object Class :  group
Group Attribute :     cn
Group GID Attribute : cn
Auth Type   :         FLAT
Admin User  :
Admin User Password :
Search Scope :        SUB
```

- 5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=symantecdomain.com

- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute  Value
galaxy       Attribute  RUNNING
nebula       Attribute  RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the VCS node using the VCS Cluster Manager (Java Console).

- 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

Accessing the VCS documentation

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

To access the VCS documentation

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

Removing permissions for communication

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

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

Installing or upgrading VCS components

This chapter includes the following topics:

- [Installing the Java Console](#)
- [Upgrading the Java Console](#)
- [Installing VCS Simulator](#)
- [Upgrading VCS Simulator](#)
- [Upgrading the VCS agents](#)

Installing the Java Console

You can administer VCS using the VCS Java-based graphical user interface, Java Console. After VCS has been installed, install the Java Console on a Windows system or Linux system. Review the software requirements for Java Console.

The system from which you run the Java Console can be a system in the cluster or a remote workstation. A remote workstation enables each system in the cluster to be administered remotely.

Review the information about using the Java Console. For more information, refer to the *Veritas Cluster Server Administrator's Guide*.

Software requirements for the Java Console

Cluster Manager (Java Console) is supported on:

- RHEL 4 Update 3, RHEL 5, SLES 9 SP3, and SLES 10
- Windows XP and Windows 2003

Note: Make sure that you are using an operating system version that supports JRE 1.5.

Hardware requirements for the Java Console

The minimum hardware requirements for the Java Console follow:

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

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

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

Symantec recommends using the following hardware:

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

Installing the Java Console on Linux

Review the procedure to install the Java console. Before you begin with the procedure, ensure that you have the `gunzip` utility installed on your system.

To install Java console on Linux

- 1 Download the Java GUI utility from <http://go.symantec.com/vcsmc> to a temporary directory.
- 2 Go to the temporary directory and unzip the compressed RPM file:

```
# cd /tmp/install
# gunzip VRTSscm.tar.gz
```

The file `VRTSscm.tar` is now present in the temporary directory.

- 3 Extract the compressed file from the tar file:

```
# tar -xvf VRTScscm.tar
```

- 4 Install the RPM using rpm -i command.

```
# rpm -i VRTScscm-5.1.00.00-GA_GENERIC.noarch.rpm
```

Installing the Java Console on a Windows system

Review the procedure to install the Java console on a Windows system.

To install the Java Console on a Windows system

- 1 Download the Java GUI utility from <http://go.symantec.com/vcsmc> to a temporary directory.
- 2 Extract the zipped file to a temporary folder.
- 3 From this extracted folder, double-click setup.exe.
- 4 The Veritas Cluster Manager Install Wizard guides you through the installation process.

Upgrading the Java Console

Use one of the following applicable procedures to upgrade Java Console.

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.

```
# rpm -e VRTScscm
```

- 3 Install the VCS Java console.

See “[Installing the Java Console on Linux](#)” on page 270.

To upgrade the Java Console on a Windows client

- 1 Stop Cluster Manager (Java Console) if it is running.
- 2 Remove Cluster Manager from the system.
 - From the Control Panel, double-click **Add/Remove Programs**
 - Select **Veritas Cluster Manager**.

- Click **Add/Remove**.
 - Follow the uninstall wizard instructions.
- 3 Install the new Cluster Manager.
- See “[Installing the Java Console on a Windows system](#)” on page 271.

Installing VCS Simulator

You can administer VCS Simulator from the Java Console or from the command line. Review the software requirements for VCS Simulator.

Software requirements for VCS Simulator

VCS Simulator is supported on:

- Windows XP and Windows 2003

Note: Make sure that you are using an operating system version that supports JRE 1.5.

Installing VCS Simulator on Windows systems

This section describes the procedure to install VCS Simulator on Windows systems.

To install VCS Simulator on Windows systems

- 1 Download VCS Simulator from the following location to a temporary directory.
<http://www.symantec.com/business/cluster-server> and click **Utilities**.
- 2 Extract the compressed files to another directory.
- 3 Navigate to the path of the Simulator installer file:
`\your_platform_architecture\cluster_server\windows\
VCSWindowsInstallers\Simulator`
- 4 Double-click the installer file.
- 5 Read the information in the Welcome screen and click **Next**.
- 6 In the Destination Folders dialog box, click **Next** to accepted the suggested installation path or click **Change** to choose a different location.
- 7 In the Ready to Install the Program dialog box, click **Back** to make changes to your selections or click **Install** to proceed with the installation.
- 8 In the Installshield Wizard Completed dialog box, click **Finish**.

Reviewing the installation

VCS Simulator installs Cluster Manager (Java Console) and Simulator binaries on the system. The Simulator installation creates the following directories:

Directory	Content
attrpool	Information about attributes associated with VCS objects
bin	VCS Simulator binaries
default_clus	Files for the default cluster configuration
sample_clus	A sample cluster configuration, which serves as a template for each new cluster configuration
templates	Various templates that are used by the Java Console
types	The types.cf files for all supported platforms
conf	Contains another directory called types. This directory contains assorted resource type definitions that are useful for the Simulator. The type definition files are present in platform-specific sub directories.

Additionally, VCS Simulator installs directories for various cluster configurations.

VCS Simulator creates a directory for every new simulated cluster and copies the contents of the sample_clus directory. Simulator also creates a log directory within each cluster directory for logs that are associated with the cluster.

Upgrading VCS Simulator

Use the following procedure to upgrade VCS Simulator.

To upgrade VCS Simulator on a Windows client

- 1 Stop all instances of VCS Simulator.
- 2 Stop VCS Simulator, if it is running.
- 3 Remove VCS Simulator from the system.
 - From the Control Panel, double-click **Add/Remove Programs**
 - Select **VCS Simulator**.
 - Click **Add/Remove**.

- Follow the uninstall wizard instructions.
- 4 Install the new Simulator.
See [“Installing VCS Simulator on Windows systems”](#) on page 272.

Upgrading the VCS agents

The `installvcs` program does not upgrade the VCS agents for DB2, Oracle, and Sybase. If previous versions of these agents are installed on your cluster, you must uninstall the previous version of the agents and manually install the agent version 5.1.

The Veritas Cluster Server product installer includes the VCS agents for DB2, Oracle, and Sybase as a bundled package.

To upgrade the VCS agents

- 1 Before you install, log in as the superuser. Mount the disc, and copy the files in a temporary folder for installation.
- 2 Use following command to install the VCS agents package:

```
# rpm -ivh VRTSvcssea-5.1.00.00-GA_dist.arch.rpm
```

where *dist* is the operating system and takes the values RHEL 5, SLES 10, or SLES 11.

arch is the architecture. It takes the values `x86_64` for RHEL5, SLES 10, and SLES11.

See the agent Installation and Configuration Guide for more information on the agent that you want to upgrade.

See *Veritas Cluster Server Release Notes* for supported versions of the agents with VCS 5.1.

Verifying the VCS installation

This chapter includes the following topics:

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

About verifying the VCS installation

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

About the LLT and GAB configuration files

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

LLT and GAB also require the initialization configuration files:

- `/etc/sysconfig/llt`

- `/etc/sysconfig/gab`

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

- The `/etc/sysconfig/llt` file

This file stores the start and stop environment variables for LLT:

- `LLT_START`—Defines the startup behavior for the LLT module after a system reboot. Valid values include:

- 1—Indicates that LLT is enabled to start up.

- 0—Indicates that LLT is disabled to start up.

- `LLT_STOP`—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include:

- 1—Indicates that LLT is enabled to shut down.

- 0—Indicates that LLT is disabled to shut down.

The installer sets the value of these variables to 1 at the end of VCS configuration.

If you manually configured VCS, make sure you set the values of these environment variables to 1.

- The `/etc/llthosts` file

The file `llthosts` is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.

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

```
0      galaxy
1      nebula
```

- The `/etc/llttab` file

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

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

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

```
set-node galaxy
set-cluster 2
```

```
link eth1 eth-00:04:23:AC:12:C4 - ether - -  
link eth2 eth-00:04:23:AC:12:C5 - ether - -
```

If you use aggregated interfaces, then the file contains the aggregated interface name instead of the `eth-MAC_address`.

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

If you configured a low priority link under LLT, the file also includes a "link-lopri" line.

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

■ The `/etc/sysconfig/gab` file

This file stores the start and stop environment variables for GAB:

- `GAB_START`—Defines the startup behavior for the GAB module after a system reboot. Valid values include:
 - 1—Indicates that GAB is enabled to start up.
 - 0—Indicates that GAB is disabled to start up.
- `GAB_STOP`—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include:
 - 1—Indicates that GAB is enabled to shut down.
 - 0—Indicates that GAB is disabled to shut down.

The installer sets the value of these variables to 1 at the end of VCS configuration.

If you manually configured VCS, make sure you set the values of these environment variables to 1.

■ The `/etc/gabtab` file

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

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

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

The `-c` option configures the driver for use. The `-nN` specifies that the cluster is not formed until at least `N` nodes are ready to form the cluster. Symantec recommends that you set `N` to be the total number of nodes in the cluster.

Note: Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. Using `-c -x` can lead to a split-brain condition.

About the cluster UUID

You can verify the existence of the cluster UUID.

To verify the cluster UUID exists

- ◆ From the prompt, run a more command.

```
more /etc/vx/.uuids/clusuuid
```

About the VCS configuration files

VCS configuration files include the following:

- `main.cf`

The installer creates the VCS configuration file in the `/etc/VRTSvcs/conf/config` folder by default during the VCS configuration. The `main.cf` file contains the minimum information that defines the cluster and its nodes.

See [“Sample main.cf file for VCS clusters”](#) on page 279.

See [“Sample main.cf file for global clusters”](#) on page 282.

- `types.cf`

The file `types.cf`, which is listed in the include statement in the `main.cf` file, defines the VCS bundled types for VCS resources. The file `types.cf` is also located in the folder `/etc/VRTSvcs/conf/config`.

Additional files similar to `types.cf` may be present if agents have been added, such as `Oracletypes.cf`.

- `/etc/sysconfig/vcs`

This file stores the start and stop environment variables for VCS engine:

- `VCS_START`—Defines the startup behavior for VCS engine after a system reboot. Valid values include:

- 1—Indicates that VCS engine is enabled to start up.

- 0—Indicates that VCS engine is disabled to start up.

- `VCS_STOP`—Defines the shutdown behavior for VCS engine during a system shutdown. Valid values include:

- 1—Indicates that VCS engine is enabled to shut down.

- 0—Indicates that VCS engine is disabled to shut down.

The installer sets the value of these variables to 1 at the end of VCS configuration.

If you manually configured VCS, make sure you set the values of these environment variables to 1.

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

- The cluster definition includes the cluster information that you provided during the configuration. This definition includes the cluster name, cluster address, and the names of users and administrators of the cluster. Notice that the cluster has an attribute `UserNames`. The `installvcs` program creates a user "admin" whose password is encrypted; the word "password" is the default password.
- If you set up the optional I/O fencing feature for VCS, then the `UseFence = SCSI3` attribute 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 if you configured the virtual IP, SMTP, SNMP, or global cluster options.

The service group also has the following characteristics:

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

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

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

Sample `main.cf` file for VCS clusters

The following sample `main.cf` file is for a cluster in secure mode.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTTypes.cf"
include "Db2udbTypes.cf"
include "SybaseTypes.cf"

cluster vcs_cluster2 (
    UserNames = { admin = cDRpdxPmHpzS, smith = dKLhKJkHLh }
    ClusterAddress = "192.168.1.16"
    Administrators = { admin, smith }
    CounterInterval = 5
    SecureClus = 1
)

system galaxy (
)

system nebula (
)

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

IP webip (
    Device = eth0
    Address = "192.168.1.16"
    NetMask = "255.255.240.0"
)

NIC csgnic (
    Device = eth0
    NetworkHosts = { "192.168.1.17", "192.168.1.18" }
)

NotifierMngr ntfr (
    SnmpConsoles = { "saturn" = Error, "jupiter" = SevereError }
    SntpServer = "smtp.example.com"
    SntpRecipients = { "ozzie@example.com" = Warning,
```

```
        "harriet@example.com" = Error }
    )

webip requires csgnic
ntfr requires csgnic

// resource dependency tree
//
//   group ClusterService
//   {
//     NotifierMgr ntfr
//     {
//       NIC csgnic
//     }
//   }
// }

group VxSS (
    SystemList = { galaxy = 0, nebula = 1 }
    Parallel = 1
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Phantom phantom_vxss (
)

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

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

Sample main.cf file for global clusters

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

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

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

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

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

In the following main.cf file example, bold text highlights global cluster specific entries.

```
include "types.cf"

cluster vcs03 (
    ClusterAddress = "10.182.13.50"
    SecureClus = 1
)

system sysA (
)

system sysB (
)

system sysC (
)
```

```
group ClusterService (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

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

IP gcoip (
    Device = eth0
    Address = "10.182.13.50"
    NetMask = "255.255.240.0"
)

NIC csgnic (
    Device = eth0
)

NotifierMngr ntfr (
    SnmpConsoles = { vcslab4079 = SevereError }
    Smtperver = "smtp.veritas.com"
    Smtperipients = { "johndoe@veritas.com" = SevereError }
)
```

gcoip requires csgnic

ntfr requires csgnic

wac requires gcoip

```
// resource dependency tree
//
//     group ClusterService
//     {
//     NotifierMngr ntfr
//     {
//     NIC csgnic
//     }
```

```
//      Application wac
//      {
//      IP gcoip
//      {
//      NIC csgnic
//      }
//      }
//      }

group VxSS (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    Parallel = 1
    AutoStartList = { sysA, sysB, sysC }
    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
//      }
```

Verifying the LLT, GAB, and VCS configuration files

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

To verify the LLT, GAB, and VCS configuration files

- 1 Navigate to the location of the configuration files:
 - LLT
/etc/llthosts

/etc/llttab

- GAB

/etc/gabtab

- VCS

/etc/VRTSvcs/conf/config/main.cf

- 2 Verify the content of the configuration files.

See [“About the LLT and GAB configuration files”](#) on page 275.

See [“About the VCS configuration files”](#) on page 278.

Verifying LLT, GAB, and cluster operation

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

To verify LLT, GAB, and cluster operation

- 1 Log in to any node in the cluster as superuser.

- 2 Make sure that the PATH environment variable is set to run the VCS commands.

See [“Setting the PATH variable”](#) on page 60.

- 3 Verify LLT operation.

See [“Verifying LLT”](#) on page 285.

- 4 Verify GAB operation.

See [“Verifying GAB”](#) on page 288.

- 5 Verify the cluster operation.

See [“Verifying the cluster”](#) on page 290.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

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

```
lltstat -n
```

The output on galaxy resembles:

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

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

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

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

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

```
lltstat -n
```

The output on nebula resembles:

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

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

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

```
lltstat -nvv | more
```

The output on galaxy resembles:

```

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

```

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

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

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

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

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

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

```
lltstat -p
```

The output resembles:

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

Verifying GAB

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

The ports indicate the following:

- Port a
 - Nodes have GAB communication.
 - gen a36e0003 is a randomly generated number
 - membership 01 indicates that nodes 0 and 1 are connected

- Port b
 - Indicates that the I/O fencing driver is connected to GAB port b.
 - Note:** After you configure VCS using the installer, port b in the `gabconfig` command output indicates that I/O fencing is up in disabled mode. You must separately configure I/O fencing to use the feature.
 - gen a23da40d is a randomly generated number
 - membership 01 indicates that nodes 0 and 1 are connected

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

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

To verify GAB

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

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

- 2 Review the output of the command:

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

```
GAB Port Memberships
=====
Port a gen a36e0003 membership 01
Port b gen a23da40d membership 01
Port h gen fd570002 membership 01
```

Note that port b in the `gabconfig` command output may not indicate that I/O fencing feature is configured. After you configure VCS using the installer, the installer starts I/O fencing in disabled mode. You can use the `vxfenadm -d` command to verify the I/O fencing configuration.

- If GAB does not operate, the command does not return any GAB port membership information:

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

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

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

Verifying the cluster

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

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

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

To verify the cluster

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

```
hastatus -summary
```

The output resembles:

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

A galaxy                 RUNNING                0
A nebula                 RUNNING                0

-- GROUP STATE
-- Group                System                Probed  AutoDisabled  State

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

- 2 Review the command output for the following information:

- The system state

If the value of the system state is `RUNNING`, the cluster is successfully started.

- The ClusterService group state

In the sample output, the group state lists the ClusterService group, which is `ONLINE` on galaxy and `OFFLINE` on nebula.

Verifying the cluster nodes

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

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

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

To verify the cluster nodes

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

```
hasys -display
```

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

```
#System      Attribute                Value
galaxy       AgentsStopped            0
galaxy       AvailableCapacity        100
galaxy       CPUUsage                 0
galaxy       CPUUsageMonitoring       Enabled 0 ActionThreshold 0
                                   ActionTimeLimit 0 Action NONE
                                   NotifyThreshold 0 NotifyTimeLimit 0

galaxy       Capacity                 100
galaxy       ConfigBlockCount         217
galaxy       ConfigChecksum           62335
galaxy       ConfigDiskState          CURRENT
galaxy       ConfigFile                /etc/VRTSvcs/conf/config
galaxy       ConfigInfoCnt            0
galaxy       ConfigModDate             Wed 14 Oct 2009 17:22:48
galaxy       ConnectorState           Down
galaxy       CurrentLimits
galaxy       DiskHbStatus
galaxy       DynamicLoad              0
galaxy       EngineRestarted          0
galaxy       EngineVersion             5.1.00.0
```

#System	Attribute	Value
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	HostUtilization	CPU 0 Swap 0
galaxy	LLTNodeId	0
galaxy	LicenseType	DEMO
galaxy	Limits	
galaxy	LinkHbStatus	link2 UP link3 UP
galaxy	LoadTimeCounter	0
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	ShutdownTimeout	600
galaxy	SourceFile	./main.cf
galaxy	SysInfo	Linux:galaxy.example.com,#1 SMP Wed Dec 17 11:41:38 EST 2008,2.6.18-128.el5,x86_64
galaxy	SysName	galaxy
galaxy	SysState	RUNNING
galaxy	SystemLocation	
galaxy	SystemOwner	
galaxy	TFrozen	0
galaxy	TRSE	0
galaxy	UpDownState	Up
galaxy	UserInt	0

#System	Attribute	Value
galaxy	UserStr	
galaxy	VCSFeatures	DR
galaxy	VCSMode	VCS

Uninstalling VCS

- [Chapter 24. Uninstalling VCS using the installer](#)
- [Chapter 25. Uninstalling VCS using response files](#)

Uninstalling VCS using the installer

This chapter includes the following topics:

- [Preparing to uninstall VCS](#)
- [Uninstalling VCS 5.1 using the script-based installer](#)
- [Uninstalling VCS with the Veritas Web-based installer](#)
- [Removing the CP server configuration using the removal script](#)

Preparing to uninstall VCS

Review the following prerequisites before you uninstall VCS:

- Before you remove VCS from any node in the cluster, shut down the applications that depend on VCS. For example, applications such as Java Console or any high availability agents for VCS.
- Before you remove VCS from fewer than all nodes in a cluster, stop the service groups on the nodes from which you uninstall VCS. You must also reconfigure VCS on the remaining nodes.
See [“About adding and removing nodes”](#) on page 311.
- If you have manually edited any of the VCS configuration files, you need to reformat them.
See [“Reformatting VCS configuration files on a stopped cluster”](#) on page 63.

Uninstalling VCS 5.1 using the script-based installer

You must meet the following conditions to use the `uninstallvcs` program to uninstall VCS on all nodes in the cluster at one time:

- Make sure that the communication exists between systems. By default, the uninstaller uses `ssh`.
- Make sure you can execute `ssh` or `rsh` commands as superuser on all nodes in the cluster.
- Make sure that the `ssh` or `rsh` is configured to operate without requests for passwords or passphrases.

If you cannot meet the prerequisites, then you must run the `uninstallvcs` program on each node in the cluster.

The `uninstallvcs` program removes all VCS RPMs.

The example demonstrates how to uninstall VCS using the `uninstallvcs` program. The `uninstallvcs` program uninstalls VCS on two nodes: `galaxy` and `nebula`. The example procedure uninstalls VCS from all nodes in the cluster.

Removing VCS 5.1 RPMs

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

To uninstall VCS

- 1 Log in as superuser from the node where you want to uninstall VCS.
- 2 Start `uninstallvcs` program.

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

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

- 3 Enter the names of the systems from which you want to uninstall VCS.

The program performs system verification checks and asks to stop all running VCS processes.

- 4 Enter `y` to stop all the VCS processes.

The program proceeds with uninstalling the software.

- 5 Review the output as the `uninstallvcs` program continues to do the following:
 - Verifies the communication between systems

- Checks the installations on each system to determine the RPMs to be uninstalled.
- 6 Review the output as the uninstaller stops processes, unloads kernel modules, and removes the RPMs.
 - 7 Note the location of summary and log files that the uninstaller creates after removing all the RPMs.

Running `uninstallvcs` from the VCS 5.1 disc

You may need to use the `uninstallvcs` program on the VCS 5.1 disc in one of the following cases:

- You need to uninstall VCS after an incomplete installation.
- The `uninstallvcs` program is not available in `/opt/VRTS/install`.

Uninstalling VCS with the Veritas Web-based installer

This section describes uninstalling VCS with the Veritas Web-based installer.

To uninstall VCS

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 152.
- 4 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 5 Select **Veritas Cluster Server** from the Product drop-down list, and click **Next**.
- 6 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Validate**.

- 7 After the validation completes successfully, click **Uninstall** to uninstall VCS on the selected system.
- 8 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 9 After the installer stops the processes, the installer removes the products from the specified system.
Click **Next**.
- 10 After the uninstall completes, the installer displays the location of the log and summary files. If required, view the files to confirm the status of the removal.
- 11 Click **Finish**. The webinstaller prompts you for another task.

Removing the CP server configuration using the removal script

This section describes how to remove the CP server configuration from a node or cluster hosting the CP server.

Warning: Ensure that no VCS cluster is using the CP server that will have its CP server configuration removed.

A configuration utility that is part of VRTScps package is used to remove the CP server configuration. When using the configuration utility, a configuration removal script is run and the following tasks are performed:

- All CP server configuration files are removed
- The VCS configuration for CP server is removed

After running the utility and script, you can then uninstall VCS from the node or cluster.

Note: The configuration script has to run only once per CP server (which can be on a single node or SFHA cluster), when removing the CP server configuration.

The configuration utility performs the following steps to remove the CP server configuration:

- Offlines the CP server service group (CPSSG), if it is online
- Removes the CPSSG service group from the VCS configuration

The following procedure describes how to remove the CP server configuration.

To remove the CP server configuration

- 1** To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2** The Veritas Coordination Point Server Configuration utility appears with an option menu.

```
VERITAS COORDINATION POINT SERVER CONFIGURATION UTILITY
=====
```

Select one of the following:

```
[1] Configure Coordination Point Server on single node VCS system
```

```
[2] Configure Coordination Point Server on SFHA cluster
```

```
[3] Unconfigure Coordination Point Server
```

- 3** Select option 3 to unconfigure the Coordination Point Server.
- 4** A warning appears and prompts you to confirm the action to unconfigure the Coordination Point Server.

Enter "y" to proceed.

```
Unconfiguring Coordination Point Server stops the vxcperv process.
VCS clusters using this server for coordination purpose
will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)[Default:n] :y
```

- 5 After entering "y" to proceed, messages appear informing you of the progress in removing the CP server configuration.

When the CP server configuration has been unconfigured, a success message appears.

For an example of the messages from a single node VCS cluster:

```
A single node VCS cluster is currently configured.
Stopping the CP server ...

Removing the CP Server from VCS configuration..

Removing resource dependencies...
Deleting the resources configured under CPSSG service group...
Deleting the CPSSG service group...

Successfully unconfigured the Veritas Coordination Point Server.
```

For an example of the messages from a CP server on an SFHA cluster:

```
A multinode CP Server cluster is currently configured.
Stopping the CP server ...

Removing the CP Server from VCS configuration..

Removing resource dependencies...
Deleting the resources configured under CPSSG service group...
Deleting the CPSSG service group...

Successfully unconfigured the Veritas Coordination Point Server.
```

- 6 You are then prompted to delete the CP server database. Enter "y" to delete the database.

For example:

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

- 7 You are then prompted to delete the CP server configuration file and log files. Enter "y" to delete these files.

For example:

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 8 Run the following `hagrps -state` command to ensure that the CPSSG resource has been removed from the node.

For example:

```
root@mycps1.symantecexample.com # hagrps -state CPSSG

VCS WARNING V-16-1-40131 Group CPSSG does not exist
in the local cluster
```


Uninstalling VCS using response files

This chapter includes the following topics:

- [Uninstalling VCS using response files](#)
- [Response file variables to uninstall VCS](#)
- [Sample response file for uninstalling VCS](#)

Uninstalling VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS uninstallation on one cluster to uninstall VCS on other clusters.

To perform automated VCS uninstallation

- 1 Make sure that you meet the pre-requisites to uninstall VCS.
- 2 Copy the response file to one of the cluster systems where you want to uninstall VCS.
See [“Sample response file for uninstalling VCS”](#) on page 307.
- 3 Edit the values of the response file variables as necessary.
See [“Response file variables to uninstall VCS”](#) on page 306.
- 4 Start the uninstallation from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installvcs -responsefile /tmp/response_file
```

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

Response file variables to uninstall VCS

[Table 25-1](#) lists the response file variables that you can define to uninstall VCS.

Table 25-1 Response file variables specific to uninstalling VCS

Variable	List or Scalar	Description
CFG{opt}{uninstall}	Scalar	Uninstalls VCS RPMs. (Required)
CFG{opt}{stopfail_allow}	Scalar	Decides whether or not to proceed if the installer fails while stopping the processes or while unloading the drivers. (Optional)
CFG{systems}	List	List of systems on which the product is to be uninstalled. (Required)
CFG{prod}	Scalar	Defines the product to be uninstalled. The value is VCS51 for VCS. (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of ssh as the communication method between systems. (Optional)

Table 25-1 Response file variables specific to uninstalling VCS (*continued*)

Variable	List or Scalar	Description
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)

Sample response file for uninstalling VCS

Review the response file variables and their definitions.

See “[Response file variables to uninstall VCS](#)” on page 306.

```
#  
# Configuration Values:  
#  
our %CFG;  
  
$CFG{opt}{uninstall}=1;  
$CFG{prod}="VCS51";  
$CFG{systems}=[ qw(galaxy nebula) ];
```


Adding and removing nodes

- [Chapter 26. Adding and removing cluster nodes](#)
- [Chapter 27. Adding a node to a single-node cluster](#)

Adding and removing cluster nodes

This chapter includes the following topics:

- [About adding and removing nodes](#)
- [Adding nodes using the VCS installer](#)
- [Manually adding a node to a cluster](#)
- [Removing a node from a cluster](#)

About adding and removing nodes

After you install VCS and create a cluster, you can add and remove nodes from the cluster. You can create a cluster of up to 32 nodes.

Adding nodes using the VCS installer

The VCS installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:

```
/etc/littab
```

```
/etc/VRTSvcs/conf/sysname
```

- Updates the following configuration files and copies them on the new node:
 - `/etc/llthosts`
 - `/etc/gabtab`
 - `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files from the existing cluster to the new node
 - `/etc/vxfenmode`
 - `/etc/vxfendg`
 - `/etc/vx/.uuids/clusuuid`
 - `/etc/sysconfig/llt`
 - `/etc/sysconfig/gab`
 - `/etc/sysconfig/vxfen`
- Configures security on the new node if the existing cluster is a secure cluster.

Warning: If the root broker system has failed, then you must recover or reconfigure the root broker system before you add a new node to the cluster.

See [“Adding a node to the secure cluster whose root broker system has failed”](#) on page 403.

- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.

At the end of the process, the new node joins the VCS cluster.

Note: If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

To add the node to an existing VCS cluster using the VCS installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the VCS installer with the `-addnode` option.

```
# cd /opt/VRTS/install
# ./installvcs -addnode
```

The installer displays the copyright message and the location where it stores the temporary installation logs.

- 3 Enter the name of a node in the existing VCS cluster. The installer uses the node information to identify the existing cluster.

Enter a node name in the VCS cluster to which
you want to add a node: **galaxy**

- 4 Review and confirm the cluster information.

- 5 Enter the name of the systems that you want to add as new nodes to the cluster.

Enter the system names separated by spaces
to add to the cluster: **saturn**

The installer checks the installed products and packages on the nodes and discovers the network interfaces.

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

Note: The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

Enter the NIC for the first private heartbeat
link on saturn: [b,q,?] **eth1**

- 7 Enter **y** to configure a second private heartbeat link.

Note: At least two private heartbeat links must be configured for high availability of the cluster.

Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)

- 8 Enter the name of the network interface that you want to configure as the second private heartbeat link.

Enter the NIC for the second private heartbeat link
on saturn: [b,q,?] **eth2**

- 9 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 10 Review and confirm the information.

- 11 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: eth3
```

- 12 If the existing cluster uses server-based fencing in secure mode, provide responses to the following installer prompts.

If you are using different root brokers for the CP server and the client VCS cluster, enter **y** to confirm the use of different root brokers. The installer attempts to establish trust between the new node being added to the cluster and the authentication broker of the CP server.

```
Are you using different Root Brokers for the CP Server(s) and the
client cluster? (If so then installer will try to establish trust
between the new node(s) being added and CP Server's
Authentication Broker) [y,n,q] (n) y
```

Enter the host name of the authentication broker used for any one of the CP servers.

```
Enter hostname of the Authentication Broker being used for any one
of the CP Server(s): [b] mycps1.symantecexample.com
```

Enter the port number where the authentication broker for the CP server listens to establish trust with the new node:

```
Enter the port where the Authentication Broker
mycps1.symantecexample.com for the CP Server(s) is listening
for establishing trust: [b] (2821)
```

Manually adding a node to a cluster

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

See “[Hardware requirements](#)” on page 33.

[Table 26-1](#) specifies the tasks that are involved in adding a cluster. The example demonstrates how to add a node saturn to already existing nodes, galaxy and nebula.

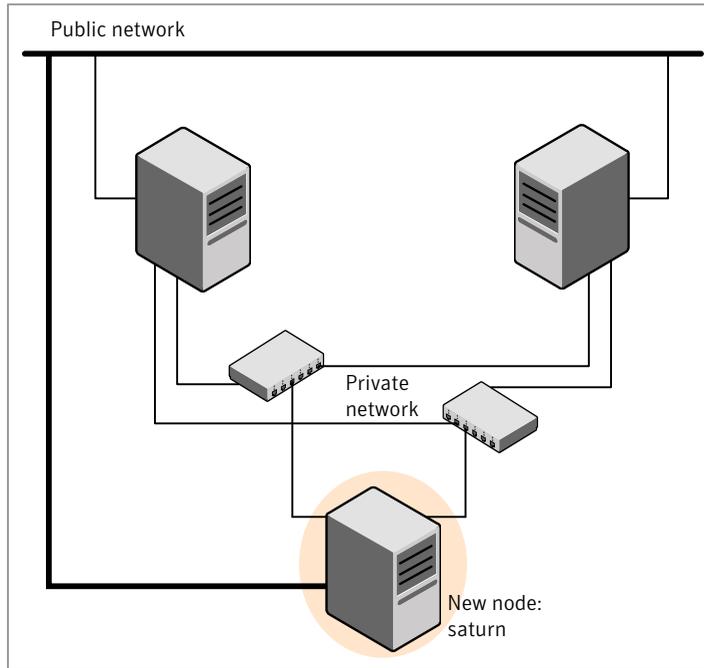
Table 26-1 Tasks that are involved in adding a node to a cluster

Task	Reference
Set up the hardware.	See “Setting up the hardware” on page 315.
Install the software manually.	See “Preparing for a manual installation” on page 187. See “Installing VCS RPMs for a manual installation” on page 191.
Add a license key.	See “Adding a license key for a manual installation” on page 193.
If the existing cluster runs in secure mode, set up the new node to run in secure mode.	See “Setting up the node to run in secure mode” on page 317. Warning: If the root broker system has failed, then you must recover or reconfigure the root broker system before you add a new node to the cluster. See “Adding a node to the secure cluster whose root broker system has failed” on page 403.
Configure LLT and GAB.	See “Configuring LLT and GAB” on page 319.
If the existing cluster is configured for I/O fencing, configure I/O fencing on the new node.	See “Configuring I/O fencing on the new node” on page 321.
Add the node to the existing cluster.	See “Adding the node to the existing cluster” on page 326.
Start VCS and verify the cluster.	See “Starting VCS and verifying the cluster” on page 327.

Setting up the hardware

[Figure 26-1](#) shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

Figure 26-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the VCS private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a two-node cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a two-node cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

[Figure 26-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.

- 2 Connect the system to the shared storage, if required.

Installing the VCS software manually when adding a node

Install the VCS 5.1 RPMs manually and add a license key.

For more information, see the following:

- See [“Adding a license key for a manual installation”](#) on page 193.

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

See [“Configuring LLT and GAB”](#) on page 319.

[Table 26-2](#) uses the following information for the following command examples.

Table 26-2 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
saturn	saturn.nodes.example.com	The new node that you are adding to the cluster.
RB1	RB1.brokers.example.com	The root broker for the cluster
RB2	RB2.brokers.example.com	Another root broker, not the cluster's RB

To verify the existing security setup on the node

- 1 If node saturn is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Else, proceed to configuring the authentication broker on node saturn.

See [“Configuring the authentication broker on node saturn”](#) on page 318.

- 2 Find out the root broker to which the node saturn belongs using the following command.

```
# vssregctl -l -q -b \  
"Security\Authentication\Authentication Broker" \  
-k "BrokerName"
```

- 3 If the node saturn already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.

See [“Setting up VCS related security configuration”](#) on page 319.

- 4 If the node saturn belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node saturn.

- Kill `/opt/VRTSat/bin/vxatd` process.
- Remove the credential that RB2 has given to AB on node saturn.

```
# vssat deletecred --domain type:domainname \  
--prplname prplname
```

For example:

```
# vssat deletecred --domain vx:root@RB2.brokers.example.com \  
--prplname saturn.nodes.example.com
```

Configuring the authentication broker on node saturn

Configure a new authentication broker (AB) on node saturn. This AB belongs to root broker RB1.

To configure the authentication broker on node saturn

- 1 Create a principal for node saturn on root broker RB1. Execute the following command on root broker RB1.

```
# vssat addprpl --pdrtype root --domain domainname \  
--prplname prplname --password password \  
--prpltype service
```

For example:

```
# vssat addprpl --pdrtype root \  
--domain root@RB1.brokers.example.com \  
--prplname saturn.nodes.example.com \  
--password flurbdicate --prpltype service
```

- 2 Ensure that there is no clock skew between the times on node saturn and RB1.
- 3 Copy the `/opt/VRTSat/bin/root_hash` file from RB1 to node saturn.

4 Configure AB on node saturn to talk to RB1.

```
# vxatd -o -a -n prplname -p password -x vx -y domainname -q \
rootbroker -z 2821 -h roothash_file_path
```

For example:

```
# vxatd -o -a -n saturn.nodes.example.com -p flurbdicate \
-x vx -y root@RB1.brokers.example.com -q RB1 \
-z 2821 -h roothash_file_path
```

5 Verify that AB is configured properly.

```
# vssat showbrokermode
```

The command should return 1, indicating the mode to be AB.

Setting up VCS related security configuration

Perform the following steps to configure VCS related security settings.

Setting up VCS related security configuration

1 Start /opt/VRTSat/bin/vxatd process.

2 Create HA_SERVICES domain for VCS.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

3 Add VCS and webserver principal to AB on node saturn.

```
# vssat addprpl --pdrtype ab --domain HA_SERVICES --prplname
webserver_VCS_prplname --password new_password --prpltype
service --can_proxy
```

4 Create /etc/VRTSvcs/conf/config/.secure file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

Configuring LLT and GAB

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

To configure LLT

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

For example, suppose you add saturn to a cluster consisting of galaxy and nebula:

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

```
0 galaxy
1 nebula
```

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

```
0 galaxy
1 nebula
2 saturn
```

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

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

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

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

- 3 Copy the following file from one of the nodes in the existing cluster to the new node:

`/etc/sysconfig/llt`

- 4 On the new system, run the command:

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

To configure GAB

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

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

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

The file on the new node should be the same. Symantec recommends that you use the `-c -nN` option, where *N* is the total number of cluster nodes.

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

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

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

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

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

- 2 Copy the following file from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/gab
```

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

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

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 288.

- 2 Run the same command on the other nodes (galaxy and nebula) 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
```

Configuring I/O fencing on the new node

If the existing cluster is configured for I/O fencing, perform the following tasks on the new node:

- Prepare to configure I/O fencing on the new node.
See [“Preparing to configure I/O fencing on the new node”](#) on page 322.
- If the existing cluster runs server-based fencing, configure server-based fencing on the new node.
See [“Configuring server-based fencing on the new node”](#) on page 323.
If the existing cluster runs disk-based fencing, you need not perform any additional step. Skip to the next task. After you copy the I/O fencing files and start I/O fencing, disk-based fencing automatically comes up.
- Copy the I/O fencing files from an existing node to the new node and start I/O fencing on the new node.
See [“Starting I/O fencing on the new node”](#) on page 325.

If the existing cluster is not configured for I/O fencing, perform the procedure to add the new node to the existing cluster.

See [“Adding the node to the existing cluster”](#) on page 326.

Preparing to configure I/O fencing on the new node

Perform the following tasks before you configure and start I/O fencing on the new node.

To prepare to configure I/O fencing on the new node

- 1 Determine whether the existing cluster runs disk-based or server-based fencing mechanism. On one of the nodes in the existing cluster, run the following command:

```
# vxfsadm -d
```

If the fencing mode in the output is SCSI3, then the cluster uses disk-based fencing.

If the fencing mode in the output is CUSTOMIZED, then the cluster uses server-based fencing.

- 2 In the following cases, install and configure Veritas Volume Manager (VxVM) on the new node.
 - The existing cluster uses disk-based fencing.
 - The existing cluster uses server-based fencing with at least one coordinator disk.

You need not perform this step if the existing cluster uses server-based fencing with all coordination points as CP servers.

See the *Veritas Storage Foundation and High Availability Installation Guide* for installation instructions.

Configuring server-based fencing on the new node

This section describes the procedures to configure server-based fencing on a new node. Depending on whether server-based fencing is configured in secure or non-secure mode on the existing cluster, perform the tasks in one of the following procedures:

- Server-based fencing in non-secure mode:
[To configure server-based fencing in non-secure mode on the new node](#)
- Server-based fencing in secure mode:
[To configure server-based fencing with security on the new node](#)

To configure server-based fencing in non-secure mode on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s mycps1.symantecexample.com \  
-a list_nodes
```

The new node must be listed in the command output.

- 4 Add the VCS user cpsclient@saturn to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e cpsclient@saturn \  
-f cps_operator -g vx
```

```
User cpsclient@saturn successfully added
```

Perform the following procedure for a secure configuration.

To configure server-based fencing with security on the new node

- 1 As the root user, create the VCS user and the domain on the new node:
 - Create a dummy configuration file `/etc/VRTSvcs/conf/config/main.cf` that resembles the following example:

```
# cat main.cf
```

```
include "types.cf"
cluster clus1 {
    SecureClus = 1
}

system saturn {
}
```

- Verify the dummy configuration file:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTSvcs/bin/hacf -verify .
```

- Start VCS in one node mode on the new node:

```
# /opt/VRTSvcs/bin/hastart -onenode
```

- 2 Verify that the VCS user and the domain are created on the new node:

```
# /opt/VRTSvcs/bin/cpsat showcred | grep _HA_VCS_
# /opt/VRTSvcs/bin/cpsat listpd -t local | grep HA_SERVICES
```

- 3 Stop VCS if the VCS user and domain are created successfully on the new node:

```
# /opt/VRTSvcs/bin/hastop -local
```

- 4 If the root broker for the CP server and the new node are different, run the following command to establish trust between the authentication broker of the CP Server and the new node:

```
# /usr/bin/echo y | /opt/VRTSvcs/bin/cpsat setuptrust \
-b mycps1.symantecexample.com -s high
```

- 5 Log in to each CP server as the root user.

- 6 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 7 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s mycps1.symantecexample.com -a list_nodes
```

The new node must be listed in the output.

- 8 Add the VCS user `_HA_VCS_saturn@HA_SERVICES@saturn.symantec.com` to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e _HA_VCS_saturn@HA_SERVICES@saturn.symantec.com \  
-f cps_operator -g vx
```

User `_HA_VCS_saturn@HA_SERVICES@saturn.symantec.com` successfully added

Adding the new node to the vxfen service group

Perform the steps in the following procedure to add the new node to the vxfen service group.

To add the new node to the vxfen group using the CLI

- 1 On one of the nodes in the existing VCS cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing vxfen group.

```
# hagrps -modify vxfen SystemList -add saturn 2
```

- 3 Save the configuration by running the following command from any node in the VCS cluster:

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

Starting I/O fencing on the new node

Copy the I/O fencing files from an existing node to the new node and start I/O fencing on the new node. This task starts I/O fencing based on the fencing mechanism that is configured in the existing cluster.

To start I/O fencing on the new node

- 1 Copy the following I/O fencing configuration files from one of the nodes in the existing cluster to the new node:

- `/etc/vxfenmode`

- `/etc/vxfendg`—This file is required only for disk-based fencing.
 - `/etc/sysconfig/vxfen`
- 2 Start I/O fencing on the new node.

```
# /etc/init.d/vxfen start
```
 - 3 Run the GAB configuration command on the new node to verify that the port b membership is formed.

```
# gabconfig -a
```

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 Copy the cluster UUID from the one of the nodes in the existing cluster to the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -copy -from_sys \  
node_name_in_running_cluster -to_sys new_sys1 ... new_sysn
```

Where you are copying the cluster UUID from a node in the cluster (*node_name_in_running_cluster*) to systems from *new_sys1* through *new_sysn* that you want to join the cluster.

- 2 Copy the following file from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/vcs
```

- 3 Enter the command:

```
# haconf -makerw
```

- 4 Add the new system to the cluster:

```
# hasys -add saturn
```

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

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

- 6 Check the VCS configuration file. No error message and a return value of zero indicates that the syntax is legal.

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

- 7 If necessary, modify any new system attributes.
- 8 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 Start VCS on the newly added system:

```
# 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
```

If the cluster uses I/O fencing, then the GAB output also shows port b membership.

Removing a node from a cluster

Table 26-3 specifies the tasks that are involved in removing a node from a cluster. In the example procedure, the cluster consists of nodes galaxy, nebula, and saturn; node saturn is to leave the cluster.

Table 26-3 Tasks that are 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. 	<p>See “Verifying the status of nodes and service groups” on page 328.</p>

Table 26-3 Tasks that are involved in removing a node (*continued*)

Task	Reference
<ul style="list-style-type: none"> ■ Switch or remove any VCS service groups on the node departing the cluster. ■ Delete the node from VCS configuration. 	<p>See “Deleting the departing node from VCS configuration” on page 329.</p>
<p>Modify the llthosts and gabtab files to reflect the change.</p>	<p>See “Modifying configuration files on each remaining node” on page 332.</p>
<p>If the existing cluster is configured to use server-based I/O fencing, remove the node configuration from the CP server.</p>	<p>See “Removing the node configuration from the CP server” on page 332.</p>
<p>For a cluster that is running in a secure mode, remove the security credentials from the leaving node.</p>	<p>See “Removing security credentials from the leaving node” on page 333.</p>
<p>On the node departing the cluster:</p> <ul style="list-style-type: none"> ■ Modify startup scripts for LLT, GAB, and VCS to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. 	<p>See “Unloading LLT and GAB and removing VCS on the departing node” on page 333.</p>

Verifying the status of nodes and service groups

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

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 galaxy      RUNNING        0  
A nebula      RUNNING        0  
A saturn      RUNNING        0  
  
-- GROUP STATE  
-- Group      System      Probed   AutoDisabled  State  
B grp1       galaxy     Y        N              ONLINE  
B grp1       nebula     Y        N              OFFLINE  
B grp2       galaxy     Y        N              ONLINE  
B grp3       nebula     Y        N              OFFLINE  
B grp3       saturn     Y        N              ONLINE  
B grp4       saturn     Y        N              ONLINE
```

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

Deleting the departing node from VCS configuration

Before you remove a node from the cluster you need to identify the service groups that run on the node.

You then need to perform the following actions:

- Remove the service groups that other service groups depend on, or
- Switch the service groups to another node that other service groups depend on.

To remove or switch service groups from the departing node

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

```
# hagrps -switch grp3 -to nebula
```

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

```
# hagrps -dep
```

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

```
# haconf -makerw  
# hagrps -unlink grp4 grp1
```

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

- 4 Stop VCS on the departing node:

```
# hastop -sys saturn
```

- 5 Check the status again. The state of the departing node should be EXITED. Make sure that any service group that you want to fail over is online on other nodes.

```
# hastatus -summary
```

```
-- SYSTEM STATE  
-- System      State          Frozen  
A galaxy      RUNNING       0  
A nebula      RUNNING       0  
A saturn      EXITED        0  
  
-- GROUP STATE  
-- Group      System        Probed   AutoDisabled  State  
B grp1       galaxy        Y        N              ONLINE  
B grp1       nebula        Y        N              OFFLINE  
B grp2       galaxy        Y        N              ONLINE  
B grp3       nebula        Y        N              ONLINE  
B grp3       saturn        Y        Y              OFFLINE  
B grp4       saturn        Y        N              OFFLINE
```

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

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

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

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

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

```
# hagrps -delete grp4
```

- 9 Check the status.

```
# hastatus -summary
-- SYSTEM STATE
-- System      State          Frozen
A galaxy      RUNNING        0
A nebula      RUNNING        0
A saturn      EXITED         0

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

- 10 Delete the node from the cluster.

```
# hasys -delete saturn
```

- 11 Save the configuration, making it read only.

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

Modifying configuration files on each remaining node

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

To modify the configuration files on a remaining node

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

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

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

Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`.

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

For example, change:

```
0 galaxy
1 nebula
2 saturn
```

To:

```
0 galaxy
1 nebula
```

Removing the node configuration from the CP server

After removing a node from a VCS cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

To remove the node configuration from the CP server

- 1 Log into the CP server as the root user.
- 2 View the list of VCS users on the CP server, using the following command:

```
# cpsadm -s cp_server -a list_users
```

Where *cp_server* is the virtual IP/ virtual hostname of the CP server.

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e _HA_VCS_saturn@HA_SERVICES@saturn.nodes.example.com \  
-f cps_operator -g vx
```

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
cpsadm -s cp_server -a rm_node -h saturn -c clus1 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
cpsadm -s cp_server -a list_nodes
```

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node saturn. Perform the following steps.

To remove the security credentials

- 1 Kill /opt/VRTSat/bin/vxatd process.
- 2 Remove the root credentials on node saturn.

```
# vssat deletecred --domain type:domainname --prplname prplname
```

Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node that is departing the cluster.

If you have configured VCS as part of the Storage Foundation and High Availability products, you may have to delete other dependent RPMs before you can delete all of the following ones.

To stop LLT and GAB and remove VCS

- 1 If you had configured I/O fencing in enabled mode, then stop I/O fencing.

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

2 Stop GAB and LLT:

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

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

3 To determine the RPMs to remove, enter:

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

4 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 VRTSvcsea
```

```
# rpm -e VRTSatServer
```

```
# rpm -e VRTSatClient
```

```
# rpm -e VRTScutil
```

```
# rpm -e VRTSvcscdr
```

```
# rpm -e VRTSvcscag
```

```
# rpm -e VRTScps
```

```
# rpm -e VRTSvcsc
```

```
# rpm -e VRTSvxfen
```

```
# rpm -e VRTSgab
```

```
# rpm -e VRTSllt
```

```
# rpm -e VRTSspt
```

```
# rpm -e VRTSperl
```

```
# rpm -e VRTSvlic
```

5 Remove the LLT and GAB configuration files.

```
# rm /etc/llttab
```

```
# rm /etc/gabtab
```

```
# rm /etc/llthosts
```

Adding a node to a single-node cluster

This chapter includes the following topics:

- [Adding a node to a single-node cluster](#)

Adding a node to a single-node cluster

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

[Table 27-1](#) specifies the activities that you need to perform to add nodes to a single-node cluster.

Table 27-1 Tasks to add a node to a single-node cluster

Task	Reference
Set up Node B to be compatible with Node A.	See “Setting up a node to join the single-node cluster” on page 336.
<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.	See “Installing and configuring Ethernet cards for private network” on page 337.
Connect both nodes to shared storage.	See “Configuring the shared storage” on page 338.

Table 27-1 Tasks to add a node to a single-node cluster (*continued*)

Task	Reference
<ul style="list-style-type: none"> ■ Bring up VCS on Node A. ■ Edit the configuration file. 	See “Bringing up the existing node” on page 338.
<p>If necessary, install VCS on Node B and add a license key.</p> <p>Make sure Node B is running the same version of VCS as the version on Node A.</p>	See “Installing the VCS software manually when adding a node to a single node cluster” on page 339.
Edit the configuration files on Node B.	See “Creating configuration files” on page 339.
Start LLT and GAB on Node B.	See “Starting LLT and GAB” on page 339.
<ul style="list-style-type: none"> ■ Start LLT and GAB on Node A. ■ Restart VCS on Node A. ■ Modify service groups for two nodes. 	See “Reconfiguring VCS on the existing node” on page 340.
<ul style="list-style-type: none"> ■ Start VCS on Node B. ■ Verify the two-node cluster. 	See “Verifying configuration on both nodes” on page 341.

Setting up a node to join the single-node cluster

The new node to join the existing single node that runs VCS must run the same operating system.

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

- 1 Do one of the following tasks:
 - If VCS is not currently running on Node B, proceed to step 2.
 - If the node you plan to add as Node B is currently part of an existing cluster, remove the node from the cluster. After you remove the node from the cluster, remove the VCS RPMs and configuration files.
See [“Removing a node from a cluster”](#) on page 327.
 - If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS.
See [“Uninstalling VCS 5.1 using the script-based installer”](#) on page 298.
 - If you renamed the LLT and GAB startup files, remove them.

See [“Renaming the LLT and GAB startup files”](#) on page 380.

- 2 If necessary, install VxVM and VxFS.

See [“Installing VxVM or VxFS if necessary”](#) on page 337.

Installing VxVM or VxFS if necessary

If you have either VxVM or VxFS with the cluster option installed on the existing node, install the same version on the new node.

Refer to the appropriate documentation for VxVM and VxFS to verify the versions of the installed products. Make sure the same version runs on all nodes where you want to use shared storage.

Installing and configuring Ethernet cards for private network

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

For high availability, use two separate NICs on each node. The two NICs provide redundancy for heartbeating.

See [“Setting up the private network”](#) on page 53.

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:

```
# shutdown -r now
```

- 3 Install the Ethernet card on Node A.

If you want to use aggregated interface to set up private network, configure aggregated interface.

- 4 Install the Ethernet card on Node B.

If you want to use aggregated interface to set up private network, configure aggregated interface.

- 5 Configure the Ethernet card on both nodes.

- 6 Make the two Ethernet cable connections from Node A to Node B for the private networks.

- 7 Restart the nodes.

Configuring the shared storage

Make the connection to shared storage from Node B. Configure VxVM on Node B and reboot the node when you are prompted.

See “[Setting up shared storage](#)” on page 57.

See “[Setting up shared storage: Fiber Channel](#)” on page 58.

Bringing up the existing node

Bring up the node.

To bring up the node

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

```
# haconf -makerw
```

- 4 Display the service groups currently configured.

```
# hagrps -list
```

- 5 Freeze the service groups.

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

Repeat this command for each service group in step 4.

- 6 Make the configuration read-only.

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

- 7 Stop VCS on Node A.

```
# hastop -local -force
```

- 8 Edit the VCS system configuration file `/etc/sysconfig/vcs`, and remove the `"-onenode"` option.

Change the line:

```
ONENODE=yes
```

To:

```
ONENODE=no
```

- 9 Enable the GAB and LLT startup files so they can be used.

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

Installing the VCS software manually when adding a node to a single node cluster

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

Refer to the following sections:

- See [“Preparing for a manual installation”](#) on page 187.
- See [“Adding a license key for a manual installation”](#) on page 193.

Creating configuration files

Create the configuration files for your cluster.

To create the configuration files

- 1 Create the file `/etc/llttab` that lists both the nodes.
See [“Setting up /etc/llttab for a manual installation”](#) on page 202.
- 2 Create the file `/etc/llthosts`. Set up `/etc/llthosts` for a two-node cluster.
See [“Setting up /etc/llthosts for a manual installation”](#) on page 202.
- 3 Create the file `/etc/gabtab`.
See [“Configuring GAB for a manual installation”](#) on page 204.

Starting LLT and GAB

On the new node, start LLT and GAB.

To start LLT and GAB

- 1 Start LLT on Node B.

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

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

Reconfiguring VCS on the existing node

Reconfigure VCS on the existing nodes.

To reconfigure VCS on existing nodes

- 1 On Node A, create the files `/etc/llttab`, `/etc/llthosts`, and `/etc/gabtab`. Use the files that are created on Node B as a guide, customizing the `/etc/llttab` for Node A.
- 2 Start LLT on Node A.

```
# /etc/init.d/llt start
```
- 3 Start GAB on Node A.

```
# /etc/init.d/gab start
```
- 4 Check the membership of the cluster.

```
# gabconfig -a
```
- 5 Copy the cluster UUID from the existing node to the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -copy -from_sys \  
node_name_in_running_cluster -to_sys new_sys1 ... new_sysn
```

Where you are copying the cluster UUID from a node in the cluster (*node_name_in_running_cluster*) to systems from *new_sys1* through *new_sysn* that you want to join the cluster.
- 6 Start VCS on Node A.

```
# hstart
```
- 7 Make the VCS configuration writable.

```
# haconf -makerw
```

8 Add Node B to the cluster.

```
# hasys -add sysB
```

9 Add Node B to the system list of each service group.

■ List the service groups.

```
# hagrp -list
```

■ For each service group that is listed, add the node.

```
# hagrp -modify group SystemList -add sysB 1
```

Verifying configuration on both nodes

Verify the configuration for the nodes.

To verify the nodes' configuration

1 On Node B, check the cluster membership.

```
# gabconfig -a
```

2 Start the VCS on Node B.

```
# 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
```


Installation reference

- [Appendix A. VCS installation packages](#)
- [Appendix B. Installation command options](#)
- [Appendix C. Changes to bundled agents in VCS 5.1](#)
- [Appendix D. Sample main.cf files](#)
- [Appendix E. Installing VCS on a single node](#)
- [Appendix F. Configuring LLT over UDP using IPv4](#)
- [Appendix G. Configuring LLT over UDP using IPv6](#)
- [Appendix H. Troubleshooting VCS installation](#)
- [Appendix I. Sample VCS cluster setup diagrams for CP server-based I/O fencing](#)

VCS installation packages

This appendix includes the following topics:

- [Veritas Cluster Server installation packages](#)

Veritas Cluster Server installation packages

[Table A-1](#) shows the package name and contents for each Veritas Cluster Server package.

Table A-1 Veritas Cluster Server packages

Package	Contents	Required/Optional
VRTSatClient	Contains the client binaries for Symantec Product Authentication Service	Optional. Required to use Symantec Product Authentication Service.
VRTSatServer	Contains the server binaries for Symantec Product Authentication Service	Optional. Required to use Symantec Product Authentication Service.
VRTScps	Contains the binaries for the Veritas Coordination Point Server.	Optional. Required to Coordination Point Server (CPS). Depends on VRTSvxfen.

Table A-1 Veritas Cluster Server packages (*continued*)

Package	Contents	Required/Optional
VRTScutil	VRTScutil contains the following components: <ul style="list-style-type: none"> ■ Contains the binaries for the Veritas Cluster Server utilities. ■ Contains the binaries for the Veritas Cluster Server configuration wizards. ■ Contains the binaries for the Veritas Cluster Server Simulator. ■ Contains the binaries for the Veritas agent for Oracle and SF Oracle RAC configuration wizards. 	Required Depends on VRTSvcsag.
VRTSgab	Contains the binaries for Veritas Cluster Server group membership and atomic broadcast services.	Required Depends on VRTSllt.
VRTSllt	Contains the binaries for Veritas Cluster Server low-latency transport.	Required
VRTSperl	Contains Perl for Veritas.	Required
VRTSspt	Contains the binaries for Veritas Software Support Tools.	Required
VRTSvcS	VRTSvcS contains the following components: <ul style="list-style-type: none"> ■ Contains the binaries for Veritas Cluster Server. ■ Contains the binaries for Veritas Cluster Server manual pages. ■ Contains the binaries for Veritas Cluster Server English message catalogs. ■ Contains the binaries for Veritas Cluster Server utilities. These utilities include security services. 	Required Depends on VRTSvxfen, VRTSgab, and VRTSllt.

Table A-1 Veritas Cluster Server packages (*continued*)

Package	Contents	Required/Optional
VRTSvcsg	Contains the binaries for Veritas Cluster Server bundled agents.	Required Depends on VRTSvcsg. Depends on VRTSvcsgdr.
VRTSvcldr	Contains the binaries for Veritas Cluster Server disk reservation.	Required
VRTSvcsea	VRTSvcsea contains the binaries for Veritas high availability agents for DB2, Sybase, and Oracle.	Optional for VCS. Required to use VCS with the high availability agents for DB2, Sybase, or Oracle.
VRTSvclic	Contains the binaries for Symantec License Utilities.	Required
VRTSvcxfen	Contains the binaries for Veritas I/O Fencing .	Optional. Required to use fencing. Depends on VRTSvcgab.

Installation command options

This appendix includes the following topics:

- [Command options for installvcs program](#)
- [Command options for uninstallvcs program](#)

Command options for installvcs program

The `installvcs` command usage takes the following form:

```
installvcs [ system1 system2... ]  
  [ -configure | -install | -license | -precheck | -requirements | -start  
    | -stop | -uninstall | -upgrade ]  
  [ -logpath log_path ]  
  [ -responsefile response_file ]  
  [ -tmppath tmp_path ]  
  [ -hostfile hostfile_path ]  
  
  [ -kickstart kickstart_path ]  
  
  [ -keyfile ssh_key_file ]  
  [ -patchpath patch_path ]  
  [ -pkgpath pkg_path ]  
  
[ -rsh | -redirect | -installminpkgs | -installrecpkgs | -installallpkgs  
  | -minpkgs | -recpkgs | -allpkgs | -ha | -pkgset | -pkginfo  
  | -serial | -makeresponsefile | -pkgtable | -security |  
  -addnode | -fencing ]
```

[Table B-1](#) lists the `installvcs` command options.

Table B-1 installvcs options

Option and Syntax	Description
-addnode	Adds a node that you specify to a cluster. The cluster must be online to use this command option to add a node.
-allpkgs	View a list of all VCS RPMs and patches. The installvcs program lists the RPMs and patches in the correct installation order. You can use the output to create scripts for command-line installation, or for installations over a network. See the -minpkgs and the -recpkgs options.
-configure	Configure VCS after using -install option to install VCS.
-fencing	Configure I/O fencing after you configure VCS. The script provides an option to configure disk-based I/o fencing or server-based I/O fencing.
-hostfile	Specifies the location of a file that contains the system names for the installer.
-install	Install product RPMs on systems without configuring VCS.
-installallpkgs	Selects all the RPMs for installation. See the -allpkgs option.
-installminpkgs	Selects the minimum RPMs for installation. See the -minpkgs option.
-installrecpkgs	Selects the recommended RPMs for installation. See the -recpkgs option.
-kickstart <i>dir_path</i>	Creates a kickstart configuration file to install VCS using the Kickstart utility for RHEL. The file contains the list of VCS RPMs in the correct installation order. The file contains the RPMs in the format that the Kickstart utility can use for installation. The <i>dir_path</i> indicates the path to an existing directory where the installer must create the file.
-keyfile <i>ssh_key_file</i>	Specifies a key file for SSH. The option passes -i <i>ssh_key_file</i> with each SSH invocation.
-license	Register or update product licenses on the specified systems. This option is useful to replace a demo license.

Table B-1 installvcs options (*continued*)

Option and Syntax	Description
<code>-logpath <i>log_path</i></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>-makeresponsefile</code>	Create a response file. This option only generates a response file and does not install VCS.
<code>-minpkgs</code>	View a list of the minimal RPMs and the patches that are required for VCS. The <code>installvcs</code> program lists the RPMs and patches in the correct installation order. The list does not include the optional RPMs. You can use the output to create scripts for command-line installation, or for installations over a network. See the <code>-allpkgs</code> and the <code>-recpkgs</code> options.
<code>-pkginfo</code>	Displays a list of packages in the order of installation in a user-friendly format. Use this option with one of the following options: <ul style="list-style-type: none"> ■ <code>-allpkgs</code> If you do not specify an option, <code>-allpkgs</code> is used by default. ■ <code>-minpkgs</code> ■ <code>-recpkgs</code>
<code>-pkgpath <i>pkg_path</i></code>	Specifies that <i>pkg_path</i> contains all RPMs that the <code>installvcs</code> program is about to install on all systems. The <i>pkg_path</i> is the complete path of a directory, usually NFS mounted.
<code>-pkgset</code>	Discovers and lists the 5.1 RPMs installed on the systems that you specify.
<code>-pkgtable</code>	Displays the VCS 5.1 RPMs in the correct installation order.
<code>-precheck</code>	Verify that systems meet the installation requirements before proceeding with VCS installation. Symantec recommends doing a precheck before you install VCS. See “Performing automated preinstallation check” on page 62.

Table B-1 installvcs options (*continued*)

Option and Syntax	Description
<p><code>-recpkgs</code></p>	<p>View a list of the recommended RPMs and the patches that are required for VCS. The <code>installvcs</code> program lists the RPMs and patches in the correct installation order. The list does not include the optional RPMs.</p> <p>You can use the output to create scripts for command-line installation, or for installations over a network.</p> <p>See the <code>-allpkgs</code> and the <code>-minpkgs</code> options.</p>
<p><code>-requirements</code></p>	<p>View a list of required operating system version, required patches, file system space, and other system requirements to install VCS.</p>
<p><code>-responsefile</code> <code>response_file</code> <code>[-enckeyfile</code> <code>encryption_key_file]</code></p>	<p>Perform automated VCS installation using the system and the configuration information that is stored in a specified file instead of prompting for information.</p> <p>The <code>response_file</code> must be a full path name. If not specified, the response file is automatically generated as <code>installernumber.response</code> where <code>number</code> is random. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.</p> <p>The <code>-enckeyfile</code> option and <code>encryption_key_file</code> name are required with the <code>-responsefile</code> option when the response file contains encrypted passwords.</p> <p>See “Installing VCS using response files” on page 161.</p> <p>See “Configuring VCS using response files” on page 167.</p> <p>See “Upgrading VCS using response files” on page 253.</p>
<p><code>-redirect</code></p>	<p>Specifies that the installer need not display the progress bar details during the installation.</p>
<p><code>-rsh</code></p>	<p>Specifies that <code>rsh</code> and <code>rscp</code> are to be used for communication between systems instead of <code>ssh</code> and <code>scp</code>. This option requires that systems be preconfigured such that <code>rsh</code> commands between systems execute without prompting for passwords or confirmations</p>

Table B-1 installvcs options (*continued*)

Option and Syntax	Description
-security	<p>Enable or disable Symantec Product Authentication Service in a VCS cluster that is running.</p> <p>See the <i>Veritas Cluster Server Administrator's Guide</i> for instructions.</p> <p>See “About Symantec Product Authentication Service (AT)” on page 29.</p>
-serial	<p>Performs the installation, uninstallation, start, and stop operations on the systems in a serial fashion. By default, the installer performs these operations simultaneously on all the systems.</p>
-start	<p>Starts the daemons and processes for VCS.</p> <p>If the installvcs program failed to start up all the VCS processes, you can use the -stop option to stop all the processes and then use the -start option to start the processes.</p> <p>See the -stop option.</p> <p>See “Starting and stopping processes for the Veritas products” on page 398.</p>
-stop	<p>Stops the daemons and processes for VCS.</p> <p>If the installvcs program failed to start up all the VCS processes, you can use the -stop option to stop all the processes and then use the -start option to start the processes.</p> <p>See the -start option.</p> <p>See “Starting and stopping processes for the Veritas products” on page 398.</p>
-tmppath <i>tmp_path</i>	<p>Specifies that <i>tmp_path</i> is the working directory for installvcs program. This path is different from the /var/tmp path. This destination is where the installvcs program performs the initial logging and where the installvcs program copies the RPMs on remote systems before installation.</p>
-upgrade	<p>Upgrades the installed RPMs on the systems that you specify.</p>
-uninstall	<p>Uninstalls VCS from the systems that you specify.</p>

Command options for uninstallvcs program

The `uninstallvcs` command usage takes the following form:

```
uninstallvcs [ system1 system2... ]
[ -logpath <log_path> ]
[ -responsefile <response_file> ]
[ -tmppath <tmp_path> ]
[ -hostfile <hostfile_path> ]
[ -keyfile <ssh_key_file> ]
[ -patchpath <patch_path> ]
[ -pkgpath <pkg_path> ]
[ -rsh | -redirect | -minpkgs | -recpkgs | -allpkgs | -ha | -pkgset |
  -pkginfo | -serial | -makeresponsefile | -pkgtable ]
```

Table B-2 lists the `uninstallvcs` command options.

Table B-2 `uninstallvcs` options

Option and Syntax	Description
<code>-hostfile</code>	Specifies the location of a file that contains the system names for the installer.
<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>-logpath</code> <i>log_path</i>	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>-makeresponsefile</code>	Use this option to create a response file or to verify that your system configuration is ready for uninstalling VCS.
<code>-pkginfo</code>	Displays a list of packages in a user-friendly format in the order that the <code>uninstallvcs</code> program uninstalls. Use this option with one of the following options: <ul style="list-style-type: none"> ■ <code>-allpkgs</code> If you do not specify an option, <code>-allpkgs</code> is used by default. ■ <code>-minpkgs</code> ■ <code>-recpkgs</code>
<code>-pkgpath</code> <i>pkg_path</i>	Specifies that <i>pkg_path</i> contains all RPMs that the <code>uninstallvcs</code> program is about to uninstall on all systems. The <i>pkg_path</i> is the complete path of a directory, usually NFS mounted.

Table B-2 uninstallocs options (*continued*)

Option and Syntax	Description
<code>-pkgset</code>	Discovers the RPM set that is installed on the systems that you specify.
<code>-pkgtable</code>	Displays VCS RPMs in the order that the uninstallocs program uninstalls.
<code>-redirect</code>	Displays progress details without showing progress bar.
<code>-responsefile response_file</code>	<p>Perform automated VCS uninstallation using the system and the configuration information that is stored in a specified file instead of prompting for information.</p> <p>The <i>response_file</i> must be a full path name. If not specified, the response file is automatically generated as <code>installer$number$.response</code> where <i>number</i> is random. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.</p> <p>See “Uninstalling VCS using response files” on page 305.</p>
<code>-rsh</code>	Specifies that <i>rsh</i> and <code>r_{cp}</code> are to be used for communication between systems instead of <code>ssh</code> and <code>s_{cp}</code> . This option requires that systems be preconfigured such that <i>rsh</i> commands between systems execute without prompting for passwords or confirmations
<code>-serial</code>	Performs the installation, uninstallation, start, and stop operations on the systems in a serial fashion. By default, the installer performs these operations simultaneously on all the systems.
<code>-tmppath tmp_path</code>	Specifies that <i>tmp_path</i> is the working directory for uninstallocs program. This path is different from the <code>/var/tmp</code> path. This destination is where the uninstallocs program performs the initial logging and where the installvcs program copies the RPMs on remote systems before installation.

Changes to bundled agents in VCS 5.1

This appendix includes the following topics:

- [Deprecated agents](#)
- [New agents](#)
- [New and modified attributes for 5.1 agents](#)
- [Manually removing deprecated resource types and modifying attributes](#)
- [Creating new VCS accounts if you used native operating system accounts](#)

Deprecated agents

The following agents are no longer supported:

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

New agents

The following new agents are in the VCS 5.1 release:

- CoordPoint—Provides server-based I/O fencing.

The following new agents were added in the 5.0 MP3 release:

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

The following new agents were added in the 5.0 release:

- **NFSRestart**—Provides high availability for NFS record locks.
- **RemoteGroup**—Monitors and manages a service group on another system.

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

New and modified attributes for 5.1 agents

[Table C-1](#) lists the attributes that VCS adds or modifies when you upgrade from VCS 5.0 MP3 or later to VCS 5.1.

Table C-1 New and modified attributes for VCS agents for upgrades from VCS 5.0 MP3

Agent	New and modified attributes	Default Value
DNS		
Modified attributes		
	Alias	
	Hostname	
DiskGroup		
Modified attributes		
	PanicSystemOnDGLoss	0
IP		
New attributes		
	IPOptions	
	IPRouteOptions	
	PrefixLen	1000
Mount		
New attributes		

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

Agent	New and modified attributes	Default Value
	OptCheck	0
	CreateMountPt	0
	ReuseMntPt	0
	MntPtPermission	
	MntPtOwner	
	MntPtGroup	
	AccessPermissionChk	
	RecursiveMnt	
MultiNICA		
New attributes		
	DualDevice	
	PrefixLen	1000
	LinkOptions	
	IPv4AddrOptions	
	IPv6AddrOptions	
	IPv4RouteOptions	
	IPv6RouteOptions	
NFS		
Modified attribute		
	Address	
Share		
New attribute		
	NFSRes	

Table C-2 lists the attributes that VCS adds or modifies when you upgrade from VCS 5.0 or later to VCS 5.0 MP3.

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

Agent	New and modified attributes	Default Value
Apache		
New attributes		
	PidFile	
	IntentionalOffline	0
DiskGroup		
New attributes		
	UmountVolumes	0
Modified attributes		
	SupportedActions	{ "license.vfd", "disk.vfd", "udid.vfd", "verifyplex.vfd", checkudid, numdisks, campusplex, joindg, splitdg, getvxvminfo, volinuse }
DNS		
New attributes		
	SupportedActions	{ "dig.vfd", "keyfile.vfd", "master.vfd" }
	ResRecord	
	CreatePTR	
	OffDelRR	
LVMVolumeGroup		
New attributes		
	SupportedActions	{ volinuse }
Mount		
New attributes		
	RegList	{ VxFSMountLock }

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

Agent	New and modified attributes	Default Value
	VxFSMountLock	0
Modified attributes		
	SupportedActions	{ "mountpoint.vfd", "mounted.vfd", "vxfslic.vfd", "chgmntlock", "mountentry.vfd" }
NFSRestart		
New attributes		
	SupportedActions	{ "lockdir.vfd", "nfsconf.vfd" }
Share		
New attributes		
	SupportedActions	{ "direxists.vfd" }

Table C-3 lists the attributes that VCS adds or modifies when you upgrade from VCS 4.1 to VCS 5.1 5.0 MP3.

Table C-3 New and modified attributes for VCS agents for upgrades from VCS 4.1

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

Table C-3 New and modified attributes for VCS agents for upgrades from VCS 4.1 (*continued*)

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

Table C-3 New and modified attributes for VCS agents for upgrades from VCS 4.1 (continued)

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

Table C-3 New and modified attributes for VCS agents for upgrades from VCS 4.1 (*continued*)

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

Table C-3 New and modified attributes for VCS agents for upgrades from VCS 4.1 (*continued*)

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

Manually removing deprecated resource types and modifying attributes

With VCS 5.1, 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.

If you use the resource type ServiceGroupHB, Symantec recommends the use of I/O fencing.

VCS 5.1 does not support gabdiskhb. So, the `installvcs` program removes the `gabdiskhb` entry from the `/etc/gabtab` file.

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

To remove the deprecated resource types and modify attributes

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

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

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

Perform the following instructions:

- Remove the resource of the deprecated resource types.
You must modify the resource dependencies to ensure that the configuration works properly.
See “[Deprecated agents](#)” on page 357.
- Modify attribute values that might have changed.
See [Table C-1](#) on page 358.
See [Table C-2](#) on page 360.
See [Table C-3](#) on page 361.
- Save the main.cf.
- Reformat the main.cf file.

```
# hacf -cftocmd config
# hacf -cmdtoef config
```

- 4 Verify the configuration.

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

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

Creating new VCS accounts if you used native operating system accounts

VCS has deprecated the AllowNativeCliUsers attribute. To use native OS accounts with VCS, use the halogin command. After you run the halogin command, VCS encrypts and stores your VCS credentials in your home directory for a specific time period. After you run the halogin command, you need not authenticate yourself every time you run a VCS command. In secure clusters, the command also sets up a trust relationship and retrieves a certificate from an authentication broker.

See the *Veritas Cluster Server Administrator’s Guide* for information on assigning user privileges to OS user groups for clusters running in secure mode and clusters not running in secure mode.

Perform the following procedure if you used the AllowNativeCliUsers attribute. Ensure that each native user running VCS commands has a home directory on the system from which the user runs VCS commands.

To set up VCS authentication for clusters running in secure mode

- 1 Set the configuration (main.cf) mode to read/write.

```
# haconf -makerw
```

- 2 Assign proper privileges to the OS users or user groups. Each operating system user must perform steps 3 and 4.

- 3 If the user executes VCS commands from a remote host, set the following environment variables:

- VCS_HOST: Name of the VCS node on which you run commands. You may specify the virtual IP address associated with the cluster.
- VCS_DOMAIN: Name of the VxSS domain to which the user belongs.
- VCS_DOMAINTYPE: Type of VxSS domain: unixpwd, nt, nis, nisplus, or vx.

- 4 Run the halogin command:

```
$ halogin vcsusername password
```

To set up VCS authentication for clusters not running in secure mode

- 1 Set the configuration (main.cf) mode to read/write.

```
# haconf -makerw
```

- 2 Create VCS user accounts for all users and assign privileges to these users.

- 3 Each VCS user must run the halogin command:

```
$ halogin vcsusername  
password
```

Creating new VCS accounts if you used native operating system accounts

Sample main.cf files

This appendix includes the following topics:

- [Sample configuration files for CP server](#)

Sample configuration files for CP server

The following are example main.cf files for a CP server that is hosted on a single node, and a CP server that is hosted on an SFHA cluster.

- The main.cf file for a CP server that is hosted on a single node:
See [“CP server hosted on a single node main.cf file”](#) on page 369.
- The main.cf file for a CP server that is hosted on an SFHA cluster:
See [“CP server hosted on an SFHA cluster main.cf file”](#) on page 371.

Note: The CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with VCS clusters. The following example main.cf files use IPv4 addresses.

CP server hosted on a single node main.cf file

The following is an example of a single CP server node main.cf.

For this CP server single node main.cf, note the following values:

- Cluster name: cps1
- Node name: mycps1

```
include "types.cf"

// cluster name: cps1
// CP server: mycps1
```

```
cluster cps1 (
    UserNames = { admin = bMnfmHmJNiNNlVNhMK, haris = fopKojNvpHouNn,
        "mycps1.symantecexample.com@root@vx" = aj,
        "root@mycps1.symantecexample.com" = hq }
    Administrators = { admin, haris,
        "mycps1.symantecexample.com@root@vx",
        "root@mycps1.symantecexample.com" }
    SecureClus = 1
    HacliUserLevel = COMMANDROOT
)

system mycps1 (
)

group CPSSG (
    SystemList = { mycps1 = 0 }
    AutoStartList = { mycps1 }
)

IP cpsvip (
    Device @mycps1 = bge0
    Address = "10.209.3.1"
    NetMask = "255.255.252.0"
)

NIC cpsnic (
    Device @mycps1 = bge0
)

Process vxcpserv (
    PathName = "/opt/VRTScps/bin/vxcpserv"
    ConfInterval = 30
    RestartLimit = 3
)

cpsvip requires cpsnic
vxcpserv requires cpsvip

// resource dependency tree
//
// group CPSSG
```

```
// {  
// Process vxcpserv  
//   {  
//     IP cpsvip  
//     {  
//       NIC cpsnic  
//     }  
//   }  
// }  
  
group VxSS (  
  SystemList = { mycps1 = 0 }  
  Parallel = 1  
  AutoStartList = { mycps1 }  
  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  
//   }
```

CP server hosted on an SFHA cluster main.cf file

The following is an example of a main.cf, where the CP server is hosted on an SFHA cluster.

For this CP server hosted on an SFHA cluster main.cf, note the following values:

- Cluster name: cps1
- Nodes in the cluster: mycps1, mycps2

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"

// cluster: cps1
// CP servers:
// mycps1
// mycps2

cluster cps1 (
    UserNames = { admin = ajkCjeJgkFkkIskEjh,
                  "mycps1.symantecexample.com@root@vx" = JK,
                  "mycps2.symantecexample.com@root@vx" = dl }
    Administrators = { admin, "mycps1.symantecexample.com@root@vx",
                       "mycps2.symantecexample.com@root@vx" }
    SecureClus = 1
)

system mycps1 (
)

system mycps2 (
)

group CPSSG (
    SystemList = { mycps1 = 0, mycps2 = 1 }
    AutoStartList = { mycps1, mycps2 } )

    DiskGroup cpsdg (
        DiskGroup = cps_dg
    )

    IP cpsvip (
        Device @mycps1 = bge0
        Device @mycps2 = bge0
        Address = "10.209.81.88"
        NetMask = "255.255.252.0"
    )
)
```

```
Mount cpsmount (
    MountPoint = "/etc/VRTScps/db"
    BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
    FSType = vxfs
    FsckOpt = "-y"
)

NIC cpsnic (
    Device @mycps1 = bge0
    Device @mycps2 = bge0
)

Process vxcpserv (
    PathName = "/opt/VRTScps/bin/vxcpserv"
)

Volume cpsvol (
    Volume = cps_volume
    DiskGroup = cps_dg
)

cpsmount requires cpsvol
cpsvip requires cpsnic
cpsvol requires cpsdg
vxcpserv requires cpsmount
vxcpserv requires cpsvip

// resource dependency tree
//
// group CPSSG
// {
// Process vxcpserv
//     {
//     Mount cpsmount
//         {
//         Volume cpsvol
//             {
//             DiskGroup cpsdg
//             }
//         }
//     }
// IP cpsvip
//     {
```

```
//          NIC cpsnic
//          }
//      }
// }

group VxSS (
    SystemList = { mycps1 = 0, mycps2 = 1 }
    Parallel = 1
    AutoStartList = { mycps1, mycps2 }
    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
// }

group cvm (
    SystemList = { mycps1 = 0, mycps2 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { mycps1, mycps2 }
)

CFSfsckd vxfscdk (
)

```

```
CVMCluster cvm_clus (
    CVMClustName = cps1
    CVMNodeId = { mycps1 = 0, mycps2 = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

cvm_clus requires cvm_vxconfigd
vxfscd requires cvm_clus

// resource dependency tree
//
// group cvm
// {
//   CFSfscd vxfscd
//   {
//     CVMCluster cvm_clus
//     {
//       CVMVxconfigd cvm_vxconfigd
//     }
//   }
// }
// }
```


Installing VCS on a single node

This appendix includes the following topics:

- [About installing VCS on a single node](#)
- [Creating a single-node cluster using the installer program](#)
- [Creating a single-node cluster manually](#)
- [Setting the path variable for a manual single node installation](#)
- [Installing VCS software manually on a single node](#)
- [Renaming the LLT and GAB startup files](#)
- [Modifying the startup files](#)
- [Configuring VCS](#)
- [Verifying single-node operation](#)

About installing VCS on a single node

You can install VCS 5.1 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 378.

See [“Creating a single-node cluster manually”](#) on page 379.

Creating a single-node cluster using the installer program

Table E-1 specifies the tasks that are involved to install VCS on a single node using the installer program.

Table E-1 Tasks to create a single-node cluster using the installer

Task	Reference
Prepare for installation.	See “ Preparing for a single node installation ” on page 378.
Install the VCS software on the system using the installer.	See “ Starting the installer for the single node cluster ” on page 378.

Preparing for a single node installation

You can use the installer program to install a cluster on a single system for either of the two following purposes:

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

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

For more information about LLT and GAB:

See “[About LLT and GAB](#)” on page 23.

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.

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,?] (y)
```

Answer `n` if you want to use the single node cluster as a stand-alone cluster.

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

Continue with the installation.

Creating a single-node cluster manually

[Table E-2](#) specifies the tasks that you need to perform to install VCS on a single node.

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

Task	Reference
Set the PATH variable	See “Setting the path variable for a manual single node installation” on page 379.
Install the VCS software manually and add a license key	See “Installing VCS software manually on a single node” on page 380.
Remove any LLT or GAB configuration files and rename LLT and GAB startup files. A single-node cluster does not require the node-to-node communication service, LLT, or the membership communication service, GAB.	See “Renaming the LLT and GAB startup files” on page 380.
Modify the VCS startup file for single-node operation.	See “Modifying the startup files” on page 380.
Create and modify the VCS configuration files.	See “Configuring VCS” on page 380.
Start VCS and verify single-node operation.	See “Verifying single-node operation” on page 381.

Setting the path variable for a manual single node installation

Set the path variable.

See [“Setting the PATH variable”](#) on page 60.

Installing VCS software manually on a single node

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

Refer to the following sections:

- See [“Installing VCS software manually”](#) on page 189.
- See [“Adding a license key for a manual installation”](#) on page 193.

Renaming the LLT and GAB startup files

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

To rename the LLT and GAB startup files

- ◆ Rename the LLT and GAB startup files.

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

Modifying the startup files

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

Change the line:

```
ONENODE=no
```

To:

```
ONENODE=yes
```

Configuring VCS

You now need to configure VCS.

See [“Configuring VCS”](#) on page 204.

Verifying single-node operation

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

To verify single-node cluster

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

```
# hastart -onenode
```

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

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


Configuring LLT over UDP using IPv4

This appendix includes the following topics:

- [Using the UDP layer for LLT](#)
- [Configuring LLT over UDP](#)

Using the UDP layer for LLT

VCS 5.1 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.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

LLT over UDP is slower than LLT over Ethernet. Use LLT over UDP only when the hardware configuration makes it necessary.

Configuring LLT over UDP

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on different physical networks.

If the LLT private links are not on different physical networks, then make sure that the links are on separate subnets. Set the broadcast address in `/etc/llttab` explicitly depending on the subnet for each link.

See [“Broadcast address in the `/etc/llttab` file”](#) on page 384.

- Make sure that each NIC has an IP address that is configured before configuring LLT.
- Make sure the IP addresses in the `/etc/llttab` files are consistent with the IP addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
See [“Selecting UDP ports”](#) on page 386.
- Set the broadcast address correctly for direct-attached (non-routed) links.
See [“Sample configuration: direct-attached links”](#) on page 387.
- For the links that cross an IP router, disable broadcast features and specify the IP address of each link manually in the `/etc/llttab` file.
See [“Sample configuration: links crossing IP routers”](#) on page 389.

Broadcast address in the `/etc/llttab` file

The broadcast address is set explicitly for each link in the following example.

- Display the content of the `/etc/llttab` file on the first node galaxy:

```
galaxy # cat /etc/llttab

set-node galaxy
set-cluster 1
link link1 udp - udp 50000 - 192.168.9.1 192.168.9.255
link link2 udp - udp 50001 - 192.168.10.1 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

- Display the content of the `/etc/llttab` file on the second node nebula:

```
nebula # cat /etc/llttab

set-node nebula
set-cluster 1
link link1 udp - udp 50000 - 192.168.9.2 192.168.9.255
link link2 udp - udp 50001 - 192.168.10.2 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

The link command in the /etc/llttab file

Review the link command information in this section for the /etc/llttab file. See the following information for sample configurations:

- See “[Sample configuration: direct-attached links](#)” on page 387.
- See “[Sample configuration: links crossing IP routers](#)” on page 389.

Table F-1 describes the fields of the link command that are shown in the /etc/llttab file examples. Note that some of the fields differ from the command for standard LLT links.

Table F-1 Field description for link command in /etc/llttab

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example udp. A place holder string. On other unix platforms like Solaris or HP, this entry points to a device file (for example, /dev/udp). Linux does not have devices for protocols. So this field is ignored.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “ Selecting UDP ports ” on page 386.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IP address</i>	IP address of the link on the local node.
<i>bcast-address</i>	<ul style="list-style-type: none"> ■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the /etc/llttab file

The `set-addr` command in the /etc/llttab file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 389.

[Table F-2](#) describes the fields of the `set-addr` command.

Table F-2 Field description for `set-addr` command in `/etc/llttab`

Field	Description
<code>node-id</code>	The ID of the cluster node; for example, 0.
<code>link tag-name</code>	The string that LLT uses to identify the link; for example link1, link2,....
<code>address</code>	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file `/etc/services`. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address      State
udp      0      0 *:32768          *:*
udp      0      0 *:956            *:*
udp      0      0 *:tftp           *:*
udp      0      0 *:sunrpc         *:*
udp      0      0 *:ipp            *:*
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# set_parms ip_address
```

For example:

- For the first network interface on the node galaxy:

```
IP address=192.168.9.1, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

For the first network interface on the node nebula:

```
IP address=192.168.9.2, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

- For the second network interface on the node galaxy:

```
IP address=192.168.10.1, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

For the second network interface on the node nebula:

```
IP address=192.168.10.2, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.

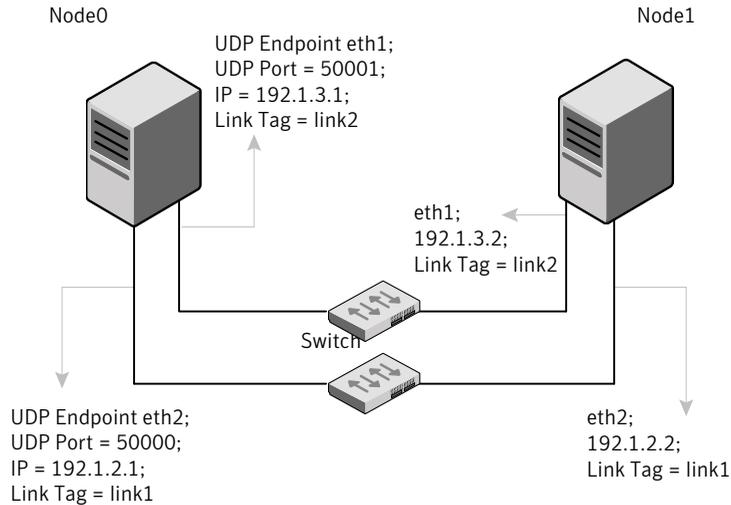
An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100  
  
link link1 udp - udp 50000 - 192.168.30.1 192.168.30.255  
link link2 udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: direct-attached links

[Figure F-1](#) depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure F-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 resembles:

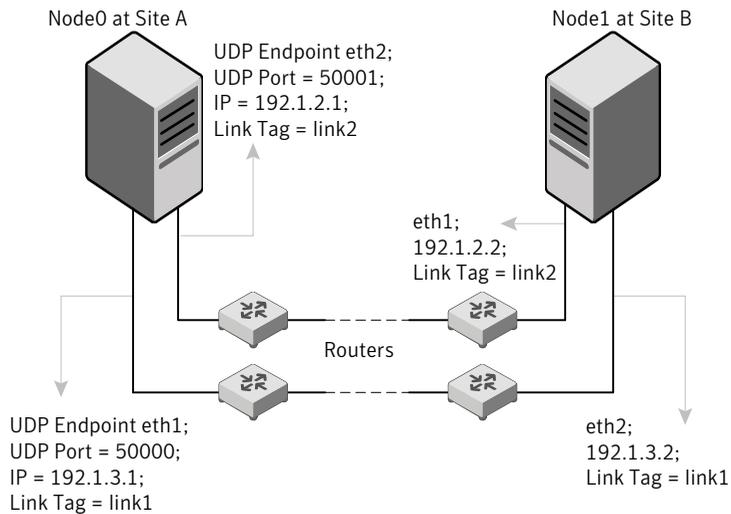
```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
```

```
link link1 udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

Figure F-2 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure F-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1

link link1 udp - udp 50000 - 192.1.3.1 -
link link2 udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr      0 link1 192.1.1.1
set-addr      0 link2 192.1.2.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
```

```
set-addr      3 link1 192.1.7.3  
set-addr      3 link2 192.1.8.3
```

```
#disable LLT broadcasts  
set-bcasthb   0  
set-arp       0
```

The `/etc/llttab` file on Node 0 resembles:

```
set-node Node0  
set-cluster 1  
  
link link1 udp - udp 50000 - 192.1.1.1 -  
link link2 udp - udp 50001 - 192.1.2.1 -  
  
#set address of each link for all peer nodes in the cluster  
#format: set-addr node-id link tag-name address  
set-addr      1 link1 192.1.3.1  
set-addr      1 link2 192.1.4.1  
set-addr      2 link1 192.1.5.2  
set-addr      2 link2 192.1.6.2  
set-addr      3 link1 192.1.7.3  
set-addr      3 link2 192.1.8.3  
  
#disable LLT broadcasts  
set-bcasthb   0  
set-arp       0
```

Configuring LLT over UDP using IPv6

This appendix includes the following topics:

- [Using the UDP layer of IPv6 for LLT](#)
- [Configuring LLT over UDP using IPv6](#)

Using the UDP layer of IPv6 for LLT

VCS 5.1 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.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

Configuring LLT over UDP using IPv6

The following checklist is to configure LLT over UDP:

- For UDP6, the multicast address is set to "-".
- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.
- Make sure the IPv6 addresses in the `/etc/llttab` files are consistent with the IPv6 addresses of the network interfaces.

- Make sure that each link has a unique not well-known UDP port.
 See [“Selecting UDP ports”](#) on page 393.
- For the links that cross an IP router, disable multicast features and specify the IPv6 address of each link manually in the `/etc/llttab` file.
 See [“Sample configuration: links crossing IP routers”](#) on page 395.

The link command in the `/etc/llttab` file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See [“Sample configuration: direct-attached links”](#) on page 394.
- See [“Sample configuration: links crossing IP routers”](#) on page 395.

Note that some of the fields in [Table G-1](#) differ from the command for standard LLT links.

[Table G-1](#) describes the fields of the link command that are shown in the `/etc/llttab` file examples.

Table G-1 Field description for link command in `/etc/llttab`

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device name of the UDP protocol; for example udp6.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp6" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “Selecting UDP ports” on page 393.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IPv6 address</i>	IPv6 address of the link on the local node.
<i>mcast-address</i>	"-" 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 multicast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 395.

[Table G-2](#) describes the fields of the `set-addr` command.

Table G-2 Field description for `set-addr` command in `/etc/llttab`

Field	Description
<code>node-id</code>	The ID of the cluster node; for example, 0.
<code>link tag-name</code>	The string that LLT uses to identify the link; for example link1, link2,....
<code>address</code>	IPv6 address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file `/etc/services`. You should also use the `netstat` command to list the UDP ports currently in use. For example:

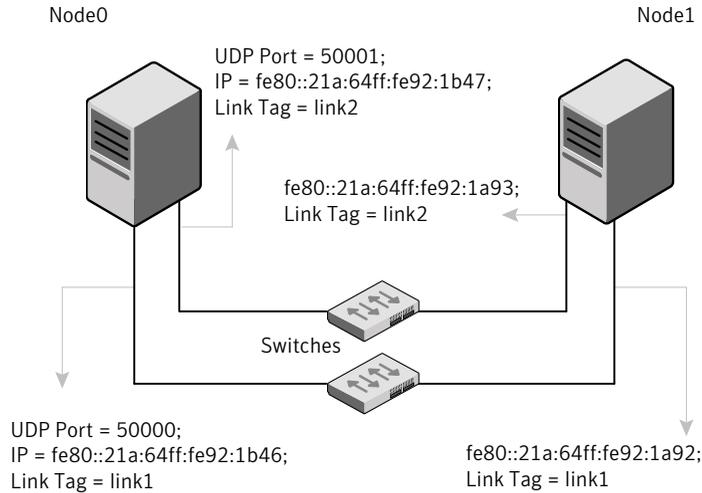
```
# netstat -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address      State
udp      0      0 *:32768          *:*
udp      0      0 *:956            *:*
udp      0      0 *:tftp           *:*
udp      0      0 *:sunrpc         *:*
udp      0      0 *:ipp            *:*
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

Sample configuration: direct-attached links

Figure G-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure G-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```

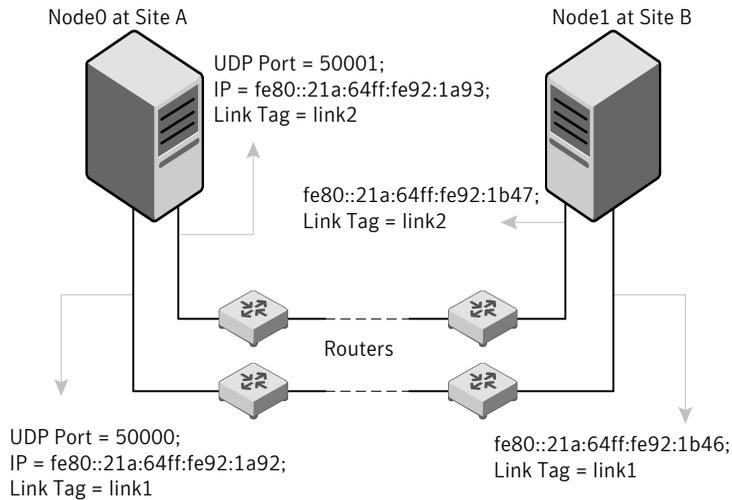
The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

Sample configuration: links crossing IP routers

Figure G-2 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure G-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

#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 fe80::21a:64ff:fe92:1b46
set-addr 0 link2 fe80::21a:64ff:fe92:1b47
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95
```

```
#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

The /etc/llttab file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

#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 fe80::21a:64ff:fe92:1a92
set-addr 1 link2 fe80::21a:64ff:fe92:1a93
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

Troubleshooting VCS installation

This appendix includes the following topics:

- [What to do if you see a licensing reminder](#)
- [Restarting the installer after a failed connection](#)
- [Starting and stopping processes for the Veritas products](#)
- [Installer cannot create UUID for the cluster](#)
- [Script-based installer cannot start LLT, GAB, VCS, or VxFEN on a newly added node](#)
- [LLT startup script displays errors](#)
- [The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails](#)
- [Issues during server-based fencing start up on VCS cluster node](#)
- [Adding a node to the secure cluster whose root broker system has failed](#)

What to do if you see a licensing reminder

In this release, you can install without a license key. In order to comply with the End User License Agreement, you must either install a license key or make the host managed by a Management Server. If you do not comply with these terms within 60 days, the following warning messages result:

```
WARNING V-365-1-1 This host is not entitled to run Veritas Storage
Foundation/Veritas Cluster Server.As set forth in the End User
License Agreement (EULA) you must complete one of the two options
```

```
set forth below. To comply with this condition of the EULA and
stop logging of this message, you have <nn> days to either:
- make this host managed by a Management Server (see
  http://go.symantec.com/sfhakeyless for details and free download),
  or
- add a valid license key matching the functionality in use on this host
  using the command 'vxlicinst' and validate using the command
  'vxkeyless set NONE'
```

To comply with the terms of the EULA, and remove these messages, you must do one of the following within 60 days:

- Install a valid license key corresponding to the functionality in use on the host. After you install the license key, you must validate the license key using the following command:

```
# vxkeyless set NONE
```

- Continue with keyless licensing by managing the server or cluster with a management server.

For more information about keyless licensing, see the following URL:
<http://go.symantec.com/sfhakeyless>

Restarting the installer after a failed connection

If an installation is killed because of a failed connection, you can restart the installer to resume the installation. The installer detects the existing installation. The installer prompts you whether you want to resume the installation. If you resume the installation, the installation proceeds from the point where the installation failed.

Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

To stop the processes

- ◆ Use the `-stop` option to the product installation script.

For example, to stop VCS processes, enter the following command:

```
# ./installvcs -stop
```

To start the processes

- ◆ Use the `-start` option to the product installation script.

For example: To start VCS processes, enter the following command:

```
# ./installvcs -start
```

Installer cannot create UUID for the cluster

The installer displays the following error message if the installer cannot find the `uuidconfig.pl` script before it configures the UUID for the cluster:

```
Couldn't find uuidconfig.pl for uuid configuration,  
please create uuid manually before start vcs
```

You may see the error message during VCS configuration, upgrade, or when you add a node to the cluster using the installer.

Workaround: To start VCS, you must run the `uuidconfig.pl` script manually to configure the UUID on each cluster node.

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

Script-based installer cannot start LLT, GAB, VCS, or VxFEN on a newly added node

After you add a node to an existing cluster and try to start the LLT, GAB, VCS, or VxFEN, the module may not start on the new node. Note that the script-based installer does not throw any error messages about the module starting or failing to start.

Workaround: To start LLT, GAB, VxFEN, and VCS modify the value of the `START` and `STOP` environment variables to 1 in the following files:

- `/etc/sysconfig/gab`
- `/etc/sysconfig/llt`

- /etc/sysconfig/vcs
- /etc/sysconfig/vxfen

You then start LLT, GAB, VxFEN, and VCS on the newly added node in existing cluster.

LLT startup script displays errors

If more than one system on the network has the same LLT node ID and the same cluster ID, then the LLT startup script displays error messages similar to the following:

```
LLT lltconfig ERROR V-14-2-15238 node 1 already exists
in cluster 8383 and has the address - 00:18:8B:E4:DE:27
LLT lltconfig ERROR V-14-2-15241 LLT not configured,
use -o to override this warning
LLT lltconfig ERROR V-14-2-15664 LLT could not
configure any link
```

Recommended action: Ensure that all systems on the network have unique clusterid-nodeid pair. You can use the `lltdump -D` command to get the list of unique clusterid-nodeid pairs connected to the network. This utility is available only for LLT-over-ethernet.

The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails

While running the vxfentsthdw utility, you may see a message that resembles as follows:

```
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 the Hitachi Data Systems 99XX arrays if bit 186 of the system mode option is not enabled.

Issues during server-based fencing start up on VCS cluster node

The following issues may occur during fencing start up on the VCS cluster node:

- `cpsadm` command on the VCS cluster gives connection error
- Authentication failure
- Authorization failure
- Preexisting split-brain

cpsadm command on the VCS cluster node gives connection error

If you receive a connection error message after issuing the `cpsadm` command on the VCS cluster, perform the following actions:

- Ensure that the CP server is reachable from all the VCS cluster nodes.
- Check that the correct CP server virtual IP/virtual hostname and port number are being used by the VCS cluster nodes.
Check the `/etc/vxfenmode` file.
- Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.

Authentication failure

If secure communication has been configured between the CP server and the VCS cluster nodes, authentication failure can occur due to the following causes:

- Symantec Product Authentication Services is not properly configured on the CP server and/or the VCS cluster.
- The CP server and the VCS cluster nodes use the same root broker but the certificate hash of the root broker is not same on the VCS cluster and the CP server. Run the following command on both the CP server and the VCS cluster to see the certificate hash:

```
# cpsat showalltrustedcreds
```

- The CP server and the VCS cluster nodes use different root brokers, and trust is not established between the authentication brokers:
- The CP server and VCS cluster do not have the same security setting.
In order to configure secure communication, both the CP server and the VCS cluster must have same security setting.

In order to have the same security setting, the security parameter must have same value in the `/etc/vxcps.conf` file on CP server and in the `/etc/vxfenmode` file on the VCS cluster nodes.

Authorization failure

Authorization failure occurs when the CP server's VCS cluster nodes or users are not added in the CP server configuration. Therefore, fencing on the VCS cluster node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points. To resolve this issue, add the VCS cluster node and user in the CP server configuration and restart fencing. Refer to the following section:

See [“Preparing the CP servers manually for use by the VCS cluster”](#) on page 215.

Preexisting split-brain

To illustrate preexisting split-brain, assume there are three CP servers acting as coordination points. One of the three CP servers then becomes inaccessible. While in this state, also one client node leaves the cluster. When the inaccessible CP server restarts, it has a stale registration from the node which left the VCS cluster. In this case, no new nodes can join the cluster. Each node that attempts to join the cluster gets a list of registrations from the CP server. One CP server includes an extra registration (of the node which left earlier). This makes the joiner node conclude that there exists a preexisting split-brain between the joiner node and the node which is represented by the stale registration. The situation is similar to that of preexisting split-brain, with coordinator disks, where the problem is solved by the administrator running the `vxfenclearpre` command. A similar solution is required using the `cpsadm` command.

The following `cpsadm` command can be used to clear a registration on a CP server:

```
# cpsadm -s cp_server -a unreg_node -c cluster_name -n nodeid
```

where `cp_server` is the virtual IP address or virtual hostname on which the CP server is listening, `cluster_name` is the VCS name for the VCS cluster, and `nodeid` specifies the node id of VCS cluster node.

After removing all stale registrations, the joiner node will be able to join the cluster.

Adding a node to the secure cluster whose root broker system has failed

If the root broker system of a cluster in secure mode has failed, you can do one of the following before you add a node to the cluster:

- If you had backed up the AT configuration files after you configured the root broker and set up the security infrastructure, you can recover from a root broker failure. Thus, you can enable the root broker to use the same broker certificates and keys for the clusters that the root broker serves. See the Symantec Product Authentication Service documentation for more information on backing up the AT configuration files and recovering the failed root broker system.
- If you did not back up the AT configuration files, then you must unconfigure the authentication brokers in the cluster and repeat the secure cluster configuration steps.

To unconfigure the authentication brokers and enable security in the cluster

- 1 In each of the nodes in the cluster, run the following command to unconfigure the authentication broker. For example,

```
galaxy> # vssregctl -l -s
-b"Security\Authentication\Authentication Broker"
-t"int" -k"Mode" -v0
nebula> # vssregctl -l -s
-b"Security\Authentication\Authentication Broker"
-t"int" -k"Mode" -v0
```

- 2 Perform the following steps to configure the cluster in secure mode.
 - If you use an external root broker, you must reconfigure the root broker. See [“Preparing to configure the clusters in secure mode”](#) on page 77. If you use one of the nodes in the cluster as root broker, proceed to the next step to enable security in the cluster.
 - Run the following command on one of the nodes in the cluster and follow the prompts to enable security in the cluster.

```
# /opt/VRTS/install/installvcs -security
```

See the *Veritas Cluster Server Administrator’s Guide* for more information.

Adding a node to the secure cluster whose root broker system has failed

Sample VCS cluster setup diagrams for CP server-based I/O fencing

This appendix includes the following topics:

- [Configuration diagrams for setting up server-based I/O fencing](#)

Configuration diagrams for setting up server-based I/O fencing

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:

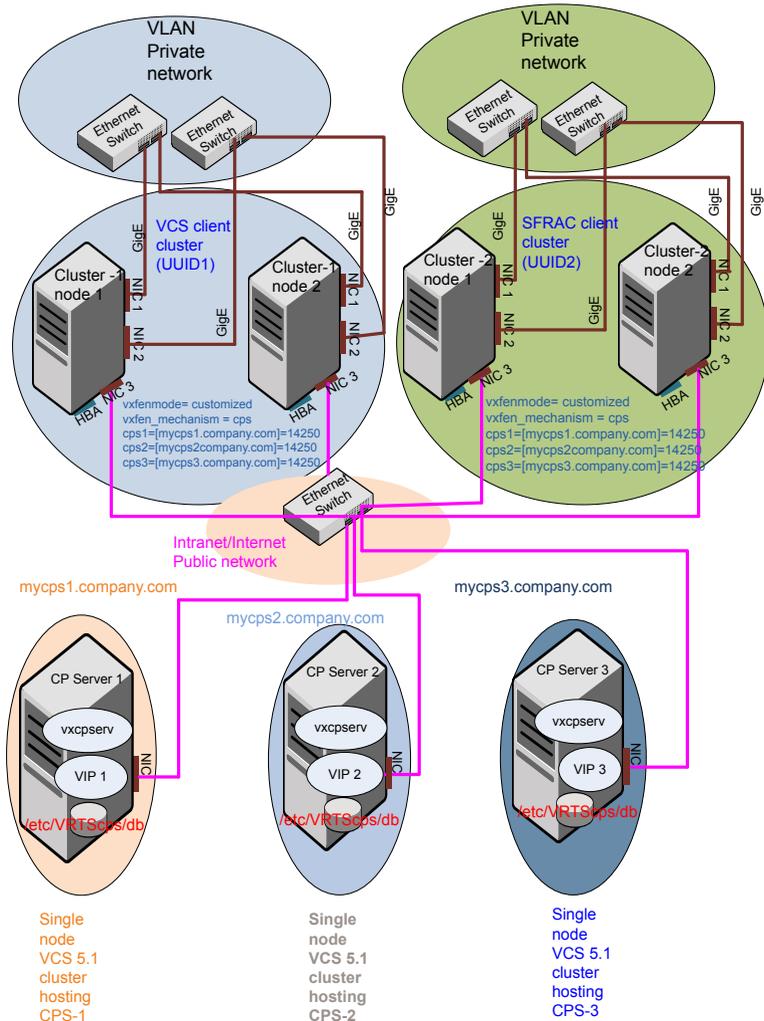
- Two unique client clusters that are served by 3 CP servers:
See [Figure I-1](#) on page 406.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
See [Figure I-2](#) on page 408.
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
See [Figure I-3](#) on page 409.
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:
See [Figure I-4](#) on page 411.

Two unique client clusters served by 3 CP servers

Figure I-1 displays a configuration where two unique client clusters are being served by 3 CP servers (coordination points). Each client cluster has its own unique user ID (UUID1 and UUID2).

In the `vxfsenmode` file on the client nodes, `vxfsenmode` is set to `customized` with `vxfsen` mechanism set to `cps`.

Figure I-1 Two unique client clusters served by 3 CP servers



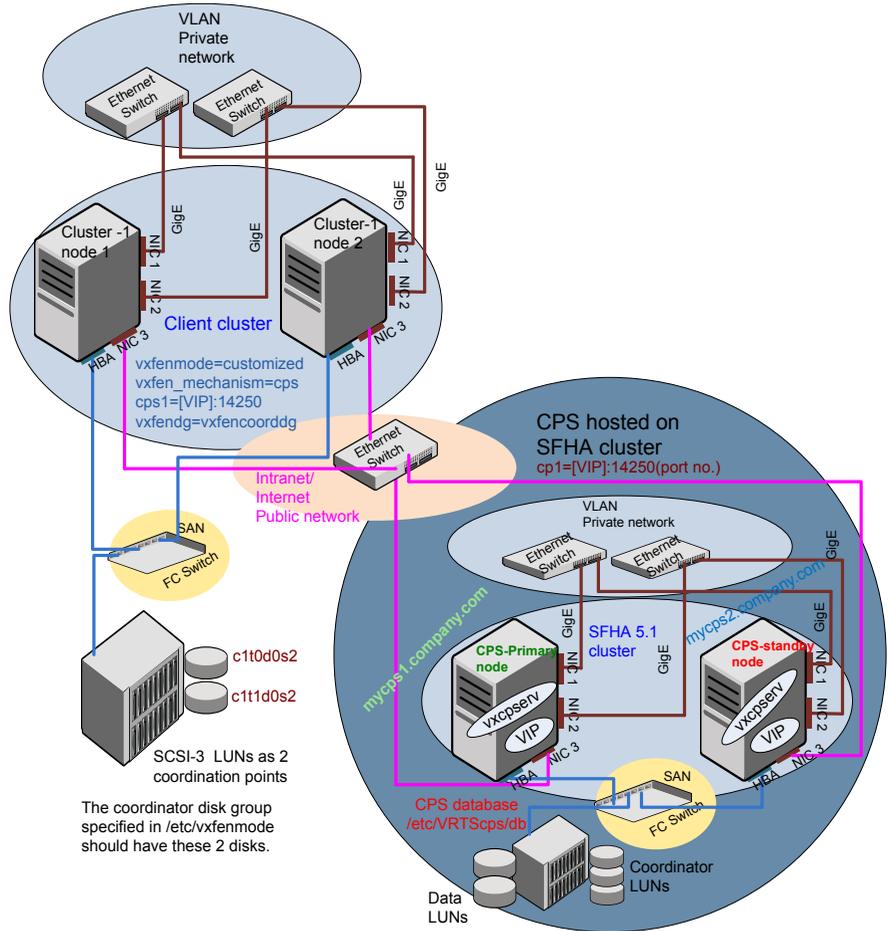
Client cluster served by highly available CPS and 2 SCSI-3 disks

[Figure I-2](#) displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure I-2 Client cluster served by highly available CP server and 2 SCSI-3 disks



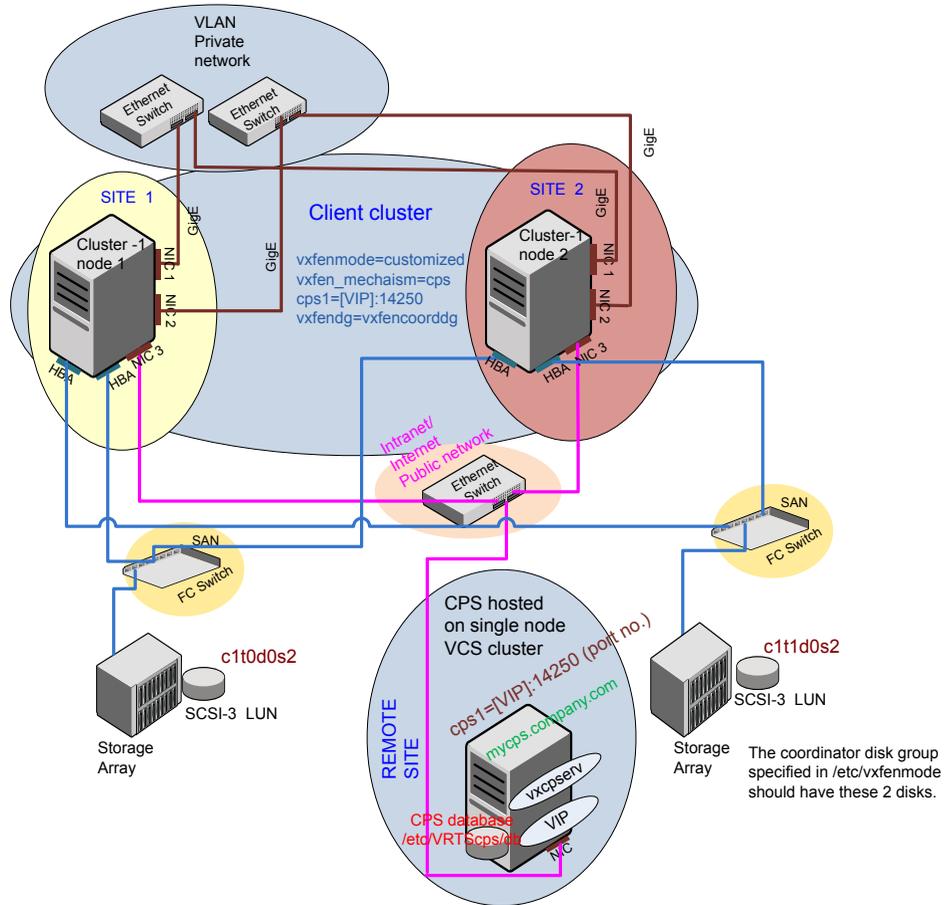
Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure I-3 displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: c1t0d0s2 and c1t1d0s2 which are part of disk group vxfencoordg. The third coordination point is a CP server on a single node VCS cluster.

Figure I-3 Two node campus cluster served by remote CP server and 2 SCSI-3



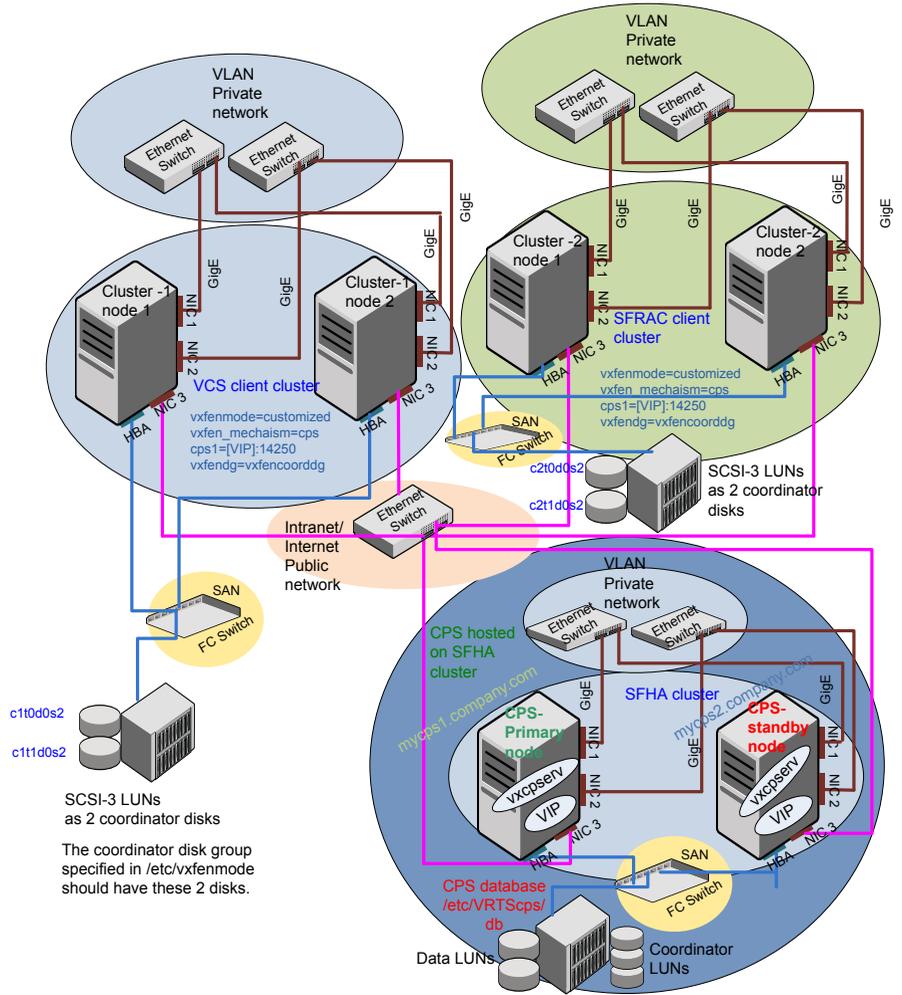
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

[Figure I-4](#) displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNS (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure I-4 Multiple client clusters served by highly available CP server and 2 SCSI-3 disks



c1t1d0s2
c1t1d0s2

SCSI-3 LUNs
as 2 coordinator disks

The coordinator disk group
specified in /etc/vxfenmode
should have these 2 disks.

Index

A

- about
 - global clusters 28
- adding
 - ClusterService group 207
 - users 118
- adding node
 - to a one-node cluster 335
- attributes
 - UseFence 212

B

- bundled agents
 - types.cf file 204

C

- cables
 - cross-over Ethernet 316
- cluster
 - creating a single-node cluster
 - installer 378
 - manual 379
 - four-node configuration 22
 - removing a node from 327
 - verifying operation 290
- Cluster Management Console 30
- Cluster Manager
 - installing Java Console 269
- ClusterService group
 - adding manually 207
- cold start
 - running VCS 24
- commands
 - gabconfig 204, 288
 - hastart 327
 - hastatus 290
 - hasys 290
 - lltconfig 275
 - lltstat 285
 - vxdisksetup (initializing disks) 129

- commands (*continued*)
 - vxlicinst 126, 193
 - vxlicrep 125, 194
- communication channels 24
- communication disk 24
- configuration files
 - types.cf 204
- configuring
 - GAB 204
 - hardware 33
 - LLT
 - manual 201
 - private network 53
 - ssh 55
 - switches 53
- configuring VCS
 - adding users 118
 - event notification 118, 120
 - global clusters 122
 - secure mode 115
 - starting 110
- controllers
 - private Ethernet 53
- coordinator disks
 - DMP devices 87
 - for I/O fencing 87
 - setting up 210

D

- data disks
 - for I/O fencing 87
- demo key 195
- directives
 - LLT 202
- disk space
 - directories 33
 - language pack 33
 - required 33
- disks
 - adding and initializing 129
 - coordinator 210

disks (*continued*)

- testing with vxfcntlsthdw 132
- verifying node access 133

documentation

- accessing 267

E

eeprom

- parameters 53

Ethernet controllers 53, 316

F

fibre channel 33

G

GAB

- description 23
- manual configuration 204
- port membership information 288
- starting 206
- verifying 288

gabconfig command 204, 288

- a (verifying GAB) 288

gatab file

- creating 204
- verifying after installation 275

global clusters 28

- configuration 122

H

hardware

- configuration 23
- configuring network and storage 33

hastart 327

hastatus -summary command 290

hasys -display command 290

hubs 53

- independent 316

I

I/O fencing

- checking disks 132
- setting up 209
- shared storage 132

installation

- required disk space 34

installing

- manual 189
- post 123
- required disk space 33
- Root Broker 81
- simulator 272

installing VCS

- required information 64

installvcs

- options 42

installvcs prompts

- b 43
- n 43
- y 43

J

Java Console

- installing 269
- installing on UNIX 269

K

kernel.panic tunable

- setting 61

L

language packages

- disk space 33

license keys

- adding with vxlicinst 126, 193
- obtaining 52
- replacing demo key 126, 195

licenses

- information about 125
- showing information 194

licensing commands

- vxlicinst 53
- vxlicrep 53
- vxlictest 53

links

- private network 275

LLT

- description 23
- directives 202
- interconnects 61
- manual configuration 201
- starting 206
- verifying 285

- LLT directives
 - link 202
 - link-lowpri 202
 - set-cluster 202
 - set-node 202
- lltconfig command 275
- llthosts file
 - verifying after installation 275
- lltstat command 285
- llttab file
 - verifying after installation 275

M

- MAC addresses 53
- main.cf file
 - contents after installation 279
- main.cf files 369
- MANPATH variable
 - setting 60
- manual installation
 - preparing 187
- media speed 61
 - optimizing 61
- membership information 288
- mounting
 - software disc 62

N

- network partition
 - preexisting 24
 - protecting against 22
- Network partitions
 - protecting against 24
- network switches 53
- NFS 21

O

- optimizing
 - media speed 61
- overview
 - VCS 21

P

- parameters
 - eeprom 53
- PATH variable
 - setting 60
 - VCS commands 285

- persistent reservations
 - SCSI-3 57
- phased 229
- phased upgrade 229
 - example 230
- pkgadd
 - command 189
- port a
 - membership 288
- port h
 - membership 288
- port membership information 288
- preparing
 - manual installation 187
- prerequisites
 - uninstalling 297
- private network
 - configuring 53

R

- RAM
 - installation requirement 33
- removing a system from a cluster 327
- remsh 111
- requirements
 - Ethernet controllers 33
 - fibre channel 33
 - hardware 33
 - RAM Ethernet controllers 33
 - SCSI host bus adapter 33
- response files 44
- Root Broker 29
 - installing 81
- rsh 55, 111

S

- SCSI host bus adapter 33
- SCSI-3
 - persistent reservations 57
- SCSI-3 persistent reservations
 - verifying 209
- seeding 24
 - automatic 24
 - manual 24
- setting
 - kernel.panic tunable 61
 - MANPATH variable 60
 - PATH variable 60

- simulator
 - installing 272
- single-node cluster
 - adding a node to 335
- single-system cluster
 - creating 378–379
 - modifying startup files 380
- SMTP email notification 118
- SNMP trap notification 120
- ssh 55, 111
 - configuring 55
- starting configuration
 - installvcs program 111
 - Veritas product installer 110
- starting VCS after manual upgrade 206
- starting VCS after rpm -i 207
- storage
 - fully shared vs. distributed 23
 - shared 23
- switches 53
- Symantec Product Authentication Service 29, 81, 115
- system communication using rsh
 - ssh 55
- system state attribute value 290

T

- types.cf 204
 - bundled agents 204
- types.cf file 204

U

- uninstalling
 - prerequisites 297
- upgrade
 - phased 229
- upgrading
 - phased 229

V

- variables
 - MANPATH 60
 - PATH 60
- VCS
 - basics 21
 - command directory path variable 285
 - configuration files
 - main.cf 278

VCS (*continued*)

- coordinator disks 210
- documentation 267
- manually installing 189
- replicated states on each system 22
- starting 206–207
- VCS installation
 - verifying
 - cluster operations 285
 - GAB operations 285
 - LLT operations 285
- verifying
 - NIC configuration 124
- vxdisksetup command 129
- vxlicinst 53
- vxlicinst command 126, 193
- vxlicrep 53
- vxlicrep command 125, 194
- vxlictest 53