Symantec™ Cluster Server 6.1 Installation Guide - Solaris

November 2014

Symantec™ 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: 6.1

Document version: 6.1 Rev 4

Legal Notice

Copyright © 2014 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate 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, whether delivered by Symantec as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

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 support 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/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

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

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support 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:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as 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 support 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

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on

page 2 of each guide. The latest product documentation is available on the Symantec website.

https://sort.symantec.com/documents

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:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation

Support agreement resources

If you want to contact Symantec regarding an existing support agreement, please contact the support 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

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

http://www.symantec.com/connect/storage-management

Contents

Technical Support 4

Section 1 Installation overview and planning 22

Chapter 1 Introducing Symantec Cluster Server 23

About Symantec™ Cluster Server 23
About VCS basics 23
About multiple nodes 24
About shared storage 24
About LLT and GAB 25
About network channels for heartbeating 25
About preexisting network partitions 26
About VCS seeding 26
About VCS features 27
About VCS notifications 27
About global clusters 27
About I/O fencing 27
About VCS optional components 28
About Veritas Operations Manager 28
About Cluster Manager (Java Console) 29
About VCS Simulator 29
About Symantec Operations Readiness Tools 30
About configuring VCS clusters for data integrity 32
About I/O fencing for VCS in virtual machines that do not support SCSI-3 PR 32
About I/O fencing components 33
About preferred fencing 35

Chapter 2 System requirements 36

Release notes 36
Important preinstallation information for VCS 37
Hardware requirements for VCS 37
Disk space requirements 38
Supported operating systems 38

Supported software for VCS 38
I/O fencing requirements 39
Coordinator disk requirements for I/O fencing 39
CP server requirements 40
Non-SCSI-3 I/O fencing requirements 43
Number of nodes supported 44
Checking installed product versions and downloading maintenance releases and hot fixes 44
Obtaining installer hot fixes 45
Disabling external network connection attempts 46

Chapter 3 Planning to install VCS 48

VCS installation methods 48
About the script-based installer 49
About the VCS installation program 51
About the web-based installer 53
About response files 53
About installation and configuration methods 55
Typical VCS cluster setup models 57
Typical configuration of two-node VCS cluster 58
Typical configuration of VCS clusters in secure mode 58
Typical configuration of VOM-managed VCS clusters 59

Chapter 4 Licensing VCS 61

About Symantec product licensing 61
Obtaining VCS license keys 62
Installing Symantec product license keys 63

Section 2 Preinstallation tasks 65

Chapter 5 Preparing to install VCS 66

About preparing to install VCS 66
Performing preinstallation tasks 66
Setting up the private network 67
About using ssh or rsh with the installer 70
Setting up shared storage 71
Creating a root user 75
Setting the PATH variable 76
Setting the MANPATH variable 76
Disabling the abort sequence on SPARC systems 76
Optimizing LLT media speed settings on private NICs 78

Guidelines for setting the media speed of the LLT interconnects 78
VCS considerations for Blade server environments 78
Preparing zone environments 79
Mounting the product disc 80
Performing automated preinstallation check 80
Reformatting VCS configuration files on a stopped cluster 81
Getting your VCS installation and configuration information ready 82
Making the IPS publisher accessible 88

Section 3 Installation using the script-based installer 90

Chapter 6 Installing VCS 91

Installing VCS using the installer 91
Installing language packages using the installer 95

Chapter 7 Preparing to configure VCS clusters for data integrity 96

About planning to configure I/O fencing 96
Typical VCS cluster configuration with disk-based I/O fencing 100
Typical VCS cluster configuration with server-based I/O fencing 101
Recommended CP server configurations 102
Setting up the CP server 105
Planning your CP server setup 105
Installing the CP server using the installer 107
Configuring the CP server cluster in secure mode 107
Setting up shared storage for the CP server database 108
Configuring the CP server using the installer program 109
Configuring the CP server using the web-based installer 121
Configuring the CP server manually 122
Configuring CP server using response files 129
Verifying the CP server configuration 133

Chapter 8 Configuring VCS 135

Overview of tasks to configure VCS using the script-based installer 136
Starting the software configuration 137

Specifying systems for configuration 137
Configuring the cluster name 138
Configuring private heartbeat links 139
Configuring the virtual IP of the cluster 142
Configuring Symantec Cluster Server in secure mode 144
Setting up trust relationships for your VCS cluster 145
Configuring a secure cluster node by node 146
Configuring the first node 146
Configuring the remaining nodes 147
Completing the secure cluster configuration 148
Adding VCS users 150
Configuring SMTP email notification 151
Configuring SNMP trap notification 152
Configuring global clusters 154
Completing the VCS configuration 155
Verifying and updating licenses on the system 155
Checking licensing information on the system 156
Updating product licenses 156

Chapter 9 Configuring VCS clusters for data integrity 158

Setting up disk-based I/O fencing using installvcs 158
Initializing disks as VxVM disks 158
Configuring disk-based I/O fencing using installvcs 159
Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installvcs 162
Checking shared disks for I/O fencing 164
Setting up server-based I/O fencing using installvcs 167
Refreshing keys or registrations on the existing coordination points for server-based fencing using the installvcs 176
Setting the order of existing coordination points for server-based fencing using the installvcs 178
Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installvcs 181
Enabling or disabling the preferred fencing policy 183

Section 4 Installation using the Web-based installer 187

Chapter 10 Installing VCS 188

Before using the web-based installer 188
Starting the web-based installer 189

Obtaining a security exception on Mozilla Firefox 189
Performing a preinstallation check with the web-based installer 190
Installing VCS with the web-based installer 190

Chapter 11 Configuring VCS 193

Configuring VCS using the web-based installer 193
Configuring VCS for data integrity using the web-based installer 198
Configuring disk-based fencing for data integrity using the web-based installer 199
Configuring server-based fencing for data integrity using the web-based installer 201
Configuring fencing in disabled mode using the web-based installer 203
Replacing, adding, or removing coordination points using the web-based installer 204
Refreshing keys or registrations on the existing coordination points using web-based installer 206
Setting the order of existing coordination points using the web-based installer 207

Section 5 Automated installation using response files 210

Chapter 12 Performing an automated VCS installation 211

Installing VCS using response files 211
Response file variables to install VCS 212
Sample response file for installing VCS 214

Chapter 13 Performing an automated VCS configuration 216

Configuring VCS using response files 216
Response file variables to configure Symantec Cluster Server 217
Sample response file for configuring Symantec Cluster Server 226

Chapter 14 Performing an automated I/O fencing configuration using response files 227

Configuring I/O fencing using response files 227
Response file variables to configure disk-based I/O fencing 228
Sample response file for configuring disk-based I/O fencing 231
Response file variables to configure server-based I/O fencing 232

Sample response file for configuring server-based I/O fencing 233
Response file variables to configure non-SCSI-3 server-based I/O fencing 234
Sample response file for configuring non-SCSI-3 server-based I/O fencing 235

Section 6 Manual installation 236

Chapter 15 Performing preinstallation tasks 237

Requirements for installing VCS 237

Chapter 16 Manually installing VCS 238

About VCS manual installation 238
Installing VCS software manually 238
Viewing the list of VCS packages 239
Installing VCS packages for a manual installation 240
Manually installing packages on Oracle Solaris 11 systems 241
Manually installing packages on Solaris brand non-global zones 242
Manually installing packages on solaris10 brand zones 243
Installing language packages in a manual installation 244
Adding a license key for a manual installation 245
Copying the installation guide to each node 247
Installing VCS on Solaris 10 using JumpStart 248
Overview of JumpStart installation tasks 248
Generating the finish scripts 248
Preparing installation resources 249
Adding language pack information to the finish file 250
Using a Flash archive to install VCS and the operating system 251
Installing VCS on Solaris 11 using Automated Installer 253
About Automated Installation 253
Using Automated Installer 254
Using AI to install the Solaris 11 operating system and SFHA products 255

Chapter 17 Manually configuring VCS 259

About configuring VCS manually 259
Configuring LLT manually 259
Setting up /etc/llthosts for a manual installation 260
Setting up /etc/lltab for a manual installation 260

About LLT directives in /etc/llttab file 261
Additional considerations for LLT for a manual installation 262
Configuring GAB manually 263
Configuring VCS manually 263
Configuring the cluster UUID when creating a cluster manually 265
Configuring VCS in single node mode 265
Disabling LLT, GAB, and I/O fencing on a single node cluster 266
Enabling LLT, GAB, and I/O fencing on a single node cluster 267
Starting LLT, GAB, and VCS after manual configuration 269
Modifying the VCS configuration 270
Configuring the ClusterService group 271

Chapter 18 Manually configuring the clusters for data integrity 272

Setting up disk-based I/O fencing manually 272
Identifying disks to use as coordinator disks 273
Setting up coordinator disk groups 273
Creating I/O fencing configuration files 274
Modifying VCS configuration to use I/O fencing 275
Verifying I/O fencing configuration 277
Setting up server-based I/O fencing manually 277
Preparing the CP servers manually for use by the VCS cluster 278
Generating the client key and certificates manually on the client nodes 281
Configuring server-based fencing on the VCS cluster manually 283
Configuring CoordPoint agent to monitor coordination points 290
Verifying server-based I/O fencing configuration 291
Setting up non-SCSI-3 fencing in virtual environments manually 292
Sample /etc/vxfenmode file for non-SCSI-3 fencing 294

Section 7 Managing your Symantec deployments 299

Chapter 19 Performing centralized installations using the Deployment Server 300

About the Deployment Server 301
How to install the Deployment Script 302
Deployment management overview 303
Setting up a Deployment Server 304
Setting deployment preferences 306
Using the Deployment Server command line option to specify a non-default repository location 307
Using the Deployment Server command line options to load and download the most recent release information 308
Viewing or downloading available release images 309
Viewing or removing repository images stored in your repository 313
Deploying Symantec product updates to your environment 316
Finding out which releases you have, and which upgrades or updates you may need 317
Deploying a specific Symantec release 319
Updating release information on systems without Internet access 320

Section 8 Upgrading VCS 322

Chapter 20 Planning to upgrade VCS 323

About upgrading to VCS 6.1 323
VCS supported upgrade paths 324
Upgrading VCS in secure enterprise environments 327
Considerations for upgrading secure VCS 5.x clusters to VCS 6.1 328
Considerations for upgrading VCS to 6.1 on systems configured with an Oracle resource 328
Considerations for upgrading secure CP servers 329
Considerations for upgrading secure CP clients 329
Setting up trust relationship between CP server and CP clients manually 329
Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes 330

Chapter 21 Performing a typical VCS upgrade using the installer 333

Before upgrading using the script-based or Web-based installer 333
Upgrading VCS using the script-based installer 333
Upgrading VCS using the web-based installer 335

Chapter 22 Performing a phased upgrade of VCS 337

About phased upgrade 337
Prerequisites for a phased upgrade 337
Planning for a phased upgrade 338
Phased upgrade limitations 338
Phased upgrade example 338
Phased upgrade example overview 339
Performing a phased upgrade using the script-based installer 340
Moving the service groups to the second subcluster 340
Upgrading the operating system on the first subcluster 345
Upgrading the first subcluster 345
Preparing the second subcluster 347
Activating the first subcluster 351
Upgrading the operating system on the second subcluster 353
Upgrading the second subcluster 354
Finishing the phased upgrade 355

Chapter 23 Performing an automated VCS upgrade using response files 359

Upgrading VCS using response files 359
Response file variables to upgrade VCS 360
Sample response file for upgrading VCS 361
Performing rolling upgrade of VCS using response files 362
Response file variables to upgrade VCS using rolling upgrade 363
Sample response file for VCS using rolling upgrade 364

Chapter 24 Performing a rolling upgrade 365

Supported rolling upgrade paths 365
About rolling upgrades 366
Performing a rolling upgrade using the installer 368
Performing a rolling upgrade using the script-based installer 368
Performing a rolling upgrade of VCS using the web-based installer 370

Chapter 25 Upgrading VCS using Live Upgrade and Boot Environment upgrade 373

About Live Upgrade 373
Symantec Cluster Server exceptions for Live Upgrade 374
About ZFS Boot Environment (BE) upgrade 374
Supported upgrade paths for Live Upgrade and Boot Environment upgrade 375
Performing Live Upgrade on Solaris 10 systems 376
Before you upgrade VCS using Solaris Live Upgrade 377
Creating a new Solaris 10 boot environment on the alternate boot disk 379
Upgrading VCS using the installer for Solaris 10 Live Upgrade 383
Completing the Solaris 10 Live Upgrade 384
Verifying the Solaris 10 Live Upgrade of VCS 385
Administering boot environments in Solaris 10 Live Upgrade 386
Performing Boot Environment upgrade on Solaris 11 systems 388
Creating a new Solaris 11 BE on the primary boot disk 388
Upgrading VCS using the installer for upgrading BE on Solaris 11 389
Completing the VCS upgrade on BE on Solaris 11 391
Verifying Solaris 11 BE upgrade 391
Administering BEs in Solaris 11 BE upgrade 392

Section 9 Post-installation tasks 395

Chapter 26 Performing post-installation tasks 396

About enabling LDAP authentication for clusters that run in secure mode 396
Enabling LDAP authentication for clusters that run in secure mode 398
Accessing the VCS documentation 403
Removing permissions for communication 403
Changing root user into root role 403

Chapter 27 Installing or upgrading VCS components 404

Installing the Java Console 404
Software requirements for the Java Console 404
Hardware requirements for the Java Console 405
Installing the Java Console on Solaris 405

Installing the Java Console on a Windows system 406
Upgrading the Java Console 406
Installing VCS Simulator 407
Software requirements for VCS Simulator 407
Installing VCS Simulator on Windows systems 407
Reviewing the installation 408
Upgrading VCS Simulator 409

Chapter 28 Verifying the VCS installation 410

About verifying the VCS installation 410
About the cluster UUID 410
Verifying the LLT, GAB, and VCS configuration files 411
Verifying LLT, GAB, and cluster operation 411
Verifying LLT 412
Verifying GAB 414
Verifying the cluster 415
Verifying the cluster nodes 416
Upgrading the disk group version 419
Performing a postcheck on a node 420
About using the postcheck option 420

Section 10 Adding and removing cluster nodes 423

Chapter 29 Adding a node to a single-node cluster 424

Adding a node to a single-node cluster 424
Setting up a node to join the single-node cluster 425
Installing and configuring Ethernet cards for private network 426
Configuring the shared storage 427
Bringing up the existing node 427
Installing the VCS software manually when adding a node to a single node cluster 428
Creating configuration files 428
Starting LLT and GAB 428
Reconfiguring VCS on the existing node 428
Verifying configuration on both nodes 430

Chapter 30 Adding a node to a multi-node VCS cluster 431

Adding nodes using the VCS installer 431
Adding a node using the web-based installer 433
Manually adding a node to a cluster 434
Setting up the hardware 435

Installing the VCS software manually when adding a node 436
Setting up the node to run in secure mode 437
Configuring LLT and GAB when adding a node to the cluster 440
Configuring I/O fencing on the new node 442
Adding the node to the existing cluster 446
Starting VCS and verifying the cluster 447
Adding a node using response files 447
Response file variables to add a node to a VCS cluster 448
Sample response file for adding a node to a VCS cluster 448

Chapter 31 Removing a node from a VCS cluster 450

Removing a node from a VCS cluster 450
Verifying the status of nodes and service groups 451
Deleting the departing node from VCS configuration 452
Modifying configuration files on each remaining node 455
Removing the node configuration from the CP server 456
Removing security credentials from the leaving node 457
Unloading LLT and GAB and removing VCS on the departing node 458

Section 11 Uninstallation of VCS 460

Chapter 32 Uninstalling VCS using the installer 461

Preparing to uninstall VCS 461
Uninstalling VCS using the script-based installer 462
Removing VCS 6.1 packages 462
Running uninstallvcs from the VCS 6.1 disc 463
Uninstalling VCS with the web-based installer 464
Removing language packages using the uninstaller program 465
Removing the CP server configuration using the installer program 465

Chapter 33 Uninstalling VCS using response files 467

Uninstalling VCS using response files 467
Response file variables to uninstall VCS 468
Sample response file for uninstalling VCS 469

Chapter 34 Manually uninstalling VCS 470

Removing VCS packages manually 470
Manually remove the CP server fencing configuration 473

Manually deleting cluster details from a CP server 474
Manually uninstalling VCS packages on non-global zones on Solaris 11 475

Section 12 Installation reference 477

Appendix A Services and ports 478

About SFHA services and ports 478

Appendix B VCS installation packages 481

Symantec Cluster Server installation packages 481

Appendix C Installation command options 486

Command options for installvcs 486
Installation script options 487
Command options for uninstallvcs 493

Appendix D Changes to bundled agents in VCS 6.1 495

Deprecated agents 495
New agents 496
New and modified attributes for VCS 6.1 agents 497
Manually removing deprecated resource types and modifying attributes 522
Creating new VCS accounts if you used native operating system accounts 524

Appendix E Configuration files 526

About the LLT and GAB configuration files 526
About the AMF configuration files 529
About the VCS configuration files 530
Sample main.cf file for VCS clusters 531
Sample main.cf file for global clusters 533
About I/O fencing configuration files 534
Sample configuration files for CP server 537
Sample main.cf file for CP server hosted on a single node that runs VCS 538
Sample main.cf file for CP server hosted on a two-node SFHA cluster 540
Sample CP server configuration (/etc/vxcps.conf) file output 543
Packaging related SMF services on Solaris 11 544

Appendix F Installing VCS on a single node 545

About installing VCS on a single node 545
Creating a single-node cluster using the installer program 546
Preparing for a single node installation 546
Starting the installer for the single node cluster 546
Creating a single-node cluster manually 547
Setting the path variable for a manual single node installation 547
Installing VCS software manually on a single node 548
Configuring VCS 548
Verifying single-node operation 548

Appendix G Configuring LLT over UDP 549

Using the UDP layer for LLT 549
When to use LLT over UDP 549
Manually configuring LLT over UDP using IPv4 549
Broadcast address in the /etc/llttab file 550
The link command in the /etc/llttab file 551
The set-addr command in the /etc/llttab file 551
Selecting UDP ports 552
Configuring the netmask for LLT 553
Configuring the broadcast address for LLT 553
Sample configuration: direct-attached links 554
Sample configuration: links crossing IP routers 556
Manually configuring LLT over UDP using IPv6 558
The link command in the /etc/llttab file 559
The set-addr command in the /etc/llttab file 560
Selecting UDP ports 560
Sample configuration: direct-attached links 561
Sample configuration: links crossing IP routers 563
LLT over UDP sample /etc/llttab 565

Appendix H Configuring the secure shell or the remote shell for communications 567

Setting up inter-system communication 567
Setting up ssh on cluster systems 567
Configuring ssh 568

Appendix I Troubleshooting VCS installation 572

What to do if you see a licensing reminder 572
Restarting the installer after a failed connection 573

Starting and stopping processes for the Symantec products 573
Installer cannot create UUID for the cluster 574
LLT startup script displays errors 575
The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails 575
Issues during fencing startup on VCS cluster nodes set up for server-based fencing 576

Appendix J Sample VCS cluster setup diagrams for CP server-based I/O fencing 577

Configuration diagrams for setting up server-based I/O fencing 577
Two unique client clusters served by 3 CP servers 577
Client cluster served by highly available CPS and 2 SCSI-3 disks 578
Two node campus cluster served by remote CP server and 2 SCSI-3 disks 580
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks 582

Appendix K Reconciling major/minor numbers for NFS shared disks 584

Reconciling major/minor numbers for NFS shared disks 584
Checking major and minor numbers for disk partitions 585
Checking the major and minor number for VxVM volumes 588

Appendix L Compatibility issues when installing Symantec Cluster Server with other products 591

Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present 591
Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present 592
Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present 592

Appendix M Upgrading the Steward process 593

Upgrading the Steward process 593

Index 596

Section 1

Installation overview and planning

- Chapter 1. Introducing Symantec Cluster Server
- Chapter 2. System requirements
- Chapter 3. Planning to install VCS
- Chapter 4. Licensing VCS

Chapter 1

Introducing Symantec Cluster Server

This chapter includes the following topics:

- About Symantec™ Cluster Server
- About VCS basics
- About VCS features
- About VCS optional components
- About Symantec Operations Readiness Tools
- About configuring VCS clusters for data integrity

About Symantec™ Cluster Server

Symantec™ Cluster Server by Symantec is a high-availability solution for applications and services configured in a cluster. Symantec 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 storage devices. When a system is part of a VCS cluster, it is called a node. VCS monitors and controls applications running in the cluster on nodes, and restarts applications in response to a variety of hardware or software faults.

Applications can continue to operate with little or no downtime. In some cases, such as NFS, this continuation is transparent to high-level applications and users. In

other cases, a user might have to retry an operation, such as a Web server reloading a page.

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

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

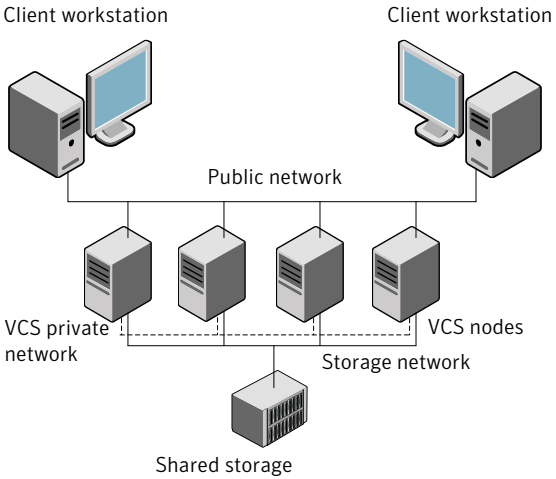

Client workstations receive service over the public network from applications running on VCS nodes. VCS monitors the nodes and their services. VCS nodes in the cluster communicate over a private network.

About multiple nodes

VCS runs in a replicated state on each node in the cluster. A private network enables the nodes to share identical state information about all resources. The private network also recognizes active nodes, 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

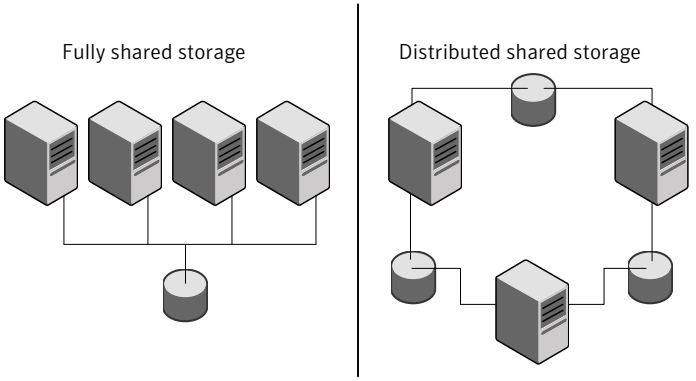

About LLT and GAB

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.

GAB (Group Membership and Atomic Broadcast) provides globally ordered message that is required to maintain a synchronized state among the nodes.

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 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 *Symantec Cluster Server Administrator's Guide*.

Figure 1-3 illustrates a two-node VCS cluster where the nodes sys1 and sys2 have two private network connections.

Figure 1-3 Two Ethernet connections connecting two nodes

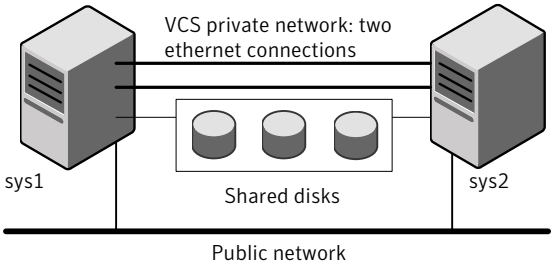

About preexisting network partitions

A preexisting network partition refers to 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 the concept of seeding. Seeding 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. However, if you have I/O fencing enabled in your cluster, you can still configure GAB to automatically seed the cluster even when some cluster nodes are unavailable.

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

About VCS features

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 27.
I/O fencing	See "About I/O fencing" on page 27.

About VCS notifications

You can configure both Simple Network Management Protocol (SNMP) and Simple Mail Transfer Protocol (SMTP) notifications for VCS. Symantec recommends you to configure at least 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 *Symantec 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 *Symantec 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.

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 installer installs the I/O fencing driver, VRTSvxfen package, when you install VCS. 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.

I/O fencing coordination points can be coordinator disks or coordination point servers (CP servers) or both. You can configure disk-based or server-based I/O fencing:

Disk-based I/O fencing	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.
Server-based I/O fencing	I/O fencing that uses at least one CP server system is referred to as server-based I/O fencing. Server-based fencing can include only CP servers, or a mix of CP servers and coordinator disks. Server-based I/O fencing ensures data integrity in clusters. In virtualized environments that do not support SCSI-3 PR, VCS supports non-SCSI-3 server-based I/O fencing.

See "About planning to configure I/O fencing" on page 96.

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

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

About VCS optional components

You can add the following optional components to VCS:

Veritas Operations Manager	See "About Veritas Operations Manager" on page 28.
Cluster Manager (Java console)	See "About Cluster Manager (Java Console)" on page 29.
VCS Simulator	See "About VCS Simulator" on page 29.

About Veritas Operations Manager

Veritas Operations Manager provides a centralized management console for Symantec Storage Foundation and High Availability products. You can use Veritas

Operations Manager to monitor, visualize, and manage storage resources and generate reports.

Symantec recommends using Veritas Operations Manager (VOM) to manage Storage Foundation and Cluster Server environments.

You can download Veritas Operations Manager at no charge at http://go.symantec.com/vom.

Refer to the Veritas Operations Manager documentation for installation, upgrade, and configuration instructions.

If you want to manage a single cluster using Cluster Manager (Java Console), a version is available for download from http://www.symantec.com/operations-manager/support. You cannot manage the new features of this release using the Java Console. Symantec Cluster Server Management Console is deprecated.

About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers administration capabilities for your cluster. Use the different views in the Java Console to monitor and manage clusters and Symantec Cluster Server (VCS) objects, including service groups, systems, resources, and resource types. You cannot manage the new features of releases 6.0 and later using the Java Console.

See *Symantec Cluster Server Administrator's Guide*.

You can download the console from http://www.symantec.com/operations-manager/support. You will need a (free) SymAccount for downloading.

The Veritas Cluster Server Management Console is deprecated. Symantec recommends using Veritas Operations Manager to manage Storage Foundation and Cluster Server environments.

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. You can install VCS Simulator only on 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 to test configurations for VCS clusters on Windows, AIX, HP-UX, Linux, and Solaris operating systems. 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/operations-manager/support.

About Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a website that automates and simplifies some of the most time-consuming administrative tasks. It helps you identify risks in your datacenters and improve operational efficiency, enabling you to manage the complexity that is associated with datacenter architectures and scale

Table 1-1 lists three major datacenter tasks and the SORT tools that can help you accomplish them.

Table 1-1 Datacenter tasks and the SORT tools

Task	SORT tools
Prepare for installations and upgrades	■ Installation and Upgrade checklists Display system requirements including memory, disk space, and architecture. ■ Installation and Upgrade custom reports Create reports that determine if you're ready to install or upgrade a Symantec enterprise product. ■ Array-specific Module Finder List the latest Array Support Libraries (ASLs) and Array Policy Modules (APMs) for UNIX servers, and Device Driver Installers (DDIs) and Device Discovery Layers (DDLs) for Windows servers. ■ High Availability Agents table Find and download the agents for applications, databases, replication, and Symantec partners.

Table 1-1 Datacenter tasks and the SORT tools *(continued)*

Task	SORT tools
Identify risks and get server-specific recommendations	■ Patch notifications Receive automatic email notifications about patch updates. (Sign in required.) ■ Risk Assessment check lists Display configuration recommendations based on your Symantec product and platform. ■ Risk Assessment custom reports Create reports that analyze your system and give you recommendations about system availability, storage use, performance, and best practices. ■ Error code descriptions and solutions Display detailed information on thousands of Symantec error codes.
Improve efficiency	■ Patch Finder List and download patches for your Symantec enterprise products. ■ License/Deployment custom reports Create custom reports that list your installed Symantec products and license keys. Display licenses by product, platform, server tier, and system. ■ Symantec Performance Value Unit (SPVU) Calculator Use the calculator to assist you with the pricing meter transition. ■ Documentation List and download Symantec product documentation, including manual pages, product guides, and support articles. ■ Related links Display links to Symantec product support, forums, customer care, and vendor information on a single page.

SORT is available at no additional charge.

To access SORT, go to:

https://sort.symantec.com

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.

See “About planning to configure I/O fencing” on page 96.

About I/O fencing for VCS in virtual machines that do not support SCSI-3 PR

In a traditional I/O fencing implementation, where the coordination points are coordination point servers (CP servers) or coordinator disks, Clustered Volume Manager (CVM) and Veritas I/O fencing modules provide SCSI-3 persistent reservation (SCSI-3 PR) based protection on the data disks. This SCSI-3 PR protection ensures that the I/O operations from the losing node cannot reach a disk that the surviving sub-cluster has already taken over.

See the *Symantec Cluster Server Administrator's Guide* for more information on how I/O fencing works.

In virtualized environments that do not support SCSI-3 PR, VCS attempts to provide reasonable safety for the data disks. VCS requires you to configure non-SCSI-3 server-based I/O fencing in such environments. Non-SCSI-3 fencing uses CP servers as coordination points with some additional configuration changes to support I/O fencing in such environments.

See "Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installvcs" on page 181.

See "Setting up non-SCSI-3 fencing in virtual environments manually" on page 292.

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 33.
- Coordination points—Act as a global lock during membership changes
  See "About coordination points" on page 33.

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. VCS prevents split-brain when vxfen races for control of the coordination points and the winner partition fences the ejected nodes from accessing the data disks.

Note: Typically, a fencing configuration for a cluster must have three coordination points. Symantec also supports server-based fencing with a single CP server as its only coordination point with a caveat that this CP server becomes a single point of failure.

The coordination points can either be disks or servers or both.

- 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's Dynamic Multi-pathing (DMP) feature. Dynamic Multi-pathing (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.

  Note: The raw disk policy supports I/O fencing only when a single hardware path from the node to the coordinator disks is available. If there are multiple hardware paths from the node to the coordinator disks then we support dmp disk policy. If few coordinator disks have multiple hardware paths and few have a single hardware path, then we support only the dmp disk policy. For new installations, Symantec recommends IO fencing with dmp disk policy even for a single hardware path.

  See the *Symantec Storage Foundation 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 active VCS cluster
  - Self-unregister from this active VCS 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 preferred fencing

The I/O fencing driver uses coordination points to prevent split-brain in a VCS cluster. By default, the fencing driver favors the subcluster with maximum number of nodes during the race for coordination points. With the preferred fencing feature, you can specify how the fencing driver must determine the surviving subcluster.

You can configure the preferred fencing policy using the cluster-level attribute PreferredFencingPolicy for the following:

- Enable system-based preferred fencing policy to give preference to high capacity systems.
- Enable group-based preferred fencing policy to give preference to service groups for high priority applications.
- Enable site-based preferred fencing policy to give preference to sites with higher priority.
- Disable preferred fencing policy to use the default node count-based race policy.

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

See “Enabling or disabling the preferred fencing policy” on page 183.

Chapter 2

System requirements

This chapter includes the following topics:

- Release notes
- Important preinstallation information for VCS
- Hardware requirements for VCS
- Disk space requirements
- Supported operating systems
- Supported software for VCS
- I/O fencing requirements
- Number of nodes supported
- Checking installed product versions and downloading maintenance releases and hot fixes
- Obtaining installer hot fixes
- Disabling external network connection attempts

Release notes

The *Release Notes* for each Symantec product contains last-minute news and important details for each product, including updates to system requirements and supported software. Review the *Release notes* for the latest information before you start installing the product.

The product documentation is available on the web at the following location:

https://sort.symantec.com/documents

Important preinstallation information for VCS

Before you install VCS, make sure that you have reviewed the following information:

- Preinstallation checklist for your configuration. Go to the SORT installation checklist tool. From the drop-down lists, select the information for the Symantec product you want to install, and click **Generate Checklist**.
- Hardware compatibility list for information about supported hardware: http://www.symantec.com/docs/TECH211575
- For important updates regarding this release, review the Late-Breaking News Technote on the Symantec Technical Support website: http://www.symantec.com/docs/TECH211540
- You can install VCS on clusters of up to 64 systems.
  Every system where you want to install VCS must meet the hardware and the software requirements.

Hardware requirements for VCS

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 64 SPARC systems running either Oracle Solaris 10 or Oracle Solaris 11 as appropriate.
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.
Disks	Typical VCS configurations require that the applications are configured to use shared disks/storage to enable migration of applications between systems in the cluster. The VCS I/O fencing feature requires that all data and coordinator disks support SCSI-3 Persistent Reservations (PR). See "About planning to configure I/O fencing" on page 96.
Disk space	See "Disk space requirements" on page 38. **Note:** VCS may require more temporary disk space during installation than the specified disk space.

Table 2-1 Hardware requirements for a VCS cluster *(continued)*

Item	Description
Ethernet controllers	In addition to the built-in public Ethernet controller, VCS requires at least one more Ethernet interface per system. Symantec recommends two additional network interfaces for private interconnects. You can also configure aggregated interfaces. Symantec recommends that you turn off the spanning tree algorithm on the switches used to connect private network interfaces..
Fibre Channel or SCSI host bus adapters	Typical VCS configuration requires at least one SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.
RAM	Each VCS node requires at least 1024 megabytes.

Disk space requirements

Before installing your products, confirm that your system has enough free disk space.

Use the **Perform a Preinstallation Check (P)** menu for the web-based installer to determine whether there is sufficient space.

Or, go to the installation directory and run the installer with the `-precheck` option.

```
# ./installer -precheck
```

See "About the script-based installer" on page 49.

Supported operating systems

For information on supported operating systems, see the *Symantec Cluster Server Release Notes*.

Supported software for VCS

VCS supports the following versions of Symantec Storage Foundation:

Symantec Storage Foundation: Veritas Volume Manager (VxVM) with Veritas File System (VxFS)

Oracle Solaris 11

- Storage Foundation 6.1

- VxVM 6.1 with VxFS 6.1
- Storage Foundation 6.0.3
  - VxVM 6.0.3 with VxFS 6.0.3

Oracle Solaris 10

- Storage Foundation 6.1
  - VxVM 6.1 with VxFS 6.1
- Storage Foundation 6.0.3
  - VxVM 6.0.3 with VxFS 6.0.3

Note: VCS supports the previous and the next versions of Storage Foundation to facilitate product upgrades.

For supported database versions of enterprise agents, refer the support matrix at http://www.symantec.com/business/support/index?page=content&id=DOC4039.

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 39.
- CP servers
  See "CP server requirements" on page 40.

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 *Symantec Storage Foundation and High Availability Installation Guide*.

If you have installed VCS in a virtual environment that is not SCSI-3 PR compliant, review the requirements to configure non-SCSI-3 server-based fencing.

See "Non-SCSI-3 I/O fencing requirements" on page 43.

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 at least three coordinator disks or there must be odd number of coordinator disks.
- The coordinator disks can be raw devices, DMP devices, or iSCSI devices.
- 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.
- The coordinator disk size must be at least 128 MB.

CP server requirements

VCS 6.1 clusters (application clusters) support coordination point servers (CP servers) which are hosted on the following VCS and SFHA versions:

- VCS 6.1 single-node cluster
  Single-node VCS clusters that hosts CP server requires you to configure LLT and GAB.
- SFHA 6.1 cluster

Warning: If you want to upgrade application clusters that use CP server based fencing to 6.1, make sure that you first upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.1. However, note that the CP server upgraded to 6.1 can support application clusters on 6.1 (HTTPS-based communication) and application clusters prior to 6.1 (IPM-based communication). When you configure the CP server, the installer asks the VIPs for HTTPS-based communication (if the clients are on release version 6.1) or VIPs for IPM-based communication (if the clients are on a release version prior to 6.1).

Make sure that you meet the basic hardware requirements for the VCS/SFHA cluster to host the CP server.

See the *Symantec Storage Foundation High Availability Installation Guide*.

See "Hardware requirements for VCS" on page 37.

Note: While Symantec recommends at least three coordination points for fencing, a single CP server as coordination point is a supported server-based fencing configuration. Such single CP server fencing configuration requires that the coordination point be a highly available CP server that is hosted on an SFHA cluster.

Make sure you meet the following additional CP server requirements which are covered in this section before you install and configure CP server:

- Hardware requirements
- Operating system requirements
- Networking requirements (and recommendations)
- Security requirements

Table 2-2 lists additional requirements for hosting the CP server.

Table 2-2 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: ■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB) ■ 300 MB in /usr ■ 20 MB in /var ■ 10 MB in /etc (for the CP server database) See "Disk space requirements" on page 38.
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the nodes of this SFHA cluster.
RAM	Each CP server requires at least 512 MB.
Network	Network hardware capable of providing TCP/IP connection between CP servers and VCS clusters (application clusters).

Table 2-3 displays the CP server supported operating systems and versions. An application cluster can use a CP server that runs any of the following supported operating systems.

Table 2-3 CP server supported operating systems and versions

CP server	Operating system and version

CP server hosted on a VCS single-node cluster or on an SFHA cluster	CP server supports any of the following operating systems:
	■ AIX 6.1 and 7.1
	■ Linux:
	  ■ RHEL 5
	  ■ RHEL 6
	  ■ SLES 11
	■ Oracle Solaris 10
	■ Oracle Solaris 11
	Review other details such as supported operating system levels and architecture for the supported operating systems.
	See the *Symantec Cluster Server Release Notes* or the *Symantec Storage Foundation High Availability Release Notes* for that platform.

Following are the CP server networking requirements and recommendations:

- Symantec recommends that network access from the application 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 application clusters by these network paths. The CP server listens for messages from the application clusters using TCP port 443 if the communication happens over the HTTPS protocol. TCP port 443 is the default port that can be changed while you configure the CP server. The CP server listens for messages from the application clusters over the IPM-based protocol using the TCP port 14250. Unlike HTTPS protocol, which is a standard protocol, IPM (Inter Process Messaging) is a VCS-specific communication protocol.
  Symantec recommends that you configure multiple network paths to access a CP server. If a network path fails, CP server does not require a restart and continues to listen on all the other available virtual IP addresses.
- The CP server supports either Internet Protocol version 4 (IPv4 addresses) when communicating with the application clusters over the IPM-based protocol. The CP server only supports Internet Protocol version 4 (IPv4) when communicating with the application clusters over the HTTPS protocol.
- When placing the CP servers within a specific network configuration, you must take into consideration the number of hops from the different application cluster nodes to the CP servers. As a best practice, Symantec recommends that the number of hops and network latency from the different application cluster nodes

to the CP servers 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 difference in number of hops or network latency between the CPS and various nodes.

For communication between the VCS cluster (application cluster) and CP server, review the following support matrix:

Table 2-4 Supported communication modes between VCS cluster (application cluster) and CP server

Communication mode	CP server (HTTPS-based communication)	CP server (IPM-based secure communication)	CP server (IPM-based non-secure communication)
VCS cluster (release version 6.1)	Yes	No	No
VCS cluster (release version prior to 6.1)	No	Yes	Yes
VCS cluster in non-secure mode (release version prior to 6.1)	No	No	Yes

For secure communications between the VCS and CP server over the IPM-based protocol, consider the following requirements and suggestions:

- In a secure communication environment, all CP servers that are used by the application cluster must be configured with security enabled. A configuration where the application cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.
- For non-secure communication between CP server and application clusters, there is no need to configure Symantec Product Authentication Service. In non-secure mode, authorization is still provided by CP server for the application 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 information about establishing secure communications between the application cluster and CP server, see the *Symantec Cluster Server Administrator's Guide*.

Non-SCSI-3 I/O fencing requirements

Supported virtual environment for non-SCSI-3 fencing:

- Oracle Solaris 11 SRU 11.1.x.y.z, and Oracle Solaris 10 Update 9, 10, and 11
  Oracle VM Server for SPARC 2.0, 2.1, 2.2, and 3.0
  Guest operating system: Oracle Solaris 10, Oracle Solaris 11

Make sure that you also meet the following requirements to configure fencing in the virtual environments that do not support SCSI-3 PR:

- VCS must be configured with Cluster attribute UseFence set to SCSI3
- All coordination points must be CP servers

Number of nodes supported

VCS supports cluster configurations with up to 64 nodes.

Checking installed product versions and downloading maintenance releases and hot fixes

Symantec provides a means to check the Symantec packages you have installed, and download any needed maintenance releases and hot fixes.

Use the `installer` command with the `-version` option to determine what is installed on your system, and download any needed maintenance releases or hot fixes. After you have installed the current version of the product, you can use the `showversion` script in the `/opt/VRTS/install` directory to find product information.

The `version` option or the `showversion` script checks the specified systems and discovers the following:

- Storage Foundation and High Availability product versions that are installed on the system
- All the required packages and the optional Symantec packages installed on the system
- Any required or optional packages (if applicable) that are not present
- Installed hot fixes
- Available base releases (major or minor)
- Available maintenance releases
- Available hot fix releases

To check your systems and download maintenance releases and hot fixes

1. Mount the media, or navigate to the installation directory.
2. Start the installer with the `-version` option.

   ```
   # ./installer -version sys1 sys2
   ```

   For each system, the installer lists all of the installed base releases, maintenance releases, and hot fixes, followed by the lists of available downloads.
3. If you have Internet access, follow the prompts to download the available maintenance releases and hot fixes to the local system.
4. If you do not have Internet access, you can download any needed maintenance releases and hot fixes from the Symantec Operations Readiness Tools (SORT) Patch Finder page at:

   https://sort.symantec.com/patch/finder

   You can obtain installer hot fixes automatically or manually.

   See "Obtaining installer hot fixes" on page 45.

Downloading maintenance releases and hot fixes requires the installer to make outbound networking calls. You can also disable external network connection attempts.

See "Disabling external network connection attempts" on page 46.

Obtaining installer hot fixes

Symantec occasionally finds issues with the Symantec Cluster Server installer, and posts public installer hot fixes on the Symantec Operations Readiness Tools (SORT) website's Patch Finder page at:

https://sort.symantec.com/patch/finder

You can access installer hot fixes automatically or manually.

To download installer hot fixes automatically

- Starting with Symantec Cluster Server version 6.1, installer hot fixes are downloaded automatically. No action is needed on your part.

  If you are running Symantec Cluster Server version 6.1 or later, and your system has Internet access, the installer automatically imports any needed installer hot fix, and begins using it.

Automatically downloading installer hot fixes requires the installer to make outbound networking calls. You can also disable external network connection attempts.

See “Disabling external network connection attempts” on page 46.

If your system does not have Internet access, you can download installer hot fixes manually.

To download installer hot fixes manually

1. Go to the Symantec Operations Readiness Tools (SORT) website's Patch Finder page, and save the most current Symantec patch on your local system.
2. Navigate to the directory where you want to unzip the file you downloaded in step 1.
3. Unzip the patch tar file. For example, run the following command:

```
# gunzip cpi-6.1P2-patches.tar.gz
```

4. Untar the file. For example, enter the following:

```
# tar -xvf cpi-6.1P2-patches.tar
patches/
patches/CPI61P2.pl
README
```

5. Navigate to the installation media or to the installation directory.
6. To start using the patch, run the `installer` command with the `-require` option. For example, enter the following:

```
# ./installer -require /target_directory/patches/CPI61P2.pl
```

Disabling external network connection attempts

When you execute the `installer` command, the installer attempts to make an outbound networking call to get information about release updates and installer hot fixes. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network connection attempts by the installer.

To disable external network connection attempts

- Disable inter-process communication (IPC).

  To do this, run the installer with the `-noipc` option.

  For example, to disable IPC for system1 (sys1) and system2 (sys2) enter the following:

  ```
  # ./installer -noipc sys1 sys2
  ```

Chapter 3

Planning to install VCS

This chapter includes the following topics:

- VCS installation methods
- About installation and configuration methods
- Typical VCS cluster setup models

VCS installation methods

Table 3-1 lists the different methods you can choose to install and configure VCS:

Table 3-1 VCS installation methods

Method	Description
Interactive installation using the script-based installer	You can use one of the following script-based installers: ■ Veritas product installer Use to install and configure multiple Symantec products. ■ `installvcs` program Use to install and configure just VCS. The script-based installer asks you a series of questions and installs and configures VCS based on the information you provide.
Interactive installation using the web-based installer	You can use a web-interface to install and configure VCS.

Table 3-1 VCS installation methods *(continued)*

Method	Description
Automated installation using the VCS response files	Use response files to perform unattended installations. You can generate a response file in one of the following ways: ■ Use the automatically generated response file after a successful installation. ■ Use the `-makeresponsefile` option to create a response file.
Manual installation using the Solaris commands and utilities	You can install VCS using the operating system commands like `pkgadd` and then manually configure VCS as described in the section on Manual installation. You can also install VCS using the JumpStart utility.

About the script-based installer

You can use the script-based installer to install Symantec products (version 6.1 and later) from a driver system that runs any supported platform to a target system that runs different supported platforms.

To install your Symantec product, use one of the following methods:

- The general product installer (`installer`). The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
  See "Installing VCS using the installer" on page 91.
- Product-specific installation scripts (`installvcs`). The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installvcs` script is identical to running the general product installer and specifying VCS from the list of products to install. Use the product-specific installation scripts to install or configure individual products you download electronically. You can find these scripts at the root of the product media. These scripts are also installed with the product.

Table 3-2 Product installation scripts

Symantec product name	Script name in the media	Script name after an installation
For all SFHA Solutions products	`installer`	N/A
Symantec ApplicationHA	`installapplicationha`	`installapplicationha` *<version>*
Symantec Cluster Server (VCS)	`installvcs`	`installvcs`*<version>*
Symantec Storage Foundation (SF)	`installsf`	`installsf`*<version>*
Symantec Storage Foundation and High Availability (SFHA)	`installsfha`	`installsfha`*<version>*
Symantec Storage Foundation Cluster File System High Availability (SFCFSHA)	`installsfcfsha`	`installsfcfsha`*<version>*
Symantec Storage Foundation for Oracle RAC (SF Oracle RAC)	`installsfrac`	`installsfrac`*<version>*
Symantec Dynamic Multi-pathing (DMP)	`installdmp`	`installdmp`*<version>*

When you install from the installation media, the script name does not include a product version.

When you configure the product after an installation, the installation scripts include the product version in the script name.

For example, for the 6.1 version:

```
# /opt/VRTS/install/installvcs61 -configure
```

Note: The general product installer (`installer`) script does not include the product version.

At most points during the installation you can type the following characters for different actions:

- Use `b` (back) to return to a previous section of the installation procedure. The back feature of the installation scripts is context-sensitive, so it returns to the beginning of a grouped section of questions.
- Use `Control+c` to stop and exit the program if an installation procedure hangs. After a short delay, the script exits.
- Use `q` to quit the installer.
- Use `?` to display help information.
- Use the Enter button to accept a default response.

See "Installation script options" on page 487.

About the VCS installation program

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

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

- Licensing VCS
- Installing VCS packages 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
- VCS configuration in secure mode
- The wide area Global Cluster Option feature
- Cluster Virtual IP address

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

See "About preparing to install VCS" on page 66.

The uninstallvcs, a companion to installvcs, uninstalls VCS packages.

See "Preparing to uninstall VCS" on page 461.

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 80.
- Upgrade VCS if a previous version of VCS currently runs on a cluster.
  See “Upgrading VCS using the script-based installer” on page 333.
- Start or stop VCS processes
  See “Starting and stopping processes for the Symantec products ” on page 573.
- Enable or disable a cluster to run in secure mode
  See the *Symantec 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” on page 158.
  See “Setting up server-based I/O fencing using installvcs” on page 167.
  See “Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installvcs” on page 181.
- Create a single-node cluster
  See “Creating a single-node cluster using the installer program” on page 546.
- Add a node to an existing cluster
  See “Adding nodes using the VCS installer” on page 431.
- Create a jumpstart finish script to install VCS using the JumpStart utility.
  See “Installing VCS on Solaris 10 using JumpStart” on page 248.
- Perform automated installations using the values that are stored in a configuration file.
  See “Installing VCS using response files” on page 211.
  See “Configuring VCS using response files” on page 216.
  See “Upgrading VCS using response files” on page 359.

Interacting with the installvcs

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

See “Preparing to uninstall VCS” on page 461.

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 does not install the VCS Java Console.

See "Installing the Java Console" on page 404.

About the web-based installer

Use the web-based installer interface to install Symantec products. The web-based installer can perform most of the tasks that the script-based installer performs.

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

When the `webinstaller` script starts the `xprtlwid` 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 operates. While the installation processes operate, 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. Symantec recommends that you keep these files for auditing, debugging, and future use.

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

See "Before using the web-based installer" on page 188.

See "Starting the web-based installer" on page 189.

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, or uninstallvcs
- *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 211.
- To upgrade VCS on multiple systems in a large VCS cluster.
  See "Upgrading VCS using response files" on page 359.
- To uninstall VCS from multiple systems in a large VCS cluster.
  See "Uninstalling VCS using response files" on page 467.

Syntax in the response file

The syntax of the Perl statements that is 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 1 ", "value 2 ", "value 3 "];
```

About installation and configuration methods

You can install and configure VCS using Symantec installation programs or using native operating system methods.

Table 3-3 shows the installation and configuration methods that VCS supports.

Table 3-3 Installation and configuration methods

Method	Description

The script-based installer	Using the script-based installer, you can install Symantec products (version 6.1 and later) from a driver system running a supported platform to target computers running any supported platform.
	
	To install your Symantec product using the installer, choose one of the following:
	
	■ The general product installer:
	`installer`
	The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
	
	■ Product-specific installation scripts:
	`installvcs`
	The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installvcs` script is identical to running the general product installer and specifying VCS from the list of products to install.
	Use the product-specific installation scripts to install or configure individual products you download electronically.
	
	See "About the script-based installer" on page 49.

The web-based installer	Using the web-based installer, you can install Symantec products (version 6.1 and later) from a driver system running a supported platform to target computers running any supported platform The web-based installer provides an interface to manage the installation and configuration from a remote site using a standard web browser. `webinstaller` See "About the web-based installer" on page 53.

Table 3-3 Installation and configuration methods *(continued)*

Method	Description
Deployment Server	Using the Deployment Server, you can store multiple release images in one central location and deploy them to systems of any supported platform. See "About the Deployment Server" on page 301.
Silent installation using response files	Response files automate installation and configuration by using the information that is stored in a specified file instead of prompting you for information. You can use any of the above options to generate a response file. You can then customize the response file for another system. Run the product installation script with the response file option to install silently on one or more systems. See "Installing VCS using response files" on page 211.
Install Bundles	Beginning with version 6.1, you can easily install or upgrade your systems directly to a base, maintenance, or hot fix level in one step using Install Bundles,. The installer installs both releases as if they were combined in the same release image. The various scripts, packages, and patch components are merged, and multiple releases are installed together as if they are one combined release. See "Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes" on page 330.
JumpStart (For Solaris 10 systems)	You can use the product installer of the product-specific installation script to generate a JumpStart script file. Use the generated script to install Symantec packages from your JumpStart server. See "Installing VCS on Solaris 10 using JumpStart" on page 248.
Flash Archive (For Solaris 10 systems)	You can use the product installer to clone the system and install the Symantec products on the master system. See "Using a Flash archive to install VCS and the operating system" on page 251.

Table 3-3 Installation and configuration methods *(continued)*

Method	Description
Manual installation and configuration	Manual installation uses the Solaris commands to install VCS. To retrieve a list of all packages and patches required for all products in the correct installation order, enter: `# installer -allpkgs` Use the Solaris commands to install VCS. Then manually or interactively configure VCS. See "Installing VCS software manually" on page 238.
Automated Installer (For Solaris 11 systems)	You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Symantec packages on multiple client systems in a network. AI performs a hands-free installation (automated installation without manual interactions) of SPARC systems. See "Installing VCS on Solaris 11 using Automated Installer" on page 253.

Typical VCS cluster setup models

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

See the *Symantec 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 "Typical configuration of two-node VCS cluster" on page 58.
- VCS clusters in secure mode
  See "Typical configuration of VCS clusters in secure mode" on page 58.
- VCS clusters centrally managed using Veritas Operations Manager (VOM)
  See "Typical configuration of VOM-managed VCS clusters" on page 59.
- VCS clusters with I/O fencing for data protection
  See "Typical VCS cluster configuration with disk-based I/O fencing" on page 100.
  See "Typical VCS cluster configuration with server-based I/O fencing" on page 101.
- VCS clusters such as global clusters, replicated data clusters, or campus clusters for disaster recovery

See the *Symantec Cluster Server Administrator's Guide* for disaster recovery cluster configuration models.

Typical configuration of two-node VCS cluster

Figure 3-1 illustrates a simple VCS cluster setup with two Solaris SPARC systems.

Figure 3-1 Typical two-node VCS cluster (Solaris SPARC systems)

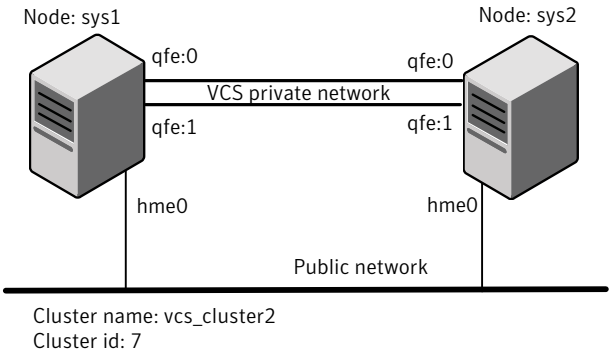

Typical configuration of VCS clusters in secure mode

Enabling secure mode for VCS guarantees that all inter-system communication is encrypted and that security credentials of users are verified.

Figure 3-2 illustrates typical configuration of VCS clusters in secure mode.

Figure 3-2 Typical configuration of VCS clusters in secure mode

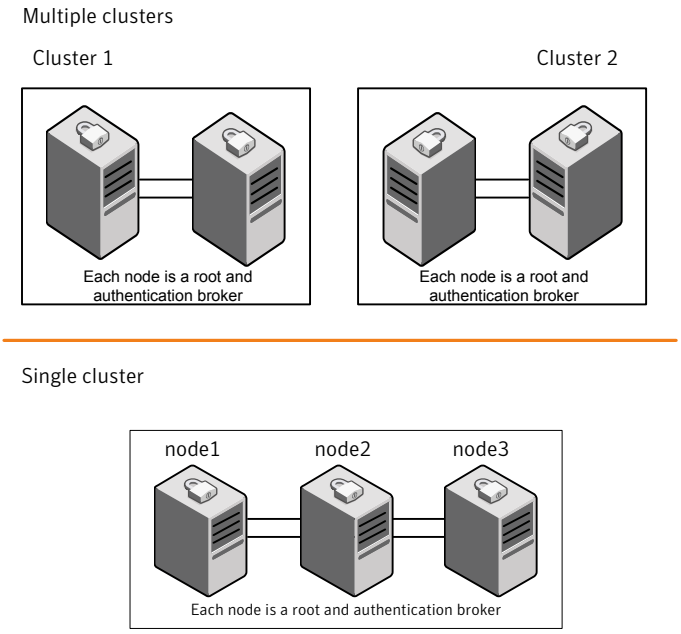

Typical configuration of VOM-managed VCS clusters

Veritas Operations Manager (VOM) provides a centralized management console for Symantec Storage Foundation and High Availability products.

See "About Veritas Operations Manager" on page 28.

Figure 3-3 illustrates a typical setup of VCS clusters that are centrally managed using Veritas Operations Manager.

Figure 3-3 Typical configuration of VOM-managed clusters

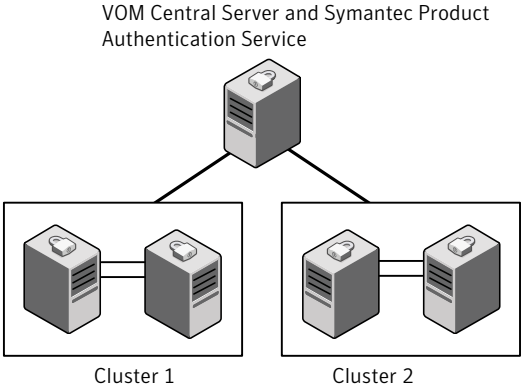

Chapter 4

Licensing VCS

This chapter includes the following topics:

- About Symantec product licensing
- Obtaining VCS license keys
- Installing Symantec product license keys

About Symantec product licensing

You have the option to install Symantec products without a license key. Installation without a license 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.

http://www.symantec.com/products-solutions/licensing/activating-software/detail.jsp?detail_id=licensing_portal

The 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. If you do not comply with the above terms, continuing to use the Symantec 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 previous release of the Symantec software, the installer asks whether you want to upgrade the key to the new version. 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 245.
  See the `vxkeyless(1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.
  See "Installing Symantec product license keys" on page 63.
  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: To change from one product group to another, you may need to perform additional steps.

Obtaining VCS license keys

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

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

http://www.symantec.com/products-solutions/licensing/activating-software/detail.jsp?detail_id=licensing_portal

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 **Get Help** link at this site for contact information and for useful links.

The VRTSvlic package enables product licensing. For information about the commands that you can use after the installing VRTSvlic:

See "Installing Symantec product license keys" on page 63.

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.

Installing Symantec product license keys

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 the currently installed licenses
`vxlictest`	Retrieves the features and their descriptions that are encoded in a license key

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

To install or change a license

1. Run the following commands. In a cluster environment, run the commands on each node in the cluster:

   ```
   # cd /opt/VRTS/bin

   # ./vxlicinst -k license key
   ```

2. Run the following Veritas Volume Manager (VxVM) command to recognize the new license:

   ```
   # vxdctl license init
   ```

   See the `vxdctl(1M)` manual page.

If you have `vxkeyless` licensing, you can view or update the keyless product licensing levels.

See "Setting or changing the product level for keyless licensing" on page 245.

Section 2

Preinstallation tasks

- Chapter 5. Preparing to install VCS

Chapter 5

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
- Making the IPS publisher accessible

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.

Performing preinstallation tasks

Table 5-1 lists the tasks you must perform before proceeding to install VCS.

Table 5-1 Preinstallation tasks

Task	Reference
Obtain license keys if you do not want to use keyless licensing.	See "Obtaining VCS license keys" on page 62.
Set up the private network.	See "Setting up the private network" on page 67.
Enable communication between systems.	See "Setting up inter-system communication" on page 567.

Table 5-1 Preinstallation tasks *(continued)*

Task	Reference
Set up ssh on cluster systems.	See "Setting up ssh on cluster systems" on page 567.
Set up shared storage for I/O fencing (optional)	See "Setting up shared storage" on page 71.
Creating root user	See "Creating a root user" on page 75.
Set the PATH and the MANPATH variables.	See "Setting the PATH variable" on page 76. See "Setting the MANPATH variable" on page 76.
Disable the abort sequence on SPARC systems.	See "Disabling the abort sequence on SPARC systems" on page 76.
Review basic instructions to optimize LLT media speeds.	See "Optimizing LLT media speed settings on private NICs" on page 78.
Review guidelines to help you set the LLT interconnects.	See "Guidelines for setting the media speed of the LLT interconnects" on page 78.
Install the compatibility/ucb additional packages from Oracle Solaris repository.	For instructions, see the Oracle documentation.
Prepare zone environments	See "Preparing zone environments" on page 79.
Mount the product disc	See "Mounting the product disc" on page 80.
Verify the systems before installation	See "Performing automated preinstallation check" on page 80.

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. However, Oracle Solaris systems assign the same MAC address to all interfaces by default. Thus, connecting two or more interfaces to a network switch can cause problems.

For example, consider the following case where:

- The IP address is configured on one interface and LLT on another
- Both interfaces are connected to a switch (assume separate VLANs)

The duplicate MAC address on the two switch ports can cause the switch to incorrectly redirect IP traffic to the LLT interface and vice versa. To avoid this issue, configure the system to assign unique MAC addresses by setting the `eeprom(1M)` parameter local-mac-address to true.

The following products make extensive use of the private cluster interconnects for distributed locking:

- Symantec Storage Foundation Cluster File System (SFCFS)
- Symantec Storage Foundation for Oracle RAC (SF Oracle RAC)

Symantec recommends network switches for the SFCFS and the SF Oracle RAC clusters due to their performance characteristics.

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

Figure 5-1 shows two private networks for use with VCS.

Figure 5-1 Private network setups: two-node and four-node clusters

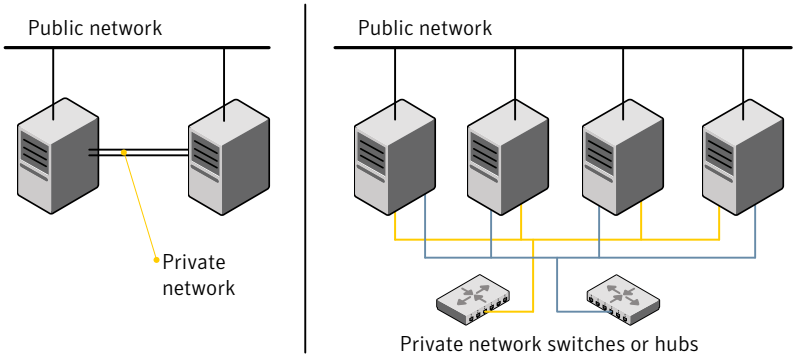

You need to configure at least two independent networks between the cluster nodes with a network switch for each network. You can also interconnect multiple layer 2 switches for advanced failure protection. Such connections for LLT are called cross-links.

Figure 5-2 shows a private network configuration with crossed links between the network switches.

Figure 5-2 Private network setup with crossed links

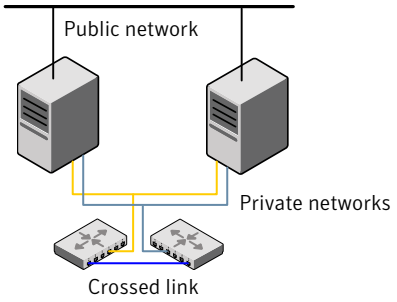

Symantec recommends one of the following two configurations:

- Use at least two private interconnect links and one public link. The public link can be a low priority link for LLT. The private interconnect link is used to share cluster status across all the systems, which is important for membership arbitration and high availability. The public low priority link is used only for heartbeat communication between the systems.
- If your hardware environment allows use of only two links, use one private interconnect link and one public low priority link. If you decide to set up only two links (one private and one low priority link), then the cluster must be configured to use I/O fencing, either disk-based or server-based fencing configuration. With only two links, if one system goes down, I/O fencing ensures that other system can take over the service groups and shared file systems from the failed node.

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 Ethernet controllers 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 Configure the Ethernet devices that are used for the private network such that the autonegotiation protocol is not used. You can achieve a more stable configuration with crossover cables if the autonegotiation protocol is not used.

To achieve this stable configuration, do one of the following:

- Edit the /etc/system file to disable autonegotiation on all Ethernet devices system-wide.
- Create a qfe.conf or bge.conf file in the /kernel/drv directory to disable autonegotiation for the individual devices that are used for private network.

Refer to the Oracle Ethernet driver product documentation for information on these methods.

5 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 installer configures the private network in the cluster during configuration.

You can also manually configure LLT.

See "Configuring LLT manually" on page 259.

About using ssh or rsh with the installer

The installer uses passwordless Secure Shell (ssh) or Remote Shell (rsh) communications among systems. The installer uses the ssh daemon or rsh daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. Or, you can run the `installer -comsetup` command to set up ssh or rsh explicitly. You then provide the installer with the superuser passwords for the systems where you plan to install. The ssh communication or rsh communication among the systems is removed when the installation process completes, unless the installation abruptly terminates. If

installation terminated abruptly, use the installation script's `-comcleanup` option to remove the ssh configuration or rsh configuration from the systems.

See "Installation script options" on page 487.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer can configure ssh or rsh on the target systems. In the following scenarios, you need to set up ssh or rsh manually, or use the `installer -comsetup` option to set up an ssh or rsh configuration from the systems.

- When you perform installer sessions using a response file.

See "Setting up inter-system communication" on page 567.

Setting up shared storage

The following sections describe how to set up the SCSI and the Fibre Channel devices that the cluster systems share.

For I/O fencing, the data disks must support SCSI-3 persistent reservations. You need to configure a coordinator disk group that supports SCSI-3 PR and verify that it works.

See "About planning to configure I/O fencing" on page 96.

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

Setting up shared storage: SCSI disks

When SCSI devices are used for shared storage, the SCSI address or SCSI initiator ID of each node must be unique. Since each node typically has the default SCSI address of "7," the addresses of one or more nodes must be changed to avoid a conflict. In the following example, two nodes share SCSI devices. The SCSI address of one node is changed to "5" by using `nvedit` commands to edit the `nvramrc` script.

If you have more than two systems that share the SCSI bus, do the following:

- Use the same procedure to set up shared storage.
- Make sure to meet the following requirements:
  - The storage devices have power before any of the systems
  - Only one node runs at one time until each node's address is set to a unique value

To set up shared storage

1. Install the required SCSI host adapters on each node that connects to the storage, and make cable connections to the storage.

   Refer to the documentation that is shipped with the host adapters, the storage, and the systems.

2. With both nodes powered off, power on the storage devices.

3. Power on one system, but do not allow it to boot. If necessary, halt the system so that you can use the ok prompt.

   Note that only one system must run at a time to avoid address conflicts.

4. Find the paths to the host adapters:

   ```
   {0} ok show-disks
   ...b) /sbus@6,0/QLGC,isp@2,10000/sd
   ```

   The example output shows the path to one host adapter. You must include the path information without the "/sd" directory, in the `nvramrc` script. The path information varies from system to system.

5. Edit the `nvramrc` script on to change the scsi-initiator-id to 5. (The *Solaris OpenBoot 3.x Command Reference Manual* contains a full list of `nvedit` commands and keystrokes.) For example:

   ```
   {0} ok nvedit
   ```

   As you edit the script, note the following points:

   - Each line is numbered, 0:, 1:, 2:, and so on, as you enter the `nvedit` commands.
   - On the line where the scsi-initiator-id is set, insert exactly one space after the first quotation mark and before scsi-initiator-id.

   In this example, edit the nvramrc script as follows:

   ```
   0: probe-all
   1: cd /sbus@6,0/QLGC,isp@2,10000
   2: 5 " scsi-initiator-id" integer-property
   3: device-end
   4: install-console
   5: banner
   6: <CTRL-C>
   ```

6 Store the changes you make to the `nvramrc` script. The changes you make are temporary until you store them.

```
{0} ok nvstore
```

If you are not sure of the changes you made, you can re-edit the script without risk before you store it. You can display the contents of the `nvramrc` script by entering:

```
{0} ok printenv nvramrc
```

You can re-edit the file to make corrections:

```
{0} ok nvedit
```

Or, discard the changes if necessary by entering:

```
{0} ok nvquit
```

7 Instruct the OpenBoot PROM Monitor to use the `nvramrc` script on the node.

```
{0} ok setenv use-nvramrc? true
```

8 Reboot the node. If necessary, halt the system so that you can use the ok prompt.

9 Verify that the scsi-initiator-id has changed. Go to the ok prompt. Use the output of the show-disks command to find the paths for the host adapters. Then, display the properties for the paths. For example:

```
{0} ok show-disks
...b) /sbus@6,0/QLGC,isp@2,10000/sd
{0} ok cd /sbus@6,0/QLGC,isp@2,10000
{0} ok .properties
scsi-initiator-id     00000005
```

Permit the system to continue booting.

10 Boot the second node. If necessary, halt the system to use the ok prompt. Verify that the scsi-initiator-id is 7. Use the output of the `show-disks` command to find the paths for the host adapters. Then, display the properties for that paths. For example:

```
{0} ok show-disks
...b) /sbus@6,0/QLGC,isp@2,10000/sd
{0} ok cd /sbus@6,0/QLGC,isp@2,10000
{0} ok .properties
scsi-initiator-id      00000007
```

Permit the system to continue booting.

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up shared storage

1 Install the required FC-AL controllers.

2 Connect the FC-AL controllers and the shared storage devices to the same hub or switch.

All systems must see all the shared devices that are required to run the critical application. If you want to implement zoning for a fibre switch, make sure that no zoning prevents all systems from seeing all these shared devices.

3 Boot each system with the reconfigure devices option:

```
ok boot -r
```

4 After all systems have booted, use the `format(1m)` command to verify that each system can see all shared devices.

If Volume Manager is used, the same number of external disk devices must appear, but device names (c#t#d#s#) may differ.

If Volume Manager is not used, then you must meet the following requirements:

- The same number of external disk devices must appear.
- The device names must be identical for all devices on all systems.

Creating a root user

On Oracle Solaris 11, you need to change the root role into a user as you cannot directly log in as root user.

To change root role into a user

1 Log in as local user and assume the root role.

```
% su  - root
```

2 Remove the root role from local users who have been assigned the role.

```
# roles admin

root

# usermod -R " " admin
```

3 Change the root role into a user.

```
# rolemod -K type=normal root
```

4 Verify the change.

- `# getent user_attr root`

  ```
  root::::auths=solaris.*;profiles=All;audit_flags=lo\
  :no;lock_after_retries=no;min_label=admin_low;clearance=admin_high
  ```

  If the `type` keyword is not present in the output or is equal to normal, the account is not a role.

- `# userattr type root`

  If the output is empty or lists normal, the account is not a role.

Note: For more information, see the Oracle documentation on Oracle Solaris 11 operating system.

Note: After installation, you may want to change root user into root role to allow local users to assume the root role.

See "Changing root user into root role" on page 403.

Setting the PATH variable

To set the PATH variable

- Do one of the following:
  - For the Bourne Shell (sh), Bourne-again Shell (bash), or Korn shell (ksh), type:

    ```
    # PATH=/opt/VRTS/bin:$PATH; export $PATH
    ```

  - For the C Shell (csh) or enhanced C Shell (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), Bourne-again Shell (bash), or Korn shell (ksh), type:

    ```
    # MANPATH=/opt/VRTS/man:$MANPATH; export MANPATH
    ```

  - For the C Shell (csh) or enhanced C Shell (tcsh), type:

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

Disabling the abort sequence on SPARC systems

Most UNIX operating systems provide a method to perform a "break" or "console abort." The inherent problem when you abort a hung system is that it ceases to heartbeat in the cluster. When other cluster members believe that the aborted node is a failed node, these cluster members may begin corrective action.

Keep the following points in mind:

- The only action that you must perform following a system abort is to reset the system to achieve the following:
  - Preserve data integrity
  - Prevent the cluster from taking additional corrective actions
- Do not resume the processor as cluster membership may have changed and failover actions may already be in progress.
- To remove this potential problem on SPARC systems, you should alias the `go` function in the OpenBoot eeprom to display a message.

To alias the go function to display a message

1. At the ok prompt, enter:

   ```
   nvedit
   ```

2. Press Ctrl+L to display the current contents of the nvramrc buffer.
3. Press Ctrl+N until the editor displays the last line of the buffer.
4. Add the following lines exactly as shown. Press Enter after adding each line.

   ```
   ." Aliasing the OpenBoot 'go' command! "
   : go ." It is inadvisable to use the 'go' command in a clustered
   environment. " cr
   ." Please use the 'power-off' or 'reset-all' commands instead. "
   cr
   ." Thank you, from your friendly neighborhood sysadmin. " ;
   ```

5. Press Ctrl+C to exit the nvramrc editor.
6. To verify that no errors exist, type the `nvrun` command. You should see only the following text:

   ```
   Aliasing the OpenBoot 'go' command!
   ```

7. Type the `nvstore` command to commit your changes to the non-volatile RAM (NVRAM) for use in subsequent reboots.
8. After you perform these commands, at reboot you see this output:

   ```
   Aliasing the OpenBoot 'go' command! go isn't unique.
   ```

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.

If you use different media speed for the private NICs, Symantec recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance.

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 use different media speed for the private NICs, Symantec recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance.
- 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), Symantec recommends that you set the media speed to the highest value common to both cards, typically 1000_Full_Duplex.

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

VCS considerations for Blade server environments

Typically, a server in the Blade environment has only two NICs. The following considerations need to be observed while configuring VCS in Blade environment in private networks:

- If your heartbeat links do not use TCP/IP and are not routable, you must use separate and dedicated physical networks. This will guard against inadvertent split brains due to inappropriate routing configurations.
- Out of the two heartbeat links, one must be dedicated and the other can be a low-priority heartbeat shared on the public IP NIC. It is assumed that the two nodes in the cluster have public IPs on the same subnet and wire.

- The size each packet of traffic on the public NIC must be 64 bytes/second and must not interfere with the public traffic.

Preparing zone environments

You need to keep the following items in mind when you install or upgrade VCS in a zone environment on an Oracle Solaris 10 operating system.

- When you install or upgrade VCS using the `installer` program, all zones are upgraded (both global and non-global) unless they are detached and unmounted.
- Make sure that all non-global zones are booted and in the running state before you install or upgrade the VCS packages in the global zone. If the non-global zones are not mounted and running at the time of upgrade, you must attach the zone with **-U** option to install or upgrade the VCS packages inside the non-globle zone.
- If you install VCS on Solaris 10 systems that run non-global zones, you need to make sure that non-global zones do not inherit the `/opt` directory. Run the following command to make sure that the `/op`t directory is not in the inherit-pkg-dir clause:

  ```
  # zonecfg -z zone_name info
  zonepath: /export/home/zone1
  autoboot: false
  pool: yourpool
  inherit-pkg-dir:
  dir: /lib
  inherit-pkg-dir:
  dir: /platform
  inherit-pkg-dir:
  dir: /sbin
  inherit-pkg-dir:
  dir: /usr
  ```

  If the /opt directory appears in the output, remove the /opt directory from the zone's configuration and reinstall the zone.

After installing packages in the global zone, you need to install the required packages in the non-global zone for Oracle Solaris 11. On Oracle Solaris 11.1, if the non-global zone has an older version of VCS packages already installed then during the upgrade of the VCS packages in global zone, packages inside non-global zone are automatically upgraded provided zone is running.

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 does not need to be part of the cluster. The systems must be in the same subnet.

2. Insert the product disc into a DVD drive that is connected to your system.

3. If Solaris volume management software is running on your system, the software disc automatically mounts as /cdrom/cdrom0.

4. If Solaris volume management software is not available to mount the DVD, you must mount it manually. After you insert the software disc, enter:

   ```
   # mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom
   ```

   Where `c0t6d0s2` is the default address for the disc drive.

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 also run the `installer -precheck` command.

See "About Symantec Operations Readiness Tools" on page 30.

To check the systems

1. Navigate to the folder that contains the installvcs.

   ```
   # cd /cdrom/cdrom0/cluster_server
   ```

2. Start the preinstallation check:

   ```
   # ./installvcs -precheck sys1 sys2
   ```

   The program proceeds in a noninteractive mode to examine the systems for licenses, packages, 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.

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 `hacf -cftocmd` and then the `hacf -cmdtocf` 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 *Symantec 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 -cmdtocf config
```

Getting your VCS installation and configuration information ready

The VCS installer prompts you for some information during the installation and configuration process. Review the following information and make sure you have made the necessary decisions and you have the required information ready before you perform the installation and configuration.

Table 5-2 lists the information you need to install the VCS packages.

Table 5-2 Information to install the VCS packages

Information	Description and sample value	Your value
System names	The system names where you plan to install VCS Example: **sys1**, **sys2**	
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 Symantec product licensing" on page 61. 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 62.	
Decide which packages to install	■ Minimum packages—provides basic VCS functionality. ■ Recommended packages—provides full functionality of VCS without advanced features. ■ All packages—provides advanced feature functionality of VCS. The default option is to install the recommended packages. See "Viewing the list of VCS packages" on page 239.	

Table 5-3 lists the information you need to configure VCS cluster name and ID.

Table 5-3 Information you need to configure VCS cluster name and ID

Information	Description and sample value	Your value
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: **my_cluster**	
A unique ID number for the cluster	A number in the range of 0-65535. If multiple distinct and separate clusters share the same network, then each cluster must have a unique cluster ID. Example: **12133**	

Table 5-4 lists the information you need to configure VCS private heartbeat links.

Table 5-4 Information you need to configure VCS private heartbeat links

Information	Description and sample value	Your value
Decide how you want to configure LLT	You can configure LLT over Ethernet or LLT over UDP. The LLT over Ethernet is the typical configuration for LLT. If the cluster nodes are across routers, then use the LLT over UDP configuration after ensuring that you meet all the prerequisites. See "Using the UDP layer for LLT" on page 549.	
Decide which configuration mode you want to choose	Installer provides you with four options: **1** Configure heartbeat links using LLT over Ethernet **2** Configure heartbeat links using LLT over UDP **3** Automatically detect configuration for LLT over Ethernet You must manually enter details for options 1, 2 whereas the installer detects the details for option 3.	
For option 1: LLT over Ethernet	■ The device names of the NICs that the private networks use among systems A network interface card or an aggregated interface. Do not use the network interface card that is used for the public network, which is typically net0 for SPARC. For example on a SPARC system: `net1, net2` ■ Choose whether to use the same NICs on all systems. If you want to use different NICs, enter the details for each system.	

Table 5-4 Information you need to configure VCS private heartbeat links *(continued)*

Information	Description and sample value	Your value
For option 2: LLT over UDP	For each system, you must have the following details: ■ The device names of the NICs that the private networks use among systems ■ IP address for each NIC ■ UDP port details for each NIC	

Table 5-5 lists the information you need to configure virtual IP address of the cluster (optional).

Table 5-5 Information you need to configure virtual IP address

Information	Description and sample value	Your value
The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. A network interface card or an aggregated interface. Example: net0	
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	

Table 5-6 lists the information you need to add VCS users.

Table 5-6 Information you need to add VCS users

Information	Description and sample value	Your value
User names	VCS usernames are restricted to 1024 characters. Example: **smith**	

Table 5-6 Information you need to add VCS users *(continued)*

Information	Description and sample value	Your value
User passwords	VCS passwords are restricted to 255 characters. Enter the password at the prompt. **Note:** VCS leverages native authentication in secure mode. Therefore, user passwords are not needed in secure mode.	
To decide user privileges	Users have three levels of privileges: Administrator, Operator, or Guest. Example: Administrator	

Table 5-7 lists the information you need to configure SMTP email notification (optional).

Table 5-7 Information you need to configure SMTP email notification (optional)

Information	Description and sample value	Your value
The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. A network interface card or an aggregated interface. Examples: **net0**	
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**	

Table 5-7 Information you need to configure SMTP email notification (optional) *(continued)*

Information	Description and sample value	Your value
To decide the minimum severity of events for SMTP email notification	Events have four levels of severity, and the severity levels are cumulative: ■ Information VCS sends notifications for important events that exhibit normal behavior. ■ Warning VCS sends notifications for events that exhibit any deviation from normal behavior. Notifications include both Warning and Information type of events. ■ Error VCS sends notifications for faulty behavior. Notifications include both Error, Warning, and Information type of events. ■ Critical VCS sends notifications for a critical error that can lead to data loss or corruption. Notifications include both Severe Error, Error, Warning, and Information type of events. Example: **Error**	

Table 5-8 lists the information you need to configure SNMP trap notification (optional).

Table 5-8 Information you need to configure SNMP trap notification (optional)

Information	Description and sample value	Your value
The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. A network interface card or an aggregated interface. Examples: **net0**	
The port number for the SNMP trap daemon	The default port number is 162.	
The system name for each SNMP console	Example: **sys5**	

Table 5-8 Information you need to configure SNMP trap notification (optional) *(continued)*

Information	Description and sample value	Your value
To decide the minimum severity of events for SNMP trap notification	Events have four levels of severity, and the severity levels are cumulative: ■ Information VCS sends notifications for important events that exhibit normal behavior. ■ Warning VCS sends notifications for events that exhibit any deviation from normal behavior. Notifications include both Warning and Information type of events. ■ Error VCS sends notifications for faulty behavior. Notifications include both Error, Warning, and Information type of events. ■ Critical VCS sends notifications for a critical error that can lead to data loss or corruption. Notifications include both Severe Error, Error, Warning, and Information type of events. Example: **Error**	

Table 5-9 lists the information you need to configure global clusters (optional).

Table 5-9 Information you need to configure global clusters (optional)

Information	Description and sample value	Your value
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. A network interface card or an aggregated interface. For example for SPARC systems: `net0`	
The virtual IP address of the NIC	You can enter either an IPv4 or an IPv6 address. 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 IPv4 address: **192.168.1.16** Example IPv6 address: **2001:454e:205a:110:203:baff:feee:10**	

Table 5-9 Information you need to configure global clusters (optional) *(continued)*

Information	Description and sample value	Your value
The netmask for the virtual IPv4 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**	
The prefix for the virtual IPv6 address	The prefix length for the virtual IPv6 address. Example: **64**	

Review the information you need to configure I/O fencing.

See "About planning to configure I/O fencing" on page 96.

Making the IPS publisher accessible

The installation of VCS 6.1 fails on Solaris 11 if the Image Packaging System (IPS) publisher is inaccessible. The following error message is displayed:

CPI ERROR V-9-20-1273 Unable to contact configured publishers on <node_name>.

Solaris 11 introduces the new Image Packaging System (IPS) and sets a default publisher (solaris) during Solaris installation. When additional packages are being installed, the set publisher must be accessible for the installation to succeed. If the publisher is inaccessible, as in the case of a private network, then package installation will fail. The following commands can be used to display the set publishers:

```
# pkg publisher
```

Example:

```
root@sol11-03:~# pkg publisher
PUBLISHER       TYPE     STATUS   URI
solaris         origin   online   http://pkg.oracle.com/solaris/release/
root@sol11-03:~# pkg publisher solaris
Publisher: solaris
                Alias:
           Origin URI: http://pkg.oracle.com/solaris/release/
              SSL Key: None
            SSL Cert: None
          Client UUID: 00000000-3f24-fe2e-0000-000068120608
```

```
 Catalog Updated: October 09:53:00 PM
         Enabled: Yes
Signature Policy: verify
```

To make the IPS publisher accessible

1. Enter the following to disable the publisher (in this case, solaris):

   ```
   # pkg set-publisher --disable solaris
   ```

2. Repeat the installation of VCS 6.1.
3. Re-enable the original publisher. If the publisher is still inaccessible (private network), then the `no-refresh` option can be used to re-enable it.

   ```
   # pkg set-publisher --enable solaris
   ```

   or

   ```
   # pkg set-publisher --enable --no-refresh solaris
   ```

Note: Unsetting the publisher will have a similar effect, except that the publisher can only be re-set if it is accessible. See pkg(1) for further information on the pkg utility.

Section 3

Installation using the script-based installer

- Chapter 6. Installing VCS
- Chapter 7. Preparing to configure VCS clusters for data integrity
- Chapter 8. Configuring VCS
- Chapter 9. Configuring VCS clusters for data integrity

Chapter 6

Installing VCS

This chapter includes the following topics:

- Installing VCS using the installer
- Installing language packages using the installer

Installing VCS using the installer

Perform the following steps to install VCS.

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 80.

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.

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: Symantec Cluster Server.

installvcs program

Perform the following steps to start the product installer:

1. Navigate to the folder that contains the installvcs.

   ```
   # cd /cdrom/cdrom0/cluster_server
   ```

2. Start the installvcs.

   ```
   # ./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 cluster_server/EULA/<lang>/EULA_VCS_Ux_6.1.pdf
   file present on media? [y,n,q,?] y
   ```

4. Choose the VCS packages that you want to install.

   See "Symantec Cluster Server installation packages" on page 481.

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

   1 Installs only the minimal required VCS packages that provides basic functionality of the product.

   2 Installs the recommended VCS packages that provides complete functionality of the product. This option does not install the optional VCS packages.

   Note that this option is the default.

   3 Installs all the VCS packages.

   You must choose this option to configure any optional VCS feature.

   4 Displays the VCS packages for each option.

   ```
   Select the packages 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 system names separated by spaces:
   [q,?] (sys1) sys1 sys2
   ```

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

See “Creating a single-node cluster using the installer program” on page 546.

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 password-less communication is not present, you can provide the location of `ssh_key` file which isused in every communication.
  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.
- Makes sure the systems install from the global zone.
- 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 6.1.
  See “About upgrading to VCS 6.1” on page 323.
- 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 packages already exists on a system.
  If the current version of any package exists, the installer removes the package from the installation list for the system. If a previous version of any package exists, the installer replaces the package with the current version.

6 Review the list of packages and patches that the installer would install on each node.

The installer installs the VCS packages and patches on the systems sys1 and sys2.

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.

The installer registers the license and completes the installation process.

8 To install the Global Cluster Option, enter y at the prompt.

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

```
Would you like to configure VCS on sys1 sys2 [y,n,q] (n) n
```

See "Overview of tasks to configure VCS using the script-based installer" on page 136.

10 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.

11 The installer checks for online updates and provides an installation summary.

12 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 packages 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 211.

Installing language packages using the installer

Before you install the language packages, do the following:

- Make sure `install_lp` command uses the *ssh* or *rsh* commands as root on all systems in the cluster.
- Make sure that permissions are granted for the system on which `install_lp` is run.

To install the language packages

1 Insert the language disc into the drive.

The Solaris volume-management software automatically mounts the disc as /cdrom/cdrom0.

2 Change to the /cdrom/cdrom0 directory.

```
# cd /cdrom/cdrom0
```

3 Install the language packages:

```
# ./install_lp
```

Chapter 7

Preparing to configure VCS clusters for data integrity

This chapter includes the following topics:

- About planning to configure I/O fencing
- Setting up the CP server

About planning to configure I/O fencing

After you configure VCS with the installer, you must configure I/O fencing in the cluster for data integrity. Application clusters on release version 6.1 (HTTPS-based communication) only support CP servers on release version 6.1 and later.

You can configure 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.

The coordination points in server-based fencing can include only CP servers or a mix of CP servers and coordinator disks.

Symantec also supports server-based fencing with a single coordination point which is a single highly available CP server that is hosted on an SFHA cluster.

Warning: For server-based fencing configurations that use a single coordination point (CP server), the coordination point becomes a single point of failure. In such configurations, the arbitration facility is not available during a failover of the CP server in the SFHA cluster. So, if a network partition occurs on any application cluster during the CP server failover, the application cluster is brought down. Symantec recommends the use of single CP server-based fencing only in test environments.

If you have installed VCS in a virtual environment that is not SCSI-3 PR compliant, you can configure non-SCSI-3 server-based fencing.

See Figure 7-2 on page 99.

Figure 7-1 illustrates a high-level flowchart to configure I/O fencing for the VCS cluster.

Figure 7-1 Workflow to configure I/O fencing

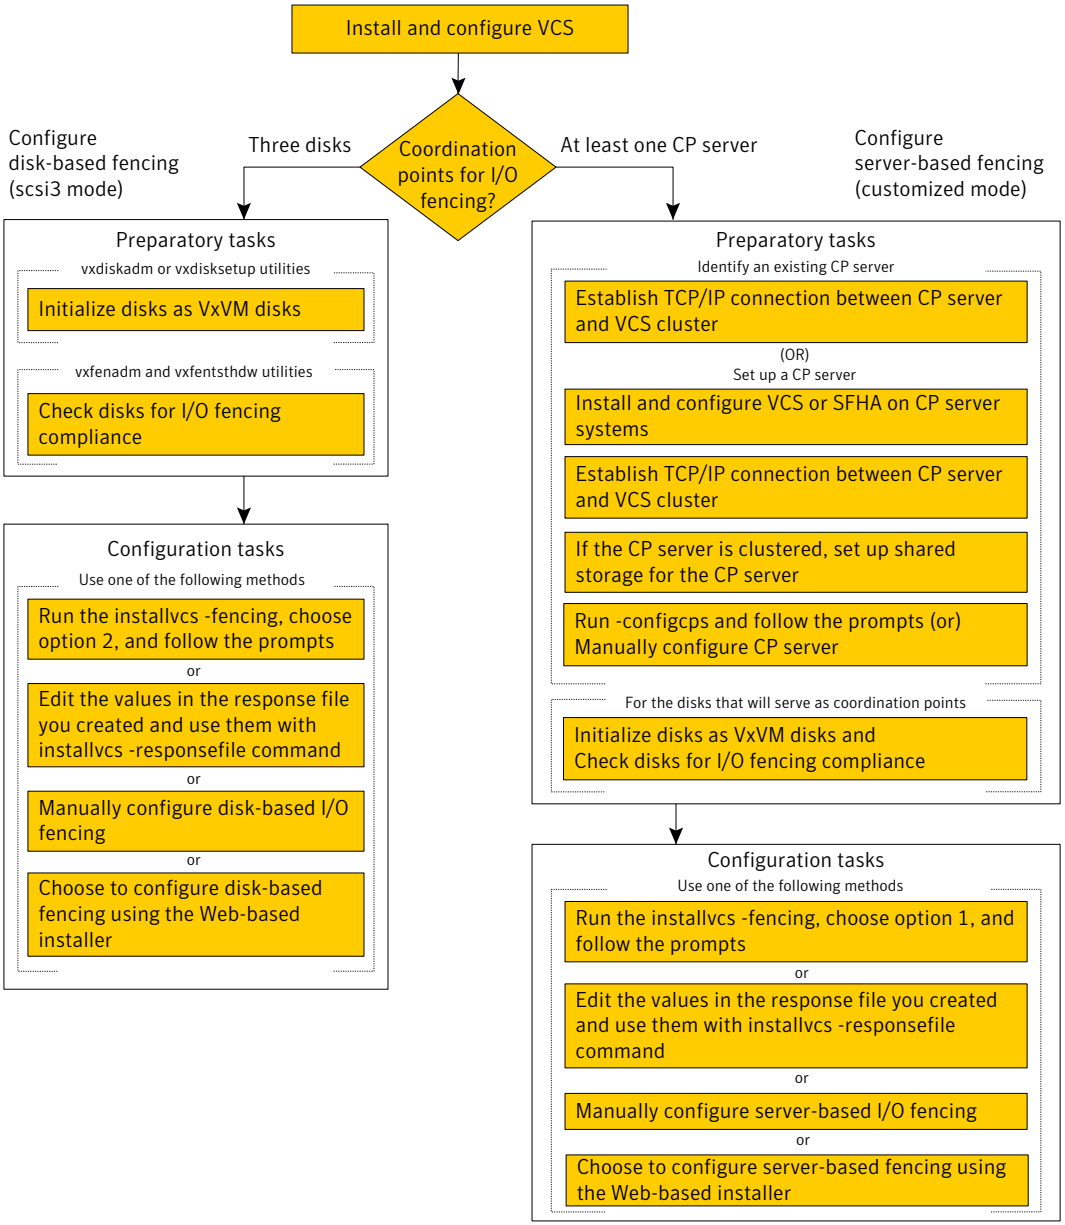

Figure 7-2 illustrates a high-level flowchart to configure non-SCSI-3 server-based I/O fencing for the VCS cluster in virtual environments that do not support SCSI-3 PR.

Figure 7-2 Workflow to configure non-SCSI-3 server-based I/O fencing

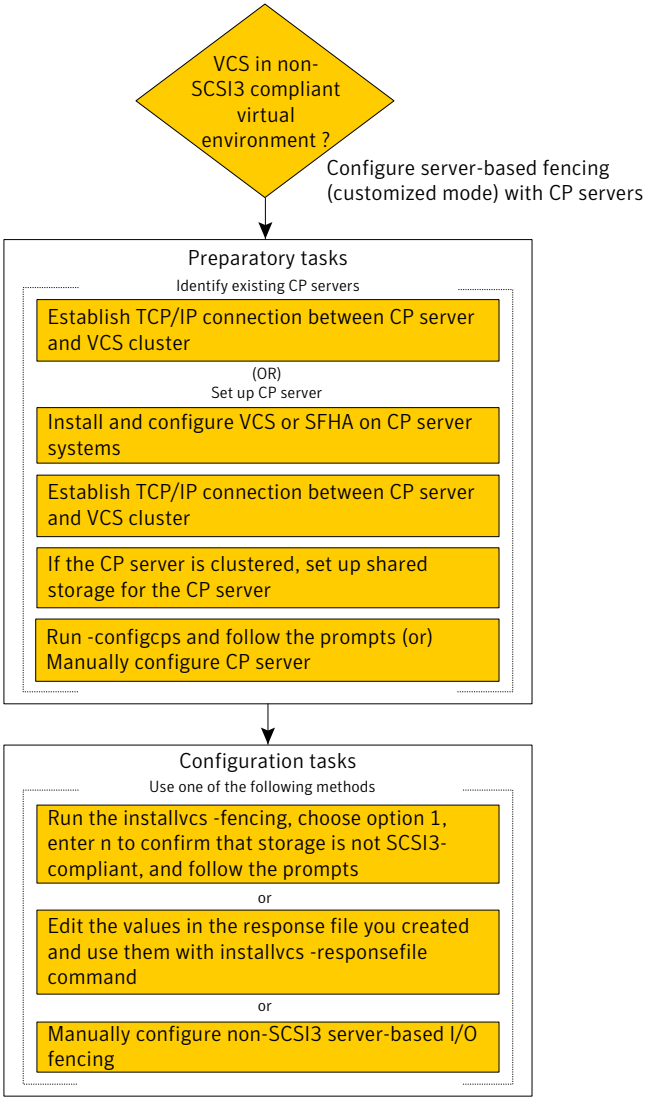

After you perform the preparatory tasks, you can use any of the following methods to configure I/O fencing:

Using the installvcs	See "Setting up disk-based I/O fencing using installvcs" on page 158. See "Setting up server-based I/O fencing using installvcs" on page 167. See "Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installvcs" on page 181.
Using the web-based installer	See "Configuring VCS for data integrity using the web-based installer" on page 198.
Using response files	See "Response file variables to configure disk-based I/O fencing" on page 228. See "Response file variables to configure server-based I/O fencing" on page 232. See "Response file variables to configure non-SCSI-3 server-based I/O fencing" on page 234. See "Configuring I/O fencing using response files" on page 227.
Manually editing configuration files	See "Setting up disk-based I/O fencing manually" on page 272. See "Setting up server-based I/O fencing manually" on page 277. See "Setting up non-SCSI-3 fencing in virtual environments manually" on page 292.

You can also migrate from one I/O fencing configuration to another.

See the *Symantec Storage foundation High Availability Administrator's Guide* for more details.

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

Figure 7-3 displays a typical VCS configuration with two nodes and shared storage. The configuration uses three coordinator disks for I/O fencing.

Figure 7-3 Typical VCS cluster configuration with disk-based I/O fencing

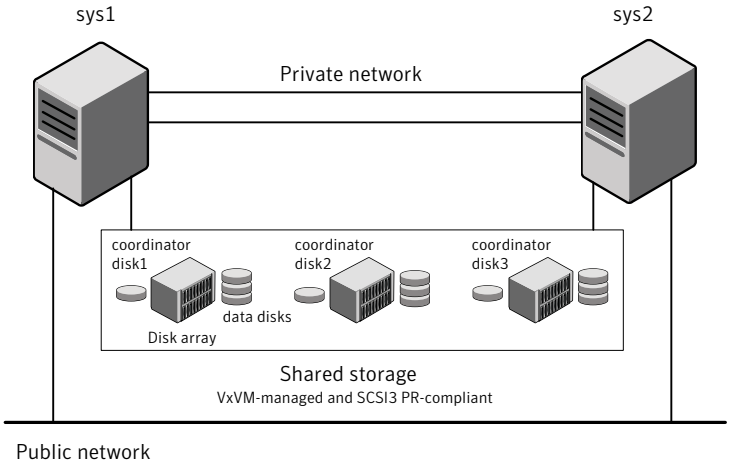

Typical VCS cluster configuration with server-based I/O fencing

Figure 7-4 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 7-4 CP server, VCS cluster, and coordinator disks

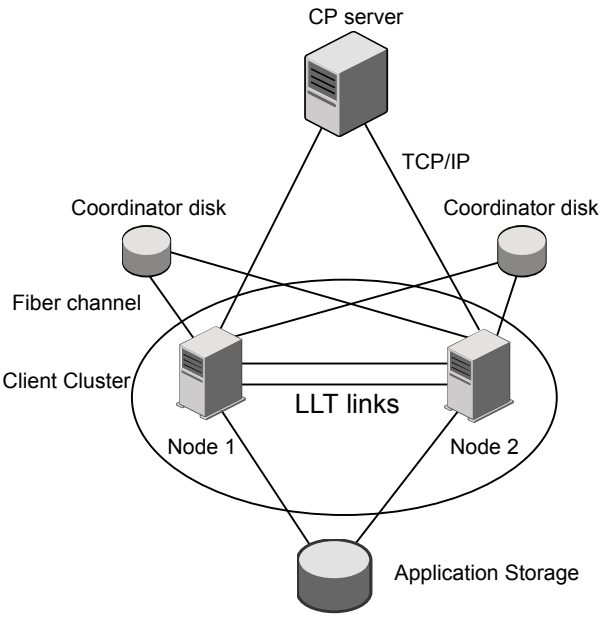

Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points
  See Figure 7-5 on page 103.
- Multiple application clusters use a single CP server and single or multiple pairs of coordinator disks (two) as their coordination points
  See Figure 7-6 on page 104.
- Multiple application clusters use a single CP server as their coordination point
  This single coordination point fencing configuration must use a highly available CP server that is configured on an SFHA cluster as its coordination point.
  See Figure 7-7 on page 104.

Warning: In a single CP server fencing configuration, arbitration facility is not available during a failover of the CP server in the SFHA cluster. So, if a network partition occurs on any application cluster during the CP server failover, the application cluster is brought down.

Although the recommended CP server configurations use three coordination points, you can use more than three coordination points for I/O fencing. Ensure that the total number of coordination points you use is an odd number. In a configuration where multiple application clusters share a common set of CP server coordination points, the application cluster as well as the CP server use a Universally Unique Identifier (UUID) to uniquely identify an application cluster.

Figure 7-5 displays a configuration using three CP servers that are connected to multiple application clusters.

Figure 7-5 Three CP servers connecting to multiple application clusters

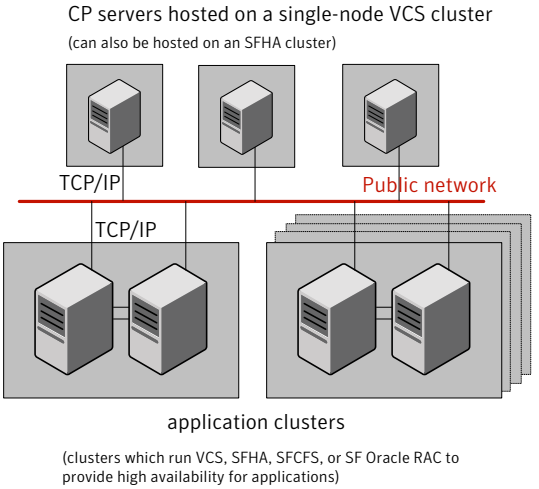

Figure 7-6 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.

Figure 7-6 Single CP server with two coordinator disks for each application cluster

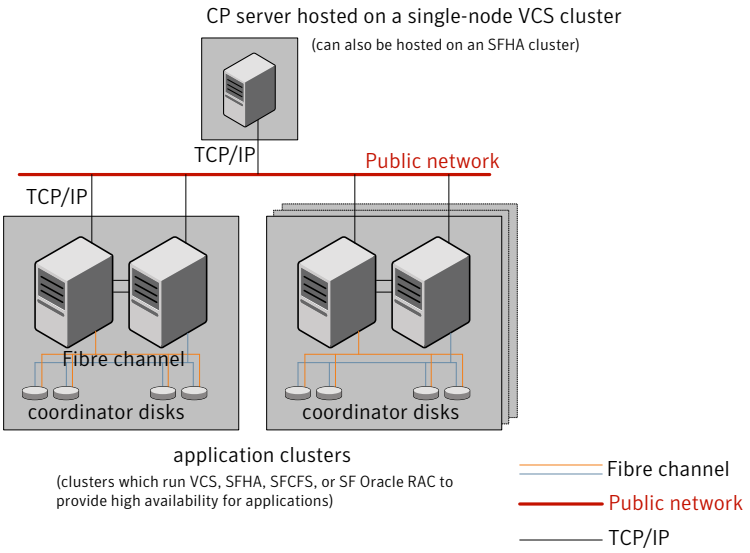

Figure 7-7 displays a configuration using a single CP server that is connected to multiple application clusters.

Figure 7-7 Single CP server connecting to multiple application clusters

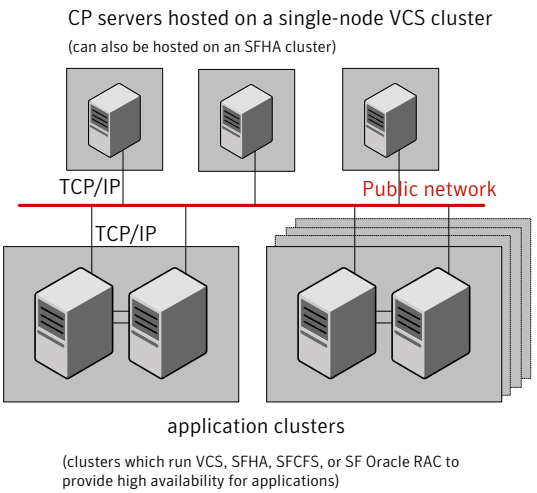

See "Configuration diagrams for setting up server-based I/O fencing" on page 577.

Setting up the CP server

Table 7-1 lists the tasks to set up the CP server for server-based I/O fencing.

Table 7-1 Tasks to set up CP server for server-based I/O fencing

Task	Reference
Plan your CP server setup	See "Planning your CP server setup" on page 105.
Install the CP server	See "Installing the CP server using the installer" on page 107.
Configure the CP server cluster in secure mode	See "Configuring the CP server cluster in secure mode" on page 107.
Set up shared storage for the CP server database	See "Setting up shared storage for the CP server database" on page 108.
Configure the CP server	See " Configuring the CP server using the installer program" on page 109. See "Configuring the CP server using the web-based installer" on page 121. See "Configuring the CP server manually" on page 122. See "Configuring CP server using response files" on page 129.
Verify the CP server configuration	See "Verifying the CP server configuration" on page 133.

Planning your CP server setup

Follow the planning instructions to set up CP server for server-based I/O fencing.

To plan your CP server setup

1. Decide whether you want to host the CP server on a single-node VCS cluster, or on an SFHA cluster.

   Symantec recommends hosting the CP server on an SFHA cluster to make the CP server highly available.

2. If you host the CP server on an SFHA cluster, review the following information. Make sure you make the decisions and meet these prerequisites when you set up the CP server:

- You must set up shared storage for the CP server database during your CP server setup.
- Decide whether you want to configure server-based fencing for the VCS cluster (application cluster) with a single CP server as coordination point or with at least three coordination points.
  Symantec recommends using at least three coordination points.

3 Decide whether you want to configure the CP server cluster for IPM-based communication or HTTPS communication.

  For IPM-based communication, the CP server on release 6.1 supports clients prior to 6.1 release. When you configure the CP server, you are required to provide VIPs for IPM-based clients.

  For HTTPS-based communication, the CP server on release 6.1 only supports clients on release 6.1.

4 Decide whether you want to configure the CP server cluster in secure mode for IPM-based communication.

  Symantec recommends configuring the CP server cluster in secure mode for IPM-based secure communication between the CP server and its clients (VCS clusters). Note that you use IPM-based communication if you want the CP server to support clients that are installed with a release version prior to 6.1 release.

5 Set up the hardware and network for your CP server.

  See "CP server requirements" on page 40.

6 Have the following information handy for CP server configuration:

- Name for the CP server
  The CP server name should not contain any special characters. CP server name can include alphanumeric characters, underscore, and hyphen.
- Port number for the CP server
  Allocate a TCP/IP port for use by the CP server.
  Valid port range is between 49152 and 65535. The default port number for HTTPS-based communication is 443 and for IPM-based secure communication is 14250.
- Virtual IP address, network interface, netmask, and networkhosts for the CP server
  You can configure multiple virtual IP addresses for the CP server.

Installing the CP server using the installer

Perform the following procedure to install and configure VCS or SFHA on CP server systems.

To install and configure VCS or SFHA on the CP server systems

- Depending on whether your CP server uses a single system or multiple systems, perform the following tasks:

CP server setup uses a single system

Install and configure VCS to create a single-node VCS cluster.

During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.

Proceed to configure the CP server.

See " Configuring the CP server using the installer program" on page 109.

See "Configuring the CP server manually" on page 122.

CP server setup uses multiple systems

Install and configure SFHA to create an SFHA cluster. This makes the CP server highly available.

Meet the following requirements for CP server:

- During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.

See the *Symantec Storage Foundation and High Availability Installation Guide* for instructions on installing and configuring SFHA.

Proceed to set up shared storage for the CP server database.

Configuring the CP server cluster in secure mode

You must configure security on the CP server only if you want IPM-based (Symantec Product Authentication Service) secure communication between the CP server and the SFHA cluster (CP server clients). However, IPM-based communication enables the CP server to support application clusters prior to release 6.1.

This step secures the HAD communication on the CP server cluster.

Note: If you already configured the CP server cluster in secure mode during the VCS configuration, then skip this section.

To configure the CP server cluster in secure mode

- Run the installer as follows to configure the CP server cluster in secure mode.

  If you have VCS installed on the CP server, run the following command:

  ```
  # /opt/VRTS/install/installvcs<version>  -security
  ```

  Where *<version>* is the specific release version.

  See "About the script-based installer" on page 49.

  If you have SFHA installed on the CP server, run the following command:

  ```
  # /opt/VRTS/install/installsfha<version>  -security
  ```

  Where *<version>* is the specific release version.

  See "About the script-based installer" on page 49.

Setting up shared storage for the CP server database

If you configured SFHA on the CP server cluster, perform the following procedure to set up shared storage for the CP server database.

The installer can set up shared storage for the CP server database when you configure CP server for the SFHA cluster.

Symantec recommends that you create a mirrored volume for the CP server database and that you use the VxFS file system type.

To set up shared storage for the CP server database

1. Create a disk group containing the disks. You require two disks to create a mirrored volume.

   For example:

   ```
   # vxdg init cps_dg  disk1 disk2
   ```

2. Create a mirrored volume over the disk group.

   For example:

   ```
   # vxassist -g cps_dg make cps_vol volume_size layout=mirror
   ```

3. 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 you must configure CP server manually.

   Depending on the operating system that your CP server runs, enter the following command:

AIX	`# mkfs -V vxfs /dev/vx/rdsk/cps_dg/cps_volume`
Linux	`# mkfs -t vxfs /dev/vx/rdsk/cps_dg/cps_volume`
Solaris	`# mkfs -F vxfs /dev/vx/rdsk/cps_dg/cps_volume`

Configuring the CP server using the installer program

Use the configcps option available in the installer program to configure the CP server.

Perform one of the following procedures:

For CP servers on single-node VCS cluster:	See "To configure the CP server on a single-node VCS cluster" on page 110.
For CP servers on an SFHA cluster:	See "To configure the CP server on an SFHA cluster" on page 115.

To configure the CP server on a single-node VCS cluster

1. Verify that the `VRTScps` package is installed on the node.
2. Run the installvcs<*version*> program with the configcps option.

   ```
   # /opt/VRTS/install/installvcs<version> -configcps
   ```

   Where `<version>` is the specific release version.

   See "About the script-based installer" on page 49.
3. Installer checks the cluster information and prompts if you want to configure CP Server on the cluster.

   Enter **y** to confirm.
4. Select an option based on how you want to configure Coordination Point server.

   ```
   1) Configure Coordination Point Server on single node VCS system
   2) Configure Coordination Point Server on SFHA cluster
   3) Unconfigure Coordination Point Server
   ```
5. Enter the option: [1-3,q] **1**.

   The installer then runs the following preconfiguration checks:

   - Checks to see if a single-node VCS cluster is running with the supported platform.
     The CP server requires VCS to be installed and configured before its configuration.

   The installer automatically installs a license that is identified as a CP server-specific license. It is installed even if a VCS license exists on the node. CP server-specific key ensures that you do not need to use a VCS license on the single-node. It also ensures that Veritas Operations Manager (VOM) identifies the license on a single-node coordination point server as a CP server-specific license and not as a VCS license.
6. Restart the VCS engine if the single-node only has a CP server-specific license.

   ```
   A single node coordination point server will be configured and
   VCS will be started in one node mode, do you want to
   continue?  [y,n,q]  (y)
   ```

7 Communication between the CP server and application clusters is secured by HTTPS from release 6.1.0 onwards. However, clusters on earlier release versions (prior to 6.1.0) that are using IPM-based communication are still supported.

Enter the name of the CP Server.

```
Enter the name of the CP Server: [b]   cps1
```

8 Enter valid virtual IP addresses for the CP Server with HTTPS-based secure communication. A CP Server can be configured with more than one virtual IP address. For HTTPS-based communication, only IPv4 addresses are supported. For IPM-based communication, both IPv4 and IPv6 addresses are supported.

```
Enter Virtual IP(s) for the CP server for HTTPS,
separated by a space:  [b]  10.200.58.231 10.200.58.232
10.200.58.233
```

Note: Ensure that the virtual IP address of the CP server and the IP address of the NIC interface on the CP server belongs to the same subnet of the IP network. This is required for communication to happen between client nodes and CP server.

---

9 Enter the corresponding CP server port number for each virtual IP address or press **Enter** to accept the default value (443).

```
Enter the default port '443' to be used for all the
virtual IP addresses for HTTPS communication or assign the
corresponding port number in the range [49152, 65535] for
each virtual IP address. Ensure that each port number is
separated by a single
space: [b]  (443) 54442 54443 54447
```

10 Decide if you want to support clusters that are on releases prior to 6.1.0. These clusters use the Symantec Product Authentication Services (AT) (secure IPM-based protocol) to securely communicate with the CP servers.

```
Do you want to support older (prior to 6.1.0)
clusters? [y,n,q,b]  (y)
```

11 Enter virtual IPs for the CP Server for IPM-based secure communication.

```
Enter Virtual IP(s) for the CP server for IPM,
separated by a space [b]  10.182.36.8 10.182.36.9
```

Note that both IPv4 and IPv6 addresses are supported.

12 Enter corresponding port number for each Virtual IP address or accept the default port.

```
Enter the default port '14250' to be used for all the
virtual IP addresses for IPM-based communication, or assign
the corresponding port number in the range [49152, 65535]
for each virtual IP address.
Ensure that each port number is separated by a single space:
[b] (14250) 54448 54449
```

13 Decide if you want to enable secure communication between the CP server and application clusters.

```
Symantec 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,q,b] (y) n
```

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

```
Enter absolute path of the database: [b] (/etc/VRTScps/db)
```

15 Verify and confirm the CP server configuration information.

```
CP Server configuration verification:
----------------------------------------------
CP Server Name:  cps1
CP Server Virtual IP(s) for HTTPS: 10.200.58.231, 10.200.58.232,
10.200.58.233
CP Server Virtual IP(s) for IPM: 10.182.36.8, 10.182.36.9
CP Server Port(s) for HTTPS: 54442, 54443, 54447
CP Server Port(s) for IPM: 54448, 54449
CP Server Security for IPM: 0
CP Server Database Dir: /etc/VRTScps/db

----------------------------------------------

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

16 The installer proceeds with the configuration process, and creates a vxcps.conf configuration file.

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

17 Configure the CP Server Service Group (CPSSG) for this cluster.

```
Enter how many NIC resources you want to configure (1 to 2): 2
```

Answer the following questions for each NIC resource that you want to configure.

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

```
Enter a valid network interface on sys1 for NIC resource - 1:  e1000g0
Enter a valid network interface on sys1 for NIC resource - 2:  e1000g1
```

19 Enter the NIC resource you want to associate with the virtual IP addresses.

```
Enter the NIC resource you want to associate with the virtual IP 10.200.58.231 (1 to 2): 1
Enter the NIC resource you want to associate with the virtual IP 10.200.58.232 (1 to 2): 2
```

20 Enter the networkhosts information for each NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure NIC resource
to be always online

Do you want to add NetworkHosts attribute for the NIC device e1000g0
on system sys1? [y,n,q] y
Enter a valid IP address to configure NetworkHosts for NIC e1000g0
on system sys1: 10.200.56.22

Do you want to add another Network Host? [y,n,q] n
```

21 Enter the netmask for virtual IP addresses. If you entered an IPv6 address, enter the prefix details at the prompt.

Note that if you are using HTTPS-based communication, only IPv4 addresses are supported.

```
Enter the netmask for virtual IP for
HTTPS 192.169.0.220: (255.255.252.0)
Enter the netmask for virtual IP for
IPM 192.169.0.221: (255.255.252.0)
```

22 Installer displays the status of the Coordination Point Server configuration. After the configuration process has completed, a success message appears.

```
For example:
Updating main.cf with CPSSG service group.. Done
Successfully added the CPSSG service group to VCS configuration.
Trying to bring CPSSG service group
ONLINE and will wait for upto 120 seconds

The Symantec coordination point server is ONLINE

The Symantec coordination point server has
been configured on your system.
```

23 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.... |ONLINE|
```

It also generates the configuration file for CP server (/etc/vxcps.conf). The vxcpserv process and other resources are added to the VCS configuration in the CP server service group (CPSSG).

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

To configure the CP server on an SFHA cluster

1 Verify that the `VRTScps` package is installed on each node.

2 Ensure that you have configured passwordless ssh or rsh on the CP server cluster nodes.

3 Run the installsfha<*version*> program with the configcps option.

```
# ./installsfha<version> -configcps
```

Where `<version>` is the specific release version.

See "About the script-based installer" on page 49.

4 Installer checks the cluster information and prompts if you want to configure CP Server on the cluster.

Enter **y** to confirm.

5 Select an option based on how you want to configure Coordination Point server.

```
1)  Configure Coordination Point Server on single node VCS system
2)  Configure Coordination Point Server on SFHA cluster
3)  Unconfigure Coordination Point Server
```

6 Enter **2** at the prompt to configure CP server on an SFHA cluster.

The installer then runs the following preconfiguration checks:

- Checks to see if an SFHA cluster is running with the supported platform. The CP server requires SFHA to be installed and configured before its configuration.

7 Communication between the CP server and application clusters is secured by HTTPS from Release 6.1.0 onwards. However, clusters on earlier release versions (prior to 6.1.0) that are using IPM-based communication are still supported.

Enter the name of the CP server.

```
Enter the name of the CP Server: [b]  cps1
```

8 Enter valid virtual IP addresses for the CP Server. A CP Server can be configured with more than one virtual IP address. For HTTPS-based communication, only IPv4 addresses are supported. For IPM-based communication, both IPv4 and IPv6 addresses are supported

```
Enter Virtual IP(s) for the CP server for HTTPS,
 separated by a space: [b] 10.200.58.231 10.200.58.232 10.200.58.233
```

9 Enter the corresponding CP server port number for each virtual IP address or press Enter to accept the default value (443).

```
Enter the default port '443' to be used for all the virtual IP addresses
for HTTPS communication or assign the corresponding port number in the range [49152,
65535] for each virtual IP address. Ensure that each port number is separated by
a single space: [b] (443) 65535 65534 65537
```

10 Decide if you want to support clusters that are on releases prior to 6.1.0. These clusters use the Symantec Product Authentication Services (AT) (secure IPM-based protocol) to securely communicate with the CP servers.

```
Do you want to support older (prior to 6.1.0) clusters? [y,n,q,b] (y)
```

11 Enter Virtual IPs for the CP Server for IPM-based secure communication. Both IPv4 and IPv6 addresses are supported.

```
Enter Virtual IP(s) for the CP server for IPM, separated by a space:
[b] 10.182.36.8 10.182.36.9
```

12 Enter corresponding port number for each Virtual IP address or accept the default port.

```
Enter the default port '14250' to be used for all
the virtual IP addresses for IPM-based communication, or
assign the corresponding port number in the range [49152, 65535]
for each virtual IP address.
Ensure that each port number is separated by a single space:
[b]  (14250) 54448 54449
```

13 Decide if you want to enable secure communication between the CP server and application clusters.

```
Symantec 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,q,b] (y)
```

14 Enter absolute path of the database.

```
CP Server uses an internal database to store the client information.
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: [b] /cpsdb
```

15 Verify and confirm the CP server configuration information.

```
CP Server configuration verification:

CP Server Name: cps1
CP Server Virtual IP(s) for HTTPS: 10.200.58.231, 10.200.58.232,
10.200.58.233
CP Server Virtual IP(s) for IPM: 10.182.36.8, 10.182.36.9
CP Server Port(s) for HTTPS: 65535, 65534, 65537
CP Server Port(s) for IPM: 54448, 54449
CP Server Security for IPM: 1
CP Server Database Dir: /cpsdb

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

16 The installer proceeds with the configuration process, and creates a vxcps.conf configuration file.

```
Successfully generated the /etc/vxcps.conf configuration file
Copying configuration file /etc/vxcps.conf to sys0....Done
Creating mount point /cps_mount_data on sys0. ... Done
Copying configuration file /etc/vxcps.conf to sys0. ... Done
Press Enter to continue.
```

17 Configure CP Server Service Group (CPSSG) for this cluster.

```
Enter how many NIC resources you want to configure (1 to 2): 2

Answer the following questions for each NIC resource that you want to configure.
```

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

```
Enter a valid network interface on sys1 for NIC resource - 1: e1000g0
Enter a valid network interface on sys1 for NIC resource - 2: e1000g1
```

19 Enter the NIC resource you want to associate with the virtual IP addresses.

```
Enter the NIC resource you want to associate with the virtual IP 10.200.58.231 (1 to 2): 1
Enter the NIC resource you want to associate with the virtual IP 10.200.58.232 (1 to 2): 2
```

20 Enter the networkhosts information for each NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure NIC resource
to be always online

Do you want to add NetworkHosts attribute for the NIC device e1000g0
on system sys1? [y,n,q] y
Enter a valid IP address to configure NetworkHosts for NIC e1000g0
on system sys1: 10.200.56.22


Do you want to add another Network Host? [y,n,q] n
Do you want to apply the same NetworkHosts for all systems? [y,n,q] (y)
```

21 Enter the netmask for virtual IP addresses. If you entered an IPv6 address, enter the prefix details at the prompt.

Note that if you are using HTTPS-based communication, only IPv4 addresses are supported.

```
Enter the netmask for virtual IP for
HTTPS 192.168.0.111: (255.255.252.0)
Enter the netmask for virtual IP for
IPM 192.168.0.112: (255.255.252.0)
```

22 Configure a disk group for CP server database. You can choose an existing disk group or create a new disk group.

```
Symantec recommends to use the disk group that has at least
two disks on which mirrored volume can be created.
Select one of the options below for CP Server database disk group:

1)  Create a new disk group
2)  Using an existing disk group

Enter the choice for a disk group: [1-2,q]  2
```

23 Select one disk group as the CP Server database disk group.

```
Select one disk group as CP Server database disk group: [1-3,q] 3
1)  mycpsdg
2)  cpsdg1
3)  newcpsdg
```

24 Select the CP Server database volume.

You can choose to use an existing volume or create new volume for CP Server database. If you chose newly created disk group, you can only choose to create new volume for CP Server database.

```
Select one of the options below for CP Server database volume:
 1)  Create a new volume on disk group newcpsdg
 2)  Using an existing volume on disk group newcpsdg
```

25 Enter the choice for a volume: [1-2,q] **2**.

26 Select one volume as CP Server database volume [1-1,q] **1**

```
1) newcpsvol
```

27 After the VCS configuration files are updated, a success message appears.

```
For example:
Updating main.cf with CPSSG service group .... Done
Successfully added the CPSSG service group to VCS configuration.
```

28 If the cluster is secure, installer creates the softlink `/var/VRTSvcs/vcsauth/data/CPSERVER` to `/cpsdb/CPSERVER` and check if credentials are already present at `/cpsdb/CPSERVER`. If not, installer creates credentials in the directory, otherwise, installer asks if you want to reuse exsting credentials.

```
Do you want to reuse these credentials? [y,n,q] (y)
```

29 After the configuration process has completed, a success message appears.

```
For example:
Trying to bring CPSSG service group ONLINE and will wait for upto 120 seconds
The Symantec Coordination Point Server is ONLINE
The Symantec Coordination Point Server has been configured on your system.
```

30 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 cps1 |ONLINE|
CPSSG State cps2 |OFFLINE|
```

It also generates the configuration file for CP server (`/etc/vxcps.conf`). The vxcpserv process and other resources are added to the VCS configuration in the CP server service group (CPSSG).

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

Configuring the CP server using the web-based installer

Perform the following steps to configure the CP server using the web-based installer.

To configure VCS on a cluster

1 Start the web-based installer.

See "Starting the web-based installer" on page 189.

2 On the Select a task and a product page, select the task and the product as follows:

Task	Configure CP server
Product	Symantec Cluster Server

Click **Next**.

3 On the Select Cluster page, enter the system names where you want to configure VCS and click **Next**.

4 In the Confirmation dialog box, verify cluster information is correct and choose whether or not to configure CP server.

- To configure CP server, click **Yes**.
- To configure CP server later, click **No**.

5 On the Select Option page, select Configure CP Server on a single-node VCS system or SFHA cluster and click **Next**.

6 On the Configure CP Server page, provide CP server information, such as, name, virtual IPs, port numbers, and absolute path of the database to store the configuration details.

Click **Next**.

7 Configure the CP Server Service Group (CPSSG), select the number of NIC resources, and associate NIC resources to virtual IPs that are going to be used to configure the CP Server.

Click **Next**.

8 Configure network hosts for the CP server.

Click **Next**.

9 Configure disk group for the CP server.

Click **Next**.

Note: This step is not applicable for a single node cluster.

10 Configure volume for the disk group associated to the CP server.

Click **Next**.

Note: This step is not applicable for a single node cluster.

11 Click **Finish** to complete configuring the CP server.

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

The CP server supports both IPM-based secure communication and HTTPS-based secure communication. CP servers that are configured for IPM-based secure communication support client nodes that are running either 6.1 or previous versions of the product. However, CP servers that are configured for HTTP-based

communication only support client nodes that are running the 6.1 version of the product. Client nodes with product versions prior to 6.1 are not supported for HTTPS-based communication.

You need to manually generate certificates for the CP server and its client nodes to configure the CP server for HTTPS-based communication.

Table 7-2 Tasks to configure the CP server manually

Task	Reference
Configure CP server manually for IPM-based secure communication	See "Configuring the CP server manually for IPM-based secure communication" on page 123.
Configure CP server manually for HTTPS-communication	See "Configuring the CP server manually for HTTPS-based communication" on page 124. See "Generating the key and certificates manually for the CP server" on page 125. See "Completing the CP server configuration" on page 129.

Configuring the CP server manually for IPM-based secure communication

Perform the following steps to manually configure the CP server in the Symantec Product Authentication Services (AT) (IPM-based) secure mode.

To manually configure the CP server

1 Stop VCS on each node in the CP server cluster using the following command:

```
# hastop -local
```

2 Edit the `main.cf` file to add the CPSSG service group on any node. Use the CPSSG service group in the sample main.cf as an example:

See "Sample configuration files for CP server" on page 537.

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

3 Verify the `main.cf` file using the following command:

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

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

4 Create the `/etc/vxcps.conf` file using the sample configuration file provided at `/etc/vxcps/vxcps.conf.sample`.

Based on whether you configured the CP server using the Symantec Product Authentication Services (AT) protocol (IPM-based) in secure mode or not, do one of the following:

- For a CP server cluster which is configured in secure mode, edit the `/etc/vxcps.conf` file to set `security`=1.
- For a CP server cluster which is not configured in secure mode, edit the `/etc/vxcps.conf` file to set `security`=0.

5 Start VCS on all the cluster nodes.

```
# hastart
```

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

```
# hagrp -state CPSSG
```

Output similar to the following appears:

```
# Group Attribute  System                      Value
  CPSSG State      cps1.symantecexample.com  |ONLINE|
```

Configuring the CP server manually for HTTPS-based communication

Perform the following steps to manually configure the CP server in the Symantec Product Authentication Services (AT) (IPM-based) secure mode.

To manually configure the CP server

1 Stop VCS on each node in the CP server cluster using the following command:

```
# hastop -local
```

2 Edit the `main.cf` file to add the CPSSG service group on any node. Use the CPSSG service group in the sample main.cf as an example:

See "Sample configuration files for CP server" on page 537.

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

3 Verify the `main.cf` file using the following command:

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

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

4 Create the `/etc/vxcps.conf` file using the sample configuration file provided at `/etc/vxcps/vxcps.conf.sample`.

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

If you configured the CP server in HTTPS mode, do the following:

- Edit the `/etc/vxcps.conf` file to set `vip_https` with the virtual IP addresses required for HTTPS communication.
- Edit the `/etc/vxcps.conf` file to set `port_https` with the ports used for HTTPS communication.

5 Manually generate keys and certificates for the CP server.

See "Generating the key and certificates manually for the CP server" on page 125.

Generating the key and certificates manually for the CP server

CP server uses the HTTPS protocol to establish secure communication with client nodes. HTTPS is a secure means of communication, which happens over a secure communication channel that is established using the SSL/TLS protocol.

HTTPS uses x509 standard certificates and the constructs from a Public Key Infrastructure (PKI) to establish secure communication between the CP server and client. Similar to a PKI, the CP server, and its clients have their own set of certificates signed by a Certification Authority (CA). The server and its clients trust the certificate.

Every CP server acts as a certification authority for itself and for all its client nodes. The CP server has its own CA key and CA certificate and a server certificate generated, which is generated from a server private key. The server certificate is issued to the Universally Unique Identifier (UUID) of the CP server. All the IP addresses or domain names that the CP server listens on are mentioned in the Subject Alternative Name section of the CP server's server certificate

The OpenSSL library must be installed on the CP server to create the keys or certificates.. If OpenSSL is not installed, then you cannot create keys or certificates. The vxcps.conf file points to the configuration file that determines which keys or certificates are used by the CP server when SSL is initialized. The configuration value is stored in the `ssl_conf_file` and the default value is `/etc/vxcps_ssl.properties`.

To manually generate keys and certificates for the CP server:

1 Create directories for the security files on the CP server.

```
# mkdir -p /var/VRTScps/security/keys /var/VRTScps/security/certs
```

2 Generate an OpenSSL config file, which includes the VIPs.

The CP server listens to requests from client nodes on these VIPs. The server certificate includes VIPs, FQDNs, and host name of the CP server. Clients can reach the CP server by using any of these values. However, Symantec recommends that client nodes use the IP address to communicate to the CP server.

The sample configuration uses the following values:

- Config file name: *https_ssl_cert.conf*
- VIP: *192.168.1.201*
- FQDN: *cpsone.company.com*
- Host name: *cpsone*

Note the IP address, VIP, and FQDN values used in the [alt_names] section of the configuration file are sample values. Replace the sample values with your configuration values. Do not change the rest of the values in the configuration file.

```
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req

[req_distinguished_name]
countryName = Country Name (2 letter code)
countryName_default = US
localityName = Locality Name (eg, city)
organizationalUnitName = Organizational Unit Name (eg, section)
commonName = Common Name (eg, YOUR name)
commonName_max = 64
emailAddress = Email Address
emailAddress_max = 40

[v3_req]
keyUsage = keyEncipherment, dataEncipherment
extendedKeyUsage = serverAuth
subjectAltName = @alt_names
```

```
[alt_names]
DNS.1 = cpsone.company.com
DNS.2 = cpsone
DNS.3 = 192.168.1.201
```

3 Generate a 4096-bit CA key that is used to create the CA certificate.

The key must be stored at `/var/VRTScps/security/keys/ca.key`. Ensure that only root users can access the CA key, as the key can be misused to create fake certificates and compromise security.

```
# /usr/bin/openssl genrsa -out /var/VRTScps/security/keys/ca.key
4096
```

4 Generate a self-signed CA certificate.

```
# /usr/bin/openssl req -new -x509 -days days -key
/var/VRTScps/security/keys/ca.key -subj \
'/C=countryname/L=localityname/OU=COMPANY/CN=CACERT' -out
/var/VRTScps/security/certs/ca.crt
```

Where, *days* is the days you want the certificate to remain valid, *countryname* is the name of the country, localityname is the city, *CACERT* is the certificate name.

5 Generate a 2048-bit private key for CP server.

The key must be stored at `/var/VRTScps/security/keys/server_private key`.

```
# /usr/bin/openssl genrsa -out
/var/VRTScps/security/keys/server_private.key 2048
```

6 Generate a Certificate Signing Request (CSR) for the server certificate.

The Certified Name (CN) in the certificate is the UUID of the CP server.

```
# /usr/bin/openssl genrsa -out
/var/VRTScps/security/keys/server_private.key 2048
```

7 Generate a Certificate Signing Request (CSR) for the server certificate.

The Certified Name (CN) in the certificate is the UUID of the CP server.

```
# /usr/bin/openssl req -new -key
/var/VRTScps/security/keys/server_private.key \

-config https_ssl_cert.conf -subj
'/C=CountryName/L=LocalityName/OU=COMPANY/CN=UUID' \

-out /var/VRTScps/security/certs/server.csr
```

Where, *countryname* is the name of the country, *localityname* is the city, *UUID* is the certificate name.

8 Generate the server certificate by using the key certificate of the CA.

```
# /usr/bin/openssl x509 -req -days days -in
/var/VRTScps/security/certs/server.csr \

-CA /var/VRTScps/security/certs/ca.crt -CAkey
/var/VRTScps/security/keys/ca.key \

-set_serial 01 -extensions v3_req -extfile https_ssl_cert.conf \

-out /var/VRTScps/security/certs/server.crt
```

Where, *days* is the days you want the certificate to remain valid, *https_ssl_cert.conf* is the configuration file name.

You successfully created the key and certificate required for the CP server.

9 Ensure that no other user except the root user can read the keys and certificates.

10 Complete the CP server configuration.

See "Completing the CP server configuration" on page 129.

Completing the CP server configuration

To verify the service groups and start VCS perform the following steps:

1 Start VCS on all the cluster nodes.

```
# hastart
```

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

```
# hagrp -state CPSSG
```

Output similar to the following appears:

```
# Group Attribute  System                      Value
  CPSSG State      cps1.symantecexample.com  |ONLINE|
```

Configuring CP server using response files

You can configure a CP server using a generated responsefile.

On a single node VCS cluster:

- Run the `installvcs<version>` command with the responsefile option to configure the CP server on a single node VCS cluster.

  ```
  # /opt/VRTS/install/installvcs<version> -responsefile
  '/tmp/sample1.res'
  ```

  Where *<version>* is the specific release version.

  See "About the script-based installer" on page 49.

On a SFHA cluster:

- Run the `installsfha<version>` command with the responsefile option to configure the CP server on a SFHA cluster.

  ```
  # /opt/VRTS/install/installsfha<version> -responsefile
  '/tmp/sample1.res'
  ```

  Where *<version>* is the specific release version.

  See "About the script-based installer" on page 49.

Response file variables to configure CP server

Table 7-3 describes the response file variables to configure CP server.

Table 7-3 describes response file variables to configure CP server

Variable	List or Scalar	Description
CFG{opt}{configcps}	Scalar	This variable performs CP server configuration task
CFG{cps_singlenode_config}	Scalar	This variable describes if the CP server will be configured on a singlenode VCS cluster
CFG{cps_sfha_config}	Scalar	This variable describes if the CP server will be configured on a SFHA cluster
CFG{cps_unconfig}	Scalar	This variable describes if the CP server will be unconfigured
CFG{cpsname}	Scalar	This variable describes the name of the CP server
CFG{cps_db_dir}	Scalar	This variable describes the absolute path of CP server database
CFG{cps_security}	Scalar	This variable describes if security is configured for the CP server
CFG{cps_reuse_cred}	Scalar	This variable describes if reusing the existing credentials for the CP server
CFG{cps_https_vips}	List	This variable describes the virtual IP addresses for the CP server configured for HTTPS-based communication
CFG{cps_ipm_vips}	List	This variable describes the virtual IP addresses for the CP server configured for IPM-based communication
CFG{cps_https_ports}	List	This variable describes the port number for the virtual IP addresses for the CP server configured for HTTPS-based communication
CFG{cps_ipm_ports}	List	This variable describes the port number for the virtual IP addresses for the CP server configured for IPM-based communication
CFG{cps_nic_list}{cpsvip<n>}	List	This variable describes the NICs of the systems for the virtual IP address

Table 7-3 describes response file variables to configure CP server *(continued)*

Variable	List or Scalar	Description
CFG{cps_netmasks}	List	This variable describes the netmasks for the virtual IP addresses
CFG{cps_prefix_length}	List	This variable describes the prefix length for the virtual IP addresses
CFG{cps_network_hosts}{cpsnic<n>}	List	This variable describes the network hosts for the NIC resource
CFG{cps_vip2nicres_map}{<vip>}	Scalar	This variable describes the NIC resource to associate with the virtual IP address
CFG{cps_diskgroup}	Scalar	This variable describes the disk group for the CP server database
CFG{cps_volume}	Scalar	This variable describes the volume for the CP server database
CFG{cps_newdg_disks}	List	This variable describes the disks to be used to create a new disk group for the CP server database
CFG{cps_newvol_volsize}	Scalar	This variable describes the volume size to create a new volume for the CP server database
CFG{cps_delete_database}	Scalar	This variable describes if deleting the database of the CP server during the unconfiguration
CFG{cps_delete_config_log}	Scalar	This variable describes if deleting the config files and log files of the CP server during the unconfiguration
CFG{cps_reconfig}	Scalar	This variable defines if the CP server will be reconfigured

Sample response file for configuring the CP server on single node VCS cluster

Review the response file variables and their definitions.

See Table 7-3 on page 130.

```
#
# Configuration Values:
#
our %CFG;

$CFG{cps_db_dir}="/etc/VRTScps/db";
$CFG{cps_https_ports}=[ qw(443) ];
$CFG{cps_https_vips}=[ qw(192.169.0.220) ];
$CFG{cps_ipm_ports}=[ qw(14250) ];
$CFG{cps_ipm_vips}=[ qw(192.169.0.221) ];
$CFG{cps_netmasks}=[ qw(255.255.252.0 255.255.252.0) ];
$CFG{cps_nic_list}{cpsvip1}=[ qw(e1000g0) ];
$CFG{cps_nic_list}{cpsvip2}=[ qw(e1000g0) ];
$CFG{cps_security}="0";
$CFG{cps_singlenode_config}=1;
$CFG{cps_vip2nicres_map}{"192.169.0.220"}=1;
$CFG{cps_vip2nicres_map}{"192.169.0.221"}=1;
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{prod}="VCS61";
$CFG{systems}=[ qw(cps1) ];
$CFG{vcs_clusterid}=64505;
$CFG{vcs_clustername}="single";

1;
```

Sample response file for configuring the CP server on SFHA cluster

Review the response file variables and their definitions.

See Table 7-3 on page 130.

```
#
# Configuration Values:
#
our %CFG;

$CFG{cps_db_dir}="/cpsdb";
$CFG{cps_diskgroup}="cps_dg1";
```

```
$CFG{cps_https_ports}=[ qw(50006 50007) ];
$CFG{cps_https_vips}=[ qw(10.198.90.6 10.198.90.7) ];
$CFG{cps_ipm_ports}=[ qw(14250) ];
$CFG{cps_ipm_vips}=[ qw(10.198.90.8) ];
$CFG{cps_netmasks}=[ qw(255.255.248.0 255.255.248.0 255.255.248.0) ];
$CFG{cps_network_hosts}{cpsnic1}=[ qw(10.198.88.18) ];
$CFG{cps_network_hosts}{cpsnic2}=[ qw(10.198.88.18) ];
$CFG{cps_newdg_disks}=[ qw(emc_clariion0_249) ];
$CFG{cps_newvol_volsize}=10;
$CFG{cps_nic_list}{cpsvip1}=[ qw(e1000g0 e1000g0) ];
$CFG{cps_nic_list}{cpsvip2}=[ qw(e1000g0 e1000g0) ];
$CFG{cps_nic_list}{cpsvip3}=[ qw(e1000g0 e1000g0) ];
$CFG{cps_security}="0";
$CFG{cps_sfha_config}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.6"}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.7"}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.8"}=1;
$CFG{cps_volume}="volcps";
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{noipc}=1;
$CFG{prod}="SFHA61";
$CFG{systems}=[ qw(cps1 cps2) ];
$CFG{vcs_clusterid}=49604;
$CFG{vcs_clustername}="sfha2233";

1;
```

Verifying the CP server configuration

Perform the following steps to verify the CP server configuration.

To verify the CP server configuration

1. Verify that the following configuration files are updated with the information you provided during the CP server configuration process:
   - /etc/vxcps.conf (CP server configuration file)
   - /etc/VRTSvcs/conf/config/main.cf (VCS configuration file)
   - /etc/VRTScps/db (default location for CP server database for a single-node cluster)

- /cps_db (default location for CP server database for a multi-node cluster)

2 Run the `cpsadm` command to check if the vxcpserv process is listening on the configured Virtual IP.

If the application cluster is configured for HTTPS-based communication, no need to provide the port number assigned for HTTP communication.

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

For IPM-based communication, you need to specify 14250 as the port number.

```
# cpsadm -s cp_server -p 14250  -a ping_cps
```

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

Chapter 8

Configuring VCS

This chapter includes the following topics:

- Overview of tasks to configure VCS using the script-based installer
- Starting the software configuration
- Specifying systems for configuration
- Configuring the cluster name
- Configuring private heartbeat links
- Configuring the virtual IP of the cluster
- Configuring Symantec Cluster Server in secure mode
- Setting up trust relationships for your VCS cluster
- Configuring a secure cluster node by node
- 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 to configure VCS using the script-based installer

Table 8-1 lists the tasks that are involved in configuring VCS using the script-based installer.

Table 8-1 Tasks to configure VCS using the script-based installer

Task	Reference
Start the software configuration	See "Starting the software configuration" on page 137.
Specify the systems where you want to configure VCS	See "Specifying systems for configuration" on page 137.
Configure the basic cluster	See "Configuring the cluster name" on page 138. See "Configuring private heartbeat links" on page 139.
Configure virtual IP address of the cluster (optional)	See "Configuring the virtual IP of the cluster" on page 142.
Configure the cluster in secure mode (optional)	See "Configuring Symantec Cluster Server in secure mode" on page 144.
Add VCS users (required if you did not configure the cluster in secure mode)	See "Adding VCS users" on page 150.
Configure SMTP email notification (optional)	See "Configuring SMTP email notification" on page 151.
Configure SNMP email notification (optional)	See "Configuring SNMP trap notification" on page 152.
Configure global clusters (optional) **Note:** You must have enabled global clustering when you installed VCS.	See "Configuring global clusters" on page 154.
Complete the software configuration	See "Completing the VCS configuration" on page 155.

Starting the software configuration

You can configure VCS using the Symantec product installer or the installvcs command.

Note: If you want to reconfigure VCS, before you start the installer you must stop all the resources that are under VCS control using the `hastop` command or the `hagrp -offline` command.

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 the corresponding number for your product:

   Symantec 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<version> -configure
   ```

   Where `<version>` is the specific release version.

   See “About the script-based installer” on page 49.

   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 configuration

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

```
Enter the operating_system system names separated
by spaces:  [q,?] (sys1) sys1 sys2
```

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. If ssh binaries cannot communicate with remote nodes, the installer tries rsh binaries. And if both ssh and rsh binaries fail, the installer prompts to help the user to setup ssh or rsh binaries.
- Makes sure that the systems are running with the supported operating system
- Makes sure the installer started from the global zone
- Checks whether VCS is installed
- Exits if VCS 6.1 is not installed

3 Review the installer output about the I/O fencing configuration and confirm whether you want to configure fencing in enabled mode.

```
Do you want to configure I/O Fencing in enabled mode? [y,n,q,?] (y)
```

See "About planning to configure I/O fencing" on page 96.

Configuring the cluster name

Enter the cluster information when the installer prompts you.

To configure the cluster

1 Review the configuration instructions that the installer presents.

2 Enter a unique cluster name.

```
Enter the unique cluster name: [q,?] clus1
```

Configuring private heartbeat links

You now configure the private heartbeat links that LLT uses.

See “Setting up the private network” on page 67.

VCS provides the option to use LLT over Ethernet or LLT over UDP (User Datagram Protocol). Symantec recommends that you configure heartbeat links that use LLT over Ethernet for high performance, unless hardware requirements force you to use LLT over UDP. If you want to configure LLT over UDP, make sure you meet the prerequisites.

You must not configure LLT heartbeat using the links that are part of aggregated links. For example, link1, link2 can be aggregated to create an aggregated link, aggr1. You can use aggr1 as a heartbeat link, but you must not use either link1 or link2 as heartbeat links.

See “Using the UDP layer for LLT” on page 549.

The following procedure helps you configure LLT heartbeat links.

To configure private heartbeat links

1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or LLT over UDP.

- Option 1: Configure the heartbeat links using LLT over Ethernet (answer installer questions)
  Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.
  Skip to step 2.
- Option 2: Configure the heartbeat links using LLT over UDP (answer installer questions)
  Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.
  Skip to step 3.
- Option 3: Automatically detect configuration for LLT over Ethernet
  Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.
  Skip to step 5.

Note: Option 3 is not available when the configuration is a single node configuration.

2 If you chose option 1, enter the network interface card details for the private heartbeat links.

The installer discovers and lists the network interface cards.

Answer the installer prompts. The following example shows different NICs based on architecture:

- For Solaris SPARC:
  You must not enter the network interface card that is used for the public network (typically net0.)

```
Enter the NIC for the first private heartbeat link on sys1:
[b,q,?] net1
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 sys1:
[b,q,?] net2
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)
```

3 If you chose option 2, enter the NIC details for the private heartbeat links. This step uses examples such as *private_NIC1* or *private_NIC2* to refer to the available names of the NICs.

```
Enter the NIC for the first private heartbeat
link on sys1: [b,q,?) private_NIC1
Do you want to use address 192.168.0.1 for the
first private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the first private heartbeat
link on sys1: [b,q,?] (50000)
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 sys1: [b,q,?] private_NIC2
Do you want to use address 192.168.1.1 for the
second private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the second private heartbeat
link on sys1: [b,q,?] (50001)
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) y
Enter the NIC for the low priority heartbeat
link on sys1: [b,q,?] (private_NIC0)
Do you want to use address 192.168.3.1 for
the low priority heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the low priority heartbeat
link on sys1: [b,q,?] (50004)
```

4 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 sys1, make sure the same NICs are available on each system. Then, enter **y** at the prompt.

For LLT over UDP, if you want to use the same NICs on other systems, you still must enter unique IP addresses on each NIC for other systems.

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.

5 If you chose option 3 , the installer detects NICs on each system and network links, and sets link priority.

If the installer fails to detect heartbeat links or fails to find any high-priority links, then choose option 1 or option 2 to manually configure the heartbeat links.

See step 2 for option 1, or step 3 for option 2 or step 5 for option 3.

6 Enter a unique cluster ID:

```
Enter a unique cluster ID number between 0-65535: [b,q,?] (60842)
```

The cluster cannot be configured if the cluster ID 60842 is in use by another cluster. Installer performs a check to determine if the cluster ID is duplicate. The check takes less than a minute to complete.

```
Would you like to check if the cluster ID is in use by another
cluster? [y,n,q] (y)
```

7 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 from the Cluster Manager (Java Console), Veritas Operations Manager (VOM), or to specify in the RemoteGroup resource.

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

See the *Symantec 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 When the system prompts whether you want to configure the virtual IP, enter `y`.

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 sys1: net0
Enter the NIC for Virtual IP of the Cluster to use on sys1:
[b,q,?](net0)
```

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 net0 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: net0
        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: net0
        IP: 2001:454e:205a:110:203:baff:feee:10
        Prefix: 64

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

If you want to set up trust relationships for your secure cluster, refer to the following topics:

See "Setting up trust relationships for your VCS cluster" on page 145.

See "Configuring a secure cluster node by node" on page 146.

Configuring Symantec Cluster Server in secure mode

Configuring VCS in secure mode ensures that all the communication between the systems is encrypted and users are verified against security credentials. VCS user names and passwords are not used when a cluster is running in secure mode.

To configure VCS in secure mode

1 To install VCS in secure mode, run the command:

```
# installsfcfsha -security
```

2 To verify the cluster is in secure mode after configuration, run the command:

```
# haclus -value SecureClus
```

The command returns 1 if cluster is in secure mode, else returns 0.

Setting up trust relationships for your VCS cluster

If you need to use an external authentication broker for authenticating VCS users, you must set up a trust relationship between VCS and the broker. For example, if Veritas Operations Manager (VOM) is your external authentication broker, the trust relationship ensures that VCS accepts the credentials that VOM issues.

Perform the following steps to set up a trust relationship between your VCS cluster and a broker.

To set up a trust relationship

1. Ensure that you are logged in as superuser on one of the nodes in the cluster.
2. Enter the following command:

   ```
   # /opt/VRTS/install/installvcs<version> -securitytrust
   ```

   Where <version> is the specific release version.

   See "About the script-based installer" on page 49.

   The installer specifies the location of the log files. It then lists the cluster information such as cluster name, cluster ID, node names, and service groups.
3. When the installer prompts you for the broker information, specify the IP address, port number, and the data directory for which you want to establish trust relationship with the broker.

   ```
   Input the broker name of IP address: 15.193.97.204

   Input the broker port: (14545)
   ```

   Specify a port number on which broker is running or press Enter to accept the default port.

   ```
   Input the data directory to setup trust with: (/var/VRTSvcs/
   vcsauth/data/HAD)
   ```

   Specify a valid data directory or press Enter to accept the default directory.
4. The installer performs one of the following actions:
   - If you specified a valid directory, the installer prompts for a confirmation.

     ```
     Are you sure that you want to setup trust for the VCS cluster
     with the broker 15.193.97.204 and port 14545? [y,n,q] y
     ```

     The installer sets up trust relationship with the broker for all nodes in the cluster and displays a confirmation.

```
Setup trust with broker 15.193.97.204 on cluster node1
........Done

Setup trust with broker 15.193.97.204 on cluster node2
........Done
```

The installer specifies the location of the log files, summary file, and response file and exits.

- If you entered incorrect details for broker IP address, port number, or directory name, the installer displays an error. It specifies the location of the log files, summary file, and response file and exits.

Configuring a secure cluster node by node

For environments that do not support passwordless ssh or passwordless rsh, you cannot use the `-security` option to enable secure mode for your cluster. Instead, you can use the `-securityonenode` option to configure a secure cluster node by node. Moreover, to enable security in fips mode, use the `-fips` option together with `-securityonenode`.

Table 8-2 lists the tasks that you must perform to configure a secure cluster.

Table 8-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See "Configuring the first node" on page 146.
Configure security on the remaining nodes	See "Configuring the remaining nodes" on page 147.
Complete the manual configuration steps	See "Completing the secure cluster configuration" on page 148.

Configuring the first node

Perform the following steps on one node in your cluster.

To configure security on the first node

1 Ensure that you are logged in as superuser.

2 Enter the following command:

```
# /opt/VRTS/install/installvcs<version> -securityonenode
```

Where *<version>* is the specific release version.

See "About the script-based installer" on page 49.

The installer lists information about the cluster, nodes, and service groups. If VCS is not configured or if VCS is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure.

```
VCS is not running on all systems in this cluster. All VCS systems
must be in RUNNING state. Do you want to continue? [y,n,q] (n) y

1) Perform security configuration on first node and export
security configuration files.

2) Perform security configuration on remaining nodes with
security configuration files.

Select the option you would like to perform [1-2,q.?] 1
```

Warning: All VCS configurations about cluster users are deleted when you configure the first node. You can use the `/opt/VRTSvcs/bin/hauser` command to create cluster users manually.

3 The installer completes the secure configuration on the node. It specifies the location of the security configuration files and prompts you to copy these files to the other nodes in the cluster. The installer also specifies the location of log files, summary file, and response file.

4 Copy the security configuration files from the location specified by the installer to temporary directories on the other nodes in the cluster.

Configuring the remaining nodes

On each of the remaining nodes in the cluster, perform the following steps.

To configure security on each remaining node

1. Ensure that you are logged in as superuser.
2. Enter the following command:

   ```
   # /opt/VRTS/install/installvcs<version> -securityonenode
   ```

   Where *<version>* is the specific release version.

   See "About the script-based installer" on page 49.

   The installer lists information about the cluster, nodes, and service groups. If VCS is not configured or if VCS is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure. Enter **2**.

   ```
   VCS is not running on all systems in this cluster. All VCS systems
   must be in RUNNING state. Do you want to continue? [y,n,q] (n) y

   1) Perform security configuration on first node and export
   security configuration files.

   2) Perform security configuration on remaining nodes with
   security configuration files.

   Select the option you would like to perform [1-2,q.?]  2
   ```

   The installer completes the secure configuration on the node. It specifies the location of log files, summary file, and response file.

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

1 On the first node, freeze all service groups except the ClusterService service group.

```
# /opt/VRTSvcs/bin/haconf -makerw

# /opt/VRTSvcs/bin/hagrp -list Frozen=0

# /opt/VRTSvcs/bin/hagrp -freeze groupname -persistent

# /opt/VRTSvcs/bin/haconf -dump -makero
```

2 On the first node, stop the VCS engine.

```
# /opt/VRTSvcs/bin/hastop -all -force
```

3 On all nodes, stop the CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

4 On the first node, edit the `/etc/VRTSvcs/conf/config/main.cf` file to resemble the following:

```
cluster clus1 (
SecureClus = 1
)
```

5 Modify `/etc/VRTSvcs/conf/config/main.cf` file on the first node, and add `-secure` to the WAC applicaton definition if GCO is configured.

For example:

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

6 On all nodes, create the `/etc/VRTSvcs/conf/config/.secure` file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

7 On the first node, start VCS. Then start VCS on the remaining nodes.

```
# /opt/VRTSvcs/bin/hastart
```

8 On all nodes, start CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer
```

9 On the first node, unfreeze the service groups.

```
# /opt/VRTSvcs/bin/haconf -makerw

# /opt/VRTSvcs/bin/hagrp -list Frozen=1

# /opt/VRTSvcs/bin/hagrp -unfreeze groupname -persistent

# /opt/VRTSvcs/bin/haconf -dump -makero
```

Adding VCS users

If you have enabled a secure VCS cluster, 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 wish to accept the default cluster credentials of
'admin/password'? [y,n,q] (y) n
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): [b,q,?] 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 *Symantec 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.

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 152.

3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the SMTP server's host name.

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

```
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 *Symantec 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.

   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 154.

3. Provide information to configure SNMP trap notification.

   Provide the following information:

   - Enter the SNMP trap daemon port.

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

   - Enter the SNMP console system name.

     ```
     Enter the SNMP console system name: [b,q,?] sys5
     ```

   - 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 sys5 [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,?] sys4
     Enter the minimum severity of events for which SNMP traps
     should be sent to sys4 [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.

```
SNMP Port: 162
Console: sys5 receives SNMP traps for Error or
higher events
Console: sys4 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 *Symantec 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.

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.

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

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 and its related processes.

To complete the VCS configuration

1. If prompted, 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 and its related processes.

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 216.

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.

See “Checking licensing information on the system” on page 156.

See “Updating product licenses” on page 156.

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:

```
# 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                  = xxxxx
License Type                   = PERMANENT
OEM ID                         = xxxxx

Features :=
Platform                       = Solaris
Version                        = 6.1
Tier                           = 0
Reserved                       = 0
Mode                           = VCS
CPU_Tier                       = 0
```

Updating product licenses

You can use the `./installer -license` command or the `vxlicinst -k` 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 157.

To update product licenses using the installer command

1 On each node, enter the license key using the command:

```
# ./installer -license
```

2 At the prompt, enter your license number.

To update product licenses using the vxlicinst command

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

```
# vxlicinst -k license key
```

Replacing a VCS demo license with a permanent license

When a VCS demo 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:

```
# hastop -all -force
```

This command does not shut down any running applications.

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

```
# vxlicinst -k license key
```

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

```
# vxlicrep
```

5 Start VCS on each node:

```
# hastart
```

Chapter 9

Configuring VCS clusters for data integrity

This chapter includes the following topics:

- Setting up disk-based I/O fencing using installvcs
- Setting up server-based I/O fencing using installvcs
- Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installvcs
- Enabling or disabling the preferred fencing policy

Setting up disk-based I/O fencing using installvcs

You can configure I/O fencing using the `-fencing` option of the installvcs.

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

1. List the new external disks or the LUNs as recognized by the operating system. On each node, enter:

   ```
   # vxdisk list
   ```

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 *Symantec Storage Foundation 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 c2t13d0
  ```

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

Configuring disk-based I/O fencing using installvcs

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

1. Start the installvcs with `-fencing` option.

   ```
   # /opt/VRTS/install/installvcs<version> -fencing
   ```

   Where `<version>` is the specific release version.

   See "About the script-based installer" on page 49.

   The installvcs starts with a copyright message and verifies the cluster information.

   Note the location of log files which you can access in the event of any problem with the configuration process.

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 6.1 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 [1-6,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 that you use three disks as coordination points for disk-based I/O fencing.
  - If the available VxVM CDS disks are less than the required, installer asks whether you want to initialize more disks as VxVM disks. Choose the disks you want to initialize as VxVM disks and then use them to create new disk group.
    If no VxVM CDS disks are available and some free exported disks exist, installer displays the following message.

```
Installer could not detect any disk that satisfies
the requirements for coordinator disks. It might happen that
if 'vxdisk -o alldisks export' was executed before
invoking the installer, in which case no disk would be listed
because exported devices cannot be used as coordinator disks.
If there are disks that have direct connectivity with all the
cluster nodes and you do not want to mark them as exported devices,
then unexport them with the command 'vxdisk unexport <disk list>'
in a separate console, and refresh the list of disks in the
installer console to continue.
```

  - Decide if you want to unexport any free exported disks to be used as a VxVM disk.

```
Do you want to unexport any disk(s)?
It should be done in a separate console.
```

Choose **Y** and run the `vxdisk unexport` *`disk list`* command in a separate console, then continue.
Choose **N** if you do not want to unexport any free exported disk.

- 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 vxfentsthdw utility and then return to this configuration program.

See "Checking shared disks for I/O fencing" on page 164.

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 disk policy for the disk(s) (raw/dmp): [b,q,?] dmp
```

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.
- Updates the I/O fencing configuration file /etc/vxfenmode.
- 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.

```
Do you want to configure Coordination Point Agent on
the client cluster? [y,n,q] (y)
```

13 Enter a name for the service group for the Coordination Point Agent.

```
Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen) vxfen
```

14 Set the level two monitor frequency.

```
Do you want to set LevelTwoMonitorFreq? [y,n,q] (y)
```

15 Decide the value of the level two monitor frequency.

```
Enter the value of the LevelTwoMonitorFreq attribute: [b,q,?] (5)
```

Installer adds Coordination Point Agent and updates the main configuration file.

See "Configuring CoordPoint agent to monitor coordination points" on page 290.

Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installvcs

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for disk-based I/O fencing using the installvcs

1 Start the installvcs with the `-fencing` option.

```
# /opt/VRTS/install/installvcs<version> -fencing
```

where, `<version>` is the specific release version.

See "About the script-based installer" on page 49.

The installvcs starts with a copyright message and verifies the cluster information.

Note down the location of log files that you can access if there is a problem with the configuration process.

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 the remote nodes and checks whether VCS 6.1 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type **5** to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q] 5
```

4 Ensure that the disk group constitution that is used by the fencing module contains the same disks that are currently used as coordination disks.

5 Verify the coordination points.

```
For example,
Disk Group: fendg
Fencing disk policy: dmp
Fencing disks:
 emc_clariion0_62
 emc_clariion0_65
 emc_clariion0_66
```

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

```
Successfully completed the vxfenswap operation
```

The keys on the coordination disks are refreshed.

6 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

7 Do you want to view the summary file? [y,n,q] **(n)**.

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 *Symantec 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 164.
- Verifying that nodes have access to the same disk
  See "Verifying that the nodes have access to the same disk" on page 165.
- Testing the shared disks for SCSI-3
  See "Testing the disks using vxfentsthdw utility" on page 166.

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
============================================================
libvx3par.so        3PARdata            VV
libvxCLARiiON.so    DGC                 All
libvxFJTSYe6k.so    FUJITSU             E6000
libvxFJTSYe8k.so    FUJITSU             All
libvxap.so          Oracle                 All
libvxatf.so         VERITAS             ATFNODES
libvxcompellent.so  COMPELNT            Compellent Vol
libvxcopan.so       COPANSYS            8814, 8818
```

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 vxfentsthdw 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/rdsk/c1t1d0s2 path on node A and the /dev/rdsk/c2t1d0s2 path on node B.

From node A, enter:

```
# vxfenadm -i /dev/rdsk/c1t1d0s2

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/rdsk/c2t1d0s2 path.

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

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

Testing the disks using vxfentsthdw utility

This procedure uses the /dev/rdsk/c1t1d0s2 disk in the steps.

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

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

```
The disk /dev/rdsk/c1t1d0s2 is ready to be configured for I/O Fencing on
node sys1
```

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

To test the disks using vxfentsthdw utility

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

See "Setting up inter-system communication" on page 567.

2 From one node, start the utility.

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: sys1
Enter the second node of the cluster: sys2
```

4 Review the output as the utility performs the checks and reports its activities.

5 If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node sys1.

```
The disk is now ready to be configured for I/O Fencing on node
sys1

ALL tests on the disk /dev/rdsk/c1t1d0s2 have PASSED
The disk is now ready to be configured for I/O fencing on node
sys1
```

6 Run the vxfentsthdw utility for each disk you intend to verify.

Setting up server-based I/O fencing using installvcs

You can configure server-based I/O fencing for the VCS cluster using the installvcs.

With server-based fencing, you can have the coordination points in your configuration as follows:

- Combination of CP servers and SCSI-3 compliant coordinator disks
- CP servers only
  Symantec also supports server-based fencing with a single highly available CP server that acts as a single coordination point.

See "About planning to configure I/O fencing" on page 96.

See “Recommended CP server configurations” on page 102.

This section covers the following example procedures:

Mix of CP servers and coordinator disks	See “To configure server-based fencing for the VCS cluster (one CP server and two coordinator disks)” on page 168.
Single CP server	See “To configure server-based fencing for the VCS cluster (single CP server)” on page 173.

To configure server-based fencing for the VCS cluster (one CP server and two coordinator disks)

1. Depending on the server-based configuration model in your setup, make sure of the following:
   - CP servers are configured and are reachable from the VCS cluster. The VCS cluster is also referred to as the application cluster or the client cluster.
     See “Setting up the CP server” on page 105.
   - The coordination disks are verified for SCSI3-PR compliance.
     See “Checking shared disks for I/O fencing” on page 164.
2. Start the installvcs with the `-fencing` option.

   ```
   # /opt/VRTS/install/installvcs<version> -fencing
   ```

   Where <version> is the specific release version. The installvcs starts with a copyright message and verifies the cluster information.

   See “About the script-based installer” on page 49.

   Note the location of log files which you can access in the event of any problem with the configuration process.
3. 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 6.1 is configured properly.
4. Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

   ```
   Select the fencing mechanism to be configured in this
   Application Cluster [1-6,b,q] 1
   ```

5 Make sure that the storage supports SCSI3-PR, and answer y at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

6 Provide the following details about the coordination points at the installer prompt:

- Enter the total number of coordination points including both servers and disks. This number should be at least 3.

```
Enter the total number of co-ordination points including both
Coordination Point servers and disks: [b] (3)
```

- Enter the total number of coordinator disks among the coordination points.

```
Enter the total number of disks among these:
[b] (0) 2
```

7 Provide the following CP server details at the installer prompt:

- Enter the total number of virtual IP addresses or the total number of fully qualified host names for each of the CP servers.

```
How many IP addresses would you like to use to communicate
 to Coordination Point Server #1?: [b,q,?] (1) 1
```

- Enter the virtual IP addresses or the fully qualified host name for each of the CP servers. The installer assumes these values to be identical as viewed from all the application cluster nodes.

```
Enter the Virtual IP address or fully qualified host name #1
for the HTTPS Coordination Point Server #1:
[b] 10.209.80.197
```

The installer prompts for this information for the number of virtual IP addresses you want to configure for each CP server.

- Enter the port that the CP server would be listening on.

```
Enter the port that the coordination point server 10.198.90.178
would be listening on or simply accept the default port
suggested: [b] (443)
```

8 Provide the following coordinator disks-related details at the installer prompt:

- Enter the I/O fencing disk policy for the coordinator disks.

```
Enter disk policy for the disk(s) (raw/dmp):
[b,q,?] raw
```

- Choose the coordinator disks from the list of available disks that the installer displays. Ensure that the disk you choose is available from all the VCS (application cluster) nodes.
  The number of times that the installer asks you to choose the disks depends on the information that you provided in step 6. For example, if you had chosen to configure two coordinator disks, the installer asks you to choose the first disk and then the second disk:

```
Select disk number 1 for co-ordination point

1) c1t1d0s2
2) c2t1d0s2
3) c3t1d0s2

Please enter a valid disk which is available from all the
cluster nodes for co-ordination point [1-3,q] 1
```

- If you have not already checked the disks for SCSI-3 PR compliance in step 1, check the disks now.
  The installer displays a message that recommends you to verify the disks in another window and then return to this configuration procedure.
  Press Enter to continue, and confirm your disk selection at the installer prompt.
- Enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):
[b] (vxfencoorddg)
```

9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 3
Coordination Point Server ([VIP or FQHN]:Port):
    1. 10.209.80.197 ([10.209.80.197]:443)
SCSI-3 disks:
    1. c1t1d0s2
    2. c2t1d0s2
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk policy used for customized fencing: raw
```

The installer initializes the disks and the disk group and deports the disk group on the VCS (application cluster) node.

10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the VCS (application cluster).

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

11 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

12 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

```
Updating client cluster information on Coordination Point Server 10.209.80.197

Adding the client cluster to the Coordination Point Server 10.209.80.197 .......... Done

Registering client node sys1 with Coordination Point Server 10.209.80.197...... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.209.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 ..Done

Updating /etc/vxfenmode file on sys1 ................................. Done
Updating /etc/vxfenmode file on sys2 ......... ....................... Done
```

See “About I/O fencing configuration files” on page 534.

13 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

14 Configure the CP agent on the VCS (application cluster). The Coordination Point Agent monitors the registrations on the coordination points.

```
Do you want to configure Coordination Point Agent on
the client cluster? [y,n,q] (y)

Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen)
```

15 Additionally the coordination point agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidently deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles.

Note that for the LevelTwoMonitorFreq attribute to be applicable there must be disks as part of the Coordinator Disk Group.

```
Enter the value of the LevelTwoMonitorFreq attribute: (5)

Adding Coordination Point Agent via sys1 .... Done
```

16 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

17 Verify the fencing configuration using:

```
# vxfenadm -d
```

18 Verify the list of coordination points.

```
# vxfenconfig -l
```

To configure server-based fencing for the VCS cluster (single CP server)

1 Make sure that the CP server is configured and is reachable from the VCS cluster. The VCS cluster is also referred to as the application cluster or the client cluster.

2 See "Setting up the CP server" on page 105.

3 Start the installvcs with `-fencing` option.

```
# /opt/VRTS/install/installvcs<version>  -fencing
```

Where <version> is the specific release version. The installvcs starts with a copyright message and verifies the cluster information.

See "About the script-based installer" on page 49.

Note the location of log files which you can access in the event of any problem with the configuration process.

4 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 6.1 is configured properly.

5 Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q] 1
```

6 Make sure that the storage supports SCSI3-PR, and answer y at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

7 Enter the total number of coordination points as **1**.

```
Enter the total number of co-ordination points including both
Coordination Point servers and disks: [b] (3) 1
```

Read the installer warning carefully before you proceed with the configuration.

8 Provide the following CP server details at the installer prompt:

- Enter the total number of virtual IP addresses or the total number of fully qualified host names for each of the CP servers.

  ```
  How many IP addresses would you like to use to communicate
  to Coordination Point Server #1? [b,q,?] (1) 1
  ```

- Enter the virtual IP address or the fully qualified host name for the CP server. The installer assumes these values to be identical as viewed from all the application cluster nodes.

  ```
  Enter the Virtual IP address or fully qualified host name
  #1 for the Coordination Point Server #1:
  [b] 10.209.80.197
  ```

  The installer prompts for this information for the number of virtual IP addresses you want to configure for each CP server.

- Enter the port that the CP server would be listening on.

  ```
  Enter the port in the range [49152, 65535] which the
  Coordination Point Server 10.209.80.197
  would be listening on or simply accept the default
  port suggested: [b] (443)
  ```

9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 1
Coordination Point Server ([VIP or FQHN]:Port):
    1. 10.209.80.197 ([10.209.80.197]:443)
```

10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the VCS (application cluster).

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

11 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

12 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

The installer also populates the `/etc/vxfenmode` file with the entry `single_cp=1` for such single CP server fencing configuration.

```
Updating client cluster information on Coordination Point Server 10.209.80.197

Adding the client cluster to the Coordination Point Server 10.209.80.197 .......... Done

Registering client node sys1 with Coordination Point Server 10.209.80.197...... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.209.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Updating /etc/vxfenmode file on sys1 ................................. Done
Updating /etc/vxfenmode file on sys2 ......... ....................... Done
```

See “About I/O fencing configuration files” on page 534.

13 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

14 Configure the CP agent on the VCS (application cluster).

```
Do you want to configure Coordination Point Agent on the
client cluster? [y,n,q] (y)

Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen)

Adding Coordination Point Agent via sys1 ... Done
```

15 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

Refreshing keys or registrations on the existing coordination points for server-based fencing using the installvcs

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss might occur because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose registrations of the cluster nodes, the cluster might panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for server-based I/O fencing using the installvcs

1 Start the installvcs with the `-fencing` option.

```
# /opt/VRTS/install/installvcs<version> -fencing
```

where `<version>` is the specific release version.

See "About the script-based installer" on page 49.

The installvcs starts with a copyright message and verifies the cluster information.

Note the location of log files that you can access if there is a problem with the configuration process.

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 the remote nodes and checks whether VCS 6.1 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type **5** to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q] 5
```

4 Ensure that the `/etc/vxfentab` file contains the same coordination point servers that are currently used by the fencing module.

Also, ensure that the disk group mentioned in the `/etc/vxfendg` file contains the same disks that are currently used by the fencing module as coordination disks.

5 Verify the coordination points.

```
For example,
Total number of coordination points being used: 3
Coordination Point Server ([VIP or FQHN]:Port):
      1. 10.198.94.146 ([10.198.94.146]:443)
      2. 10.198.94.144 ([10.198.94.144]:443)
SCSI-3 disks:
      1. emc_clariion0_61
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk policy used for customized fencing: dmp
```

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

```
Updating client cluster information on Coordination Point Server
 IPaddress

Successfully completed the vxfenswap operation
```

The keys on the coordination disks are refreshed.

7 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

8 Do you want to view the summary file? [y,n,q] **(n)**.

Setting the order of existing coordination points for server-based fencing using the installvcs

This section describes the reasons, benefits, considerations, and the procedure to set the order of the existing coordination points for server-based fencing.

About deciding the order of existing coordination points

You can decide the order in which coordination points can participate in a race during a network partition. In a network partition scenario, I/O fencing attempts to contact coordination points for membership arbitration based on the order that is set in the `vxfentab` file.

When I/O fencing is not able to connect to the first coordination point in the sequence it goes to the second coordination point and so on. To avoid a cluster panic, the surviving subcluster must win majority of the coordination points. So, the order must begin with the coordination point that has the best chance to win the race and must end with the coordination point that has the least chance to win the race.

For fencing configurations that use a mix of coordination point servers and coordination disks, you can specify either coordination point servers before coordination disks or disks before servers.

Note: Disk-based fencing does not support setting the order of existing coordination points.

Considerations to decide the order of coordination points

- Choose the coordination points based on their chances to gain membership on the cluster during the race and hence gain control over a network partition. In effect, you have the ability to save a partition.

- First in the order must be the coordination point that has the best chance to win the race. The next coordination point you list in the order must have relatively lesser chance to win the race. Complete the order such that the last coordination point has the least chance to win the race.

Setting the order of existing coordination points using the installvcs

To set the order of existing coordination points

1 Start the installvcs with `-fencing` option.

```
# /opt/VRTS/install/installvcs<version> -fencing
```

where `<version>` is the specific release version.

See "About the script-based installer" on page 49.

The installvcs starts with a copyright message and verifies the cluster information.

Note the location of log files that you can access if there is a problem with the configuration process.

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 6.1 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type **6** to set the order of existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q] 6


Installer will ask the new order of existing coordination points.
Then it will call vxfenswap utility to commit the
coordination points change.
```

Warning: The cluster might panic if a node leaves membership before the coordination points change is complete.

4 Review the current order of coordination points.

```
Current coordination points order:
(Coordination disks/Coordination Point Server)
Example,
1)  /dev/vx/rdmp/emc_clariion0_65,/dev/vx/rdmp/emc_clariion0_66,
/dev/vx/rdmp/emc_clariion0_62
2)  [10.198.94.144]:443
3)  [10.198.94.146]:443
b)  Back to previous menu
```

5 Enter the new order of the coordination points by the numbers and separate the order by space [1-3,b,q] **3 1 2**.

```
New coordination points order:
(Coordination disks/Coordination Point Server)
Example,
1) [10.198.94.146]:443
2) /dev/vx/rdmp/emc_clariion0_65,/dev/vx/rdmp/emc_clariion0_66,
/dev/vx/rdmp/emc_clariion0_62
3) [10.198.94.144]:443
```

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

```
Preparing vxfenmode.test file on all systems...
Running vxfenswap...
Successfully completed the vxfenswap operation
```

7 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

8 Do you want to view the summary file? [y,n,q] **(n)**.

9 Verify that the value of `vxfen_honor_cp_order` specified in the `/etc/vxfenmode` file is set to **1**.

```
For example,
vxfen_mode=customized
vxfen_mechanism=cps
port=443
scsi3_disk_policy=dmp
cps1=[10.198.94.146]
vxfendg=vxfencoorddg
cps2=[10.198.94.144]
vxfen_honor_cp_order=1
```

10 Verify that the coordination point order is updated in the output of the `vxfenconfig -l` command.

```
For example,
I/O Fencing Configuration Information:
======================================

single_cp=0
[10.198.94.146]:443 {e7823b24-1dd1-11b2-8814-2299557f1dc0}
/dev/vx/rdmp/emc_clariion0_65  60060160A38B1600386FD87CA8FDDD11
/dev/vx/rdmp/emc_clariion0_66  60060160A38B1600396FD87CA8FDDD11
/dev/vx/rdmp/emc_clariion0_62  60060160A38B16005AA00372A8FDDD11
[10.198.94.144]:443 {01f18460-1dd2-11b2-b818-659cbc6eb360}
```

Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installvcs

If you have installed VCS in virtual environments that do not support SCSI-3 PR-compliant storage, you can configure non-SCSI-3 fencing.

To configure I/O fencing using the installvcs in a non-SCSI-3 PR-compliant setup

1 Start the installvcs with `-fencing` option.

```
# /opt/VRTS/install/installvcs<version> -fencing
```

Where `<version>` is the specific release version.

See "About the script-based installer" on page 49.

The installvcs 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 6.1 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster
[1-6,q] 1
```

4 Enter **n** to confirm that your storage environment does not support SCSI-3 PR.

```
Does your storage environment support SCSI3 PR?
[y,n,q] (y) n
```

5 Confirm that you want to proceed with the non-SCSI-3 I/O fencing configuration at the prompt.

6 Enter the number of CP server coordination points you want to use in your setup.

7 Enter the following details for each CP server:

- Enter the virtual IP address or the fully qualified host name.
- Enter the port address on which the CP server listens for connections. The default value is 443. You can enter a different port address. Valid values are between 49152 and 65535.

The installer assumes that these values are identical from the view of the VCS cluster nodes that host the applications for high availability.

8 Verify and confirm the CP server information that you provided.

9 Verify and confirm the VCS cluster configuration information.

Review the output as the installer performs the following tasks:

- Updates the CP server configuration files on each CP server with the following details:
  - Registers each node of the VCS cluster with the CP server.
  - Adds CP server user to the CP server.
  - Adds VCS cluster to the CP server user.
- Updates the following configuration files on each node of the VCS cluster
  - `/etc/vxfenmode` file
  - `/etc/default/vxfen` file
  - `/etc/vxenviron` file
  - `/etc/llttab` file
  - `/etc/vxfentab`

10 Review the output as the installer stops VCS on each node, starts I/O fencing on each node, updates the VCS configuration file main.cf, and restarts VCS with non-SCSI-3 server-based fencing.

Confirm to configure the CP agent on the VCS cluster.

11 Confirm whether you want to send the installation information to Symantec.

12 After the installer configures I/O fencing successfully, note the location of summary, log, and response files that installer creates.

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

Enabling or disabling the preferred fencing policy

You can enable or disable the preferred fencing feature for your I/O fencing configuration.

You can enable preferred fencing to use system-based race policy or group-based race policy. If you disable preferred fencing, the I/O fencing configuration uses the default count-based race policy.

See “About preferred fencing” on page 35.

To enable preferred fencing for the I/O fencing configuration

1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

3 To enable system-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as System.

```
# haclus -modify PreferredFencingPolicy System
```

- Set the value of the system-level attribute FencingWeight for each node in the cluster.
  For example, in a two-node cluster, where you want to assign sys1 five times more weight compared to sys2, run the following commands:

```
# hasys -modify sys1 FencingWeight 50
# hasys -modify sys2 FencingWeight 10
```

- Save the VCS configuration.

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

- Verify fencing node weights using:

```
# vxfenconfig -a
```

4 To enable group-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Group.

```
# haclus -modify PreferredFencingPolicy Group
```

- Set the value of the group-level attribute Priority for each service group.

For example, run the following command:

```
# hagrp -modify service_group Priority 1
```

Make sure that you assign a parent service group an equal or lower priority than its child service group. In case the parent and the child service groups are hosted in different subclusters, then the subcluster that hosts the child service group gets higher preference.

- Save the VCS configuration.

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

5 To enable site-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Site.

```
# haclus -modify PreferredFencingPolicy Site
```

- Set the value of the site-level attribute Preference for each site.

```
For example,
# hasite -modify Pune Preference 2
```

- Save the VCS configuration.

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

6 To view the fencing node weights that are currently set in the fencing driver, run the following command:

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

3 To disable preferred fencing and use the default race policy, set the value of the cluster-level attribute PreferredFencingPolicy as Disabled.

```
# haconf -makerw
# haclus -modify PreferredFencingPolicy Disabled
# haconf -dump -makero
```

Section 4

Installation using the Web-based installer

- Chapter 10. Installing VCS
- Chapter 11. Configuring VCS

Chapter 10

Installing VCS

This chapter includes the following topics:

- Before using the web-based installer
- Starting the web-based installer
- Obtaining a security exception on Mozilla Firefox
- Performing a preinstallation check with the web-based installer
- Installing VCS with the web-based installer

Before using the web-based installer

The web-based installer requires the following configuration.

Table 10-1 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Symantec products.	Must be a supported platform for VCS 6.1.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must be at one of the supported operating system update levels.
Administrative system	The system where you run the web browser to perform the installation.	Must have a web browser. Supported browsers: ■ Internet Explorer 6, 7, and 8 ■ Firefox 3.x and later

Starting the web-based installer

This section describes starting the web-based installer.

To start the web-based installer

1 Start the Veritas XPortal Server process `xprtlwid`, on the installation server:

```
# ./webinstaller start
```

The `webinstaller` script displays a URL. Note this URL.

Note: If you do not see the URL, please check your firewall and iptables settings. If you have configured a firewall, ensure that the firewall settings allow access to the port 14172. You can alternatively use the `-port` option to use a free port instead.

You can use the following command to display the details about ports used by webinstaller and its status:

```
# ./webinstaller status
```

2 On the administrative server, start the web browser.

3 Navigate to the URL that the script displayed.

4 Certain browsers may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

When you are prompted, enter `root` and root's password of the installation server.

5 Log in as superuser.

Obtaining a security exception on Mozilla Firefox

You may need to get a security exception on Mozilla Firefox.

The following instructions are general. They may change because of the rapid release cycle of Mozilla browsers.

To obtain a security exception

1 Click **Or you can add an exception** link.

2 Click **I Understand the Risks**, or **You can add an exception**.

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 preinstallation check with the web-based installer

This section describes performing a preinstallation check with the web-based installer.

To perform a preinstallation check

1 Start the web-based installer.

See "Starting the web-based installer" on page 189.

2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.

3 Select **Symantec Cluster Server** 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 **Next**.

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 **Yes** to install on the selected system. Click **No** to install later.

7 If the validation completes successfully, click **Next**. The installer prompts you to begin the installation. Click **Yes** to install on the selected system. Click **No** to install later.

8 Click **Finish**. The installer prompts you for another task.

Installing VCS with the web-based installer

This section describes installing VCS with the Symantec web-based installer.

To install VCS using the web-based installer

1. Perform preliminary steps.

   See "Performing a preinstallation check with the web-based installer" on page 190.

2. Start the web-based installer.

   See "Starting the web-based installer" on page 189.

3. Select **Install a Product** from the **Task** drop-down list.

4. Select **Symantec Cluster Server** from the Product drop-down list, and click **Next**.

5. On the License agreement page, read 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 where you want to install. Separate multiple system names with spaces. Click **Next**.

8. If you have not yet configured a communication mode among systems, you have the option to let the installer configure ssh or rsh. If you choose to allow this configuration, select the communication mode and provide the superuser passwords for the systems.

9. After the validation completes successfully, click **Next** to install VCS on the selected system.

10. After the installation completes, you must choose your licensing method.

    On the license page, select one of the following radio buttons:

    - Enable keyless licensing and complete system licensing later

      Note: The keyless license option enables you to install without entering a key. However, 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

      Click **Next**

      Complete the following information:

      - Choose whether you want to enable Global Cluster option.
      - Click **Next**.

- Enter a valid license key
  If you have a valid license key, input the license key and click **Next**.

11 The installer prompts you to configure the cluster. Select **Yes** to continue with configuring the product.

   If you select **No**, 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.

12 If you are prompted, enter the option to specify whether you want to send your installation information to Symantec.

```
Installation procedures and diagnostic information were saved in
the log files under directory
/var/tmp/installer-<platform>-<uuid>. Analyzing this information
helps Symantec discover and fix failed operations performed by
the installer. Would you like to send the information about this
installation to Symantec to help improve installation in the
future? [y,n,q,?]
```

   Click **Finish**. The installer asks if you want to read the summary file. Select **Yes** to read the summary file. If you select **No**, the installer prompts you for another task.

Chapter 11

Configuring VCS

This chapter includes the following topics:

- Configuring VCS using the web-based installer
- Configuring VCS for data integrity using the web-based installer

Configuring VCS using the web-based installer

Before you begin to configure VCS using the web-based installer, review the configuration requirements.

See "Getting your VCS installation and configuration information ready" on page 82.

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.

You can click **Quit** to quit the web-installer at any time during the configuration process.

To configure VCS on a cluster

1 Start the web-based installer.

See "Starting the web-based installer" on page 189.

2 On the Select a task and a product page, select the task and the product as follows:

Task	Configure CP Server
Product	Symantec Cluster Server

Click **Next**.

3 On the Select Systems page, enter the system names where you want to configure VCS, and click **Next**.

Example: **sys1 sys2**

The installer performs the initial system verification. It checks for the system communication. 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 In the Confirmation dialog box that appears, choose whether or not to configure CP Server.

Would you like to configure CP Server on the cluster?, click **Yes**.

Would you like to configure CP Server on the cluster later?, click **No**. You can configure I/O fencing later using the web-based installer.

See "Configuring VCS for data integrity using the web-based installer" on page 198.

You can also configure I/O fencing later using the `installvcs<version> -fencing` command, the response files, or manually configure.

Where `<version>` is the specific release version.

See "About the script-based installer" on page 49.

5 On the Set Cluster Name/ID page, specify the following information for the cluster.

Cluster Name	Enter a unique cluster name.
Cluster ID	Enter a unique cluster ID. Note that you can have the installer check to see if the cluster ID is unique. Symantec recommends that you use the installer to check for duplicate cluster IDs in multi-cluster environments.
Check duplicate cluster ID	Select the check box if you want the installer to verify if the given cluster ID is unique in your private network. The verification is performed after you specify the heartbeat details in the following pages. The verification takes some time to complete.
LLT Type	Select an LLT type from the list. You can choose to configure LLT over UDP or LLT over Ethernet.
Number of Heartbeats	Choose the number of heartbeat links you want to configure. See "Setting up the private network" on page 67.
Additional Low Priority Heartbeat NIC	Select the check box if you want to configure a low priority link. The installer configures one heartbeat link as low priority link. See "Setting up the private network" on page 67.
Unique Heartbeat NICs per system	For LLT over Ethernet, select the check box if you do not want to use the same NIC details to configure private heartbeat links on other systems. For LLT over UDP, this check box is selected by default.

Click **Next**.

6 On the Set Cluster Heartbeat page, select the heartbeat link details for the LLT type you chose on the Set Cluster Name/ID page.

For **LLT over Ethernet**: Do the following:

- If you are using the same NICs on all the systems, select the NIC for each private heartbeat link.
- If you had selected **Unique Heartbeat NICs per system** on the Set Cluster Name/ID page, provide the NIC details for each system.

For **LLT over UDP**: Select the NIC, Port, and IP address for each private heartbeat link. You must provide these details for each system.

Click **Next**.

7 On the Optional Configuration page, decide the optional VCS features that you want to configure. Click the corresponding tab to specify the details for each option:

Security

To configure a secure VCS cluster, select the **Configure secure cluster** check box.

If you want to perform this task later, do not select the **Configure secure cluster** check box. You can use the `-security` option of the installvcs.

Virtual IP

- Select the **Configure Virtual IP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select the interface on which you want to configure the virtual IP.
- Enter a virtual IP address and value for the netmask. You can use an IPv4 or an IPv6 address.

VCS Users

- Reset the password for the Admin user, if necessary.
- Select the **Configure VCS users** option.
- Click **Add** to add a new user.
  Specify the user name, password, and user privileges for this user.

SMTP

- Select the **Configure SMTP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- 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.
- In the **SMTP Server** box, enter the domain-based hostname of the SMTP server. Example: smtp.yourcompany.com
- In the **Recipient** box, enter the full email address of the SMTP recipient. Example: user@yourcompany.com.
- In the **Event** list box, select the minimum security level of messages to be sent to each recipient.
- Click **Add** to add more SMTP recipients, if necessary.

SNMP

- Select the **Configure SNMP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- 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.
- In the **SNMP Port** box, enter the SNMP trap daemon port: (162).
- In the **Console System Name** box, enter the SNMP console system name.
- In the **Event** list box, select the minimum security level of messages to be sent to each console.
- Click **Add** to add more SNMP consoles, if necessary.

GCO

If you installed a valid HA/DR license, you can now enter the wide-area heartbeat link details for the global cluster that you would set up later.

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

- Select the **Configure GCO** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select a NIC.
- Enter a virtual IP address and value for the netmask. You can use an IPv4 or an IPv6 address.

Click **Next**.

8 On the Stop Processes page, click **Next** after the installer stops all the processes successfully.

9 On the Start Processes page, click **Next** after the installer performs the configuration based on the details you provided and starts all the processes successfully.

If you did not choose to configure I/O fencing in step 4, then skip to step 11. Go to step 10 to configure fencing.

10 On the Select Fencing Type page, choose the type of fencing configuration:

Configure Coordination Point client based fencing	Choose this option to configure server-based I/O fencing.
Configure disk based fencing	Choose this option to configure disk-based I/O fencing.

Based on the fencing type you choose to configure, follow the installer prompts.

See "Configuring VCS for data integrity using the web-based installer" on page 198.

11 Click **Next** to complete the process of configuring VCS.

On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.

12 Select the checkbox to specify whether you want to send your installation information to Symantec.

Click **Finish**. The installer prompts you for another task.

Configuring VCS for data integrity using the web-based installer

After you configure VCS, you must configure the cluster for data integrity. Review the configuration requirements.

See "Configuring VCS using the web-based installer" on page 193.

See "About planning to configure I/O fencing" on page 96.

Ways to configure I/O fencing using the web-based installer:

- See "Configuring disk-based fencing for data integrity using the web-based installer" on page 199.
- See "Configuring server-based fencing for data integrity using the web-based installer" on page 201.
- See "Configuring fencing in disabled mode using the web-based installer" on page 203.
- See "Replacing, adding, or removing coordination points using the web-based installer" on page 204.
- See "Refreshing keys or registrations on the existing coordination points using web-based installer" on page 206.

- See “Setting the order of existing coordination points using the web-based installer” on page 207.

Configuring disk-based fencing for data integrity using the web-based installer

After you configure VCS, you must configure the cluster for data integrity. Review the configuration requirements.

See “Configuring VCS using the web-based installer” on page 193.

See “About planning to configure I/O fencing” on page 96.

To configure VCS for data integrity

1. Start the web-based installer.

   See “Starting the web-based installer” on page 189.

2. On the Select a task and a product page, select the task and the product as follows:

Task	I/O fencing configuration
Product	Symantec Cluster Server

   Click **Next**.

3. Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4. On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

   The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

5. On the Select Fencing Type page, select the `Configure disk-based fencing` option.

6. In the Confirmation dialog box that appears, confirm whether your storage environment supports SCSI-3 PR.

   You can configure non-SCSI-3 server-based fencing in a virtual environment that is not SCSI-3 PR compliant.

7. On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.

   Click **Next**.

8 On the Configure Fencing page, specify the following information:

Select a Disk Group Select the **Create a new disk group** option or select one of the disk groups from the list.

- If you selected one of the disk groups that is listed, choose the fencing disk policy for the disk group.
- If you selected the **Create a new disk group** option, make sure you have SCSI-3 PR enabled disks, and click **Yes** in the confirmation dialog box.

Click **Next**.

9 On the Create New DG page, specify the following information:

New Disk Group Name Enter a name for the new coordinator disk group you want to create.

Select Disks Select at least three disks to create the coordinator disk group.

If you want to select more than three disks, make sure to select an odd number of disks.

Fencing Disk Policy Choose the fencing disk policy for the disk group.

10 Verify and confirm the I/O fencing configuration information.

The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

11 If you want to configure the Coordination Point agent on the client cluster, do the following:

- At the prompt for configuring the Coordination Point agent on the client cluster, click **Yes** and enter the Coordination Point agent service group name.
- If you want to set the LevelTwoMonitorFreq attribute, click Yes at the prompt and enter a value (0 to 65535).
- Follow the rest of the prompts to complete the Coordination Point agent configuration.

12 Click **Next** to complete the process of configuring I/O fencing.

On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.

13 Select the checkbox to specify whether you want to send your installation information to Symantec.

Click **Finish**. The installer prompts you for another task.

Configuring server-based fencing for data integrity using the web-based installer

After you configure VCS, you must configure the cluster for data integrity. Review the configuration requirements.

See "Configuring VCS using the web-based installer" on page 193.

See "About planning to configure I/O fencing" on page 96.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 189.

2 On the Select a task and a product page, select the task and the product as follows:

Task	I/O fencing configuration
Product	Symantec Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

5 On the Select Fencing Type page, select the `Configure Coordination Point client based fencing` option.

6 In the Confirmation dialog box that appears, confirm whether your storage environment supports SCSI-3 PR.

You can configure non-SCSI-3 server-based fencing in a virtual environment that is not SCSI-3 PR compliant.

7 On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.

Click **Next**.

8 Provide the following details for each of the CP servers:

- Enter the virtual IP addresses or host names of the virtual IP address. The installer assumes these values to be identical as viewed from all the application cluster nodes.
- Enter the port that the CP server must listen on.
- Click **Next**.

9 If your server-based fencing configuration also uses disks as coordination points, perform the following steps:

- If you have not already checked the disks for SCSI-3 PR compliance, check the disks now, and click OK in the dialog box.
- If you do not want to use the default coordinator disk group name, enter a name for the new coordinator disk group you want to create.
- Select the disks to create the coordinator disk group.
- Choose the fencing disk policy for the disk group.
  The default fencing disk policy for the disk group is dmp.

10 In the Confirmation dialog box that appears, confirm whether the coordination points information you provided is correct, and click **Yes**.

11 Verify and confirm the I/O fencing configuration information.

The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

12 If you want to configure the Coordination Point agent on the client cluster, do the following:

- At the prompt for configuring the Coordination Point agent on the client cluster, click **Yes** and enter the Coordination Point agent service group name.
- Follow the rest of the prompts to complete the Coordination Point agent configuration.

13 Click **Next** to complete the process of configuring I/O fencing.

On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.

14 Select the checkbox to specify whether you want to send your installation information to Symantec.

Click **Finish**. The installer prompts you for another task.

Configuring fencing in disabled mode using the web-based installer

After you configure VCS, you must configure the cluster for data integrity. Review the configuration requirements.

See "Configuring VCS using the web-based installer" on page 193.

See "About planning to configure I/O fencing" on page 96.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 189.

2 On the Select a task and a product page, select the task and the product as follows:

Task	I/O fencing configuration
Product	Symantec Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

5 Fencing may be enabled, installer may prompt whether you want to reconfigure it.

Click **Yes**.

6 On the Select Fencing Type page, select the `Configure fencing in disabled mode` option.

7 Installer stops VCS before applying the selected fencing mode to the cluster.

Note: Unfreeze any frozen service group and unmount any file system that is mounted in the cluster.

Click **Yes**.

8 Installer restarts VCS on all systems of the cluster. I/O fencing is disabled.

9 Verify and confirm the I/O fencing configuration information.

On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.

10 Select the checkbox to specify whether you want to send your installation information to Symantec.

Click **Finish**. The installer prompts you for another task.

Replacing, adding, or removing coordination points using the web-based installer

After you configure VCS, you must configure the cluster for data integrity. Review the configuration requirements.

See "Configuring VCS using the web-based installer" on page 193.

See "About planning to configure I/O fencing" on page 96.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 189.

2 On the Select a task and a product page, select the task and the product as follows:

Task	I/O Fencing configuration
Product	Symantec Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O Fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

5 Fencing may be enabled, installer may prompt whether you want to reconfigure it.

Click **Yes**.

6 On the Select Fencing Type page, select the `Replace/Add/Remove coordination points` option.

7 The installer prompts to select the coordination points you want to remove from the currently configured coordination points.

Click **Next**.

8 Provide the number of Coordination point server and disk coordination points to be added to the configuration.

Click **Next**.

9 Provide the number of virtual IP addresses or Fully Qualified Host Name (FQHN) used for each coordination point server.

Click **Next**.

10 Provide the IP or FQHN and port number for each coordination point server.

Click **Next**.

11 Installer prompts to confirm the online migration coordination point servers.

Click **Yes**.

12 Installer proceeds with migration of the new coordination point servers. VCS is restarted during configuration.

Click **Next**.

13 You can add a Coordination Point agent to the client cluster and also provide name to the agent.

14 Click **Next**.

15 On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.

16 Select the check box to specify whether you want to send your installation information to Symantec.

Click **Finish**. The installer prompts you for another task.

Refreshing keys or registrations on the existing coordination points using web-based installer

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points using web-based installer

1. Start the web-based installer.

   See "Starting the web-based installer" on page 189.

2. On the **Select a task and a product** page, select the task and the product as follows:

Task	I/O Fencing configuration
Product	Symantec Cluster Server

   Click **Next**.

3. Verify the cluster information that the installer presents and click **Yes** to confirm whether you want to configure I/O fencing on the cluster.

4. On the **Select Cluster** page, enter the system name and click **Yes** to confirm cluster information.

5. On the **Select Cluster** page, click **Next** when the installer completes the cluster verification successfully.

   The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

6. The installer may prompt you to reconfigure fencing if it is already enabled. Click **Yes** to reconfigure fencing.

7 On the **Select Fencing Type** page, select the `Refresh keys/registrations on the existing coordination points` option.

8 Ensure that the `/etc/vxfenmode` file contains the same coordination point servers that are currently used by the fencing module.

9 Ensure that the disk group mentioned in the `/etc/vxfenmode` file contains the same disks that are currently used by the fencing module as coordination disks.

10 Installer lists the reasons for the loss of registrations.

Click **OK**.

11 Verify the coordination points.

Click **Yes** if the information is correct.

12 Installer updates the client cluster information on the coordination point servers.

Click **Next**.

Installer prepares the `vxfenmode` file on all nodes and runs the `vxfenswap` utility to refresh registrations on the coordination points.

13 On the **Completion** page, view the `summary` file, `log` file, or `response` file to confirm the configuration.

14 Select the check box to specify whether you want to send your installation information to Symantec.

Click **Finish**.

Setting the order of existing coordination points using the web-based installer

This section describes the reasons, benefits, considerations, and the procedure to set the order of the existing coordination points using the web-based installer.

About deciding the order of existing coordination points

You can decide the order in which coordination points can participate in a race during a network partition. In a network partition scenario, I/O fencing attempts to contact coordination points for membership arbitration based on the order that is set in the `vxfenmode` file.

When I/O fencing is not able to connect to the first coordination point in the sequence it goes to the second coordination point and so on. To avoid a cluster panic, the surviving subcluster must win majority of the coordination points. So, the order must begin with the coordination point that has the best chance to win the race and must end with the coordination point that has the least chance to win the race

For fencing configurations that use a mix of coordination point servers and coordination disks, you can either specify coordination point servers before coordination point disks or disks before servers.

Note: Disk-based fencing does not support setting the order of existing coordination points.

Considerations to decide the order of coordination points

- Choose coordination points based on their chances gain membership on the cluster during the race and hence gain control over a network partition. In effect, you have the ability to save a partition.
- First in the order must be the coordination point that has the best chance to win the race. The next coordination point you list in the order must have relatively lesser chance to win the race. Complete the order such that the last coordination point has the least chance to win the race.

Setting the order of existing coordination points using the web-based installer

To set the order of existing coordination points for server-based fencing using the web-based installer

1. Start the web-based installer.

   See "Starting the web-based installer" on page 189.

2. On the **Select a task and a product** page, select the task and the product as follows:

Task	I/O Fencing configuration
Product	Symantec Cluster Server

   Click **Next**.

3. Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.
4. On the **Select Cluster** page, enter the system name and click **Yes**.
5. On the **Select Cluster** page, click **Next** if the installer completes the cluster verification successfully.

   The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

6 The installer may prompt you to reconfigure fencing if it is already enabled. Click **Yes** to reconfigure fencing.

Click **Yes**.

7 On the **Select Fencing Type** page, select the `Set the order of existing coordination points` option.

8 Confirm **OK** at the installer message about the procedure.

9 Decide the new order by moving the existing coordination points to the box on the window in the order you want. If you want to change the current order of coordination points, click **Reset** and start again.

10 Click **Next** if the information is correct.

11 On the **Confirmation** window, click **Yes**.

Installer prepares the `vxfenmode` file on all nodes and runs the `vxfenswap` utility to update the new order of coordination points.

12 On the **Completion** page, view the summary file, log file, or response file to confirm the configuration.

13 Select the check box to specify whether you want to send your installation information to Symantec.

Click **Finish**.

Section 5

Automated installation using response files

- Chapter 12. Performing an automated VCS installation
- Chapter 13. Performing an automated VCS configuration
- Chapter 14. Performing an automated I/O fencing configuration using response files

Chapter 12

Performing an 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.
2. Make sure that the preinstallation tasks are completed.

   See "Performing preinstallation tasks" on page 66.
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 214.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to install VCS" on page 212.

5 Mount the product disc and navigate to the directory 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.

See "About the script-based installer" on page 49.

7 Complete the VCS post-installation tasks.

For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.

Response file variables to install VCS

Table 12-1 lists the response file variables that you can define to install VCS.

Table 12-1 Response file variables specific to installing VCS

Variable	List or Scalar	Description
CFG{opt}{install}	Scalar	Installs VCS packages. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{systems}	List	List of systems on which the product is to be installed, uninstalled or configured. Required
CFG{prod}	Scalar	Defines the product to be installed. The value is VCS61 for VCS. (Required)

Table 12-1 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 packages based on the variable that has the value set to 1: ■ installallpkgs: Installs all packages ■ installrecpkgs: Installs recommended packages ■ installminpkgs: Installs minimum packages **Note:** The installer requires only one of these variable values to be set to 1. (Required)
CFG{opt}{rsh}	Scalar	Defines that *rsh* must be used instead of ssh 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 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)

Table 12-1 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 packages. 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 packages 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 *logpath* 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 212.

```
#
# Configuration Values:
#
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installrecpkgs}=1;
$CFG{prod}="VCS61";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{uuid} = "16889f4e-1dd2-11b2-a559-afce02598e1b";
1;
```

Chapter 13

Performing an automated VCS configuration

This chapter includes the following topics:

- Configuring VCS using response files
- Response file variables to configure Symantec Cluster Server
- Sample response file for configuring Symantec Cluster Server

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 packages 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 Symantec Cluster Server" on page 226.

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 Symantec Cluster Server" on page 217.

4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installvcs<version>
-responsefile /tmp/response_file
```

Where *<version>* is the specific release version, and */tmp/response_file* is the response file's full path name.

See "About the script-based installer" on page 49.

Response file variables to configure Symantec Cluster Server

Table 13-1 lists the response file variables that you can define to configure VCS.

Table 13-1 Response file variables specific to configuring Symantec Cluster Server

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the packages are already installed. (Required) Set the value to 1 to configure VCS.
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
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 VCS61 for VCS. (Required)

Table 13-1 Response file variables specific to configuring Symantec Cluster Server *(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 *rsh* 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 *logpath* location. (Optional)
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec website. The value 0 indicates that the installation logs are not uploaded to the Symantec website. (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 smtprsev), the SNMP trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Table 13-2 lists the response file variables that specify the required information to configure a basic VCS cluster.

Table 13-2 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)
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{fencingenabled}	Scalar	In a VCS configuration, defines if fencing is enabled. Valid values are 0 or 1. (Required)

Table 13-3 lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 13-3 Response file variables specific to configuring private LLT over Ethernet

Variable	List or Scalar	Description
CFG{vcs_lltlink#} {"system"}	Scalar	Defines the NIC to be used for a private heartbeat link on each system. Atleast two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. See "Setting up the private network" on page 67. You must enclose the system name within double quotes. (Required)

Table 13-3 Response file variables specific to configuring private LLT over Ethernet *(continued)*

Variable	List or Scalar	Description
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. If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on. You must enclose the system name within double quotes. (Optional)

Table 13-4 lists the response file variables that specify the required information to configure LLT over UDP.

Table 13-4 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltoverudp}=1	Scalar	Indicates whether to configure heartbeat link using LLT over UDP. (Required)
CFG{vcs_udplink<n>_address} {<sys1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)

Table 13-4 Response file variables specific to configuring LLT over UDP *(continued)*

Variable	List or Scalar	Description
CFG {vcs_udplinklowpri<n>_address} {<sys1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<sys1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_port} {<sys1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_netmask} {<sys1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)

Table 13-4 Response file variables specific to configuring LLT over UDP *(continued)*

Variable	List or Scalar	Description
CFG {vcs_udplinklowpri<n>_netmask} {<sys1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 13-5 lists the response file variables that specify the required information to configure virtual IP for VCS cluster.

Table 13-5 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)
CFG{vcs_csgnetmask}	Scalar	Defines the Netmask of the virtual IP address for the cluster. (Optional)

Table 13-6 lists the response file variables that specify the required information to configure the VCS cluster in secure mode.

Table 13-6 Response file variables specific to configuring VCS cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonenode}	Scalar	Specifies that the securityonenode option is being used.
CFG{securityonenode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. ■ 1—Configure the first node ■ 2—Configure the other node
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.
CFG{vcs_eat_security_fips}	Scalar	Specifies that the enabled security is FIPS compliant.

Table 13-7 lists the response file variables that specify the required information to configure VCS users.

Table 13-7 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)

Table 13-7 Response file variables specific to configuring VCS users *(continued)*

Variable	List or Scalar	Description
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 13-8 lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 13-8 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_smtprsev}	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 13-9 lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 13-9 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)
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 13-10 lists the response file variables that specify the required information to configure VCS global clusters.

Table 13-10 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 Symantec Cluster Server

Review the response file variables and their definitions.

See "Response file variables to configure Symantec Cluster Server" on page 217.

```
#
# Configuration Values:
#
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{gco}=1;
$CFG{prod}="VCS61";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_csgnetmask}="255.255.255.0";
$CFG{vcs_csgnic}{all}="net0";
$CFG{vcs_csgvip}="10.10.12.1";
$CFG{vcs_gconetmask}="255.255.255.0";
$CFG{vcs_gcovip}="10.10.12.1";
$CFG{vcs_lltlink1}{sys1}="net1";
$CFG{vcs_lltlink1}{sys2}="net1";
$CFG{vcs_lltlink2}{sys1}="net2";
$CFG{vcs_lltlink2}{sys2}="net2";

$CFG{vcs_smtprecp}=[ qw(earnie@symantecexample.com) ];
$CFG{vcs_smtprsev}=[ qw(SevereError) ];
$CFG{vcs_smtpserver}="smtp.symantecexample.com";
$CFG{vcs_snmpcons}=[ qw(neptune) ];
$CFG{vcs_snmpcsev}=[ qw(SevereError) ];
$CFG{vcs_snmpport}=162;
1;
```

Chapter 14

Performing an automated I/O fencing configuration using response files

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
- Response file variables to configure non-SCSI-3 server-based I/O fencing
- Sample response file for configuring non-SCSI-3 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.

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 96.

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 231.

See "Sample response file for configuring server-based I/O fencing" on page 233.

4 Edit the values of the response file variables as necessary.

See "Response file variables to configure disk-based I/O fencing" on page 228.

See "Response file variables to configure server-based I/O fencing" on page 232.

5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installvcs<version>
-responsefile /tmp/response_file
```

Where *<version>* is the specific release version, and */tmp/response_file* is the response file's full path name.

See "About the script-based installer" on page 49.

Response file variables to configure disk-based I/O fencing

Table 14-1 lists the response file variables that specify the required information to configure disk-based I/O fencing for VCS.

Table 14-1 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)

Table 14-1 Response file variables specific to configuring disk-based I/O fencing *(continued)*

Variable	List or Scalar	Description
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode ■ 4—Replace/Add/Remove coordination points ■ 5—Refresh keys/registrations on the existing coordination points ■ 6—Set the order of existing coordination points (Required)
CFG {fencing_scsi3_disk_policy}	Scalar	Specifies the I/O fencing mechanism. This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the fencing_scsi3_disk_policy variable and either the fencing_dgname variable or the fencing_newdg_disks variable. (Optional)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) **Note:** You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.

Table 14-1 Response file variables specific to configuring disk-based I/O fencing *(continued)*

Variable	List or Scalar	Description
CFG{fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) **Note:** You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.
CFG{fencing_cpagent_monitor_freq}	Scalar	Specifies the frequency at which the Coordination Point Agent monitors for any changes to the Coordinator Disk Group constitution. **Note:** Coordination Point Agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidently deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles. If LevelTwoMonitorFreq attribute is not set, the agent will not monitor any changes to the Coordinator Disk Group. 0 means not to monitor the Coordinator Disk Group constitution.
CFG {fencing_config_cpagent}	Scalar	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 14-1 Response file variables specific to configuring disk-based I/O fencing *(continued)*

Variable	List or Scalar	Description
CFG {fencing_cpagentgrp}	Scalar	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the **fencing_config_cpagent** field is given a value of '0'.

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 228.

```
#
# Configuration Values:
#
our %CFG;

$CFG{fencing_config_cpagent}=1;
$CFG{fencing_cpagent_monitor_freq}=5;
$CFG{fencing_cpagentgrp}="vxfen";
$CFG{fencing_dgname}="fencingdg1";
$CFG{fencing_newdg_disks}=[ qw(emc_clariion0_155
 emc_clariion0_162 emc_clariion0_163) ];
$CFG{fencing_option}=2;
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{fencing_cpagent_monitor_freq}=5;

$CFG{prod}="VCS61";

$CFG{systems}=[ qw(pilot25) ];
$CFG{vcs_clusterid}=32283;
$CFG{vcs_clustername}="clus1";
1;
```

Response file variables to configure server-based I/O fencing

You can use a coordination point server-based fencing response file to configure server-based customized I/O fencing.

Table 14-2 lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table 14-2 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_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.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the `fencing_config_cpagent` field is given a value of '0'.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_reusedg}	This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks). Enter either a "1" or "0". Entering a "1" indicates reuse, and entering a "0" indicates do not reuse. When reusing an existing DG name for the mixed mode fencing configuration. you need to manually add a line of text , such as "$CFG{fencing_reusedg}=0" or "$CFG{fencing_reusedg}=1" before proceeding with a silent installation.

Table 14-2 Coordination point server (CP server) based fencing response file definitions *(continued)*

Response file field	Definition
CFG {fencing_dgname}	The name of the disk group to be used in the customized fencing, where at least one disk is being used.
CFG {fencing_disks}	The disks being used as coordination points if any.
CFG {fencing_ncp}	Total number of coordination points being used, including both CP servers and disks.
CFG {fencing_ndisks}	The number of disks being used.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_cps_ports}	The port that the virtual IP address or the fully qualified host name of the CP server listens on.
CFG {fencing_scsi3_disk_policy}	The disk policy that the customized fencing uses. The value for this field is either "raw" or "dmp"
CFG{fencing_option}	Specifies the I/O fencing configuration mode. ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode ■ 4—Replace/Add/Remove coordination points ■ 5—Refresh keys/registrations on the existing coordination points ■ 6—Set the order of existing coordination points

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_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cps_vips}{"10.200.117.145"}=[ qw(10.200.117.145) ];
$CFG{fencing_dgname}="vxfencoorddg";
$CFG{fencing_disks}=[ qw(emc_clariion0_37 emc_clariion0_13) ];
$CFG{fencing_scsi3_disk_policy}="raw";
```

```
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=2;
$CFG{fencing_cps_ports}{"10.200.117.145"}=443;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="VCS61";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Response file variables to configure non-SCSI-3 server-based I/O fencing

Table 14-3 lists the fields in the response file that are relevant for non-SCSI-3 server-based customized I/O fencing.

See "About I/O fencing for VCS in virtual machines that do not support SCSI-3 PR" on page 32.

Table 14-3 Non-SCSI-3 server-based I/O fencing response file definitions

Response file field	Definition
CFG{non_scsi3_fencing}	Defines whether to configure non-SCSI-3 server-based I/O fencing. Valid values are 1 or 0. Enter 1 to configure non-SCSI-3 server-based I/O fencing.
CFG {fencing_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.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the `fencing_config_cpagent` field is given a value of '0'.

Table 14-3 Non-SCSI-3 server-based I/O fencing response file definitions *(continued)*

Response file field	Definition
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_ncp}	Total number of coordination points (CP servers only) being used.
CFG {fencing_cps_ports}	The port of the CP server that is denoted by *cps* .

Sample response file for configuring non-SCSI-3 server-based I/O fencing

The following is a sample response file used for non-SCSI-3 server-based I/O fencing :

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.198.89.251 10.198.89.252 10.198.89.253) ];
$CFG{fencing_cps_vips}{"10.198.89.251"}=[ qw(10.198.89.251) ];
$CFG{fencing_cps_vips}{"10.198.89.252"}=[ qw(10.198.89.252) ];
$CFG{fencing_cps_vips}{"10.198.89.253"}=[ qw(10.198.89.253) ];
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=0;
$CFG{fencing_cps_ports}{"10.198.89.251"}=443;
$CFG{fencing_cps_ports}{"10.198.89.252"}=443;
$CFG{fencing_cps_ports}{"10.198.89.253"}=443;
$CFG{non_scsi3_fencing}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="VCS61";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Section 6

Manual installation

- Chapter 15. Performing preinstallation tasks
- Chapter 16. Manually installing VCS
- Chapter 17. Manually configuring VCS
- Chapter 18. Manually configuring the clusters for data integrity

Chapter 15

Performing preinstallation tasks

This chapter includes the following topics:

- Requirements for installing VCS

Requirements for installing VCS

Review requirements before you install.

See “Important preinstallation information for VCS” on page 37.

Chapter 16

Manually installing VCS

This chapter includes the following topics:

- About VCS manual installation
- Installing VCS software manually
- Installing VCS on Solaris 10 using JumpStart
- Installing VCS on Solaris 11 using Automated Installer

About VCS manual installation

You can manually install and configure VCS instead of using the installvcs.

A manual installation takes a lot of time, patience, and care. Symantec recommends that you use the installvcs instead of the manual installation when possible.

Installing VCS software manually

If you manually install VCS software to upgrade your cluster, make sure to back up the previous VCS configuration files before you start the installation. The configuration files that you must back up are as follows:

- All the files from `/etc/VRTSvcs/conf/config` directory.
- /etc/llttab
- /etc/gabtab
- /etc/llthosts
- /etc/amftab
- /etc/default/vcs
- /etc/default/amf

- /etc/default/vxfen
- /etc/default/gab
- /etc/default/llt

Table 16-1 lists the tasks that you must perform when you manually install and configure VCS 6.1.

Table 16-1 Manual installation tasks for VCS 6.1

Task	Reference
Install VCS software manually on each node in the cluster.	See "Installing VCS packages for a manual installation" on page 240.
Install VCS language pack software manually on each node in the cluster.	See "Installing language packages in a manual installation" on page 244.
Add a license key.	See "Adding a license key for a manual installation" on page 245.
Copy the installation guide to each node.	See "Copying the installation guide to each node" on page 247.
Configure LLT and GAB.	■ See "Configuring LLT manually" on page 259. ■ See "Configuring GAB manually" on page 263.
Configure VCS.	See "Configuring VCS manually" on page 263.
Start LLT, GAB, and VCS services.	See "Starting LLT, GAB, and VCS after manual configuration" on page 269.
Modify the VCS configuration.	See "Modifying the VCS configuration" on page 270.
Replace demo license with a permanent license.	See "Replacing a VCS demo license with a permanent license for manual installations" on page 247.

Viewing the list of VCS packages

During the VCS installation, the installer prompts you with an option to choose the VCS packages to install. You can view the list of packages 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 packages in a specified order. For example, you must install some packages before other packages because of various product dependencies. The following installer

command options list the packages in the order in which you must install these packages.

Table 16-2 describes the VCS package installation options and the corresponding command to view the list of packages.

Table 16-2 Installer command options to view VCS packages

Option	Description	Command option to view the list of packages
1	Installs only the minimal required VCS packages that provide basic functionality of the product.	`installvcs -minpkgs`
2	Installs the recommended VCS packages that provide complete functionality of the product. This option does not install the optional VCS packages.	`installvcs -recpkgs`
3	Installs all the VCS packages. You must choose this option to configure any optional VCS feature.	`installvcs -allpkgs`

To view the list of VCS packages

1 Navigate to the directory from where you can start the installvcs.

```
# cd cluster_server
```

2 Run the following command to view the list of packages. Based on what packages you want to install, enter the appropriate command option:

```
# ./installvcs -minpkgs
```

Or

```
# ./installvcs -recpkgs
```

Or

```
# ./installvcs -allpkgs
```

Installing VCS packages for a manual installation

All packages are installed into the `/opt` directory and a few files are installed into the `/etc` and `/var` directories.

You can create lists of the packages to install.

See "Viewing the list of VCS packages" on page 239.

If you copied the Symantec packages to `/tmp/install`, navigate to the directory and perform the following on each system:

To install VCS packages on a node

- Install the following required packages on a Solaris 10 node in the order shown:

```
# pkgadd -d VRTSperl.pkg
# pkgadd -d VRTSvlic.pkg
# pkgadd -d VRTSspt.pkg
# pkgadd -d VRTSllt.pkg
# pkgadd -d VRTSgab.pkg
# pkgadd -d VRTSvxfen.pkg
# pkgadd -d VRTSamf.pkg
# pkgadd -d VRTSvcs.pkg
# pkgadd -d VRTScps.pkg
# pkgadd -d VRTSvcsag.pkg
# pkgadd -d VRTSvcsea.pkg
# pkgadd -d VRTSsfmh.pkg
# pkgadd -d VRTSvbs.pkg
# pkgadd -d VRTSvcswiz.pkg
# pkgadd -d VRTSsfcpi6.1.pkg
```

Note: To configure an Oracle VM Server logical domain for disaster recovery, install the following required package inside the logical domain:

```
# pkgadd -d VRTSvcsnr.pkg
```

See "Symantec Cluster Server installation packages" on page 481.

Manually installing packages on Oracle Solaris 11 systems

To install packages on Solaris 11 system

1. Copy the VRTSpkgs.p5p package from the pkgs directory from the installation media to the system at `/tmp/install` directory.
2. Disable the publishers that are not reachable as package install may fail if any of the already added repositories are unreachable.

```
# pkg set-publisher --disable <publisher name>
```

3 Add a file-based repository in the system.

```
# pkg set-publisher -g /tmp/install/VRTSpkgs.p5p Symantec
```

4 Install the required packages.

```
# # pkg install --accept VRTSperl VRTSvlic VRTSspt
VRTSllt VRTSgab VRTSvxfenVRTSamf VRTSvcs VRTScps VRTSvcsag VRTSvcsea VRTS
VRTSvcswiz VRTSsfcpi61
```

5 To configure an OracleVMServer logical domain for disaster recovery, install the following required package inside the logical domain:

```
# pkg install --accept VRTSvcsnr
```

6 Remove the publisher from the system.

```
# pkg unset-publisher Symantec
```

7 Clear the state of the SMF service if non-global zones are present in the system. In presence of non-global zones, setting the file-based repository causes SMF service `svc:/application/pkg/system-repository:default` to go into maintenance state.

```
# svcadm clear svc:/application/pkg/system-repository:default
```

8 Enable the publishers that were disabled earlier.

```
# pkg set-publisher --enable <publisher name>
```

Manually installing packages on Solaris brand non-global zones

With Oracle Solaris 11, you must manually install VCS packages inside non-global zones. The native non-global zones are called Solaris brand zones.

To install packages manually on Solaris brand non-global zones:

1 Ensure that the SMF service svc:/application/pkg/system-repository:default and svc:/application/pkg/zones-proxyd:default is online on the global zone.

```
# svcs svc:/application/pkg/system-repository:default
# svcs svc:/application/pkg/zones-proxyd:default
```

2 Log on to the non-global zone as a superuser.

3 Ensure that the SMF service svc:/application/pkg/zones-proxy-client:default is online inside non-global zone:

```
# svcs svc:/application/pkg/zones-proxy-client:default
```

4 Copy the VRTSpkgs.p5p package from the pkgs directory from the installation media to the non-global zone (for example at `/tmp/install` directory).

5 Disable the publishers that are not reachable, as package install may fail if any of the already added repositories are unreachable.

```
#pkg set-publisher --disable <publisher name>
```

6 Add a file-based repository in the non-global zone.

```
# pkg set-publisher -g/tmp/install/VRTSpkgs.p5p Symantec
```

7 Install the required packages.

```
# pkg install --accept VRTSperl VRTSvlic VRTSvcs
VRTSvcsag  VRTSvcsea
```

8 Remove the publisher on the non-global zone.

```
#pkg unset-publisher Symantec
```

9 Clear the state of the SMF service, as setting the file-based repository causes SMF service svc:/application/pkg/system-repository:default to go into maintenance state.

```
# svcadm clear svc:/application/pkg/system-repository:default
```

10 Enable the publishers that were disabled earlier.

```
# pkg set-publisher --enable <publisher>
```

Note: Perform steps 2 through 10 on each non-global zone.

Manually installing packages on solaris10 brand zones

You need to manually install VCS 6.1 packages inside the solaris10 brand zones.

1 Boot the zone.

2 Logon to the solaris10 brand zone as a super user.

3 Copy the Solaris 10 packages from the pkgs directory from the installation media to the non-global zone (such as `/tmp/install` directory).

4 Install the following VCS packages on the brand zone.

```
# cd /tmp/install
# pkgadd -d VRTSperl.pkg
# pkgadd -d VRTSvlic.pkg
# pkgadd -d VRTSvcs.pkg
# pkgadd -d VRTSvcsag.pkg
# pkgadd -d VRTSvcsea.pkg
```

Note: Perform all the above steps on each Solaris 10 brand zone.

For more information on the support for Branded Zones, refer the *Symantec Storage Foundation and High Availability Solutions Virtualization Guide*.

Installing language packages in a manual installation

Install the language packages that VCS requires after you install the base VCS packages.

See "Symantec Cluster Server installation packages" on page 481.

Before you install, make sure that you are logged on as superuser and that you have mounted the language disc.

See "Mounting the product disc" on page 80.

Perform the steps on each node in the cluster to install the language packages.

To install the language packages on a Solaris 10 node

1 Copy the package files from the software disc to the temporary directory.

```
# cp -r pkgs/* /tmp
```

2 Install the following required and optional VCS packages from the compressed files:

- Install the following required packages in the order shown for Japanese language support:

```
# pkgadd -d VRTSjacse.pkg
# pkgadd -d VRTSjacs.pkg
```

To install the language packages on a Solaris 11 node:

1. Copy the VRTSpkgs.p5p package from the pkgs directory from the installation media to the system at `/tmp/install` directory.
2. Add a file-based repository in the system.

```
# pkg set-publisher -p /tmp/install/VRTSpkgs.p5p Symantec
```

3. Install the following required packages in the order shown for Japanese language support:

```
# pkg install --accept VRTSjacse
# pkg install --accept VRTSjacs
```

Adding a license key for a manual installation

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

```
# 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 Symantec products and functionality that are licensed.

For more information to use keyless licensing 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.

When you upgrade from a previous release, the product installer prompts you to update the `vxkeyless` license product level to the current release level. If you update the `vxkeyless` license product level during the upgrade process, no further action is required. If you do not update the `vxkeyless` license product level, the output you see when you run the `vxkeyless display` command includes the previous release's `vxkeyless` license product level. Each `vxkeyless` license product level name includes the suffix *_previous_release_version*. For example, DMP_6.0, or SFENT_VR_5.1SP1, or VCS_GCO_5.1. If there is no suffix, it is the current release version.

You would see the suffix *_previous_release_version* if you did not update the `vxkeyless` product level when prompted by the product installer. Symantec highly recommends that you always use the current release version of the product levels. To do so, use the `vxkeyless set` command with the desired product levels. If you see SFENT_60, VCS_60, use the `vxkeyless set SFENT,VCS` command to update the product levels to the current release.

After you install or upgrade, 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 Change your current working directory:

```
# export PATH=$PATH:/opt/VRTSvlic/bin
```

2 View the current setting for the product level.

```
# vxkeyless -v display
```

3 View the possible settings for the product level.

```
# vxkeyless displayall
```

4 Set the desired product level.

```
# vxkeyless set prod_levels
```

where *prod_levels* is a comma-separated list of keywords. The keywords are the product levels as shown by the output of step 3.

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 Symantec products until you install a new key or set a new product level.

See "Installing Symantec product license keys" on page 63.

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 Symantec licenses on a system. For example, enter:

```
# 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 156.

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 to the /opt/VRTS/docs directory on each node to make it available for reference. The PDF is located at `docs/cluster_server/vcs_install_`*`version_platform`*`.pdf` where *version* is the release version and *platform* is the name of the operating system.

Installing VCS on Solaris 10 using JumpStart

This installation method applies only to Solaris 10. These JumpStart instructions assume a working knowledge of JumpStart. See the JumpStart documentation that came with your operating system for details on using JumpStart.

Upgrading is not supported. The following procedure assumes a standalone configuration.

For the language pack, you can use JumpStart to install packages. You add the language packages in the script, and put those files in the JumpStart server directory.

You can use a Flash archive to install VCS and the operating system with JumpStart.

See "Using a Flash archive to install VCS and the operating system" on page 251.

Overview of JumpStart installation tasks

Review the summary of tasks before you perform the JumpStart installation.

Summary of tasks

1. Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.
2. Read the JumpStart installation instructions.
3. Generate the finish scripts.

   See "Generating the finish scripts" on page 248.
4. Prepare shared storage installation resources.

   See "Preparing installation resources" on page 249.
5. Modify the rules file for JumpStart.

   See the JumpStart documentation that came with your operating system for details.
6. Install the operating system using the JumpStart server.
7. When the system is up and running, run the installer command from the installation media to configure the Symantec software.

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

   See "About the script-based installer" on page 49.

Generating the finish scripts

Perform these steps to generate the finish scripts to install VCS.

To generate the script

1 Run the product installer program to generate the scripts for all products.

```
./installer -jumpstart directory_to_generate_scripts
```

Or

```
./install<productname> -jumpstart directory_to_generate_script
```

where **<*productname*>** is the product's installation command, and *directory_to_generate_scripts* is where you want to put the product's script.

For example:

```
# ./installvcs -jumpstart /js_scripts
```

2 Modify the JumpStart script according to your requirements. You must modify the *BUILDSRC* and *ENCAPSRC* values. Keep the values aligned with the resource location values.

```
BUILDSRC="hostname_or_ip:/path_to_pkgs"
// If you don't want to encapsulate the root disk automatically
// comment out the following line.
ENCAPSRC="hostname_or_ip:/path_to_encap_script"
```

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1 Copy the `pkgs` directory of the installation media to the shared storage.

```
# cd /path_to_installation_media
# cp -r pkgs BUILDSRC
```

2 Generate the response file with the list of packages.

```
# cd BUILDSRC/pkgs/
# pkgask -r package_name.response -d /
BUILDSRC/pkgs/packages_name.pkg
```

3 Create the `adminfile` file under *BUILDSRC*`/pkgs/` directory.

```
mail=
instance=overwrite
partial=nocheck
runlevel=quit
idepend=quit
rdepend=nocheck
space=quit
setuid=nocheck
conflict=nocheck
action=nocheck
basedir=default
```

Adding language pack information to the finish file

To add the language pack information to the finish file, perform the following procedure.

To add the language pack information to the finish file

1 For the language pack, copy the language packages from the language pack installation disc to the shared storage.

```
# cd /cdrom/cdrom0/pkgs
# cp -r * BUILDSRC/pkgs
```

If you downloaded the language pack:

```
# cd /path_to_language_pack_installation_media/pkgs
# cp -r * BUILDSRC/pkgs
```

2 In the finish script, copy the product package information and replace the product packages with language packages.

3 The finish script resembles:

```
. . .
for PKG in product_packages
do
...
done. . .
for PKG in language_packages
do
...
done. . .
```

Using a Flash archive to install VCS and the operating system

You can only use Flash archive on the Solaris 10 operating system. In the following outline, refer to Solaris documentation for Solaris-specific tasks.

Note: Symantec does not support Flash Archive installation if the root disk of the master system is encapsulated.

The following is an overview of the creation and installation of a Flash archive with Symantec software.

- If you plan to start flar (flash archive) creation from bare metal, perform step 1 through step 10.
- If you plan to start flar creation from a system where you have installed, but not configured the product, perform step 1 through step 4. Skip step 5 and finish step 6 through step 10.

- If you plan to start flar creation from a system where you have installed and configured the product, perform step 5 through step 10.

Flash archive creation overview

1. Ensure that you have installed Solaris 10 on the master system.
2. Use JumpStart to create a clone of a system.
3. Restart the cloned system.
4. Install the Symantec products on the master system.

   Perform one of the installation procedures from this guide.
5. If you have configured the product on the master system, create the `vrts_deployment.sh` file and the `vrts_deployment.cf` file and copy them to the master system.

   See "Creating the Symantec post-deployment scripts" on page 252.
6. Use the `flarcreate` command to create the Flash archive on the master system.
7. Copy the archive back to the JumpStart server.
8. Use JumpStart to install the Flash archive to the selected systems.
9. Configure the Symantec product on all nodes in the cluster.

   The scripts that are installed on the system include the product version in the script name. For example, to install the SF script from the install media, run the `installsf` command. However, to run the script from the installed binaries, run the `installsf<version>` command. For example, for the 6.1 version:

   ```
   # /opt/VRTS/install/installvcs61 -configure
   ```

   See "About the script-based installer" on page 49.
10. Perform post-installation and configuration tasks.

    See the product installation guide for the post-installation and configuration tasks.

Creating the Symantec post-deployment scripts

The generated files vrts_deployment.sh and vrts_post-deployment.cf are customized Flash archive post-deployment scripts. These files clean up Symantec product settings on a cloned system before you reboot it for the first time. Include these files in your Flash archives.

To create the post-deployment scripts

1. Mount the product disc.
2. From the prompt, run the `-flash_archive` option for the installer. Specify a directory where you want to create the files.

   ```
   # ./installer -flash_archive /tmp
   ```

3. Copy the vrts_postedeployment.sh file and the vrts_postedeployment.cf file to the golden system.
4. On the golden system perform the following:
   - Put the vrts_postdeployment.sh file in the /etc/flash/postdeployment directory.
   - Put the vrts_postdeployment.cf file in the /etc/vx directory.
5. Make sure that the two files have the following ownership and permissions:

   ```
   # chown root:root /etc/flash/postdeployment/vrts_postdeployment.sh
   # chmod 755 /etc/flash/postdeployment/vrts_postdeployment.sh
   # chown root:root /etc/vx/vrts_postdeployment.cf
   # chmod 644 /etc/vx/vrts_postdeployment.cf
   ```

   Note that you only need these files in a Flash archive where you have installed Symantec products.

Installing VCS on Solaris 11 using Automated Installer

You can use the Oracle Solaris Automated Installer (AI) to install the Solaris 11 operating system and Storage Foundation product on multiple client systems in a network. AI performs a hands-free installation (automated installation without manual interactions) of SPARC systems. You can also use AI media to install the Oracle Solaris OS on a single SPARC platform. Oracle provides the AI bootable image and it can be downloaded from the Oracle website. All cases require access to a package repository on the network to complete the installation.

About Automated Installation

AI automates the installation of the Oracle Solaris 11 OS on one or more SPARC clients in a network. Automated Installation applies to Solaris 11 only. You can install the Oracle Solaris OS on many different types of clients. The clients can differ in:

- Architecture

- Memory characteristics
- MAC address
- IP address
- CPU

The installations can differ depending on specifications including network configuration and packages installed.

An automated installation of a client in a local network consists of the following high-level steps:

1. A client system boots and gets IP information from the DHCP server
2. Characteristics of the client determine which AI service and which installation instructions are used to install the client.
3. The installer uses the AI service instructions to pull the correct packages from the package repositories and install the Oracle Solaris OS on the client.

Using Automated Installer

To use Automated Installer to install systems over the network, set up DHCP and set up an AI service on an AI server. The DHCP server and AI server can be the same system or two different systems.

Make sure that the systems can access an Oracle Solaris Image Packaging System (IPS) package repository. The IPS package repository can reside on the AI server, on another server on the local network, or on the Internet.

An AI service is associated with a SPARC AI install image and one or more sets of installation instructions. The installation instructions specify one or more IPS package repositories from where the system retrieves the packages that are needed to complete the installation. The installation instructions also include the names of additional packages to install and information such as target device and partition information. You can also specify instructions for post-installation configuration of the system.

Consider the operating systems and packages you want to install on the systems. Depending on your configuration and needs, you may want to do one of the following:

- If two systems have different architectures or need to be installed with different versions of the Oracle Solaris OS, create two AI services. Then, associate each AI service with a different AI image
- If two systems need to be installed with the same version of the Oracle Solaris OS but need to be installed differently in other ways, create two sets of installation

instructions for the AI service. The different installation instructions can specify different packages to install or a different slice as the install target.

The installation begins when you boot the system. DHCP directs the system to the AI install server, and the system accesses the install service and the installation instructions within that service.

For more information, see the *Oracle® Solaris 11 Express Automated Installer Guide*.

Using AI to install the Solaris 11 operating system and SFHA products

Use the following procedure to install the Solaris 11 operating system and SFHA products using AI.

To use AI to install the Solaris 11 operating system and SFHA products

1. Follow the Oracle documentation to set up a Solaris AI server and DHCP server.

   You can find the documentation at http://docs.oracle.com.

2. Set up the Symantec package repository.

   Run the following commands to startup necessary SMF services and create directories:

   ```
   # svcadm enable svc:/network/dns/multicast:default
   # mkdir /ai
   # zfs create -o compression=on -o mountpoint=/ai rpool/ai
   ```

3 Run the following commands to set up IPS repository for Symantec SPARC packages:

```
# mkdir -p /ai/repo_symc_sparc
# pkgrepo create /ai/repo_symc_sparc
# pkgrepo add-publisher -s /ai/repo_symc_sparc Symantec
# pkgrecv -s <media_sparc>/pkgs/VRTSpkgs.p5p -d
/ai/repo_symc_sparc '*'
# svccfg -s pkg/server list
# svcs -a | grep pkg/server
# svccfg -s pkg/server add symcsparc
# svccfg -s pkg/server:symcsparc addpg pkg application
# svccfg -s pkg/server:symcsparc setprop pkg/port=10003
# svccfg -s pkg/server:symcsparc setprop pkg/inst_root=
/ai/repo_symc_sparc
# svccfg -s pkg/server:symcsparc addpg general framework
# svccfg -s pkg/server:symcsparc addpropvalue general/complete
astring: symcsparc
# svccfg -s pkg/server:symcsparc addpropvalue general/enable
boolean: true
# svcs -a | grep pkg/server
# svcadm refresh application/pkg/server:symcsparc
# svcadm enable application/pkg/server:symcsparc
```

Or run the following commands to set up the private depot server for testing purposes:

```
# /usr/lib/pkg.depotd -d /ai/repo_symc_sparc -p 10003 > /dev/null &
```

Check the following URL on IE or Firefox browser:

http://<host>:10003

4 Set up the install service on the AI server.

Run the following command:

```
# mkdir /ai/iso
```

Download the AI image from the Oracle website and place the `iso` in the `/ai/iso` directory.

Create an install service.

For example:

To set up the AI install service for SPARC platform::

```
# # installadm create-service -n sol11sparc -s\
 /ai/iso/sol-11-1111-ai-sparc.iso -d /ai/aiboot/
```

5 Run the installer to generate manifest XML files for all the SFHA products that you plan to install.

```
# mkdir /ai/manifests
# <media>/installer -ai /ai/manifests
```

6 For each system, generate the system configuration and include the host name, user accounts, and IP addresses. For example, enter one of the following:

```
# mkdir /ai/profiles
# sysconfig create-profile -o /ai/profiles/profile_client.xml
```

or

```
# cp /ai/aiboot/auto-install/sc_profiles/sc_sample.xml
/ai/profiles/profile_client.xml
```

7 Add a system and match it to the specified product manifest and system configuration.

Run the following command to add a SPARC system, for example:

```
# installadm create-client -e "<client_MAC>" -n sol11sparc
# installadm add-manifest -n sol11sparc -f \
/ai/manifests/vrts_manifest_sfha.xml
# installadm create-profile -n sol11sparc -f \
/ai/profiles/profile_client.xml -p profile_sc
# installadm set-criteria -n sol11sparc -m \
vrts_sfha -p profile_sc -c mac="<client_MAC>"
# installadm list -m -c -p -n sol11sparc
```

8 For SPARC system, run the following command to restart the system and install the operating system and Storage Foundation products:

```
# boot net:dhcp - install
```

Chapter 17

Manually configuring VCS

This chapter includes the following topics:

- About configuring VCS manually
- Configuring LLT manually
- Configuring GAB manually
- Configuring VCS manually
- Configuring VCS in single node mode
- Starting LLT, GAB, and VCS after manual configuration
- Modifying the VCS configuration

About configuring VCS manually

This section describes the procedures to manually configure VCS.

Note: For manually configuring VCS in single node mode, you can skip steps about configuring LLT manually and configuring GAB manually.

Configuring LLT manually

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 over Ethernet, 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 260.
- Set up the file /etc/lltab.
  See "Setting up /etc/lltab for a manual installation" on page 260.
- 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/default/llt

You can also configure LLT over UDP.

See "Using the UDP layer for LLT" on page 549.

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 sys1
1 sys2
```

Setting up /etc/lltab for a manual installation

The /etc/lltab 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 lltab file in /opt/VRTSllt.

See "About LLT directives in /etc/lltab file" on page 261.

Run the `dladm show-dev` command to query all NICs.

Use vi or another editor to create the file /etc/lltab that contains the entries that resemble the following:

- For SPARC:

```
set-node sys1
set-cluster 2
```

```
link net1 net:0 - ether - -
link net2 net:1 - ether - -
```

The first line must identify the system where the file exists. In the example, the value for `set-node` can be: sys1, 0, or the file name /etc/nodename. The file needs to contain the name of the system (sys1 in this example). The next line, beginning with the `set-cluster` command, identifies the cluster number, which must be a unique number when more than one cluster is configured on the same physical network connection.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 `lltab` file in `/opt/VRTSllt`.

If you use different media speed for the private NICs, Symantec recommends that you configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example:

Use vi or another editor to create the file /etc/lltab that contains the entries that resemble the following:

- For SPARC:

```
set-node sys1
set-cluster 2
link net1 net:0 - ether - -
link net2 net:1 - ether - -
link-lowpri qfe2 qfe:2 - ether - -
```

See “Setting up the private network” on page 67.

About LLT directives in /etc/lltab file

Table 17-1 lists the LLT directives in /etc/lltab file for LLT over Ethernet.

Table 17-1 LLT directives

Directive	Description
`set-node`	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 0-63. The symbolic name corresponds to the system ID, which is in /etc/llthosts file. Note that LLT fails to operate if any systems share the same ID.
`set-cluster`	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.

Table 17-1 LLT directives *(continued)*

Directive	Description
`link`	Attaches LLT to a network interface. At least one link is required, and up to eight are supported. LLT distributes network traffic evenly across all available network connections unless you mark the link as low-priority using the link-lowpri directive or you configured LLT to use destination-based load balancing. The first argument to link is a user-defined tag shown in the `lltstat(1M)` output to identify the link. It may also be used in `llttab` to set optional static MAC addresses. The second argument to link is the device name of the network interface. Its format is device_name:device_instance_number. The remaining four arguments to link 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 `llttab(4)` 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 in LLT over Ethernet mode.
`link-lowpri`	Use this directive in place of `link` 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. If you use private NICs with different speed, use "link-lowpri" directive in place of "link" for all links with lower speed. Use the "link" directive only for the private NIC with higher speed to enhance LLT performance. LLT uses low-priority network links for VCS communication only when other links fail.

For more information about the 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.

By default, Oracle systems assign the same MAC address to all interfaces. Thus, connecting two or more interfaces to a network switch can cause problems. Consider the following example. You configure an IP on one public interface and LLT on another. Both interfaces are connected to a switch. The duplicate MAC address on

the two switch ports can cause the switch to incorrectly redirect IP traffic to the LLT interface and vice versa. To avoid this issue, configure the system to assign unique MAC addresses by setting the eeprom(1M) parameter local-mac-address? to true.

Configuring GAB manually

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 `-n`N 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 or `-x` option for `/sbin/gabconfig`. Using `-c -x` or `-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/default/gab

Configuring VCS manually

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 resource type definitions.
- 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 resource type definitions. 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 *Symantec 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 sys1 ( )
system sys2 ( )
```

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

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

3 Save and close the main.cf file.

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.

Configuring VCS in single node mode

In addition to the steps mentioned in the manual configuration section, complete the following steps to configure VCS in single node mode.

See "Configuring VCS manually" on page 263.

To configure VCS in single node mode

1. Disable the VCS SMF service imported by VRTSvcs package.

```
# svcadm disable -s system/vcs:default
```

2. Delete the VCS SMF service configuration.

```
# svccfg delete -f  system/vcs:default
```

3. Edit the following file to change the value of the ONENODE environment variable to **yes**.

```
/etc/default/vcs
```

4. Import the SMF service for vcs-onenode.

```
# svccfg import /etc/VRTSvcs/conf/vcs-onenode.xml
```

5. If the single node is intended only to manage applications, you can disable LLT, GAB, I/O fencing kernel modules.

Note: Disabling VCS kernel modules means that you cannot make the applications highly available across multiple nodes.

See "Disabling LLT, GAB, and I/O fencing on a single node cluster" on page 266.

See "Enabling LLT, GAB, and I/O fencing on a single node cluster" on page 267.

Disabling LLT, GAB, and I/O fencing on a single node cluster

This section discusses how to disable kernel modules on a single node VCS cluster.

Typically, LLT, GAB, and I/O fencing kernel modules are loaded on a node when you install VCS. However, you can disable LLT, GAB, and I/O fencing modules if you do not require high availability for the applications. You can continue to manage applications on the single node and use the application restart capabilities of VCS.

If you later decide to extend the cluster to multiple nodes, you can enable these modules and make the applications highly available across multiple nodes.

Note: If VCS engine hangs on the single node cluster with GAB disabled, GAB cannot detect the hang state and cannot take action to restart VCS. For such a condition, you need to detect that VCS engine has hung and take corrective action. For more information, refer to the 'About GAB client process failure' section in the *Symantec Cluster Server Administrator's Guide*.

See "Disabling LLT, GAB, and I/O fencing on Oracle Solaris 10 and 11" on page 266.

Disabling LLT, GAB, and I/O fencing on Oracle Solaris 10 and 11

Complete the following procedures to disable the kernel modules.

To disable I/O fencing

1 Edit the following file to set the value of `VXFEN_START` and `VXFEN_STOP` to **0**.

```
/etc/default/vxfen
```

2 Run the following commands.

```
# svcadm disable -s system/vxfen
# svccfg delete -f system/vxfen
# rem_drv vxfen
# modinfo | grep -w vxfen
# modunload -i <mod id>
```

Where, *mod id* is the module ID

To disable GAB

1 Edit the following file to set the value of `GAB_START` and `GAB_STOP` to **0**.

```
/etc/default/gab
```

2 Run the following commands.

```
# svcadm disable -s system/gab
# svccfg delete -f system/gab
# rem_drv gab
# modinfo | grep -w gab
# modunload -i <mod id>
```

Where, *mod id* is the module ID.

To disable LLT

1 Edit the following file to set the value of `LLT_START` and `LLT_STOP` to **0**.

```
/etc/default/llt
```

2 Run the following commands.

```
# svcadm disable -s system/llt
# svccfg delete -f system/llt
# rem_drv llt
# modinfo | grep -w llt
# modunload -i <mod id>
```

Where, *mod id* is the module ID.

Enabling LLT, GAB, and I/O fencing on a single node cluster

This section provides the steps to enable the kernel modules on a single node cluster.

Enabling LLT, GAB, and I/O fencing on Solaris 11

Complete the following procedures to enable the kernel modules.

To enable LLT:

1 In the `/etc/default/llt` file, ensure LLT_START=1.

2 Run the following commands.

```
# /usr/sbin/add_drv -v -f -m '* 0600 root sys' llt
# svccfg -s system/llt delcust
# svccfg disable system/llt
# svccfg enable system/llt
```

To enable GAB:

1 In the `/etc/default/gab` file, ensure GAB_START=1.

2 Run the following commands:

```
# /usr/sbin/add_drv -v -f -m '* 0600 root sys' gab
# svccfg -s system/gab delcust
# svccfg disable system/gab
# svccfg enable system/gab
```

To enable I/O fencing:

1 In the `/etc/default/vxfen` file, ensure VXFEN_START=1.

2 Run the following commands:

```
# /usr/sbin/add_drv -v -f -m '* 0600 root sys' vxfen
# svccfg -s system/vxfen delcust
# svccfg disable system/vxfen
# svccfg enable system/vxfen
```

3 Reboot the nodes.

Enabling LLT, GAB, and I/O fencing on Solaris 10

Complete the following procedures to enable the kernel modules.

To enable LLT:

1 In the `/etc/default/llt` file, ensure LLT_START=1.

2 Run the following commands.

```
#/usr/sbin/add_drv -v -f -m '* 0600 root sys' llt
#/usr/sbin/svccfg  import /var/svc/manifest/system/llt.xml
#/usr/sbin/svcadm enable -s system/llt
```

To enable GAB:

1 In the `/etc/default/gab` file, ensure GAB_START=1.

2 Run the following commands:

```
#/usr/sbin/add_drv -v -f -m '* 0600 root sys' gab
#/usr/sbin/svccfg  import /var/svc/manifest/system/gab.xml
#/usr/sbin/svcadm enable -s system/gab
```

To enable I/O fencing:

1 In the `/etc/default/vxfen` file, ensure VXFEN_START=1.

2 Run the following commands:

```
#/usr/sbin/add_drv -v -f -m '* 0600 root sys' vxfen
#/usr/sbin/svccfg  import /var/svc/manifest/system/vxfen.xml
#/usr/sbin/svcadm enable -s system/vxfen
```

3 Reboot the nodes.

Starting LLT, GAB, and VCS after manual configuration

After you have configured LLT, GAB, and VCS, use the following procedures to start LLT, GAB, and VCS.

To start LLT

1 On each node, run the following command to start LLT:

```
# svcadm enable llt
```

If LLT is configured correctly on each node, the console output resembles:

```
Jun 26 19:04:24 sys1 kernel: [1571667.550527] LLT INFO V-14-1-10009 LLT 6
```

2 On each node, run the following command to verify that LLT is running:

```
# /sbin/lltconfig
LLT is running
```

To start GAB

1 On each node, run the following command to start GAB:

```
# svcadm enable gab
```

If GAB is configured correctly on each node, the console output resembles:

```
Jun 26 19:10:34 sys1 kernel: [1572037.501731] GAB INFO
V-15-1-20021 GAB 6.0.100.000-SBLD available
```

2 On each node, run the following command to verify that GAB is running:

```
# /sbin/gabconfig -a
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
```

To start VCS

- On each node, type:

```
# svcadm enable vcs
```

If VCS is configured correctly on each node, the console output resembles:

```
Apr  5 14:52:02 sys1 gab: GAB:20036: Port h gen 3972a201
      membership 01
```

See "Verifying the cluster" on page 415.

To start VCS as single node

- Run the following command:

```
# svcadm enable vcs-onenode
```

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 Operations Manager, or the Cluster Manager (Java Console). For information on management tools, refer to the *Symantec 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 *Symantec 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. You can refer to the configuration examples of a system with a ClusterService group. See the *Symantec Cluster Server Administrator's Guide* for more information.

See "Sample main.cf file for VCS clusters" on page 531.

Chapter 18

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 non-SCSI-3 fencing in virtual environments manually

Setting up disk-based I/O fencing manually

Table 18-1 lists the tasks that are involved in setting up I/O fencing.

Table 18-1

Task	Reference
Initializing disks as VxVM disks	See "Initializing disks as VxVM disks" on page 158.
Identifying disks to use as coordinator disks	See "Identifying disks to use as coordinator disks" on page 273.
Checking shared disks for I/O fencing	See "Checking shared disks for I/O fencing" on page 164.
Setting up coordinator disk groups	See "Setting up coordinator disk groups" on page 273.
Creating I/O fencing configuration files	See "Creating I/O fencing configuration files" on page 274.
Modifying VCS configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 275.

Table 18-1 *(continued)*

Task	Reference
Configuring CoordPoint agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 290.
Verifying I/O fencing configuration	See "Verifying I/O fencing configuration" on page 277.

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See "Initializing disks as VxVM disks" on page 158.

Review the following procedure to identify disks to use as coordinator disks.

To identify the coordinator disks

1 List the disks on each node.

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

```
# 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 164.

Setting up coordinator disk groups

From one node, create a disk group named vxfencoorddg. 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 *Symantec Storage Foundation Administrator's Guide* for details on how to create disk groups.

The following example procedure assumes that the disks have the device names c1t1d0s2, c2t1d0s2, and c3t1d0s2.

To create the vxfencoorddg disk group

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

```
# vxdg init vxfencoorddg c1t1d0s2 c2t1d0s2 c3t1d0s2
```

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 Ensure that you 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/default/vxfen

Modifying VCS configuration to use I/O fencing

After you add coordination points 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 To ensure High Availability has stopped cleanly, run `gabconfig -a`.

In the output of the commans, check that Port h is not present.

4 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

```
# svcadm disable -t vxfen
```

5 Make a backup of the main.cf file on all the nodes:

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

6 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
)
```

Regardless of whether the fencing configuration is disk-based or server-based, the value of the cluster-level attribute UseFence is set to SCSI3.

7 Save and close the file.

8 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

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

9 Using rcp or another utility, copy the VCS configuration file from a node (for example, sys1) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp sys1:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

10 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 coordination points that are listed in /etc/vxfentab.

  ```
  # svcadm enable vxfen
  ```

- Start VCS on the node where main.cf is modified.

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

- Start VCS on all other nodes once VCS on first node reaches RUNNING state.

  ```
  # /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

1 On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is SCSI3 and the SCSI3 disk policy is dmp:

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

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

   * 0 (sys1)
   1 (sys2)

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

2 Verify that the disk-based I/O fencing is using the specified disks.

```
# vxfenconfig -l
```

Setting up server-based I/O fencing manually

Tasks that are involved in setting up server-based I/O fencing manually include:

Table 18-2 Tasks to set up server-based I/O fencing manually

Task	Reference
Preparing the CP servers for use by the VCS cluster	See "Preparing the CP servers manually for use by the VCS cluster" on page 278.
Generating the client key and certificates on the client nodes manually	See "Generating the client key and certificates manually on the client nodes " on page 281.
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 283.
Modifying VCS configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 275.
Configuring Coordination Point agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 290.
Verifying the server-based I/O fencing configuration	See "Verifying server-based I/O fencing configuration" on page 291.

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 18-3 displays the sample values used in this procedure.

Table 18-3 Sample values in procedure

CP server configuration component	Sample name
CP server	cps1
Node #1 - VCS cluster	sys1
Node #2 - VCS cluster	sys2
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 (sys1):

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster clus1

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2-bb31-00306eea460a}
```

2 Use the `cpsadm` command to check whether the VCS cluster and nodes are present in the CP server.

For example:

```
# cpsadm -s cps1.symantecexample.com -a list_nodes

ClusName  UUID                                    Hostname(Node ID) Registered
clus1  {f0735332-1dd1-11b2-bb31-00306eea460a} sys1(0)         0
clus1  {f0735332-1dd1-11b2-bb31-00306eea460a} sys2(1)         0
```

If the output does not show the cluster and nodes, then add them as described in the next step.

For detailed information about the `cpsadm` command, see the *Symantec Cluster Server Administrator's Guide*.

3 Add the VCS cluster and nodes to each CP server.

For example, issue the following command on the CP server (cps1.symantecexample.com) to add the cluster:

```
# cpsadm -s cps1.symantecexample.com -a add_clus\
 -c clus1  -u {f0735332-1dd1-11b2}

Cluster clus1 added successfully
```

Issue the following command on the CP server (cps1.symantecexample.com) to add the first node:

```
# cpsadm -s cps1.symantecexample.com -a add_node\
 -c clus1 -u {f0735332-1dd1-11b2} -h sys1 -n0

Node 0 (sys1) successfully added
```

Issue the following command on the CP server (cps1.symantecexample.com) to add the second node:

```
# cpsadm -s cps1.symantecexample.com -a add_node\
 -c clus1 -u {f0735332-1dd1-11b2} -h sys2 -n1

Node 1 (sys2) successfully added
```

4 If security is to be disabled, then add the user name "cpsclient@hostname" to the server.

5 Add the users to the CP server.

Issue the following commands on the CP server (cps1.symantecexample.com):

```
# cpsadm -s cps1.symantecexample.com -a add_user -e\
cpsclient@hostname\
 -f cps_operator -g vx

User cpsclient@hostname
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 (cps1.symantecexample.com) for VCS cluster clus1 with two nodes sys1 and sys2:

```
# cpsadm -s cps1.symantecexample.com -a\
add_clus_to_user -c clus1\
 -u {f0735332-1dd1-11b2}\
 -e cpsclient@hostname\
 -f cps_operator -g vx

Cluster successfully added to user
 cpsclient@hostname  privileges.
```

See “Generating the client key and certificates manually on the client nodes ” on page 281.

Generating the client key and certificates manually on the client nodes

The client node that wants to connect to a CP server using HTTPS must have a private key and certificates signed by the Certificate Authority (CA) on the CP server

The client uses its private key and certificates to establish connection with the CP server. The key and the certificate must be present on the node at a predefined location. Each client has one client certificate and one CA certificate for every CP server, so, the certificate files must follow a specific naming convention. Distinct certificate names help the `cpsadm` command to identify which certificates have to be used when a client node connects to a specific CP server.

The certificate names must be as follows: ca_*cps-vip*.crt and client _*cps-vip*.crt

Where, *cps-vip* is the VIP or FQHN of the CP server listed in the `/etc/vxfenmode` file. For example, for a sample VIP, *192.168.1.201*, the corresponding certificate name is ca_*192.168.1.201*.

To manually set up certificates on the client node

1 Create the directory to store certificates.

```
# mkdir -p /var/VRTSvxfen/security/keys
/var/VRTSvxfen/security/certs
```

Note: Since the openssl utility might not be available on client nodes, Symantec recommends that you access the CP server using SSH to generate the client keys or certificates on the CP server and copy the certificates to each of the nodes.

2 Generate the private key for the client node.

```
# /usr/bin/openssl genrsa -out client_private.key 2048
```

3 Generate the client CSR for the cluster. CN is the UUID of the client's cluster.

```
# /usr/bin/openssl req -new -key client_private.key\
-subj '/C=countryname/L=localityname/OU=COMPANY/CN=CLUS_UUID\
-out client_192.168.1.201.csr
```

Where, *countryname* is the name of the country, *localityname* is the city, *COMPANY* is the name of the company, and *CLUS_UUID* is the certificate name.

4 Generate the client certificate by using the CA key and the CA certificate.

```
# /usr/bin/openssl x509 -req -days days -in
client_192.168.1.201.csr\
-CA /var/VRTScps/security/certs/ca.crt -CAkey\
/var/VRTScps/security/keys/ca.key -set_serial 01 -out
client_192.168.10.1.crt
```

Where, *days* is the days you want the certificate to remain valid, *192.168.1.201* is the VIP or FQHN of the CP server.

5 Copy the client key, client certificate, and CA certificate to each of the client nodes at the following location.

Copy the client key at `/var/VRTSvxfen/security/keys/client_private.key`. The client is common for all the client nodes and hence you need to generate it only once.

Copy the client certificate at `/var/VRTSvxfen/security/certs/client_192.168.1.201.crt`.

Copy the CA certificate at `/var/VRTSvxfen/security/certs/ca_192.168.1.201.crt`

Note: Copy the certificates and the key to all the nodes at the locations that are listed in this step.

6 If the client nodes need to access the CP server using the FQHN and or the host name, make a copy of the certificates you generated and replace the VIP with the FQHN or host name. Make sure that you copy these certificates to all the nodes.

7 Repeat the procedure for every CP server.

8 After you copy the key and certificates to each client node, delete the client keys and client certificates on the CP server.

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)
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)
- Set the order of coordination points

Note: Whenever coordinator disks are used as coordination points in your I/O fencing configuration, you must create a disk group (vxfencoorddg). You must specify this disk group in the `/etc/vxfenmode` file.

See "Setting up coordinator disk groups" on page 273.

The customized fencing framework also generates the `/etc/vxfentab` file which has coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

To configure server-based fencing on the VCS cluster manually

1. Use a text editor to edit the following file on each node in the cluster:

   `/etc/default/vxfen`

   You must change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1.

2. Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

   - If your server-based fencing configuration uses a single highly available CP server as its only coordination point, make sure to add the `single_cp=1` entry in the `/etc/vxfenmode` file.
   - If you want the `vxfen` module to use a specific order of coordination points during a network partition scenario, set the `vxfen_honor_cp_order` value to be **1**. By default, the parameter is disabled.

   The following sample file output displays what the `/etc/vxfenmode` file contains:

   See "Sample vxfenmode file output for server-based fencing" on page 284.

3. After editing the `/etc/vxfenmode` file, run the vxfen init script to start fencing.

   For example:

   ```
   # svcadm enable vxfen
   ```

Sample vxfenmode file output for server-based fencing

The following is a sample vxfenmode file for server-based fencing:

```
#
# 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 parameter is deprecated release 6.1 onwards
# since communication with CP server will always happen
# over HTTPS which is inherently secure. In pre-6.1 releases,
# it was used to configure 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

#
# vxfen_honor_cp_order determines the order in which vxfen
# should use the coordination points specified in this file.
#
```

```
# available options:
# 0 - vxfen uses a sorted list of coordination points specified
# in this file,
# the order in which coordination points are specified does not matter.
#   (default)
# 1 - vxfen uses the coordination points in the same order they are
#    specified in this file

# Specify 3 or more odd number of coordination points in this file,
# each one in its own line. 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 follows:
#
#  cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames
# over different subnets, all of the IPs/names can be specified
# in a comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,
...,[<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#  is the serial number of the CPS as a coordination point; must
#  start with 1.
# <vip>
#  is the virtual IP address of the CPS, must be specified in
#  square brackets ("[]").
# <vhn>
#  is the virtual hostname of the CPS, must be specified in square
#  brackets ("[]").
# <port>
#  is the port number bound to a particular <vip/vhn> of the CPS.
#  It is optional to specify a <port>. However, if specified, it
#  must follow a colon (":") after <vip/vhn>. If not specified, the
#  colon (":") must not exist after <vip/vhn>.
#
```

```
# For all the <vip/vhn>s which do not have a specified <port>,
# a default port can be specified as follows:
#
# port=<default_port>
#
#  Where <default_port> is applicable to all the <vip/vhn>s for
#  which a <port> is not specified. In other words, specifying
#  <port> with a <vip/vhn> overrides the <default_port> for that
#  <vip/vhn>. If the <default_port> is not specified, and there
#  are <vip/vhn>s for which <port> is not specified, then port
#  number 14250 will be used for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[cps1.company.com]
#  cps2=[192.168.0.25]
# cps3=[cps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [cps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
# - if default port 57777 were not specified, port 14250
# would be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
#
# 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
# cps1=[cps1.company.com]
# cps2=[cps2.company.com]
# cps3=[cps3.company.com]
# port=443
```

Table 18-4 defines the vxfenmode parameters that must be edited.

Table 18-4 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	Deprecated from release 6.1 onwards. Security parameter is deprecated release 6.1 onwards as communication between CP servers and application clusters happens over the HTTPS protocol which is inherently secure. In releases prior to 6.1, the security parameter was used to configure secure communication to the CP server using the VxAT (Veritas Authentication Service) options. The options are: ■ 0 - Do not use Veritas Authentication Service for CP server communication ■ 1 - Use Veritas Authentication Service for CP server communication

Table 18-4 vxfenmode file parameters *(continued)*

vxfenmode File Parameter	Description

cps1, cps2, or vxfendg	Coordination point parameters.
	Enter either the virtual IP address or the FQHN (whichever is accessible) of the CP server.
	`cps<number>=[virtual_ip_address/virtual_host_name]:port`
	Where *port* is optional. The default port value is 443.
	If you have configured multiple virtual IP addresses or host names over different subnets, you can specify these as comma-separated values. For example:
	`cps1=[192.168.0.23],[192.168.0.24]:58888, [cps1.company.com]`
	Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfencoorddg) 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).

port	Default port for the CP server to listen on. If you have not specified port numbers for individual virtual IP addresses or host names, the default port number value that the CP server uses for those individual virtual IP addresses or host names is 443. You can change this default port value using the port parameter.
single_cp	Value 1 for single_cp parameter indicates that the server-based fencing uses a single highly available CP server as its only coordination point. Value 0 for single_cp parameter indicates that the server-based fencing uses at least three coordination points.
vxfen_honor_cp_order	Set the value to 1 for vxfen module to use a specific order of coordination points during a network partition scenario. By default the parameter is disabled. The default value is 0.

Configuring CoordPoint agent to monitor coordination points

The following procedure describes how to manually configure the CoordPoint agent to monitor coordination points.

The CoordPoint agent can monitor CP servers and SCSI-3 disks.

See the *Symantec Cluster Server Bundled Agents Reference Guide* for more information on the agent.

To configure CoordPoint 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 using the following commands:

```
# haconf -makerw
# hagrp -add vxfen
# hagrp -modify vxfen SystemList sys1 0 sys2 1
# hagrp -modify vxfen AutoFailOver 0
# hagrp -modify vxfen Parallel 1
# hagrp -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 0
# hares -override coordpoint LevelTwoMonitorFreq
# hares -modify coordpoint LevelTwoMonitorFreq 5
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

3. Configure the Phantom resource for the vxfen disk group.

```
# haconf -makerw
# hares -add RES_phantom_vxfen Phantom vxfen
# hares -modify RES_phantom_vxfen Enabled 1
# haconf -dump -makero
```

4 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 coordpoint

# Resource    Attribute   System   Value
coordpoint   State        sys1   ONLINE
coordpoint   State        sys2   ONLINE
```

5 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed CoordPoint agent monitoring information; including information about whether the CoordPoint agent is able to access all the coordination points, information to check on which coordination points the CoordPoint agent is reporting missing keys, etc.

To view the debug logs in the engine log, change the dbg level for that node using the following commands:

```
# haconf -makerw

# hatype -modify Coordpoint LogDbg 10

# haconf -dump -makero
```

The agent log can now be viewed at the following location:

/var/VRTSvcs/log/engine_A.log

Verifying server-based I/O fencing configuration

Follow the procedure described below to verify your server-based I/O fencing configuration.

To verify the server-based I/O fencing configuration

1 Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

Note: For troubleshooting any server-based I/O fencing configuration issues, refer to the *Symantec Cluster Server Administrator's Guide*.

2 Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```

If the output displays `single_cp=1`, it indicates that the application cluster uses a CP server as the single coordination point for server-based fencing.

Setting up non-SCSI-3 fencing in virtual environments manually

To manually set up I/O fencing in a non-SCSI-3 PR compliant setup

1 Configure I/O fencing in customized mode with only CP servers as coordination points.

See "Setting up server-based I/O fencing manually" on page 277.

2 Make sure that the VCS cluster is online and check that the fencing mode is customized.

```
# vxfenadm -d
```

3 Make sure that the cluster attribute UseFence is set to SCSI3.

```
# haclus -value UseFence
```

4 On each node, edit the /etc/vxenviron file as follows:

```
data_disk_fencing=off
```

5 On each node, edit the /kernel/drv/vxfen.conf file as follows:

```
vxfen_vxfnd_tmt=25
```

6 On each node, edit the /etc/vxfenmode file as follows:

```
loser_exit_delay=55
vxfen_script_timeout=25
```

Refer to the sample /etc/vxfenmode file.

7 On each node, set the value of the LLT sendhbcap timer parameter value as follows:

- Run the following command:

```
lltconfig -T sendhbcap:3000
```

- Add the following line to the /etc/llttab file so that the changes remain persistent after any reboot:

```
set-timer senhbcap:3000
```

8 On any one node, edit the VCS configuration file as follows:

- Make the VCS configuration file writable:

```
# haconf -makerw
```

- For each resource of the type DiskGroup, set the value of the MonitorReservation attribute to 0 and the value of the Reservation attribute to NONE.

```
# hares -modify <dg_resource> MonitorReservation 0

# hares -modify <dg_resource> Reservation "NONE"
```

- Run the following command to verify the value:

```
# hares -list Type=DiskGroup MonitorReservation!=0

# hares -list Type=DiskGroup Reservation!="NONE"
```

The command should not list any resources.

- Modify the default value of the Reservation attribute at type-level.

```
# haattr -default DiskGroup Reservation "NONE"
```

- Make the VCS configuration file read-only

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

9 Make sure that the UseFence attribute in the VCS configuration file main.cf is set to SCSI3.

10 To make these VxFEN changes take effect, stop and restart VxFEN and the dependent modules

- On each node, run the following command to stop VCS:

  ```
  # svcadm disable -t vcs
  ```

- After VCS takes all services offline, run the following command to stop VxFEN:

  ```
  # svcadm disable -t vxfen
  ```

- On each node, run the following commands to restart VxFEN and VCS:

  ```
  # svcadm enable vxfen
  ```

Sample /etc/vxfenmode file for non-SCSI-3 fencing

```
#
# 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

#
# Seconds for which the winning sub cluster waits to allow for the
# losing subcluster to panic & drain I/Os. Useful in the absence of
# SCSI3 based data disk fencing loser_exit_delay=55
#
# Seconds for which vxfend process wait for a customized fencing
# script to complete. Only used with vxfen_mode=customized
# vxfen_script_timeout=25

# security parameter is deprecated release 6.1 onwards since
# communication with CP server will always happen over HTTPS
# which is inherently secure. In pre-6.1 releases, it was used
# to configure 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

#
# vxfen_honor_cp_order determines the order in which vxfen
# should use the coordination points specified in this file.
#
# available options:
# 0 - vxfen uses a sorted list of coordination points specified
# in this file, the order in which coordination points are specified
# does not matter.
#   (default)
# 1 - vxfen uses the coordination points in the same order they are
#    specified in this file

# Specify 3 or more odd number of coordination points in this file,
# each one in its own line. 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 follows:
#
#  cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames
# over different subnets, all of the IPs/names can be specified
# in a comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,
# ...,[<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#  is the serial number of the CPS as a coordination point; must
#  start with 1.
# <vip>
#  is the virtual IP address of the CPS, must be specified in
#  square brackets ("[]").
# <vhn>
#  is the virtual hostname of the CPS, must be specified in square
#  brackets ("[]").
# <port>
#  is the port number bound to a particular <vip/vhn> of the CPS.
#  It is optional to specify a <port>. However, if specified, it
#  must follow a colon (":") after <vip/vhn>. If not specified, the
#  colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>,
# a default port can be specified as follows:
#
# port=<default_port>
#
#  Where <default_port> is applicable to all the <vip/vhn>s for which a
#  <port> is not specified. In other words, specifying <port> with a
#  <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
```

```
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[cps1.company.com]
#  cps2=[192.168.0.25]
# cps3=[cps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [cps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
# - if default port 57777 were not specified, port 14250 would be
#  used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
#
# 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
# cps1=[cps1.company.com]
# cps2=[cps2.company.com]
```

```
# cps3=[cps3.company.com]
# port=443
```

Section 7

Managing your Symantec deployments

- Chapter 19. Performing centralized installations using the Deployment Server

Chapter 19

Performing centralized installations using the Deployment Server

This chapter includes the following topics:

- About the Deployment Server
- How to install the Deployment Script
- Deployment management overview
- Setting up a Deployment Server
- Setting deployment preferences
- Using the Deployment Server command line option to specify a non-default repository location
- Using the Deployment Server command line options to load and download the most recent release information
- Viewing or downloading available release images
- Viewing or removing repository images stored in your repository
- Deploying Symantec product updates to your environment
- Finding out which releases you have, and which upgrades or updates you may need
- Deploying a specific Symantec release
- Updating release information on systems without Internet access

About the Deployment Server

The Deployment Server is a script that makes it easier to install or upgrade SFHA releases. The Deployment Server lets you store multiple release images in one central location and deploy them to systems of any supported UNIX or Linux operating system (6.1 or later). Prior to 6.1, releases still require the same platform, architecture, distribution, and version of the operating system. You can use the Deployment Server if you want to install or upgrade multiple releases and or multiple platforms.

The Deployment Server lets you do the following as described in Table 19-1.

Table 19-1 Deployment Server functionality

Feature	Description
Manage release images	■ View available SFHA releases. ■ Download maintenance and hot fix release images from the Symantec Operations Readiness Tools (SORT) website into a repository. ■ Load the downloaded release image files from FileConnect and SORT into the repository. ■ View and remove release image files stored in the repository.
Check versions	■ Discover packages and patches installed on designated systems and informs you of the product and version installed, including installed hot fixes. ■ Identify base, maintenance, and hot fix level upgrades to your system and download maintenance and hot fix releases. ■ Query SORT for the most recent updates.

Install or upgrade systems	■ Install or upgrade a release stored in the repository on selected systems.
	■ In release 6.1 and later:
	  ■ Install hot fix level releases.
	  ■ Install SFHA from any supported UNIX or Linux operating system to any other supported UNIX or Linux operating system.
	  ■ Automatically load the script-based installer hot fixes that apply to that release.

Update metadata and preferences	■ Download, load the release matrix updates, and script-based installer updates for systems behind a firewall. ■ Define or reset program settings.

Note: The Deployment Server is available only for the script-based installer, not the web-based installer.

How to install the Deployment Script

The Deployment Script is the utility for managing your Deployment Server.

You can obtain the Deployment Script by either:

- Installing the Deployment Script manually.
- Running the Deployment Script after installing at least one Symantec 6.1 product.

Note: The `VRTSperl` and the `VRTSsfcpi<version>`packages are included in all Storage Foundation (SF) products, so installing any Symantec 6.1 product lets you access the Deployment Script.

To install the Deployment Script manually

1. Log in as superuser.
2. Mount the installation media.

   See "Mounting the product disc" on page 80.
3. For Solaris 10, move to the top-level directory on the disc.

   ```
   # cd /cdrom/cdrom0
   ```
4. For Solaris 10, navigate to the following directory:

   ```
   # cd pkgs
   ```
5. For Solaris 10, run the following commands to install the `VRTSperl` and the `VRTSsfcpi<version>` packages:

   ```
   # pkgadd -d ./VRTSperl.pkg VRTSperl
   # pkgadd -d ./VRTSsfcpi<version>.pkg VRTSsfcpi<version>
   ```
6. For Solaris 11, move to the top-level directory on the disc.

   ```
   # cd /cdrom/cdrom0
   ```

7 For Solaris 11, navigate to the following directory:

```
# cd pkgs
```

8 For Solaris 11, run the following commands to install the `VRTSperl` and the `VRTSsfcpi<version>` packages:

```
# pkg install --accept -g ./VRTSpkgs.p5p VRTSperl VRTSsfcpi<version>
```

To run the Deployment Script

1 Log in as superuser.

2 Navigate to the following directory:

```
# cd /opt/VRTS/install
```

3 Run the Deployment Script.

```
# ./deploy_sfha
```

Deployment management overview

After obtaining and installing the Deployment Server and defining a central repository, you can begin managing your deployments from that repository. You can load and store product images for Symantec products up to version 5.1 in your Deployment Server. The Deployment Server is a central installation server for storing and managing your product updates.

See "How to install the Deployment Script" on page 302.

Setting up and managing your repository involves the following tasks:

- Setting up a Deployment Server.
  See "Setting up a Deployment Server" on page 304.
- Finding out which products you have installed, and which upgrades or updates you may need.
  See "Viewing or downloading available release images" on page 309.
- Adding release images to your Deployment Server.
  See "Viewing or downloading available release images" on page 309.
- Removing release images from your Deployment Server.
  See "Viewing or removing repository images stored in your repository" on page 313.

Later, when your repository is set up, you can use it to deploy Symantec products to other systems in your environment.

See "Deploying Symantec product updates to your environment" on page 316.

Setting up a Deployment Server

For large deployments, Symantec recommends that you create a dedicated Deployment Server to manage your product updates.

A Deployment Server is useful for doing the following tasks:

- Downloading and storing release images for the latest upgrades and updates from Symantec in a central repository directory.
- Installing and updating systems directly by accessing the release images that are stored within a central repository (direct installation).
- Performing heterogeneous push installations (installing Symantec products from the Deployment Server to systems running any supported platform).

  Note: The script-based installer for version 6.1 and higher supports installations from one operating system node onto a different operating system. Therefore, heterogeneous push installations are supported for 6.1 and higher releases only. To perform push installations for product versions 5.1, 6.0, or 6.0.1 releases, you must have a separate Deployment Server for each operating system.

- Creating a file share on the repository directory provides a convenient, central location from which systems running any supported platform can install the latest Symantec products and updates.

Create a central repository on the Deployment Server to store and manage the following types of Symantec releases:

- Base releases. These major releases and minor releases are available for all Symantec products. They contain new features, and you can download them from FileConnect.
- Maintenance releases. These releases are available for all Symantec products. They contain bug fixes and a limited number of new features, and you can download them from the Symantec Operations Readiness Tools (SORT) website.
- Hot fixes. These releases contain fixes for specific products, and you can download them from the SORT website.

Note: All VCS base releases and maintenance releases can be deployed using the install scripts that are included in the release. Hot fixes are typically installed manually, however, from the 6.1 release and onwards, install scripts are included with hot fix releases.

You can set up a Deployment Server with or without Internet access.

- If you set up a Deployment Server that has Internet access, you can download VCS maintenance releases and hot fixes from Symantec directly. Then, you can deploy them to your systems.
  Setting up a Deployment Server that has Internet access
- If you set up a Deployment Server that does not have Internet access, you can download VCS maintenance releases and hot fixes from Symantec on another system that has Internet access. Then, you can load the images onto the Deployment Server separately.
  Setting up a Deployment Server that does not have Internet access

Setting up a Deployment Server that has Internet access

Figure 19-1 shows a Deployment Server that can download product images directly from Symantec using the Deployment Server.

Figure 19-1 Example Deployment Server that has Internet access

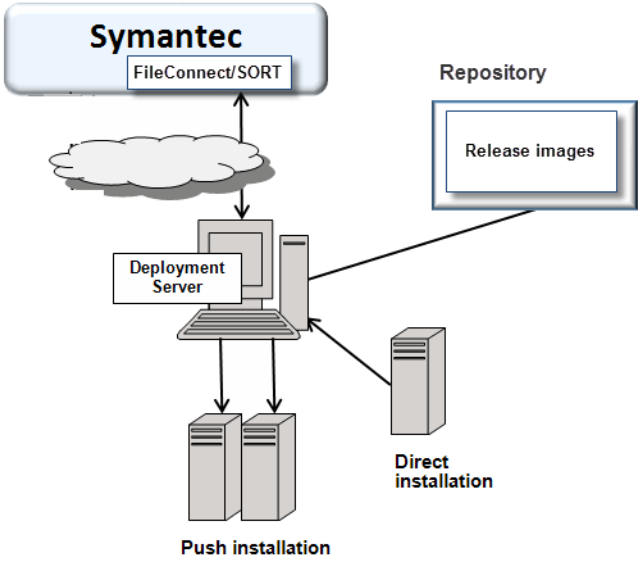

Setting up a Deployment Server that does not have Internet access

Figure 19-2 shows a Deployment Server that does not have Internet access. In this scenario, release images and metadata updates are downloaded from another system. Then, they are copied to a file location available to the Deployment Server, and loaded.

Figure 19-2 Example Deployment Server that does not have Internet access

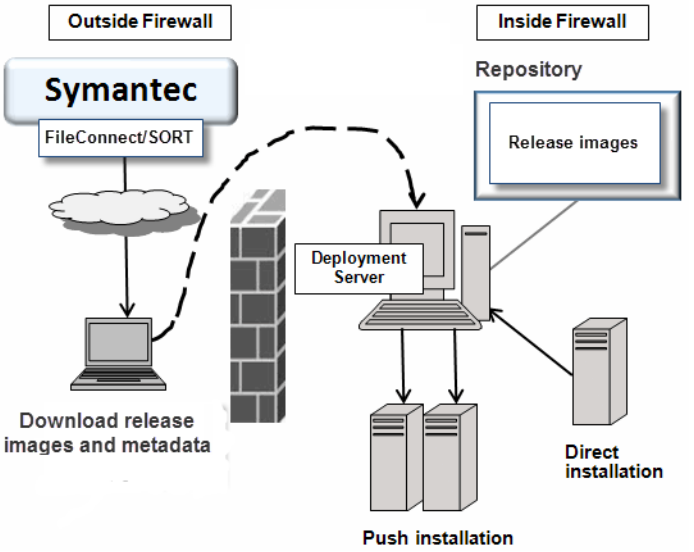

Setting deployment preferences

You can set preferences for managing the deployment of products dating back to version 5.1.

Note: You can select option **T (Terminology and Usage)** to obtain more information about Deployment Server terminology and usage.

To set deployment preferences

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
U) Upgrade/Install Systems               T) Terminology and Usage
?) Help                                  Q) Quit

Enter a Task: [R,M,V,S,U,T,?,Q]
```

2 Select option **S**, **Set Preferences**.

3 In the current preferences page, select option **S**, **Set Preferences**.

4 Do one of the following:

- To set the default repository, enter **1**. Then enter the name of the repository in which you want to store your downloads. For example, enter the following:

  ```
  # /opt/VRTS/install/ProductDownloads
  ```

  If the specified repository replaces a previous repository, the installer asks if you want to move all your files into the new repository. To move your files to the new repository, enter **y**.

- To set the option for saving or removing tar files, enter **2**. At the prompt, if you want to save the tar files after untarring them, enter **y**. Or, if you want to remove tar files after untarring them, enter **n**.
  By default, the installer does not remove tar files after the releases have been untarred.

Using the Deployment Server command line option to specify a non-default repository location

You can specify a repository location other than the default that has been set within the system preferences by using the command line option. The command line option is mainly used to install a release image from a different repository location. When you use the command line option, the designated repository folder is used instead

of the default for the execution time of the script. Using the command line option does not override the repository preference set by the **Set Preference** menu item.

To use the command line option to specify a non-default repository location

- At the command line, to specify a non-default repository location, enter the following:

  ```
  # ./deploy_sfha -repository repository_path
  ```

  where *repository_path* is the preferred location of the repository.

Using the Deployment Server command line options to load and download the most recent release information

You can use the Deployment Server command line options to perform the following operations:

- Load the most recent SF release information and installer hot fixes on systems without Internet access using the `deploy_sfha` command.
- Download a `.tar` file containing the most recent SF release information and installer hot fixes from the SORT site. This `.tar` file is used to update release information on systems without Internet access using the `deploy_sfha -load_metadata` command.

To use the Deployment Server command line option to load the most recent SF release information and installer hot fixes without Internet access

- At the command line, enter the following:

  ```
  # ./deploy_sfha -load_metadata metadata_tar_file
  ```

  where *metadata_tar_file* is the name of the metadata tar file where you want to load the most recent SF release information and installer hot fixes.

  You can download the `.tar` file from the SORT site at:

  https://sort.symantec.com/support/related_links/offline-release-updates

  Or, you can create it by running the `deploy_sfha -download_metadata` command from a system that does have Internet access and can connect to the SORT site.

To use the Deployment Server command line option to download the most recent SF release information and installer hot fixes from the SORT site

- At the command line, enter the following:

```
# ./deploy_sfha -download_metadata
```

Viewing or downloading available release images

You can use the Deployment Server to conveniently view or download available VCS release images to be deployed on other systems in your environment.

Note: If you have Internet access, communication with the Symantec Operations Readiness Tools (SORT) provides the latest release information. If you do not have Internet access, static release matrix files are referenced, and the most recent updates may not be included.

See "Updating release information on systems without Internet access" on page 320.

To view or download available release images

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images      M) Update Metadata
V) Version Check Systems         S) Set Preferences
U) Upgrade/Install Systems       T) Terminology and Usage
?) Help                          Q) Quit

Enter a Task: [R,M,V,S,U,T,?,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

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

3 Select option **1**, **View/Download Available Releases**, to view or download what is currently installed on your system.

You see a list of platforms and release levels.

```
To view or download available releases, the platform type and release
level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  SLES10 x86_64             8)  SLES11 x86_64
9)  Solaris 9 Sparc          10)  Solaris 10 Sparc
11)  Solaris 10 x64          12)  Solaris 11 Sparc
13)  Solaris 11 x64           b)  Back to previous menu

Select the platform of the release to view/download [1-13,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels.

```
1) Base
2) Maintenance
3) Hot Fix
b) Back to previous menu

Select the level of the <platform> releases to view/download
[1-3,b,q,?]
```

5 Select the number corresponding to the type of release you want to view (Base, Maintenance, or Hot Fix).

You see a list of releases available for download.

```
Available Maintenance releases for sol10_sparc:

release_version SORT_release_name           DL OBS AI rel_date   size_KB
================================================================================
5.1SP1PR2RP2  sfha-sol10_sparc-5.1SP1PR2RP2 -  Y   Y  2011-09-28 145611
5.1SP1PR2RP3  sfha-sol10_sparc-5.1SP1PR2RP3 -  Y   Y  2012-10-02 153924
5.1SP1PR2RP4  sfha-sol10_sparc-5.1SP1PR2RP4 -  -   -  2013-08-21 186859
5.1SP1PR3RP2  sfha-sol10_sparc-5.1SP1PR3RP2 -  Y   Y  2011-09-28 145611
5.1SP1PR3RP3  sfha-sol10_sparc-5.1SP1PR3RP3 -  Y   Y  2012-10-02 153924
5.1SP1PR3RP4  sfha-sol10_sparc-5.1SP1PR3RP4 -  -   -  2013-08-21 186859
6.0RP1        sfha-sol10_sparc-6.0RP1       Y  -   -  2012-03-22 245917
6.0.3         sfha-sol10_sparc-6.0.3        Y  -   -  2013-02-01 212507

Enter the release_version to view details about a release or press 'Enter' to continue [b,q,?]
```

The following are the descriptions for the column headers:

- release_version: The version of the release.
- SORT_release_name: The name of the release, used when accessing SORT (https://sort.symantec.com).
- DL: An indicator that the release is present in your repository.
- OBS: An indicator that the release has been obsoleted by another higher release.
- AI: An indicator that the release has scripted install capabilities. All base and maintenance releases have auto-install capabilities. Hot Fix releases

with auto-install capabilities are available beginning with version 6.1. Otherwise the hot fix will require a manual installation.

- rel_date: The date the release is available.
- size_KB: The file size of the release in kilobytes.

6 If you are interested in viewing more details about any release, type the release version. For example, enter the following:

```
6.0.3
```

You see the following output:

```
release_version: 6.0.3
release_name: sfha-sol10_sparc-6.0.3
release_type: MR
release_date: 2013-02-01
downloaded: Y
install_path: sol10_sparc/installmr
upload_location: ftp://ftp.veritas.com/pub/support/patchcentral
/Solaris/6.0.3/sfha/sfha-sol10_sparc-6.0.3-patches.tar.gz
obsoletes: 6.0.1.200-fs,6.0.1.200-vm,6.0.1.300-fs
obsoleted_by: None
Would you like to download this patch? [y,n,q] (y) n

Enter the release_version to view the details about a release or press
'Enter' to continue [b,q,?]
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to download a sol10_sparc Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all releases that are not currently in the repository.

```
1)  5.1SP1PR1RP2
2)  5.1SP1PR1RP4
3)  6.0RP1
4)  All non-obsoleted releases
5)  All releases
b)  Back to previous menu

Select the patch release to download, 'All non-obsoleted releases' to
download all non-obsoleted releases, or 'All releases' to download
all releases [1-5,b,q] 3
```

8 Select the number corresponding to the release that you want to download. It is possible to download a single release, all non-obsoleted releases, or all releases.

The selected release images are downloaded to the Deployment Server.

```
Downloading sfha-sol10_sparc-6.0RP1 from SORT - https://sort.symantec.com
Downloading 215118373 bytes (Total 215118373 bytes [205.15 MB]): 100%
Untarring sfha-sol10_sparc-6.0RP1 ..................................... Done

sfha-sol10_sparc-6.0RP1 has been downloaded successfully.
```

9 From the menu, select option **2**, **View/Remove Repository Images**, and follow the prompts to check that the release images are loaded.

See "Viewing or downloading available release images" on page 309.

Viewing or removing repository images stored in your repository

You can use the Deployment Server to conveniently view or remove the release images that are stored in your repository.

To view or remove release images stored in your repository

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images      M) Update Metadata
V) Version Check Systems         S) Set Preferences
U) Upgrade/Install Systems       T) Terminology and Usage
?) Help                          Q) Quit

Enter a Task: [R,M,V,S,U,T,?,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

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

3 Select option **2**, **View/Remove Repository Images**, to view or remove the release images currently installed on your system.

You see a list of platforms and release levels if you have downloaded the corresponding Base, Maintenance, or Hot Fix release on that platform.

```
To view or remove repository images, the platform type and
release level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  SLES10 x86_64             8)  SLES11 x86_64
9)  Solaris 9 Sparc          10)  Solaris 10 Sparc
11)  Solaris 10 x64          12)  Solaris 11 Sparc
13)  Solaris 11 x64           b)  Back to previous menu

Select the platform of the release to view/remove [1-13,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels if you have downloaded the corresponding Base, Maintenance, or Hot Fix release.

```
1) Base
2) Maintenance
3) Hot Fix
b) Back to previous menu

Select the level of the <platform> releases to view/remove
[1-3,b,q]
```

5 Select the number corresponding to the type of release you want to view or remove (Base, Maintenance, or Hot Fix).

You see a list of releases that are stored in your repository.

```
Stored Repository Releases:

release_version SORT_release_name        OBS AI
================================================
6.0RP1          sfha-sol10_sparc-6.0RP1   -   Y
6.0.3           sfha-sol10_sparc-6.0.3    -   Y
```

6 If you are interested in viewing more details about a release image stored in your repository, type the release version. For example, enter the following:

```
6.0.3
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to remove a sol10_sparc Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all releases stored in your repository that match the selected platform and release level.

```
1)  6.0RP1
2)  6.0.3
b)  Back to previous menu

Select the patch release to remove [1-2,b,q] 1
```

8 Type the number corresponding to the release version you want to remove.

The release images are removed from the Deployment Server.

```
Removing sfha-sol10_sparc-6.0RP1-patches ..................... Done
sfha-sol10_sparc-6.0RP1-patches has been removed successfully.
```

9 From the menu, select option **2**, **View/Remove Repository Images**, and follow the prompts to check that the release images are removed.

See "Viewing or downloading available release images" on page 309.

Deploying Symantec product updates to your environment

After you install at least one Symantec 6.1 product on a server, you can use the Deployment Server to deploy release images to the systems in your environment as follows:

- If you are not sure what to deploy, perform a version check. A version check tells you if there are any Symantec products installed on your systems. It suggests patches and maintenance releases, and gives you the option to install updates.
  See "Finding out which releases you have, and which upgrades or updates you may need" on page 317.

- If you know which update you want to deploy on your systems, use the Upgrade/Install Systems script to deploy a specific Symantec release.
  See "Deploying a specific Symantec release" on page 319.

Finding out which releases you have, and which upgrades or updates you may need

Use the Version Check script to determine which Symantec product you need to deploy. The Version Check script is useful if you are not sure which releases you already have installed, or want to know about available releases.

The Version Check script gives the following information:

- Installed products and their versions (base, maintenance releases, and hot fixes)
- Installed packages (required and optional)
- Available releases (base, maintenance releases, and hot fixes) relative to the version which is installed on the system

To determine which Symantec product updates to deploy

1 Launch the Deployment Server. For example, enter the following:

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images                M) Update Metadata
V) Version Check Systems                   S) Set Preferences
U) Upgrade/Install Systems                 T) Terminology and Usage
?) Help                                    Q) Quit

Enter a Task: [R,M,V,S,U,T,?,Q]
```

2 Select option **V**, **Version Check Systems**.

3 At the prompt, enter the system names for the systems you want to check. For example, enter the following:

sys1

You see output for the installed packages (required, optional, or missing).

You see a list of releases available for download.

```
Available Base Releases for Veritas Storage Foundation HA 6.0.1:
None

Available Maintenance Releases for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name      DL OBS AI rel_date   size_KB
=====================================================================
6.0.3           sfha-sol10_sparc-6.0.3 Y  -   -  2013-02-01 212507

Available Public Hot Fixes for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name        DL OBS AI rel_date   size_KB
=========================================================================
6.0.1.200-fs    fs-sol10_sparc-6.0.1.200 -  Y   -  2012-09-20 14346
6.0.1.200-vm    vm-sol10_sparc-6.0.1.200 -  Y   -  2012-10-10 47880

Would you like to download the available Maintenance or Public Hot
Fix releases which cannot be found in the repository? [y,n,q] (n) y
```

4 If you want to download any of the available maintenance releases or hot fixes, enter **y**.

5 If you have not set a default repository for releases you download, the installer prompts you for a directory. (You can also set the default repository in **Set Preferences**).

See "Setting deployment preferences" on page 306.

You can also specify a non-default repository location using the command line.

See "Using the Deployment Server command line option to specify a non-default repository location" on page 307.

6 Select an option for downloading products.

The installer downloads the releases you specified and stores them in the repository.

Deploying a specific Symantec release

After you install at least one Symantec 6.1 product on a server, you can use the Deployment Server to deploy your licensed Symantec products dating back to version 5.1. If you know which product version you want to install, follow the steps in this section to install it.

To deploy a specific Symantec release

1 From the directory in which you installed your Symantec product (version 6.1 or later), launch the Deployment Server with the upgrade and install systems option. For example, enter the following:

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images             M) Update Metadata
V) Version Check Systems                S) Set Preferences
U) Upgrade/Install Systems              T) Terminology and Usage
?) Help                                 Q) Quit

Enter a Task: [R,M,V,S,U,T,?,Q]
```

2 Select option **U**, **Upgrade/Install Systems**.

You see the following output:

```
    1)  AIX 5.3
    2)  AIX 6.1
    3)  AIX 7.1
    4)  RHEL5 x86_64
    b)  Back to previous menu

Select the platform of the available release(s) to be
upgraded/installed [1-4,b,q,?]
```

3 Select the number corresponding to the platform for the release you want to deploy. For example, select the number for the **RHEL5 x86_64** release or the **AIX 6.1** release.

4 Select a Symantec product release.

The installation script is executed and the release is deployed on the specified server.

Updating release information on systems without Internet access

When you install the Deployment Server, the installation includes product metadata that includes information about Symantec, all prior base, maintenance, and hot fix releases across all the products and the platforms. If your system has Internet access, release matrix information is automatically updated from the Symantec Operations Readiness Tools (SORT) site with each use. If your system is behind a firewall, these updates are not possible and the release matrices eventually get out of date.

To update release information on systems without Internet access, you can download a `.tar` file (`deploy_sfha.tar`) containing all the latest release matrices. Then, load it on to your Deployment Server. The `deploy_sfha.tar` file available from the SORT site is updated on a daily basis, and there are typically several release updates every week captured in the updates.

Downloading a .tar file from the SORT site

Loading releases and hot fixes onto your Deployment Server

Downloading a .tar file from the SORT site

To obtain a `.tar` file with release updates, the easiest method is to download a copy from the SORT website.

To download a .tar file from the SORT site

1 Navigate to the following link:

https://sort.symantec.com/support/related_links/offline-release-updates

2 Click on **deploy_sfha.tar [Download]**.

3 Save the file to your desktop.

Loading releases and hot fixes onto your Deployment Server

In this procedure, the Internet-enabled system is the system to which you downloaded the `deploy_sfha.tar` file.

To load releases and hot fixes onto your Deployment Server

1. On the Internet-enabled system, copy the `deploy_sfha.tar` file you downloaded to a location accessible by the Deployment Server.
2. On the Deployment Server, change to the installation directory. For example, enter the following:

   ```
   # cd /opt/VRTS/install/
   ```

3. Run the Deployment Script. Enter the following:

   ```
   # ./deploy_sfha
   ```

4. Select option **M**, **Update Metadata**, and enter the location of the `deploy_sfha.tar` file (the installer calls it a "meta-data tar file").

   ```
   Enter the location of the meta-data tar file: [b]
   (/opt/VRTS/install/deploy_sfha.tar)
   ```

   For example, enter the location of the meta-data tar file:

   ```
   /tmp/deploy_sfha.tar
   ```

Section 8

Upgrading VCS

- Chapter 20. Planning to upgrade VCS
- Chapter 21. Performing a typical VCS upgrade using the installer
- Chapter 22. Performing a phased upgrade of VCS
- Chapter 23. Performing an automated VCS upgrade using response files
- Chapter 24. Performing a rolling upgrade
- Chapter 25. Upgrading VCS using Live Upgrade and Boot Environment upgrade

Chapter 20

Planning to upgrade VCS

This chapter includes the following topics:

- About upgrading to VCS 6.1
- VCS supported upgrade paths
- Upgrading VCS in secure enterprise environments
- Considerations for upgrading secure VCS 5.x clusters to VCS 6.1
- Considerations for upgrading VCS to 6.1 on systems configured with an Oracle resource
- Considerations for upgrading secure CP servers
- Considerations for upgrading secure CP clients
- Setting up trust relationship between CP server and CP clients manually
- Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes

About upgrading to VCS 6.1

When you upgrade to VCS 6.1, you need not reconfigure application monitoring with VCS. All existing monitoring configurations are preserved.

You can upgrade VCS using one of the following methods:

- Typical upgrade using Veritas product installer or the installvcs
  See "VCS supported upgrade paths" on page 324.
  See "Upgrading VCS using the script-based installer" on page 333.
- Typical upgrade using Veritas web installer
  See "VCS supported upgrade paths" on page 324.

See "Upgrading VCS using the web-based installer" on page 335.

- Phased upgrade to reduce downtime
  See "Performing a phased upgrade using the script-based installer" on page 340.
- Automated upgrade using response files
  See "VCS supported upgrade paths" on page 324.
  See "Upgrading VCS using response files" on page 359.
- Upgrade using supported native operating system utility
  Live Upgrade
  See "About Live Upgrade" on page 373.
- Rolling upgrade to minimize downtime
  See "Performing a rolling upgrade of VCS using the web-based installer" on page 370.

You can upgrade VCS 6.1 to Storage Foundation High Availability 6.1 using Veritas product installer or response files.

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

Note: In a VMware virtual environment, you can use the vSphere Client to directly install VCS and supported high availability agents (together called guest components) on the guest virtual machines. For details, see the *Symantec High Availability Solution Guide for VMware*.

If zones are present on the system, make sure that all non-global zones are in the running state before you use the Symantec product installer to upgrade the Storage Foundation products in the global zone, so that any packages present inside non-global zones also gets updated automatically. For Oracle Solaris 10, if the non-global zones are in configured state at the time of the upgrade, you must attach the zone with `-U` option to upgrade the SFHA packages inside non-global zone. For Oracle Solaris 11.1, If the non-global zone has previous version of VCS packages (VRTSperl, VRTSvlic, VRTSvcs, VRTSvcsag, VRTSvcsea) already installed, then during upgrade of the VCS packages in global zone, packages inside non-global zone are automatically upgraded if the zone is in running state. If non-global zones are not in running state, you must set the publisher inside the global zone and also attach the zone with `-u` option to upgrade the SFHA packages inside non-global zone.

VCS supported upgrade paths

Table 20-1 lists the supported upgrade paths for Solaris SPARC.

Table 20-1 Supported upgrade paths for Solaris SPARC

Current version of VCS	Solaris 8 or older	Solaris 9	Solaris 10	Solaris 11
3.5 3.5 MP4 4.0 4.0 MP1 4.0 MP2	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	Not applicable.	Not applicable.
4.1 4.1 MP1 4.1 MP2	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	No upgrade path exists. Uninstall VCS and then use the installer to perform a full installation of VCS 6.1.	Not applicable.
5.0 5.0 MP1	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	No upgrade path exists. Uninstall VCS and then use the installer to perform a full installation of VCS 6.1.	Not applicable.

Table 20-1 Supported upgrade paths for Solaris SPARC *(continued)*

Current version of VCS	Solaris 8 or older	Solaris 9	Solaris 10	Solaris 11
5.0 MP3 5.0 MP3RPs	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	Upgrade directly to VCS 6.1 using the installer script.	Not applicable.
5.1 5.1RPs 5.1SP1 5.1SP1RPs	Not applicable.	No upgrade path exists. Uninstall VCS. Upgrade the operating system to at least Solaris 10. Use the installer to perform a full installation of VCS 6.1.	Upgrade directly to VCS 6.1 using the installer script.	Not applicable.
6.0 6.0RPs	Not applicable.	Not applicable.	Upgrade directly to VCS 6.1 using the installer script.	Upgrade directly to VCS 6.1 using the installer script.
6.0PR1	Not applicable.	Not applicable.	Not applicable.	Upgrade directly to VCS 6.1 using the installer script.
6.0.1	Not applicable.	Not applicable.	Upgrade directly to VCS 6.1 using the installer script.	Upgrade directly to VCS 6.1 using the installer script.

Table 20-1 Supported upgrade paths for Solaris SPARC *(continued)*

Current version of VCS	Solaris 8 or older	Solaris 9	Solaris 10	Solaris 11
6.0.3	Not applicable.	Not applicable	Upgrade directly to VCS 6.1 using the installer script.	Upgrade directly to VCS 6.1 using the installer script.

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).

To upgrade VCS in secure enterprise environments with no rsh or ssh communication

1. Run the installvcs program on each node to upgrade the cluster to VCS 6.1.

   On each node, the installvcs program updates the configuration, stops the cluster, and then upgrades VCS on the node. The program also generates a cluster UUID on the node. Each node may have a different cluster UUID at this point.

2. Start VCS on the first node.

   ```
   # hastart
   ```

   VCS generates the cluster UUID on this node. Run the following command to display the cluster UUID on the local node:

   ```
   # /opt/VRTSvcs/bin/uuidconfig.pl -clus -display systemname
   ```

3. On each of the other nodes, perform the following steps:

   - Set the value of the VCS_HOST environment variable to the name of the first node.
   - Display the value of the CID attribute that stores the cluster UUID value:

     ```
     # haclus -value CID
     ```

   - Copy the output of the CID attribute to the file /etc/vx/.uuids/clusuuid.
   - Update the VCS_HOST environment variable to remove the set value.
   - Start VCS.
     The node must successfully join the already running nodes in the cluster.

See "Verifying LLT, GAB, and cluster operation" on page 411.

Considerations for upgrading secure VCS 5.x clusters to VCS 6.1

When you upgrade a secure VCS 5.x cluster to VCS 6.1, the upgrade does not migrate the old broker configuration to the new broker because of the change in architecture. Both the old broker (`/opt/VRTSat/bin/vxatd`) and new broker (`/opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver`) continue to run. In such a scenario, you must consider the following:

- The HA commands that you run in VCS 6.1 are processed by the new broker by default. To ensure that the HA commands are processed by the old broker, set the VCS_REMOTE_BROKER environment variable as follows:

  ```
  # export VCS_REMOTE_BROKER=localhost IP,2821
  ```

  See "About enabling LDAP authentication for clusters that run in secure mode" on page 396.

- VCS 6.1 does not prompt non-root users who run HA commands for passwords. In 5.x, non-root users required a password to run HA commands. If you want non-root users to enter passwords before they run HA commands, set the VCS_DOMAINTYPE environment variable to unixpwd.

- Trust relationships are not migrated during the upgrade. If you had configured secure GCO or secure steward, ensure that trust relationships are recreated between the clusters and the steward.
  See "Setting up trust relationships for your VCS cluster" on page 145.

- For Zones, the HA commands run within the container and use credentials that were deployed by the old broker. However, you can migrate to the newer credentials from the new broker by running `hazonesetup` again.

When the old broker is not used anymore, you can delete the old VRTSat package.

Considerations for upgrading VCS to 6.1 on systems configured with an Oracle resource

If you plan to upgrade VCS running on systems configured with an Oracle resource, set the `MonitorOption` attribute to 0 (zero) before you start the upgrade. If you use the product installer for the rolling upgrade, it sets the `MonitorOption` to 0 through its scripts. In a manual upgrade, the `MonitorOption` value must be set to 0 using

the `hares` command. When the upgrade is complete, invoke the `build_oraapi.sh` script, and then set the `MonitorOption` to 1 to enable the Oracle health check.

For more information on enabling the Oracle health check, see the *Symantec Cluster Server Agent for Oracle Installation and Configuration Guide*.

Considerations for upgrading secure CP servers

CP server supports Symantec Product Authentication Services (AT) (IPM-based protocol) and HTTPS communication to securely communicate with clusters. For HTTPS communication, you do not need to consider setting up trust relationships.

When you upgrade the CP Server that supports IPM-based protocol, trust relationships are not migrated.

If you upgrade the CP clients after you upgrade the CP server that supports IPM-based protocol, the installer recreates the trust relationships that are established by the client. You do not need to establish the trust relationships manually. However, the CP server and CP clients cannot communicate with each other till trust relationships are established.

If you do not upgrade the CP clients after you upgrade the CP server that supports IPM-based protocol, you must recreate the trust relationships between the CP server and CP clients.

Considerations for upgrading secure CP clients

Passwordless communication from CP clients to CP server must exist for the installer to reconfigure fencing. If passwordless communication does not exist, you must reconfigure fencing manually.

See “Setting up disk-based I/O fencing manually” on page 272.

See “Setting up server-based I/O fencing manually” on page 277.

Setting up trust relationship between CP server and CP clients manually

You need to set up trust relationship only if you use the Symantec Product Authentication Services (AT) (IPM-based protocol) for communication between CP servers and CP server clients.

For each client cluster on release version 6.0 and later, run the following command on the CP server:

```
EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSERVER \
/opt/VRTSvcs/bin/vcsat setuptrust -b client_ip_addres:14149 -s high
```

For each client cluster on release version prior to 6.0, run the following command on the CP server:

```
EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSERVER \
/opt/VRTSvcs/bin/vcsat setuptrust -b client_ip_addres:2821 -s high
```

For each client node on release version 6.0 and later, run the following command:

```
EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSADM \
/opt/VRTSvcs/bin/vcsat setuptrust -b cpserver_ip_address:14149 -s high
```

For each client node on release version prior to 6.0, run the following command:

```
/opt/VRTScps/bin/cpsat setuptrust -b cpserver_ip_address:14149 -s high
```

Using Install Bundles to simultaneously install or upgrade base releases, maintenance patches, and hot fixes

Beginning with version 6.1, Symantec offers you a method to easily install or upgrade your systems directly to a base, maintenance, or hot fix level in one step using Install Bundles. With Install Bundles, the installers have the ability to merge so that customers can install or upgrade directly to maintenance or hot fix levels in one execution. The various scripts, packages, and patch components are merged, and multiple releases are installed together as if they are one combined release. You do not have to perform two or more install actions to install or upgrade systems to maintenance levels or hot fix levels.

Releases are divided into the following categories:

Table 20-2 Release Levels

Level	Content	Form factor	Applies to	Release types	Download location
Base	Features	packages	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect

Table 20-2 Release Levels *(continued)*

Level	Content	Form factor	Applies to	Release types	Download location
Maintenance	Fixes, new features	packages, patches	All products	Maintenance Release (MR)	Symantec Operations Readiness Tools (SORT)
Hot fix	Fixes	packages	Single product	P-Patch, Public hot fix, Private hot fix	SORT, Support site

When you install or upgrade using Install Bundles:

- SFHA products are discovered and assigned as a single version to the maintenance level. Each system can also have one or more hot fixes applied.
- Base releases are accessible from FileConnect that requires customer serial numbers. Maintenance and hot fix releases can be automatically downloaded from SORT. You can download them from the SORT website manually or use the `deploy_sfha` script.
- Public hot fix releases can be installed using automated installers from the 6.1 version or later.
- Private hot fixes can now be detected to prevent upgrade conflict. Private hot fix releases are not offered as a combined release. They are only available from Symantec Technical Support on a need basis.

You can use the `-base_path` and `-hotfix_path` options to import installation code from multiple releases. You can find packages and patches from different media paths, and merge package and patch definitions for multiple releases. You can use these options to use new task and phase functionality to correctly perform required operations for each release component. You can install the packages and patches in defined phases using these options, which helps you when you want to perform a single start or stop process and perform pre and post operations for all level in a single operation.

Four possible methods of integration exist. All commands must be executed from the highest base or maintenance level install script.

For example:

1. Base + maintenance:

   This integration method can be used when you install or upgrade from a lower version to 6.1.1.

Enter the following command:

```
# installmr -base_path <path_to_base>
```

2. Base + hot fix:

   This integration method can be used when you install or upgrade from a lower version to 6.1.0.100.

   Enter the following command:

   ```
   # installer -hotfix_path <path_to_hotfix>
   ```

3. Maintenance + hot fix:

   This integration method can be used when you upgrade from version 6.1 to 6.1.1.100.

   Enter the following command:

   ```
   # installmr -hotfix_path <path_to_hotfix>
   ```

4. Base + maintenance + hot fix:

   This integration method can be used when you install or upgrade from a lower version to 6.1.1.100.

   Enter the following command:

   ```
   # installmr -base_path <path_to_base>
   -hotfix_path <path_to_hotfix>
   ```

   Note: For the 6.1 release, you can add a maximum of five hot fixes using *-hotfix_path <path_to_hotfix> -hotfix2_path <path_to_hotfix> ... -hotfix5_path <path_to_hotfix>*

Chapter 21

Performing a typical VCS upgrade using the installer

This chapter includes the following topics:

- Before upgrading using the script-based or Web-based installer
- Upgrading VCS using the script-based installer
- Upgrading VCS using the web-based installer

Before upgrading using the script-based or Web-based installer

As a result of OS upgrade, if VCS is not in running state before upgrade, the installer does not start VCS after the upgrade is completed. You need to manually start it or restart the cluster nodes.Before you upgrade VCS, perform the following steps. You first need to remove deprecated resource types and modify changed values.

To prepare to upgrade to VCS 6.1

- Make sure that all non-global zones are booted and in the running state before you install or upgrade the VCS packages in the global zone. If the non-global zones are not mounted and running at the time of upgrade, you must upgrade each package in each non-global zone manually.

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. Choose **1** for Full Upgrade.
5. 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.
6. When the verification checks are complete, the installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.

   The installer lists the packages to upgrade.
7. The installer asks if you want to stop VCS processes. Press the Enter key to continue.

   The installer stops VCS processes, uninstalls packages, installs or upgrades packages, and configures VCS.

   The installer lists the nodes that Symantec recommends you restart.
8. The installer asks if you would like to send the information about this installation to Symantec to help improve installation in the future. Enter your response.

   The installer displays the location of log files, summary file, and response file.
9. If you want to upgrade CP server systems that use VCS or SFHA to VCS 6.1, make sure that you first upgrade all application clusters to version VCS 6.1. Then, upgrade VCS or SFHA on the CP server systems.

   For instructions to upgrade VCS or SFHA, see the *Symantec Cluster Server Installation Guide* or the *Storage Foundation and High Availability Installation Guide*.

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 524.

Upgrading VCS using the web-based installer

This section describes upgrading VCS with the web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems.

To upgrade VCS

1. Perform the required steps to save any data that you want to preserve. For example, make configuration file backups.
2. If you want to upgrade a high availability (HA) product, take all service groups offline. List all service groups:

   ```
   # /opt/VRTSvcs/bin/hagrp -list
   ```

   For each service group listed, take it offline:

   ```
   # /opt/VRTSvcs/bin/hagrp -offline service_group -any
   ```

3. Start the web-based installer.

   See "Starting the web-based installer" on page 189.
4. On the Select a task and a product page, select **Upgrade a Product** from the Task drop-down menu.

   The product is discovered once you specify the system. Click **Next**.
5. Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.
6. Installer detects the product that is installed on the specified system. It shows the cluster information and lets you confirm if you want to perform upgrade on the cluster. Select **Yes** and click **Next**.
7. 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**.
8. 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.
9. If you are prompted to restart the systems, enter the following restart command:

   ```
   # /usr/sbin/shutdown -y -i6 -g0
   ```

10 After the upgrade, if the product is not configured, the web-based installer asks: "Do you want to configure this product?" If the product is already configured, it does not ask any questions.

11 Click **Finish**. The installer prompts you for another task.

12 If you want to upgrade CP server systems that use VCS or SFHA to 6.1, make sure that you upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.1. For instructions to upgrade VCS or SFHA, see the *VCS or SFHA Installation Guide*.

If you want to upgrade 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 524.

Chapter 22

Performing a phased upgrade of VCS

This chapter includes the following topics:

- About phased upgrade
- Performing a phased upgrade using the script-based installer

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:

Table 22-1

Fail over condition	Downtime
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 into two sub-clusters of equal or near equal size.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.
- Before you start the upgrade, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.

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.
- After you upgrade the first half of your cluster (the first subcluster), you need to set up password-less SSH or RSH. Create the connection between an upgraded node in the first subcluster and a node from the other subcluster. The node from the other subcluster is where you plan to run the installer and also plan to upgrade.
- 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 a secure cluster that you have configured to run on 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.

Figure 22-1 Example of phased upgrade set up

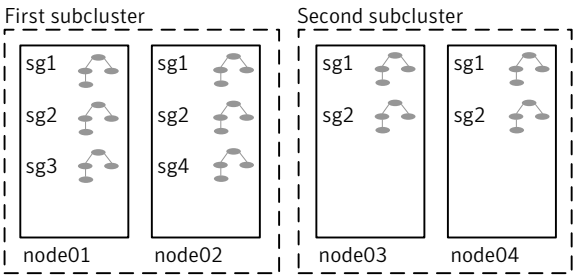

Each service group is running on the nodes as follows:

- 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.

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.

Phased upgrade example overview

This example's upgrade path follows:

- Move all the failover service groups from the first subcluster to the second subcluster.
- Take all the parallel service groups offline on the first 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. After activating the first cluster, switch the service groups online on the second subcluster to 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 using the script-based installer” on page 340.

Performing a phased upgrade using the script-based installer

This section explains how to perform a phased upgrade of VCS on four nodes with four service groups. Note that in this scenario, VCS and the service groups cannot stay online on the second subcluster 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.

An example of a phased upgrade follows. It illustrates the steps to perform a phased upgrade. The example makes use of a secure VCS cluster.

You can perform a phased upgrade from VCS 5.1 or other supported previous versions to VCS 6.1.

See “About phased upgrade” on page 337.

See “Phased upgrade example” on page 338.

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.

```
# hagrp -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|
```

2 Offline the parallel service groups (sg1 and sg2) 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). For SFHA, vxfen sg is the parallel service group.

```
# hagrp -offline sg1 -sys node01
# hagrp -offline sg2 -sys node01
# hagrp -offline sg1 -sys node02
# hagrp -offline sg2 -sys node02
# hagrp -switch sg3 -to node03
# hagrp -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 -k
```

```
Filesystem          kbytes    used   avail capacity  Mounted on
/dev/dsk/c1t0d0s0   66440242 10114415 55661425  16%  /
/devices                0       0       0     0%  /devices
ctfs                    0       0       0     0%  /system/contract
proc                    0       0       0     0%  /proc
mnttab                  0       0       0     0%  /etc/mnttab
swap                 5287408    1400 5286008     1%  /etc/svc/volatile
objfs                   0       0       0     0%  /system/object
sharefs                 0       0       0     0%  /etc/dfs/sharetab
/platform/sun4u-us3/lib/libc_psr/libc_psr_hwcap1.so.1
                  66440242 10114415 55661425    16%  /platform/sun4u-us3/lib/
libc_psr.so.1
/platform/sun4u-us3/lib/sparcv9/libc_psr/libc_psr_hwcap1.so.1
                    66440242 10114415 55661425    16%    /platform/sun4u-us3/lib/
sparcv9/libc_psr.so.1
fd                      0       0       0     0%    /dev/fd
swap                 5286064      56 5286008     1%    /tmp
swap                 5286056      48 5286008     1%    /var/run
swap                 5286008       0 5286008     0%    /dev/vx/dmp
swap                 5286008       0 5286008     0%    /dev/vx/rdmp
                      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.

```
# hagrp -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|
```

9 Back up the lltab, llthosts, gabtab, types.cf, main.cf and AT configuration files on the first subcluster.

```
# cp /etc/lltab /etc/lltab.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.

Before performing operating system upgrade, it is better to prevent LLT from starting automatically when the node starts. For example, you can do the following:

```
# mv /etc/llttab /etc/llttab.save
```

or you can change the `/etc/default/llt` file by setting LLT_START = **0**.

After you finish upgrading the OS, remember to change the LLT configuration to its original configuration.

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. Make sure that you can ssh or rsh from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
3. Navigate to the folder that contains installvcs.

   ```
   # cd cluster_server
   ```

4. 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.

5. 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
   cluster_server/EULA/<lang>/EULA_SFHA_Ux_<version>.pdf
   file present on media? [y,n,q,?] y
   ```

6 Review the available installation options.

See “Symantec Cluster Server installation packages” on page 481.

1 Installs only the minimal required VCS packages that provides basic functionality of the product.

2 Installs the recommended VCS packages that provide complete functionality of the product. This option does not install the optional VCS packages.

Note that this option is the default.

3 Installs all the VCS packages.

You must choose this option to configure any optional VCS feature.

4 Displays the VCS packages for each option.

For this example, select 3 for all packages.

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

7 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility.

8 When you are prompted, reply **y** to continue with the upgrade.

```
Do you want to continue? [y,n,q] (y)
```

9 When you are prompted, reply **y** to stop appropriate processes.

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

10 The installer ends for the first subcluster with the following output:

```
    Configuring VCS: 100%

    Estimated time remaining: 0:00

    Performing VCS upgrade configuration ..................... Done

Symantec Cluster Server Configure completed successfully


You are performing phased upgrade (Phase 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.

11 In the `/etc/default/llt` file, set LLT_START = 0.

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
```

2 Unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -k
```

```
Filesystem            kbytes    used   avail capacity  Mounted on
/dev/dsk/c1t0d0s0    66440242 10114415 55661425    16%    /
/devices                   0       0       0     0%    /devices
ctfs                       0       0       0     0%    /system/contract
proc                       0       0       0     0%    /proc
mnttab                     0       0       0     0%    /etc/mnttab
swap                 5287408    1400 5286008     1%    /etc/svc/volatile
objfs                      0       0       0     0%    /system/object
sharefs                    0       0       0     0%    /etc/dfs/sharetab
/platform/sun4u-us3/lib/libc_psr/libc_psr_hwcap1.so.1
                     66440242 10114415 55661425    16%    /platform/sun4u-us3/
lib/libc_psr.so.1
/platform/sun4u-us3/lib/sparcv9/libc_psr/libc_psr_hwcap1.so.1
                     66440242 10114415 55661425    16%    /platform/sun4u-us3/
lib/sparcv9/libc_psr.so.1
fd                         0       0       0     0%    /dev/fd
swap                 5286064      56 5286008     1%    /tmp
swap                 5286056      48 5286008     1%    /var/run
swap                 5286008       0 5286008     0%    /dev/vx/dmp
swap                 5286008       0 5286008     0%    /dev/vx/rdmp
                      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 Make the configuration writable on the second subcluster.

```
# haconf -makerw
```

4 Unfreeze the service groups.

```
# hagrp -unfreeze sg1 -persistent
# hagrp -unfreeze sg2 -persistent
# hagrp -unfreeze sg3 -persistent
# hagrp -unfreeze sg4 -persistent
```

5 Dump the configuration and make it read-only.

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

6 Take the service groups offline on node03 and node04.

```
# hagrp -offline sg1 -sys node03
# hagrp -offline sg1 -sys node04
# hagrp -offline sg2 -sys node03
# hagrp -offline sg2 -sys node04
# hagrp -offline sg3 -sys node03
# hagrp -offline sg4 -sys node04
```

7 Verify the state of the service groups.

```
# hagrp -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|
```

8 Stop all VxVM volumes (for each disk group) that VCS does not manage.

9 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

- Solaris 9:

```
# /opt/VRTSvcs/bin/hastop -local
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- Solaris 10 and 11:

```
# svcadm disable -t /system/vcs
# svcadm disable -t /system/vxfen
```

```
# svcadm disable -t /system/gab
# svcadm disable -t /system/llt
```

10 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 are not configured.

- Solaris 9:

```
# /etc/init.d/vxfen status
VXFEN: loaded
# /etc/init.d/gab status
GAB: module not configured
# /etc/init.d/llt status
LLT: is loaded but not configured
```

- Solaris 10 and 11:

```
# /lib/svc/method/vxfen status
VXFEN: loaded

# /lib/svc/method/gab status
GAB: module not configured

# /lib/svc/method/llt status
LLT: is loaded but not configured
```

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 10.

To activate the first subcluster

1 Reboot the node01 and node02 in the first subcluster.

```
# /usr/sbin/shutdown -y -i6 -g0
```

2 In the `/etc/default/llt` file, change the value of the LLT_START attribute.

```
LLT_START = 1
```

3 Start LLT and GAB.

```
#  svcadm enable system/llt
# svcadm enable system/gab
```

4 Seed node01 and node02 in the first subcluster.

```
# gabconfig -x
```

5 On the first half of the cluster, start VCS:

```
# cd /opt/VRTS/install
```

```
# ./installvcs<version> -start sys1 sys2
```

Where <*version*> is the specific release version.

See “About the script-based installer” on page 49.

6 Start VCS in first half of the cluster:

```
# svcadm enable system/vcs
```

7 Start VCS on node01 and node02. On each node run:

```
# hastart
```

8 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

9 Unfreeze the nodes in the first subcluster.

```
# hasys -unfreeze -persistent node01
# hasys -unfreeze -persistent node02
```

10 Dump the configuration and make it read-only.

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

11 Bring the service groups online on node01 and node02.

```
# hagrp -online sg1 -sys node01
# hagrp -online sg1 -sys node02
# hagrp -online sg2 -sys node01
# hagrp -online sg2 -sys node02
# hagrp -online sg3 -sys node01
# hagrp -online sg4 -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

1 On the second subcluster, disable VCS so that it does not start after reboot. Edit the `vcs` file in `/etc/default`. Open the `vcs` file in an editor, and change the line that reads VCS_START=1 to VCS_START=**0**. Save and close the file.

2 On the second subcluster, disable VXFEN so that it does not start after reboot. Edit the `vxfen` file in `/etc/default`. Open the `vxfen` file in an editor, and change the line that reads VXFEN_START=1 to VXFEN_START=**0**. Save and close the file.

3 On the second subcluster, disable GAB so that it does not start after reboot. Edit the `gab` file in `/etc/default`. Open the `gab` file in an editor, and change the line that reads GAB_START=1 to GAB_START=**0**. Save and close the file.

4 On the second subcluster, disable LLT so that it does not start after reboot. Edit the `llt` file in `/etc/default`. Open the `llt` file in an editor, and change the line that reads LLT_START=1 to LLT_START=**0**. Save and close the file.

5 For a cluster that uses secure mode, create a password-less SSH connection. The connection is from the node where you plan to run the installer to one of the nodes that you have already upgraded.

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
   cluster_server/EULA/<lang>/EULA_VCS_Ux_<version>.pdf
   file present on media? [y,n,q,?] y
   ```

5. Review the available installation options.

   See “Symantec Cluster Server installation packages” on page 481.

   1. Installs only the minimal required VCS packages that provides basic functionality of the product.
   2. Installs the recommended VCS packages that provide complete functionality of the product. This option does not install the optional VCS packages.

      Note that this option is the default.
   3. Installs all the VCS packages.

      You must choose this option to configure any optional VCS feature.
   4. Displays the VCS packages for each option.

      For this example, select 3 for all packages.

      Select the packages 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

Complete the following procedure to complete the upgrade.

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
-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 Reboot the node03 and node04 in the second subcluster.

```
# /usr/sbin/shutdown -y -i6 -g0
```

The nodes in the second subcluster join the nodes in the first subcluster.

3 In the /etc/default/llt file, change the value of the LLT_START attribute.

In the /etc/default/gab file, change the value of the gab_START attribute.

In the /etc/default/vxfen file, change the value of the vxfen_START attribute.

In the /etc/default/vcs file, change the value of the vcs_START attribute.

```
LLT_START = 1
GAB_START = 1
VXFEN_START =1
VCS_START =1
```

4 Start LLT and GAB.

```
# svcadm enable system/llt

# svcadm enable system/gab
```

5 Seed node03 and node04 in the second subcluster.

```
# gabconfig -x
```

6 On the second half of the cluster, start VCS:

```
# cd /opt/VRTS/install

# ./installvcs<version> -start sys3 sys4
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 49.

7 For nodes that use Solaris 10, start VCS in first half of the cluster:

```
# svcadm enable system/vcs
```

8 Check to see if VCS and its components are up.

```
# gabconfig -a
GAB Port Memberships
===============================================================
Port a gen   nxxxnn membership 0123
Port b gen   nxxxnn membership 0123
Port h gen   nxxxnn membership 0123
```

9 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  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              ONLINE
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
```

10 After the upgrade is complete, start the VxVM volumes (for each disk group) and mount the VxFS file systems.

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.

Note: If you want to upgrade application clusters that use CP server based fencing to 6.1, make sure that you first upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.1. However, note that the CP server upgraded to 6.1 can support application clusters on 6.1 (HTTPS-based communication) and application clusters prior to 6.1 (IPM-based communication). When you configure the CP server, the installer asks the VIPs for HTTPS-based communication (if the clients are on release version 6.1) or VIPs for IPM-based communication (if the clients are on a release version prior to 6.1).

For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.

Chapter 23

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
- Performing rolling upgrade of VCS using response files
- Response file variables to upgrade VCS using rolling upgrade
- Sample response file for VCS using rolling upgrade

Upgrading VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS upgrade on one system to upgrade VCS on other systems. 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 systems where you want to upgrade VCS.

   See "Sample response file for upgrading VCS" on page 361.

4 Edit the values of the response file variables as necessary.

See “Response file variables to upgrade VCS” on page 360.

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 23-1 lists the response file variables that you can define to upgrade VCS.

Table 23-1 Response file variables specific to upgrading VCS

Variable	List or Scalar	Description
CFG{opt}{upgrade}	Scalar	Upgrades VCS packages. (Required)
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{systems}	List	List of systems on which the product is to be upgraded. (Required)
CFG{prod}	Scalar	Defines the product to be upgraded. The value is VCS61 for VCS. (Optional)
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)

Table 23-1 Response file variables specific to upgrading 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}{pkgpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product packages. 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 packages 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 *logpath* location. (Optional)
CFG{opt}{rsh}	Scalar	Defines that *rsh* must be used instead of ssh 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 360.

```
#
# Configuration Values:
```

```
#
our %CFG;

$CFG{accepteula}=1;
$CFG{vcs_allowcomms}=1;
$CFG{opt}{upgrade}=1;
$CFG{prod}="VCS61";
$CFG{systems}=[ qw( sys1 sys2) ];
1;
```

Performing rolling upgrade of VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS upgrade on one system to upgrade VCS on other systems.

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

To perform automated VCS rolling 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 the systems where you want to launch the installer.

   See "Sample response file for VCS using rolling upgrade" on page 364.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to upgrade VCS using rolling upgrade" on page 363.
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 using rolling upgrade

Table 23-2 lists the response file variables that you can define to upgrade VCS using rolling upgrade.

Table 23-2 Response file variables for upgrading VCS using rolling upgrade

Variable	Description
CFG{phase1}{0}	A series of $CFG{phase1}{N} items define sub-cluster division. The index N indicatse the order to do RU phase1. The index starts from 0. Each item has a list of node(at least 1). List or scalar: list Optional or required: conditional required Required if rolling upgrade phase1 needs to be performed.
CFG{rollingupgrade_phase2}	The `CFG{rollingupgrade_phase2}` option is used to perform rolling upgrade Phase 2. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version. List or scalar: scalar Optional or required: conditional required Required if rolling upgrade phase 2 needs to be performed.
CFG{rolling_upgrade}	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade Phase 1 or Phase 2 explicitly.
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required

Table 23-2 Response file variables for upgrading VCS using rolling upgrade *(continued)*

Variable	Description
CFG{opt}{upgrade}	Upgrades all packages installed, without configuration. List or scalar: scalar Optional or required: optional
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required

Sample response file for VCS using rolling upgrade

The following example shows a response file for VCS using Rolling Upgrade.

```
our %CFG;
$CFG{accepteula}=1;
$CFG{client_vxfen_warning}=1;
$CFG{fencing_cps}=[ qw(10.198.90.6) ];
$CFG{fencing_cps_ports}{"10.198.90.6"}=50006;
$CFG{fencing_cps_vips}{"10.198.90.6"}=[ qw(10.198.90.6) ];
$CFG{opt}{gco}=1;
$CFG{opt}{noipc}=1;
$CFG{opt}{rolling_upgrade}=1;
$CFG{opt}{rollingupgrade_phase2}=1;
$CFG{opt}{updatekeys}=1;
$CFG{opt}{upgrade}=1;
$CFG{opt}{vr}=1;
$CFG{phase1}{"0"}=[ qw(sys3 sys2) ];
$CFG{phase1}{"1"}=[ qw(sys1) ];
$CFG{systems}=[ qw(sys1 sys2 sys3) ];
$CFG{vcs_allowcomms}=1;
1;
```

Chapter 24

Performing a rolling upgrade

This chapter includes the following topics:

- Supported rolling upgrade paths
- About rolling upgrades
- Performing a rolling upgrade using the installer
- Performing a rolling upgrade of VCS using the web-based installer

Supported rolling upgrade paths

You can perform a rolling upgrade of VCS with the script-based installer, the web-based installer, or manually.

The rolling upgrade procedures support only minor operating system upgrades.

Table 24-1 shows the versions of VCS for which you can perform a rolling upgrade to VCS 6.1.

Table 24-1 Supported rolling upgrade paths

Platform	VCS version
Solaris 10 SPARC	5.1, 5.1RPs 5.1SP1, 5.1SP1RPs 6.0, 6.0RP1 6.0.1, 6.0.3

Table 24-1 Supported rolling upgrade paths *(continued)*

Platform	VCS version
Solaris 11 SPARC	6.0PR1 6.0.1, 6.0.3

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.1, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in Patch Lookup on the SORT website.

About rolling upgrades

The rolling upgrade minimizes downtime for highly available clusters to the amount of time that it takes to perform a service group failover. The rolling upgrade has two main phases where the installer upgrades kernel packages in phase 1 and VCS agent packages in phase 2.

Note: You need to perform a rolling upgrade on a completely configured cluster.

The following is an overview of the flow for a rolling upgrade:

1. The installer performs prechecks on the cluster.
2. The installer moves service groups to free nodes for the first phase of the upgrade as is needed.

   Application downtime occurs during the first phase as the installer moves service groups to free nodes for the upgrade. The only downtime that is incurred is the normal time required for the service group to fail over. The downtime is limited to the applications that are failed over and not the entire cluster.
3. The installer performs the second phase of the upgrade on all of the nodes in the cluster. The second phase of the upgrade includes downtime of the Symantec Cluster Server (VCS) engine HAD, but does not include application downtime.

Figure 24-1 illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 24-1 Example of the installer performing a rolling upgrade

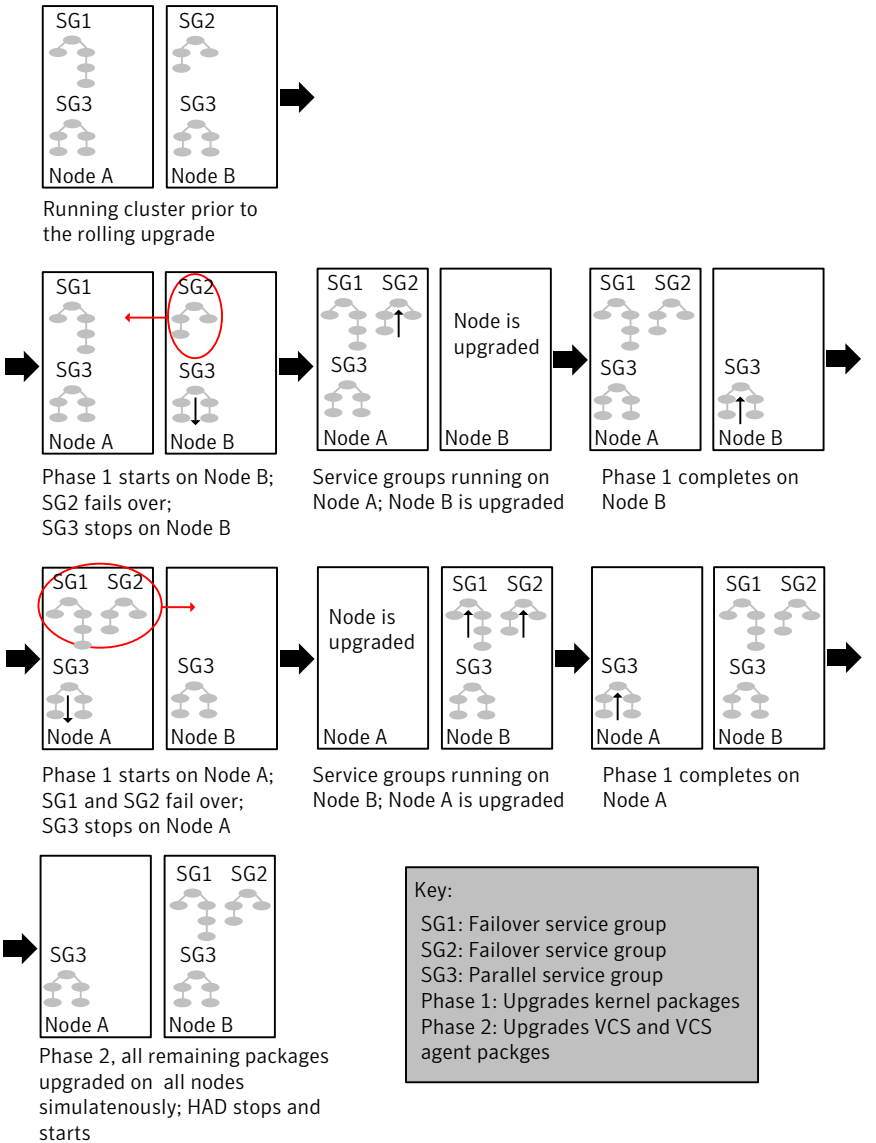

The following limitations apply to rolling upgrades:

- Rolling upgrades are not compatible with phased upgrades. Do not mix rolling upgrades and phased upgrades.
- You can perform a rolling upgrade from 5.1 and later versions.

Performing a rolling upgrade using the installer

Use a rolling upgrade to upgrade Symantec Cluster Server to the latest release with minimal application downtime.

Performing a rolling upgrade using the script-based installer

Before you start the rolling upgrade, make sure that Symantec Cluster Server (VCS) is running.

To perform a rolling upgrade

1. Complete the preparatory steps on the first sub-cluster.
2. Log in as superuser and mount the VCS 6.1 installation media.
3. From root, start the installer.

   ```
   # ./installer
   ```

4. From the menu, select `Upgrade a Product` and from the sub menu, select `Rolling Upgrade`.
5. The installer suggests system names for the upgrade. Press **Enter** to upgrade the suggested systems, or enter the name of any one system in the cluster on which you want to perform a rolling upgrade and then press**Enter**.
6. The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
7. The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.
8. The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.
9. Review the end-user license agreement, and type **y** if you agree to its terms.
10. After the installer detects the online service groups, the installer checks for the installed packages and displays them. Then, the installer displays the list of packages that would be installed and prompts the user to do one of the following:
    - Manually switch service groups
    - Use the CPI to automatically switch service groups

The downtime is the time that it normally takes for the service group's failover.

Note: It is recommended that you manually switch the service groups. Automatic switching of service groups does not resolve dependency issues.

---

11 The installer prompts you to stop the applicable processes. Type **y** to continue.

The installer evacuates all service groups to the node or nodes that are not upgraded at this time. The installer stops parallel service groups on the nodes that are to be upgraded.

12 The installer stops relevant processes, uninstalls old kernel packages, and installs the new packages. The installer asks if you want to update your licenses to the current version. Select **Yes** or **No**. Symantec recommends that you update your licenses to fully use the new features in the current release.

13 If the cluster has configured Coordination Point Server based fencing, then during upgrade, installer asks the user to provide the new HTTPS Coordination Point Server.

The installer performs the upgrade configuration and restarts the nodes, if required. Then, the installer starts the processes.

14 Complete the preparatory steps on the nodes that you have not yet upgraded.

15 The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade.

If the installer prompts to restart nodes, restart the nodes.

Restart the installer.

The installer repeats step 7 through step 12.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.

16 When phase 1 of the rolling upgrade completes, begin phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.

17 The installer determines the remaining packages to upgrade. Press **Enter** to continue.

18 The installer stops Symantec Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs prestop, uninstalls old packages, and installs the new packages. It performs post-installation tasks, and the configuration for the upgrade.

19 If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

20 A prompt message appears to ask if the user wants to read the summary file. You can choose **y** if you want to read the install summary file.

21 Upgrade application to the supported version.

22 If you want to upgrade application clusters that use CP server based fencing to 6.1, make sure that you upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.1. However, note that the CP server upgraded to 6.1 can support application clusters on 6.1 (HTTPS-based communication) and application clusters prior to 6.1 (IPM-based communication). When you configure the CP server, the installer asks the VIPs for HTTPS-based communication (if the clients are on release version 6.1) or VIPs for IPM-based communication (if the clients are on a release version prior to 6.1).

For instructions to upgrade VCS or SFHA on the CP server systems, refer to the appropriate installation guide.

Note: You have to configure `vset` under VCS if you want the new configuration changes to take effect.

Performing a rolling upgrade of VCS using the web-based installer

This section describes using the web-based installer to perform a rolling upgrade. 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.

See “About rolling upgrades” on page 366.

To start the rolling upgrade—phase 1

1. Perform the required steps to save any data that you want to preserve. For example, take backups of configuration files.
2. Start the web-based installer.

   See "Starting the web-based installer" on page 189.
3. In the Task pull-down menu, select `Rolling Upgrade`.

   The option `Phase-1: Upgrade Kernel packages` is displayed and selected by default.

   Click **Next** to proceed.
4. Enter the name of any one system in the cluster on which you want to perform a rolling upgrade.The installer identifies the cluster information of the system and displays the information.

   Click **Yes** to confirm the cluster information. The installer now displays the nodes in the cluster that will be upgraded during phase 1 of the upgrade.
5. Review the systems that the installer has chosen for phase 1 of the rolling upgrade. These systems are chosen to minimize downtime during the upgrade.

   Click **Yes** to proceed.

   The installer validates systems.
6. Review the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
7. If you have online failover service groups, the installer prompts you to choose to switch these service groups either manually or automatically. Choose any option and follow the steps to switch all the failover service groups to the other subcluster.
8. The installer stops all processes. Click **Next** to proceed.

   The installer removes old software and upgrades the software on the systems that you selected.
9. The installer asks if you want to update your licenses to the current version. Select **Yes** or **No**. Symantec recommends that you update your licenses to fully use the new features in the current release.
10. If the cluster has configured Coordination Point Server-based fencing, then during upgrade, installer asks the user to provide the new HTTPS Coordination Point Server. If you are prompted, restart the node.

    The installer starts all the relevant processes and brings all the service groups online if the nodes do not require a restart.

11 Restart the nodes, if required.

Restart the installer.

12 Repeat step 5 through step 11 until the kernel packages of all the nodes are upgraded. For clusters with larger number of nodes, this process may get repeated several times. Service groups come down and are brought up to accommodate the upgrade.

13 When prompted, perform step 3 through step 11 on the nodes that you have not yet upgraded.

14 When prompted, start phase 2. Click **Yes** to continue with the rolling upgrade.

You may need to restart the web-based installer to perform phase 2.

See "Starting the web-based installer" on page 189.

To upgrade the non-kernel components—phase 2

1 In the Task pull-down menu, make sure that **Rolling Upgrade** is selected.

Click the **Next** button to proceed.

2 The installer detects the information of cluster and the state of rolling upgrade.

The installer validates systems. Click **Next**. If it throws an error, address the error and return to the installer.

3 Review the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.

4 The installer stops the `HAD` and `CmdServer` processes in phase 2 of the rolling upgrade process but the applications continue to run. Click **Next** to proceed.

5 The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted. Register the software and click **Next** to proceed. The installer starts all the relevant processes and brings all the service groups online.

6 If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

7 A prompt message appears to ask if the user wants to read the summary file. You can choose **y** if you want to read the install summary file.

The upgrade is complete.

Chapter 25

Upgrading VCS using Live Upgrade and Boot Environment upgrade

This chapter includes the following topics:

- About Live Upgrade
- About ZFS Boot Environment (BE) upgrade
- Supported upgrade paths for Live Upgrade and Boot Environment upgrade
- Performing Live Upgrade on Solaris 10 systems
- Performing Boot Environment upgrade on Solaris 11 systems

About Live Upgrade

Solaris Live Upgrade provides a method of upgrading a system while the system continues to operate. This is done by creating an alternate boot environment (ABE) from the current boot environment and then upgrading the ABE. Once the ABE is upgraded, you can activate the ABE and then reboot the system.

On Solaris 10 or previous releases, you can use Live Upgrade technology to reduce downtime associated with the OS upgrade and VCS product upgrade by creating a boot environment on a alternate boot disk.

- See "Performing Live Upgrade on Solaris 10 systems" on page 376.

Figure 25-1 illustrates an example of an upgrade of Symantec products from 5.1 SP1 to 6.1, and the operating system from Solaris 9 to Solaris 10 using Live Upgrade.

Figure 25-1 Live Upgrade process

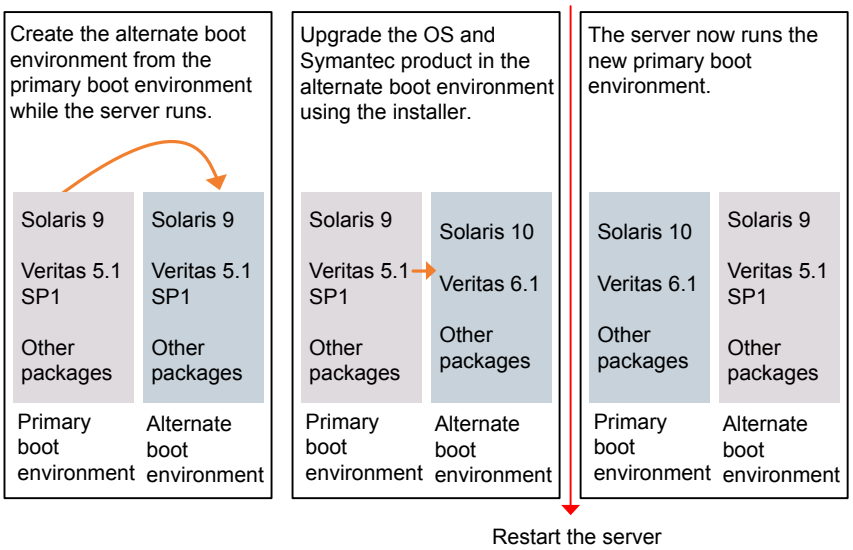

Some service groups (failover and parallel) may be online in this cluster and the Live Upgrade process does not affect them. Downtime is experienced only when the server is restarted to boot into the alternate boot environment.

Symantec Cluster Server exceptions for Live Upgrade

If you have configured Veritas File System or Veritas Volume Manager, use the Live Upgrade instructions in the *Storage Foundation and High Availability Installation Guide*.

About ZFS Boot Environment (BE) upgrade

A Boot Environment (BE) is a bootable instance of the Oracle Solaris operating system image along with any other application software packages installed into that image. System administrators can maintain multiple BEs on their systems, and each BE can have different software versions installed. Upon the initial installation of the Oracle Solaris 11 release onto a system, a BE is created.

On Solaris 11, you can use the `beadm` utility to create and administer additional BEs on your system.

Figure 25-2 Boot Environment upgrade process

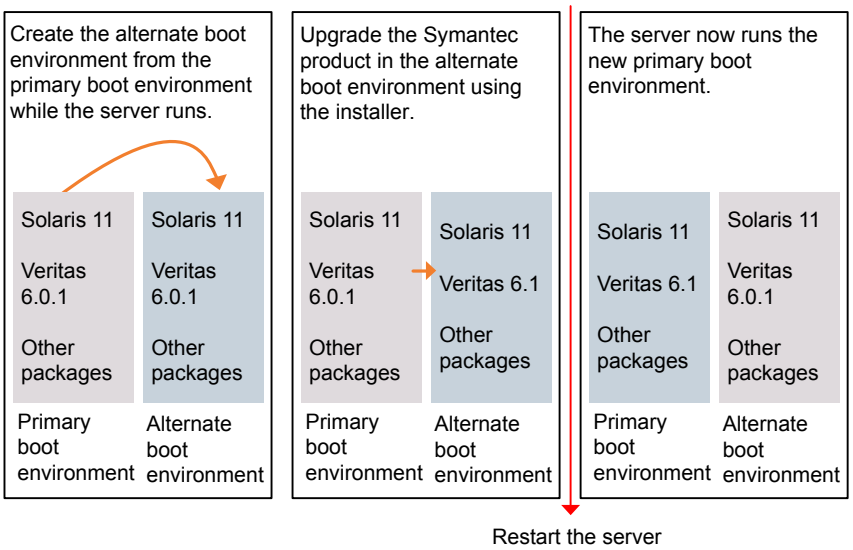

Supported upgrade paths for Live Upgrade and Boot Environment upgrade

The systems where you plan to use Live Upgrade must run Solaris 9 or Solaris 10. Boot environment upgrade can be used on Solaris 11 system only. You can upgrade from those systems that run Solaris 9, but VCS 6.1 is not supported on Solaris 9.

For Live Upgrade method, existing VCS version must be at least 5.0 MP3. For Boot Environment upgrade method, the VCS version you are upgrading to must be at least 6.1.0.

Symantec requires that both global and non-global zones run the same version of Symantec products.

You can use Live Upgrade or Boot environment upgrade in the following virtualized environments:

Table 25-1 Live Upgrade or Boot environment upgrade support in virtualized environments

Environment	Procedure
Solaris native zones	Perform Live Upgrade or Boot Environment upgrade to upgrade both global and non-global zones. See "Performing Live Upgrade on Solaris 10 systems" on page 376. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 388.
Solaris branded zones (BrandZ)	Perform Live Upgrade or Boot Environment upgrade to upgrade the global zone. See "Performing Live Upgrade on Solaris 10 systems" on page 376. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 388. VCS 6.1 does not support Branded zones on Solaris 10 operating system. You must migrate applications running on Solaris 8/9 branded zones to Solaris 10 non-global zones if you need to manage the applications by VCS.
Oracle VM Server for SPARC	Use Live upgrade or Boot Environment upgrade procedure for Control domain as well as guest domains. See "Performing Live Upgrade on Solaris 10 systems" on page 376. See "Performing Boot Environment upgrade on Solaris 11 systems" on page 388.

Performing Live Upgrade on Solaris 10 systems

Perform the Live Upgrade using the installer.

For VCS, the nodes do not form a cluster until all of the nodes are upgraded. At the end of the Live Upgrade of the last node, all the nodes must boot from the alternate boot environment and join the cluster

Table 25-2 Upgrading VCS using Solaris 10 Live Upgrade

Step	Description
Step 1	Prepare to upgrade using Solaris Live Upgrade. See "Before you upgrade VCS using Solaris Live Upgrade" on page 377.
Step 2	Create a new boot environment on the alternate boot disk. See "Creating a new Solaris 10 boot environment on the alternate boot disk" on page 379.
Step 3	Upgrade VCS using the installer. See "Upgrading VCS using the installer for Solaris 10 Live Upgrade" on page 383.
	To upgrade only Solaris See the Oracle documentation on Solaris 10 operating system **Note:** If you choose to upgrade the operating system, the Solaris operating system on the alternate boot environment is upgraded. A new boot environment is created on the alternate boot disk by cloning the primary boot environment.
Step 4	Switch the alternate boot environment to be the new primary. See "Completing the Solaris 10 Live Upgrade" on page 384.
Step 5	Verify Live Upgrade of VCS. See "Verifying the Solaris 10 Live Upgrade of VCS" on page 385.

Before you upgrade VCS using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

To prepare for the Live Upgrade

1. Make sure that the VCS installation media and the operating system installation images are available and on hand.
2. On the nodes to be upgraded, select an alternate boot disk that is at least the same size as the root partition of the primary boot disk

   If the primary boot disk is mirrored, you need to break off the mirror for the alternate boot disk.

3 Before you perform the Live Upgrade, take offline any services that involve non-root file systems. This prevents file systems from being copied to the alternate boot environment that can potentially cause a root file system to run out of space.

4 On the primary boot disk, patch the operating system for Live Upgrade.

For upgrade from Solaris 9 to 10:

- SPARC system: Patch 137447-01 or later is required.

Verify that the patches are installed.

5 The version of the Live Upgrade packages must match the version of the operating system to which you want to upgrade on the alternate boot disk. If you upgrade the Solaris operating system, do the following steps:

- Remove the installed Live Upgrade packages for the current operating system version:
  All Solaris versions: `SUNWluu`, `SUNWlur` packages.
  Solaris 10 update 7 or later also requires: `SUNWlucfg` package.
  Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.
- From the new Solaris installation image, install the new versions of the following Live Upgrade packages:
  All Solaris versions: `SUNWluu`, `SUNWlur`, and `SUNWlucfg` packages.
  Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.

Solaris installation media comes with a script for this purpose named liveupgrade20. Find the script at /*cdrom*/*solaris_release*/Tools/Installers/liveupgrade20. If scripting, you can use:

```
# /cdrom/solaris_release/Tools/Installers/liveupgrade20 \
-nodisplay -noconsole
```

If the specified image has some missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you have to install any missing patches on the alternate boot disk.

6 If the specified image has some missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you need to install any missing patches on the alternate boot disk.

Creating a new Solaris 10 boot environment on the alternate boot disk

Symantec provides the `vxlustart` script that runs a series of commands to create the alternate boot environment for the upgrade.

To preview the commands, specify the `vxlustart` script with the `-V` option.

Symantec recommends that you preview the commands with `-V` option to ensure there are no problems before beginning the Live Upgrade process. The vxlustart script is located in the scripts directory on the distribution media.

Note: This step can take several hours to complete. Do not interrupt the session as it may leave the boot environment unstable.

```
# cd /cdrom/scripts

# ./vxlustart -V -u targetos_version -s osimage_path -d diskname
```

Table 25-3

vxlustart option	Usage
-V	Lists the commands to be executed during the upgrade process without executing them and pre-checks the validity of the command. If the operating system is upgraded, the user is prompted to compare the patches that are installed on the image with the patches installed on the primary boot disk. This determines if any critical patches are not present from the new operating system image.
-v	Indicates verbose, print commands before executing them.
-f	Forces the vtoc creation on the disk.
-Y	Indicates a default yes with no questions asked.
-m	Uses the already existing vtoc on the disk.
-D	Prints with debug option on, and is for debugging.

Table 25-3 *(continued)*

vxlustart option	Usage
-U	Specifies that only the Storage Foundation products are upgraded. The operating system is cloned from the primary boot disk.
-g	Specifies the DG to which the rootdisk belongs. Optional.
-d	Indicates the name of the alternate boot disk `c#t#d#s2` on which you intend to upgrade. The default disk is mirrordisk.
-u	Specifies the operating system version for the upgrade on the alternate boot disk. For example, use `5.9` for Solaris 9 and `5.10` for Solaris 10. If you want to upgrade only SF products, specify the current OS version.
-F	Specifies the root disk's file system, where the default is `ufs`.
-S	Specifies the path to the Solaris image. It can be a network/directory path. If the installation uses the CD, this option must not be specified. See *Solaris Live Upgrade installation guide* for more information about the path.
-r	Specifies that if the computer crashes or restarts before the `vxlufinish` command is run, the alternate disk is remounted using this option.
-k	Specifies the location of file containing auto-registration information. This file is required by `luupgrade(1M)` for OS upgrade to Solaris 10 9/10 or a later release.
-x	Excludes file from newly created BE. (`lucreate -x` option)
-X	Excludes file list from newly created BE. (`lucreate -f` option

Table 25-3 *(continued)*

vxlustart option	Usage
-i	Includes file from newly created BE. (`lucreate -y` option)
-I	Includes file list from newly created BE. (`lucreate -Y` option)
-z	Filters file list from newly created BE. (`lucreate -z` option)
-w	Specifies additional mount points. (`lucreate -m` option)
-W	Specifies additional mount points in a file (`lucreate -M` option

If the `-U` option is specified, you can omit the `-s` option. The operating system is cloned from the primary boot disk.

For example, to preview the commands to upgrade only the Symantec product:

```
# ./vxlustart -V -u 5.10 -U -d disk_name
```

For example, to preview the commands for an upgrade to Solaris 10 update 6:

```
# ./vxlustart -V -u 5.10 -s /mnt/Solaris_10u6 -d c0t1d0s2
```

In the procedure examples, the primary or current boot environment resides on Disk0 (c0t0d0s2) and the alternate or inactive boot environment resides on Disk1 (c0t1d0s2).

At the end of the process:

- A new boot environment is created on the alternate boot disk by cloning the primary boot environment.
- The Solaris operating system on the alternate boot disk is upgraded, if you have chosen to upgrade the operating system.

To create a new boot environment on the alternate boot disk

Perform the steps in this procedure on each node in the cluster.

1 Navigate to the install media for the Symantec products:

```
# cd /cdrom/scripts
```

2 Before you upgrade, make sure that you exclude the file system mount points on a shared storage that applications use from getting copied to the new boot environment. To prevent these shared mount points from being copied to the new boot environment, create a temporary file containing the file system mountpoints that need to be excluded.

```
# cat /var/tmp/file_list
- /ora_mnt
- /sap_mnt
```

Where `/var/tmp/file_list` is a temporary file that contains the list of mount points to be excluded from the new boot environment. The items in the file list are preceded either by a *'+'* or *'-'* symbol. The *'+'* symbol indicates that the mount point is included in the new boot environment. The *'-'* symbol indicates that the mount point is excluded from the new boot environment.

Apart from file system mount points, you may choose to include or exclude other files. If you have non-global zone in running state in the current boot environment and zone root path is on a shared storage, setup another disk of same or more size for each zone root in alternate boot environment.

3 Run one of the following commands to create the alternate boot environment.

For example: To upgrade the operating system:

```
# ./vxlustart -v -u 5.10 -s /mnt/sol10u9 -d
c0t1d0s2 -z /var/tmp/file_list
```

Where `/mnt/sol10u9` is the path to the operating system image that contains the `.cdtoc` file.

To clone the operating system of current boot environment:

```
# ./vxlustart -v -u 5.10 -U -d c0t1d0s2 -z /var/tmp/file_list
```

If you have non-global zone with zone root path on shard storage, then to upgrade the OS:

```
# ./vxlustart -v -u 5.10 -U -d c0t1d0s2 -z /var/tmp/file_list -w
/zone1-rootpath:/dev/dsk/<new disk>/:<type of fs>
```

Where `zone1-rootpath` is root path of zone in present boot environment.

4 Update the permissions, user name, and group name of the mount points (created on the ABE) to match that of the existing directories on the primary boot environment.

5 If zone root path is on shared storage, update the `/altroot.5.10/etc/VRTSvcs/conf/config/main.cf` file with new block device created in step 2 for all zones to reflect the ABE zone root paths.

6 Review the output and note the new mount points. If the system is restarted before completion of the upgrade or if the mounts become unmounted, you may need to remount the disks.

If you need to remount, run the command:

```
# vxlustart -r -u targetos_version -d disk_name
```

7 After the alternate boot disk is created and mounted on */altroot.5.10*, install any operating system patches or packages on the alternate boot disk that are required for the Symantec product installation.

```
# pkgadd -R /altroot.5.10 -d pkg_dir
```

Upgrading VCS using the installer for Solaris 10 Live Upgrade

You can use the Symantec product installer to upgrade VCS as part of the Live Upgrade.

On a node in the cluster, run the installer on the alternate boot disk to upgrade VCS on all the nodes in the cluster. The program uninstalls the existing version of VCS on the alternate boot disk during the process.

At the end of the process, VCS 6.1 is installed on the alternate boot disk.

To perform Live Upgrade of VCS using the installer

1 Insert the product disc with VCS 6.1 or access your copy of the software on the network.

2 Run the installer script specifying the root path as the alternate boot disk:

```
# ./installer -upgrade -rootpath /altroot.5.10
```

3 Enter the names of the nodes that you want to upgrade to VCS 6.1.

The installer displays the list of packages to be installed or upgraded on the nodes.

4 Press **Return** to continue with the installation.

During Live Upgrade, if the OS of the alternate boot disk is upgraded, the installer does not update the VCS configurations for Oracle, Netlsnr, and Sybase resources. If cluster configurations include these resources, you are prompted to run a list of commands to manually update the configurations after the cluster restarts from the alternate boot disks.

5 Verify that the version of the Veritas packages on the alternate boot disk is 6.1.

```
# pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSvcs
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/logs`.

Completing the Solaris 10 Live Upgrade

At the end of the process:

- The alternate boot environment is activated.
- The system is booted from the alternate boot disk.

To complete the Live Upgrade

1 Complete the Live upgrade process. Enter the following command on all nodes in the cluster.

```
# ./vcslufinish -u target_os_version
Live Upgrade finish on the Solaris release <5.10>
```

2 After the successful completion of `vxlustart`, if the system crashes or restarts before Live Upgrade completes successfully, you can remount the alternate disk using the following command:

```
# ./vxlustart -r -u target_os_version
```

Then, rerun the `vcslufinish` command from step 1

```
# ./vcslufinish -u target_os_version
```

3 **Note:** Do not use the `reboot`, `halt`, or `uadmin` commands to restart the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate boot environment.

You can ignore the following error if it appears: Error: boot environment *<dest.13445>* already mounted on *</altroot.5.10>*.

```
# shutdown -g0 -y -i6
```

4 If you want to upgrade the CP server systems that use VCS or SFHA to this version, make sure that you have upgraded all application clusters to this version. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the relevant Installation Guide.

Verifying the Solaris 10 Live Upgrade of VCS

To ensure that Live Upgrade has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

To verify that Live Upgrade completed successfully

1 Verify that the alternate boot environment is active.

```
# lustatus
```

If the alternate boot environment fails to be active, you can revert to the primary boot environment.

See "Reverting to the primary boot environment on a Solaris 10 system" on page 386.

2 Make sure that GAB ports a and h are up.

```
# gabconfig -a
Port a gen   39d901 membership 01
Port h gen   39d909 membership 01
```

3 Perform other verification as required to ensure that the new boot environment is configured correctly.

4 In a zone environment, verify the zone configuration.

Administering boot environments in Solaris 10 Live Upgrade

Use the following procedures to perform relevant administrative tasks for boot environments.

Reverting to the primary boot environment on a Solaris 10 system

If the alternate boot environment fails to start, you can revert to the primary boot environment.

On each node, start the system from the primary boot environment in the PROM monitor mode.

```
ok> boot disk0
```

where *disk0* is the primary boot disk.

Failure to perform this step can result in the operating system booting from the alternate boot environment after the restart.

The `vcslufinish` script displays the way to revert to primary boot environment. Here is a sample output.

```
Notes:
 ****************************************************************
 In case of a failure while booting to the target BE, the following
 process needs to be followed to fallback to the currently working
 boot environment:
 1. Enter the PROM monitor (ok prompt).
 2. Change the boot device back to the original boot environment
 by typing:
 setenv boot-device /pci@1c,600000/scsi@2/disk@0,0:a
 3. Boot to the original boot environment by typing:
 boot
 ****************************************************************
```

Switching the boot environment for Solaris SPARC

You do not have to perform the following procedures to switch the boot environment when you use the `vxlustart` and `vcslufinish` scripts to process Live Upgrade. You must perform the following procedures when you perform a manual Live Upgrade.

To switch the boot environment

1 Display the status of Live Upgrade boot environments.

```
# lustatus

Boot Environment  Is       Active Active    Can    Copy
Name              Complete Now    On Reboot Delete Status
----------------  -------- ------ --------- ------ ------
source.2657       yes      yes    yes       no     -
dest.2657         yes      no     no        yes    -
```

In this example, the primary boot environment is currently (source.2657). You want to activate the alternate boot environment (dest.2657).

2 Unmount any file systems that are mounted on the alternate boot environment (dest.2657).

```
# lufslist dest.2657

               boot environment name: dest.2657

Filesystem        fstype device size  Mounted on Mount Options
----------------- ------ ------------ ---------- -------------
/dev/dsk/c0t0d0s1 swap     4298342400 -          -
/dev/dsk/c0t0d0s0 ufs     15729328128 /          -
/dev/dsk/c0t0d0s5 ufs      8591474688 /var       -
/dev/dsk/c0t0d0s3 ufs      5371625472 /vxfs      -

# luumount dest.2657
```

3 Activate the Live Upgrade boot environment.

```
# luactivate dest.2657
```

4 Restart the system.

```
# shutdown -g0 -i6 -y
```

The system automatically selects the boot environment entry that was activated.

Performing Boot Environment upgrade on Solaris 11 systems

Perform the BE upgrade manually or use the installer. For VCS, the nodes do not form a cluster until all of the nodes are upgraded. At the end of the BE upgrade of the last node, all the nodes must boot from the alternate BE and join the cluster.

Table 25-4 Upgrading VCS using BE upgrade

Step	Description
Step 1	Create a new BE on the primary boot disk. See "Creating a new Solaris 11 BE on the primary boot disk" on page 388.
Step 2	Upgrade VCS using the installer. See "Upgrading VCS using the installer for upgrading BE on Solaris 11" on page 389.
	To upgrade only Solaris See the Oracle documentation on Oracle Solaris 11 operating system.
Step 3	Switch the alternate BE to be the new primary. See "Completing the VCS upgrade on BE on Solaris 11" on page 391.
Step 4	Verify BE upgrade of VCS. See "Verifying Solaris 11 BE upgrade" on page 391.

Creating a new Solaris 11 BE on the primary boot disk

At the end of the process, a new BE is created on the primary boot disk by cloning the primary BE.

To create a new BE on the primary boot disk

Perform the steps in this procedure on each node in the cluster.

1. View the list of BE in the primary disk.

   ```
   # beadm list
   ```

2. If you have solaris brand zones in running state for which zone root is on shared storage, set `AutoStart` to 0 for the service group containing zone resource.

   ```
   # hagrp -modify <group> AutoStart 0

   # haconf -dump
   ```

3. Create a new BE in the primary boot disk.

   ```
   # beadm create beName

   # beadm mount beName mountpoint
   ```

4. Reset `AutoStart` to 1 for the service group containing zone resource in step 2

   ```
   # hagrp -modify <group> AutoStart 1

   # haconf -dump
   ```

If VVR is configured, it is recommended that <*beName*> should have the value *altroot.5.11* and <*mountpoint*> should have the value `/altroot.5.11`.

Upgrading VCS using the installer for upgrading BE on Solaris 11

You can use the Symantec product installer to upgrade VCS on a BE.

On a node in the cluster, run the installer on the primary boot disk to upgrade VCS on all the nodes in the cluster.

At the end of the process, the VCS 6.1 is installed on the alternate BE.

To perform BE upgrade of VCS using the installer

1. Insert the product disc with VCS 6.1 or access your copy of the software on the network.
2. If you had solaris brand zones in running state in the present BE when you created alternate BE, set the publisher for package repository for BEs of each of the zones.

   ```
   # /usr/bin/pkg -R /altrootpath/zone-root/root
   set-publisher -g  /<path>/VRTSpkgs.p5p Symantec
   ```

   For example:

   ```
   # /usr/bin/pkg -R /altroot.5.11/export/home/zone1/root
   set-publisher -g /mnt/VRTSpkgs.p5p Symantec
   ```

3. Run the installer script specifying the root path as the alternate BE:

   ```
   # ./installer -upgrade -rootpath /altroot.5.11
   ```

4. Enter the names of the nodes that you want to upgrade to VCS 6.1.

   The installer displays the list of packages to be installed or upgraded on the nodes.

5. Press **Return** to continue with the installation.

   During BE upgrade, if the OS of the alternate BE is upgraded, the installer does not update the VCS configurations for Oracle, Netlsnr, and Sybase resources. If cluster configurations include these resources, you are prompted to run a list of commands to manually update the configurations after the cluster restarts from the alternate BE.

6. Verify that the version of the Veritas packages on the alternate BE is 6.1.

   ```
   # pkg -R /altroot.5.11 list VRTS\*
   ```

   For example:

   ```
   # pkg -R /altroot.5.11 list VRTSvcs
   ```

   Review the installation logs at `/altroot.5.11/opt/VRTS/install/logs`.

7. Unset the publisher set in step 2.

   ```
   # /usr/bin/pkg -R /altrootpath/zone-root/root
   unset-publisher Symantec
   ```

Completing the VCS upgrade on BE on Solaris 11

At the end of the process:

- The alternate BE is activated.
- The system is booted from the alternate BE.

To complete the BE upgrade

1. Activate the alternate BE.

   ```
   # beadm activate altroot.5.11
   ```

2. Stop application and VCS on all nodes.

   ```
   # hastop -all
   ```

   If you have enabled VVR,

3. Restart all the nodes in the cluster. The BE on the alternate disk is activated when you restart the nodes.

   Note: Do not use the `reboot`, `halt`, or `uadmin` commands to restart the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate BE.

   ```
   # shutdown -g0 -y -i6
   ```

4. If you want to upgrade the CP server systems that use VCS or SFHA to this version, make sure that you upgrade all application clusters to this version. Then, upgrade VCS or SFHA on the CP server systems.

   For instructions to upgrade VCS or SFHA on the CP server systems, see the *VCS or SFHA Installation Guide*.

Verifying Solaris 11 BE upgrade

To ensure that BE upgrade has completed successfully, verify that all the nodes have booted from the alternate BE and joined the cluster.

To verify that BE upgrade completed successfully

1. Verify that the alternate BE is active.

```
# beadm list
```

   If the alternate BE fails to be active, you can revert to the primary BE.

   See "Reverting to the primary BE on a Solaris 11 system" on page 393.

2. Make sure that GAB ports a and h are up.

```
# gabconfig -a
Port a gen   39d901 membership 01
Port h gen   39d909 membership 01
```

3. Perform other verification as required to ensure that the new BE is configured correctly.
4. In a zone environment, verify the zone configuration.

If you have set `AutoStart` to 0 for the service group containing zone resource earlier, perform the following steps:

- Verify whether the zpool on which the root file system of the zone is residing is imported

```
# zpool list
```

  If not imported, `online` the zpool resource.

- Attach the zone.

```
# zoneadm -z <zone> attach
```

- Reset `AutoStart` to 1 for the service group containing zone resource.

```
# hagrp -modify <group> AutoStart 1
```

If you have solaris10 brand zone on your system, you must manually upgrade the packages inside the solaris10 brand zone with packages from Solaris 10 install media.

If you have installed `VRTSvxfs` or `VRTSodm` packages inside the zones, you need to manually upgrade these packages inside the zone.

Administering BEs in Solaris 11 BE upgrade

Use the following procedures to perform relevant administrative tasks for BEs.

Switching the BE for Solaris SPARC

1 Display the status of BEs.

```
# beadm list

Boot Emvironment Active Mountpoint Space  Policy Created
--               ------ ---------  -----  ------ -------
solaris          NR        /       13.08G static 2012-11-14 10:22
altroot.5.11     -         -        3.68G static 2013-01-06 18:41
```

In this example, the primary boot disk is currently *solaris*. You want to activate the alternate boot disk *altroot.5.11*.

2 Activate the BE.

```
# beadm activate altroot.5.11
```

3 Restart the system to complete the BE activation.

```
# shutdown -g0 -i6 -y
```

The system automatically selects the BE entry that was activated.

4 You can destroy an existing BE.

```
# beadm destroy altroot.5.11
```

Reverting to the primary BE on a Solaris 11 system

Boot the system to `ok` prompt.

View the available BEs.

To view the BEs, enter the following:

```
ok> boot -L
```

Select the option of the original BE to which you need to boot.

To boot to the BE

```
# boot -Z <path to boot env>
```

For example:

```
{0} ok boot -L
Boot device: /virtual-devices@100/channel-devices@200/disk@0:a
File and args: -L
1 Oracle Solaris 11 11/11 SPARC
```

```
2 solaris-backup-1
Select environment to boot: [ 1 - 2 ]: 1
```

To boot the selected entry, enter the following:

```
boot [<root-device>] -Z rpool/ROOT/solaris

Program terminated
{0} ok boot -Z rpool/ROOT/solaris
```

Section 9

Post-installation tasks

- Chapter 26. Performing post-installation tasks
- Chapter 27. Installing or upgrading VCS components
- Chapter 28. Verifying the VCS installation

Chapter 26

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
- Changing root user into root role

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 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.

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

See the *Symantec Cluster Server Administrator's Guide* for instructions to add VCS users.

Figure 26-1 depicts the VCS cluster communication with the LDAP servers when clusters run in secure mode.

Figure 26-1 Client communication with LDAP servers

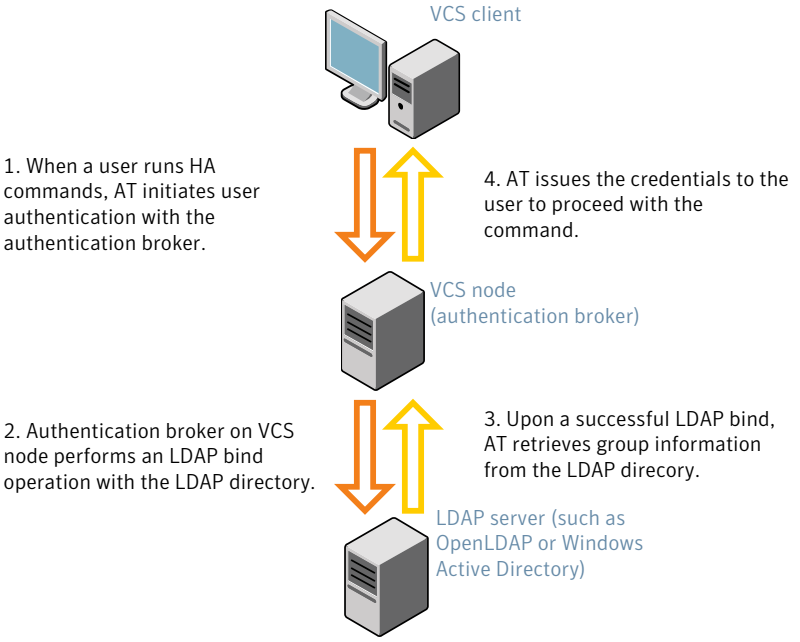

The LDAP schema and syntax for LDAP commands (such as, ldapadd, ldapmodify, and ldapsearch) vary based on your LDAP implementation.

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 6.1.6.0 or later.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion
vssat version: 6.1.6.0
```

To enable OpenLDAP 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/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-d -s domain_controller_name_or_ipaddress -u domain_user

Attribute list file name not provided, using AttributeList.txt

Attribute file created.
```

You can use the `cat` command to view the entries in the attributes file.

2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-c -d LDAP_domain_name

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/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -x

Using default broker port 14149

CLI file not provided, using default CLI.txt

Looking for AT installation...

AT found installed at ./vssat

Successfully added LDAP domain.
```

4 Check the AT version and list the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion

vssat version: 6.1.12.8

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat listldapdomains

Domain Name : mydomain.com

Server URL : ldap://192.168.20.32:389

SSL Enabled : No

User Base DN : CN=people,DC=mydomain,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 Check the other domains in the cluster.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showdomains -p vx
```

The command output lists the number of domains that are found, with the domain names and domain types.

6 Generate credentials for the user.

```
# unset EAT_LOG

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat authenticate \
-d ldap:LDAP_domain_name -p user_name -s user_password -b \
localhost:14149
```

7 Add non-root users as applicable.

```
# useradd user1

# passwd pw1

Changing password for "user1"

user1's New password:

Re-enter user1's new password:

# su user1

# bash

# id

uid=204(user1) gid=1(staff)

# pwd

# mkdir /home/user1

# chown user1 /home/ user1
```

8 Log in as non-root user and run VCS commands as LDAP user.

```
# cd /home/user1

# ls

# cat .vcspwd

101 localhost mpise LDAP_SERVER ldap

# unset VCS_DOMAINTYPE

# unset VCS_DOMAIN

# /opt/VRTSvcs/bin/hasys -state

   #System        Attribute          Value

cluster1:sysA     SysState           FAULTED

cluster1:sysB     SysState           FAULTED

cluster2:sysC     SysState           RUNNING

cluster2:sysD     SysState           RUNNING
```

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.

Changing root user into root role

On Oracle Solaris 11, you need to create root user to perform installation. This means that a local user cannot assume the root role. After installation, you may want to turn root user into root role for a local user, who can log in as root.

1. Log in as root user.
2. Change the root account into role.

```
# rolemod -K type=role root

# getent user_attr root

root::::type=role;auths=solaris.*;profiles=All;audit_flags=lo\
:no;lock_after_retries=no;min_label=admin_low;clearance=admin_high
```

3. Assign the root role to a local user who was unassigned the role.

```
# usermod -R root admin
```

For more information, see the Oracle documentation on Oracle Solaris 11 operating system.

Chapter 27

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

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 Solaris system with X-Windows. 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.

When you install the Java Console on the Solaris system, make sure a printer is configured to that system. If you print the online JavaHelp on a system that does not have a printer that is configured, the Java Console might hang.

Review the information about using the Java Console. For more information, refer to the *Symantec Cluster Server Administrator's Guide*.

Software requirements for the Java Console

Cluster Manager (Java Console) is supported on:

- Solaris SPARC 2.10
- Windows XP and Windows 2003

Note: Make sure that you are using an operating system version that supports JRE 1.6.

Hardware requirements for the Java Console

The minimum hardware requirements for the Java Console are as follows:

- 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 or higher, 256MB RAM or higher, and 800x600 display resolution or higher.

The version of the Java™ 2 Runtime Environment (JRE) requires 32 megabytes of RAM.

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 Solaris

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 Solaris

1 Create a directory for installation of the Java Console:

```
# mkdir /tmp/install
```

2 Download the Java GUI utility from http://www.symantec.com/operations-manager/support to a temporary directory.

3 Go to the temporary directory and unzip the compressed package file using the gunzip utility:

```
# cd /tmp/install
# gunzip VRTScscm.tar.gz
```

The file VRTScscm.tar is now present in the temporary directory.

4 Extract the compressed file from the tar file:

```
# tar -xvf VRTScscm.tar
```

5 Install the software:

```
# pkgadd -d . VRTScscm
```

6 Answer `Yes` if prompted.

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://www.symantec.com/operations-manager/support to a temporary directory.

2 Extract the zipped file to a temporary folder.

3 From this extracted folder, double-click setup.exe.

4 The Symantec 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 Solaris

1. Log in as superuser on the node where you intend to install the package.
2. Remove the GUI from the previous installation.

   ```
   # pkgrm VRTScscm
   ```

3. Install the VCS Java console.

   See "Installing the Java Console on Solaris" on page 405.

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 406.

Installing VCS Simulator

You can administer VCS Simulator from the Java Console or from the command line. For more information, see the *Symantec Cluster Server Administrator's Guide*.

Review the software requirements for VCS Simulator.

Software requirements for VCS Simulator

VCS Simulator is supported on:

- Windows XP SP3, Windows 2008, Windows Vista, and Windows 7

Note: Make sure that you are using an operating system version that supports JRE 1.6 or later.

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:

   \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 407.

Chapter 28

Verifying the VCS installation

This chapter includes the following topics:

- About verifying the VCS installation
- About the cluster UUID
- Verifying the LLT, GAB, and VCS configuration files
- Verifying LLT, GAB, and cluster operation
- Upgrading the disk group version
- Performing a postcheck on a node

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 cluster UUID

You can verify the existence of the cluster UUID.

To verify that the cluster UUID exists

- From the prompt, run a cat command.

```
cat /etc/vx/.uuids/clusuuid
```

To display UUID of all the nodes in the cluster

- From the prompt, run the command from any node.

```
/opt/VRTSvcs/bin/uuidconfig.pl -rsh -clus -display -use_llthost
```

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/lltab
   - 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 526.

   See "About the VCS configuration files" on page 530.

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 76.
3. Verify LLT operation.

   See "Verifying LLT" on page 412.
4. Verify GAB operation.

   See "Verifying GAB" on page 414.

5 Verify the cluster operation.

See “Verifying the cluster” on page 415.

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 sys1.

2 Run the `lltstat` command on the node sys1 to view the status of LLT.

```
lltstat -n
```

The output on sys1 resembles:

```
LLT node information:
    Node             State           Links
    *0 sys1          OPEN           2
     1 sys2          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 sys1             OPEN       2
     1 sys2             OPEN       2
     2 sys5             OPEN       1
```

3 Log in as superuser on the node sys2.

4 Run the `lltstat` command on the node sys2 to view the status of LLT.

```
lltstat -n
```

The output on sys2 resembles:

```
LLT node information:
    Node             State           Links
     0 sys1          OPEN            2
    *1 sys2          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 sys1 in a two-node cluster:

```
lltstat -nvv active
```

The output on sys1 resembles the following:

- For Solaris SPARC:

```
Node          State      Link    Status      Address
*0 sys1       OPEN
                         net:0 UP      08:00:20:93:0E:34
                         net:1 UP      08:00:20:93:0E:38
1 sys2        OPEN
                         net:0 UP      08:00:20:8F:D1:F2
                         net:1 DOWN
```

The command reports the status on the two active nodes in the cluster, sys1 and sys2.

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

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

However, the output in the example shows different details for the node sys2. The private network connection is possibly broken or the information in the /etc/lltab 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 sys1 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 ... 60 61 62 63
        connects:    0 1
  7     gab          0x7
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
        connects:    0 1
  31    gab          0x1F
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
        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: Port b appears in the `gabconfig` command output only if you had configured I/O fencing after you configured VCS.
- 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 *Symantec 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:

For a cluster where I/O fencing is not configured:

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

For a cluster where I/O fencing is configured:

```
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 appears in the `gabconfig` command output only if you had configured I/O fencing. You can also 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 *Symantec 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  sys1                 RUNNING                 0
A  sys2                 RUNNING                 0

-- GROUP STATE
-- Group           System      Probed  AutoDisabled   State

B  ClusterService  sys1          Y       N            ONLINE
B  ClusterService  sys2          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 sys1 and OFFLINE on sys2.

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 *Symantec 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 in the following procedure is for SPARC and it shows the output when the command is run on the node sys1. The list continues with similar information for sys2 (not shown) and any other nodes in the cluster.

```
#System    Attribute             Value

sys1       AgentsStopped         0

sys1       AvailableCapacity     100

sys1       CPUBinding            BindTo None CPUNumber 0

sys1       CPUThresholdLevel     Critical 90 Warning 80 Note 70
                                 Info 60

sys1       CPUUsage              0

sys1       CPUUsageMonitoring    Enabled 0 ActionThreshold 0
                                 ActionTimeLimit 0 Action NONE
                                 NotifyThreshold 0 NotifyTimeLimit 0

sys1       Capacity              100

sys1       ConfigBlockCount      130

sys1       ConfigCheckSum        46688

sys1       ConfigDiskState       CURRENT

sys1       ConfigFile            /etc/VRTSvcs/conf/config

sys1       ConfigInfoCnt         0

sys1       ConfigModDate         Mon Sep 03 07:14:23 CDT 2012

sys1       ConnectorState        Up

sys1       CurrentLimits

sys1       DiskHbStatus

sys1       DynamicLoad           0

sys1       EngineRestarted       0

sys1       EngineVersion         6.1.00.0

sys1       FencingWeight         0

sys1       Frozen                0

sys1       GUIIPAddr

sys1       HostUtilization       CPU 0 Swap 0
```

```
sys1        LLTNodeId               0
sys1        LicenseType             PERMANENT_SITE
sys1        Limits
sys1        LinkHbStatus            net:0 UP net:1 UP
sys1        LoadTimeCounter         0
sys1        LoadTimeThreshold       600
sys1        LoadWarningLevel        80
sys1        NoAutoDisable           0
sys1        NodeId                  0
sys1        OnGrpCnt                7
sys1        PhysicalServer
sys1        ShutdownTimeout         600
sys1        SourceFile              ./main.cf
sys1        SwapThresholdLevel      Critical 90 Warning 80 Note 70
                                    Info 60
sys1        SysInfo                 Solaris:sys1,Generic_
                                    118558-11,5.9,SUN4u
sys1        SysName                 sys1
sys1        SysState                RUNNING
sys1        SystemLocation
sys1        SystemOwner
sys1        SystemRecipients
sys1        TFrozen                 0
sys1        TRSE                    0
sys1        UpDownState             Up
sys1        UserInt                 0
sys1        UserStr
```

```
sys1       VCSFeatures            DR

sys1       VCSMode                VCS
```

Upgrading the disk group version

After you upgrade from previous versions to 6.1, you have to upgrade the disk group version manually.

To upgrade disk group version, you have to first upgrade the cluster protocol version using the `vxdctl upgrade` command.

```
# vxdctl list
Volboot file
version: 3/1
seqno:   0.1
cluster protocol version: 120
hostid:  sys1
hostguid:  {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
#
# vxdctl upgrade
#

# vxdctl list

Volboot file
version: 3/1
seqno:   0.2
cluster protocol version: 130
hostid:  sys1
hostguid:  {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
```

Verify if the cluster protocol version shows 130 and disk group version is upgraded to 190.

```
# vxdg list dg_name |grep version

version:   130
#
# vxdg upgrade dg_name
#
# vxdg list dg_name |grep version
```

```
version:    190
```

Performing a postcheck on a node

The installer's `postcheck` command can help you to determine installation-related problems and provide troubleshooting information.

See "About using the postcheck option" on page 420.

To run the postcheck command on a node

1 Run the installer with the `-postcheck` option.

```
# ./installer -postcheck system_name
```

2 Review the output for installation-related information.

About using the postcheck option

You can use the installer's post-check to determine installation-related problems and to aid in troubleshooting.

Note: This command option requires downtime for the node.

When you use the `postcheck` option, it can help you troubleshoot the following VCS-related issues:

- The heartbeat link does not exist.
- The heartbeat link cannot communicate.
- The heartbeat link is a part of a bonded or aggregated NIC.
- A duplicated cluster ID exists (if LLT is not running at the check time).
- The VRTSllt pkg version is not consistent on the nodes.
- The llt-linkinstall value is incorrect.
- The llthosts(4) or lltttab(4) configuration is incorrect.
- the `/etc/gabtab` file is incorrect.
- The incorrect GAB linkinstall value exists.
- The VRTSgab pkg version is not consistent on the nodes.
- The `main.cf` file or the `types.cf` file is invalid.

- The `/etc/VRTSvcs/conf/sysname` file is not consistent with the hostname.
- The cluster UUID does not exist.
- The `uuidconfig.pl` file is missing.
- The VRTSvcs pkg version is not consistent on the nodes.
- The `/etc/vxfenmode` file is missing or incorrect.
- The `/etc/vxfendg file` is invalid.
- The vxfen link-install value is incorrect.
- The VRTSvxfen pkg version is not consistent.

The `postcheck` option can help you troubleshoot the following SFHA or SFCFSHA issues:

- Volume Manager cannot start because the `/etc/vx/reconfig.d/state.d/install-db` file has not been removed.
- Volume Manager cannot start because the `volboot` file is not loaded.
- Volume Manager cannot start because no license exists.
- Cluster Volume Manager cannot start because the CVM configuration is incorrect in the `main.cf` file. For example, the Autostartlist value is missing on the nodes.
- Cluster Volume Manager cannot come online because the node ID in the `/etc/llthosts` file is not consistent.
- Cluster Volume Manager cannot come online because Vxfen is not started.
- Cluster Volume Manager cannot start because gab is not configured.
- Cluster Volume Manager cannot come online because of a CVM protocol mismatch.
- Cluster Volume Manager group name has changed from "cvm", which causes CVM to go offline.

You can use the installer's post-check option to perform the following checks:

General checks for all products:

- All the required packages are installed.
- The versions of the required packages are correct.
- There are no verification issues for the required packages.

Checks for Volume Manager (VM):

- Lists the daemons which are not running (`vxattachd`, `vxconfigbackupd`, `vxesd`, `vxrelocd` ...).

- Lists the disks which are not in 'online' or 'online shared' state (`vxdisk list`).
- Lists the diskgroups which are not in 'enabled' state (`vxdg list`).
- Lists the volumes which are not in 'enabled' state (`vxprint -g <dgname>`).
- Lists the volumes which are in 'Unstartable' state (`vxinfo -g <dgname>`).
- Lists the volumes which are not configured in `/etc/vfstab`.

Checks for File System (FS):

- Lists the VxFS kernel modules which are not loaded (`vxfs/fdd/vxportal`.).
- Whether all VxFS file systems present in `/etc/vfstab` file are mounted.
- Whether all VxFS file systems present in `/etc/vfstab` are in disk layout 6 or higher.
- Whether all mounted VxFS file systems are in disk layout 6 or higher.

Checks for Cluster File System:

- Whether FS and ODM are running at the latest protocol level.
- Whether all mounted CFS file systems are managed by VCS.
- Whether cvm service group is online.

See “Performing a postcheck on a node” on page 420.

Section 10

Adding and removing cluster nodes

- Chapter 29. Adding a node to a single-node cluster
- Chapter 30. Adding a node to a multi-node VCS cluster
- Chapter 31. Removing a node from a VCS cluster

Chapter 29

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 29-1 specifies the activities that you need to perform to add nodes to a single-node cluster.

Table 29-1 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 425.
■ 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 426.
Connect both nodes to shared storage.	See "Configuring the shared storage" on page 427.

Table 29-1 Tasks to add a node to a single-node cluster *(continued)*

Task	Reference
■ Bring up VCS on Node A. ■ Edit the configuration file.	See "Bringing up the existing node" on page 427.
If necessary, install VCS on Node B and add a license key. Make sure Node B is running the same version of VCS as the version on Node A.	See "Installing the VCS software manually when adding a node to a single node cluster" on page 428.
Edit the configuration files on Node B.	See "About the VCS configuration files" on page 530.
Start LLT and GAB on Node B.	See "Starting LLT and GAB" on page 428.
■ Start LLT and GAB on Node A. ■ Copy UUID from Node A to Node B. ■ Restart VCS on Node A. ■ Modify service groups for two nodes.	See "Reconfiguring VCS on the existing node" on page 428.
■ Start VCS on Node B. ■ Verify the two-node cluster.	See "Verifying configuration on both nodes" on page 430.

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 packages and configuration files.
  See "Removing a node from a VCS cluster" on page 450.
- If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS.
- If you renamed the LLT and GAB startup files, remove them.

2 If necessary, install VxVM and VxFS.

See “Installing VxVM or VxFS if necessary” on page 426.

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 67.

To install and configure Ethernet cards for private network

1 Shut down VCS on Node A.

```
# hastop -local
```

2 Shut down the node to get to the `OK` prompt:

```
# sync;sync;init 0
```

3 Install the Ethernet card on Node A.

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 71.

Bringing up the existing node

Bring up the node.

To bring up the node

1. Start the operating system. On a SPARC node (Node A) enter the command:

   ```
   ok boot -r
   ```

2. Log in as superuser.
3. Make the VCS configuration writable.

   ```
   # haconf -makerw
   ```

4. Display the service groups currently configured.

   ```
   # hagrp -list
   ```

5. Freeze the service groups.

   ```
   # hagrp -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. If you have configured I/O Fencing, GAB, and LLT on the node, stop them.

   ```
   # /usr/sbin/svcadm disable -t gab
   ```

   ```
   # /usr/sbin/svcadm disable -t llt
   ```

Installing the VCS software manually when adding a node to a single node cluster

Install the VCS 6.1 packages manually and install the license key.

Refer to the following sections:

- See “Adding a license key for a manual installation” on page 245.

Creating configuration files

Create the configuration files for your cluster.

To create the configuration files

1. Create the file /etc/llttab for a two-node cluster

   See “Setting up /etc/llttab for a manual installation” on page 260.
2. Create the file /etc/llthosts that list both the nodes.

   See “Setting up /etc/llthosts for a manual installation” on page 260.
3. Create the file /etc/gabtab.

   See “Configuring GAB manually” on page 263.

Starting LLT and GAB

On the new node, start LLT and GAB.

To start LLT and GAB

1. Start LLT on Node B.

   ```
   # /usr/sbin/svcadm enable llt
   ```
2. Start GAB on Node B

   ```
   # /usr/sbin/svcadm enable gab
   ```

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/lltab, /etc/llthosts, and /etc/gabtab. Use the files that are created on Node B as a guide, customizing the /etc/lltab for Node A.
2. Start LLT on Node A.

   ```
   # /usr/sbin/svcadm enable llt
   ```

3. Start GAB on Node A.

   ```
   # /usr/sbin/svcadm enable gab
   ```

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. Delete the VCS one-node SMF configuration from SMF respository.

   ```
   # svccfg -f delete vcs-onenode
   ```

7. Import the VCS SMF configuration to SMF respository.

   ```
   # svccfg import /etc/VRTSvcs/conf/vcs.xml
   ```

   Note: To start VCS using SMF service, use the `svcadm enable vcs` command.

8. Start VCS on Node A.

   ```
   # hastart
   ```

9. Make the VCS configuration writable.

   ```
   # haconf -makerw
   ```

10 Add Node B to the cluster.

```
# hasys -add sysB
```

11 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 Save the new two-node configuration.

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

Chapter 30

Adding a node to a multi-node VCS cluster

This chapter includes the following topics:

- Adding nodes using the VCS installer
- Adding a node using the web-based installer
- Manually adding a node to a cluster
- Adding a node using response files

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/llttab`

  `/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/default/llt
/etc/default/gab
/etc/default/vxfen

- 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<version> -addnode
   ```

   Where <version> is specific to the release version.

   See "About the script-based installer" on page 49.

   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: sys1
   ```

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: sys5
   ```

   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 LLT configuration for the new node must be the same as that of the existing cluster. If your existing cluster uses LLT over UDP, the installer asks questions related to LLT over UDP for the new node.

See "Configuring private heartbeat links" on page 139.

```
Enter the NIC for the first private heartbeat
link on sys5: [b,q,?] net:0
```

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 sys5: [b,q,?] net:1
```

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 sys5: net:2
```

Adding a node using the web-based installer

You can use the web-based installer to add a node to a cluster.

To add a node to a cluster using the web-based installer

1. From the Task pull-down menu, select **Add a Cluster node**.

   From the product pull-down menu, select the product.

   Click the **Next** button.

2. Click **OK** to confirm the prerequisites to add a node.

3. In the System Names field enter a name of a node in the cluster where you plan to add the node and click **OK**.

   The installer program checks inter-system communications and compatibility. If the node fails any of the checks, review the error and fix the issue.

   If prompted, review the cluster's name, ID, and its systems. Click the **Yes** button to proceed.

4. In the System Names field, enter the names of the systems that you want to add to the cluster as nodes. Separate system names with spaces. Click the **Next** button.

   The installer program checks inter-system communications and compatibility. If the system fails any of the checks, review the error and fix the issue.

   Click the **Next** button. If prompted, click the **Yes** button to add the system and to proceed.

5. From the heartbeat NIC pull-down menus, select the heartbeat NICs for the cluster. Click the **Next** button.

6. Once the addition is complete, review the log files. Optionally send installation information to Symantec. Click the **Finish** button to complete the node's addition to the cluster.

Manually adding a node to a cluster

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

See "Hardware requirements for VCS" on page 37.

Table 30-1 specifies the tasks that are involved in adding a cluster. The example demonstrates how to add a node saturn to already existing nodes, sys1 and sys2.

Table 30-1 Tasks that are involved in adding a node to a cluster

Task	Reference
Set up the hardware.	See "Setting up the hardware" on page 435.

Table 30-1 Tasks that are involved in adding a node to a cluster *(continued)*

Task	Reference
Install the software manually.	See "Installing VCS packages for a manual installation" on page 240.
Add a license key.	See "Adding a license key for a manual installation" on page 245.
Configure LLT and GAB.	See "Configuring LLT and GAB when adding a node to the cluster" on page 440.
Copy the UUID.	See "Reconfiguring VCS on the existing node" on page 428.
Add the node to the existing cluster.	See "Adding the node to the existing cluster" on page 446.
Start VCS and verify the cluster.	See "Starting VCS and verifying the cluster" on page 447.

Setting up the hardware

Figure 30-1 shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

Figure 30-1 Adding a node to a two-node cluster using two switches

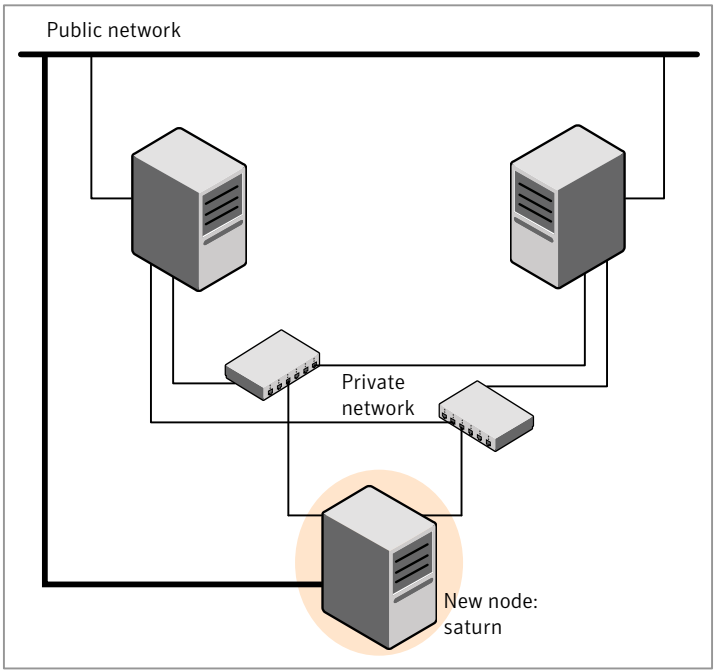

To set up the hardware

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 30-1 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 6.1 packages manually and add a license key.

For more information, see the following:

- See “Installing VCS software manually” on page 238.
- See “Adding a license key for a manual installation” on page 245.

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 when adding a node to the cluster” on page 440.

Table 30-2 uses the following information for the following command examples.

Table 30-2 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
sys5	sys5.nodes.example.com	The new node that you are adding to the cluster.

Configuring the authentication broker on node sys5

To configure the authentication broker on node sys5

1 Extract the embedded authentication files and copy them to temporary directory:

```
# mkdir -p /var/VRTSvcs/vcsauth/bkup

# cd /tmp; gunzip -c /opt/VRTSvcs/bin/VxAT.tar.gz | tar xvf -
```

2 Edit the setup file manually:

```
# cat /etc/vx/.uuids/clusuuid 2>&1
```

The output is a string denoting the UUID. This UUID (without { and }) is used as the ClusterName for the setup file.

```
{UUID}

# cat /tmp/eat_setup 2>&1
```

The file content must resemble the following example:

```
AcceptorMode=IP_ONLY

BrokerExeName=vcsauthserver

ClusterName=UUID

DataDir=/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER

DestDir=/opt/VRTSvcs/bin/vcsauth/vcsauthserver

FipsMode=0

IPPort=14149

RootBrokerName=vcsroot_uuid

SetToRBPlusABorNot=0

SetupPDRs=1

SourceDir=/tmp/VxAT/version
```

3 Set up the embedded authentication file:

```
# cd /tmp/VxAT/version/bin/edition_number; \
./broker_setup.sh/tmp/eat_setup

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssregctl -s -f
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER/root/.VRTSat/profile \
/VRTSatlocal.conf -b 'Security\Authentication \
\Authentication Broker' -k UpdatedDebugLogFileName \
-v /var/VRTSvcs/log/vcsauthserver.log -t string
```

4 Copy the broker credentials from one node in the cluster to sys5 by copying the entire `bkup` directory.

The `bkup` directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/

# ls

CMDSERVER  HAD  VCS_SERVICES  WAC
```

5 Import the VCS_SERVICES domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER -f /var/VRTSvcs/vcsauth/bkup \
/VCS_SERVICES -p password
```

6 Import the credentials for HAD, CMDSERVER, and WAC.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCS_SERVICES -f /var/VRTSvcs/vcsauth/bkup \
/HAD -p password
```

7 Start the vcsauthserver process on sys5.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh
```

8 Perform the following tasks:

```
# mkdir /var/VRTSvcs/vcsauth/data/CLIENT

# mkdir /var/VRTSvcs/vcsauth/data/TRUST

# export EAT_DATA_DIR='/var/VRTSvcs/vcsauth/data/TRUST'

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat setuptrust -b \
localhost:14149 -s high
```

9 Create the `/etc/VRTSvcs/conf/config/.secure` file:

```
# touch /etc/VRTSvcs/conf/config/.secure
```

Configuring LLT and GAB when adding a node to the cluster

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

To configure LLT when adding a node to the cluster

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 sys5 to a cluster consisting of sys1 and sys2:

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

```
0 sys1
1 sys2
```

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

```
0 sys1
1 sys2
2 sys5
```

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

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

The following example describes a system where node sys2 is the new node on cluster ID number 2:

- For Solaris SPARC:

```
set-node sys2
set-cluster 2
link net1 net:0 - ether - -
link net2 net:1 - ether - -
```

3 Copy the following file from one of the nodes in the existing cluster to the new node:

/etc/default/llt

4 On the new system, run the command:

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

In a setup that uses LLT over UDP, new nodes automatically join the existing cluster if the new nodes and all the existing nodes in the cluster are not separated by a router. However, if you use LLT over UDP6 link with IPv6 address and if the new node and the existing nodes are separated by a router, then do the following:

- Edit the `/etc/lltab` file on each node to reflect the link information about the new node.
- Specify the IPv6 address for UDP link of the new node to all existing nodes. Run the following command on each existing node for each UDP link:

```
# /sbin/lltconfig -a set systemid device_tag address
```

To configure GAB when adding a node to the cluster

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/default/gab

3 On the new node, to configure GAB run the command:

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

To verify GAB

1 On the new node, run the command:

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

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

```
GAB Port Memberships
===================================
Port a gen a3640003 membership 012
```

See "Verifying GAB" on page 414.

2 Run the same command on the other nodes (sys1 and sys2) 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 443.
- 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 443.

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 445.

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 446.

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:

   ```
   # vxfenadm -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 *Symantec 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 cps1.symantecexample.com \
-a add_node -c clus1 -h sys5 -n2

Node 2 (sys5) successfully added
```

3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s cps1.symantecexample.com \
-a list_nodes
```

The new node must be listed in the command output.

4 Add the VCS user cpsclient@sys5 to each CP server:

```
# cpsadm -s cps1.symantecexample.com \
-a add_user -e cpsclient@sys5 \
-f cps_operator -g vx

User cpsclient@sys5 successfully added
```

To configure server-based fencing with security 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 cps1.symantecexample.com \
-a add_node -c clus1 -h sys5 -n2

Node 2 (sys5) successfully added
```

3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s cps1.symantecexample.com -a list_nodes
```

The new node must be listed in the output.

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 sys5 to the existing vxfen group.

```
# hagrp -modify vxfen SystemList -add sys5 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/default/vxfen

2 Start I/O fencing on the new node.

```
# svcadm enable vxfen
```

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 Enter the command:

```
# haconf -makerw
```

2 Add the new system to the cluster:

```
# hasys -add sys1
```

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

```
# rcp /etc/VRTSvcs/conf/config/main.cf \
  sys5:/etc/VRTSvcs/conf/config/
```

4 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/
```

5 If necessary, modify any new system attributes.

6 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 To start VCS service using SMF, use the following command:

```
# svcadm enable vcs
```

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
```

Adding a node using response files

Typically, you can use the response file that the installer generates on one system to add nodes to an existing cluster.

To add nodes using response files

1 Make sure the systems where you want to add nodes meet the requirements.

2 Make sure all the tasks required for preparing to add a node to an existing VCS cluster are completed.

3 Copy the response file to one of the systems where you want to add nodes.

See "Sample response file for adding a node to a VCS cluster" on page 448.

4 Edit the values of the response file variables as necessary.

See "Response file variables to add a node to a VCS cluster" on page 448.

5 Mount the product disc and navigate to the folder that contains the installation program.

6 Start adding nodes 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.

Depending on the fencing configuration in the existing cluster, the installer configures fencing on the new node. The installer then starts all the required Symantec processes and joins the new node to cluster. The installer indicates the location of the log file and summary file with details of the actions performed.

Response file variables to add a node to a VCS cluster

Table 30-3 lists the response file variables that you can define to add a node to an VCS cluster.

Table 30-3 Response file variables for adding a node to an VCS cluster

Variable	Description
$CFG{opt}{addnode}	Adds a node to an existing cluster. List or scalar: scalar Optional or required: required
$CFG{newnodes}	Specifies the new nodes to be added to the cluster. List or scalar: list Optional or required: required

Sample response file for adding a node to a VCS cluster

The following example shows a response file for upgrading VCS.

```
our %CFG;

$CFG{clustersystems}=[ qw(sys1) ];
$CFG{newnodes}=[ qw(sys5) ];
$CFG{nic_add_ip_to_files}=1;
$CFG{opt}{addnode}=1;
```

```
$CFG{opt}{configure}=1;
$CFG{opt}{vr}=1;
$CFG{prod}="VCS60";
$CFG{sys5}{net:0}{haipip}="192.168.12.3";
$CFG{sys5}{net:0}{hostname_for_haip}="sys5-haip1";
$CFG{sys5}{net:1}{haipip}="192.168.13.3";
$CFG{sys5}{net:1}{hostname_for_haip}="sys5-haip2";
$CFG{systems}=[ qw(sys1 sys5) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_lltlink1}{sys5}="net:0";
$CFG{vcs_lltlink2}{sys5}="net:1";

1;
```

Chapter 31

Removing a node from a VCS cluster

This chapter includes the following topics:

- Removing a node from a VCS cluster

Removing a node from a VCS cluster

Table 31-1 specifies the tasks that are involved in removing a node from a cluster. In the example procedure, the cluster consists of nodes sys1, sys2, and sys5; node sys5 is to leave the cluster.

Table 31-1 Tasks that are involved in removing a node

Task	Reference
■ Back up the configuration file. ■ Check the status of the nodes and the service groups.	See "Verifying the status of nodes and service groups" on page 451.
■ Switch or remove any VCS service groups on the node departing the cluster. ■ Delete the node from VCS configuration.	See "Deleting the departing node from VCS configuration" on page 452.
Modify the llthosts(4) and gabtab(4) files to reflect the change.	See "Modifying configuration files on each remaining node" on page 455.
For a cluster that is running in a secure mode, remove the security credentials from the leaving node.	See "Removing security credentials from the leaving node " on page 457.

Table 31-1 Tasks that are involved in removing a node *(continued)*

Task	Reference
On the node departing the cluster: ■ Modify startup scripts for LLT, GAB, and VCS to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. ■ Remove the VCS packages.	See "Unloading LLT and GAB and removing VCS on the departing node" on page 458.

Verifying the status of nodes and service groups

Start by issuing the following commands from one of the nodes to remain in the cluster node sys1 or node sys2 in our example.

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  sys1    RUNNING        0
A  sys2    RUNNING        0
A  sys5      RUNNING        0

-- GROUP STATE
-- Group    System      Probed   AutoDisabled   State
B  grp1     sys1      Y        N              ONLINE
B  grp1     sys2      Y        N              OFFLINE
B  grp2     sys1      Y        N              ONLINE
B  grp3     sys2      Y        N              OFFLINE
B  grp3     sys5    Y        N              ONLINE
B  grp4     sys5    Y        N              ONLINE
```

The example output from the `hastatus` command shows that nodes sys1, sys2, and sys5 are the nodes in the cluster. Also, service group grp3 is configured to run on node sys2 and node sys5, the departing node. Service group grp4 runs only on node sys5. Service groups grp1 and grp2 do not run on node sys5.

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 sys5 to node sys2.

```
# hagrp -switch grp3 -to sys2
```

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

```
# hagrp -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
# hagrp -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 sys5
```

To stop VCS using SMF, run the following command:

```
# svcadm disable vcs
```

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  sys1    RUNNING        0
    A  sys2      RUNNING        0
    A  sys5      EXITED         0

    -- GROUP STATE
    -- Group    System      Probed   AutoDisabled   State
    B  grp1     sys1      Y        N              ONLINE
    B  grp1     sys2      Y        N              OFFLINE
    B  grp2     sys1      Y        N              ONLINE
    B  grp3     sys2      Y        N              ONLINE
    B  grp3     sys5    Y        Y              OFFLINE
    B  grp4     sys5    Y        N              OFFLINE
```

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

```
# haconf -makerw
# hagrp -modify grp3 SystemList -delete sys5
# hagrp -modify grp4 SystemList -delete sys5
```

Note: If sys5 was in the autostart list, then you need to manually add another system in the autostart list so that after reboot, the group comes online automatically.

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

```
# hagrp -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.

```
# hagrp -delete grp4
```

9 Check the status.

```
# hastatus -summary
    -- SYSTEM STATE
    -- System      State          Frozen
    A  sys1      RUNNING        0
    A  sys2      RUNNING        0
    A  sys5    EXITED         0

    -- GROUP STATE
    -- Group    System      Probed   AutoDisabled   State
    B  grp1     sys1      Y        N              ONLINE
    B  grp1     sys2      Y        N              OFFLINE
    B  grp2     sys1      Y        N              ONLINE
    B  grp3     sys2      Y        N              ONLINE
```

10 Delete the node from the cluster.

```
# hasys -delete sys5
```

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 `-n`*N* option, where *N* is the number of cluster systems.

If the command has the form `/sbin/gabconfig -c -n`*N*, 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/llthosts file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 sys1
1 sys2
2 sys5
```

To:

```
0 sys1
1 sys2
```

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.

Note: The `cpsadm` command is used to perform the steps in this procedure. For detailed information about the `cpsadm` command, see the *Symantec Cluster Server Administrator's Guide*.

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.

   If the CP server is configured to use HTTPS-based communication, run the following command:

   ```
   # cpsadm -s cp_server -a list_users
   ```

   If the CP server is configured to use IPM-based communication, run the following command:

   ```
   # cpsadm -s cp_server -p 14250 -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 non-secure mode:

   ```
   # cpsadm -s cp_server  -a rm_user \
   -e cpsclient@sys5  -f cps_operator  -g vx
   ```

4. Remove the node entry from the CP server:

   ```
   # cpsadm -s cp_server -a rm_node  -h sys5 -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 sys5. Perform the following steps.

To remove the security credentials

1. Stop the AT process.

   ```
   # /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh \
   stop
   ```

2. Remove the credentials.

   ```
   # rm -rf /var/VRTSvcs/vcsauth/data/
   ```

Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node that is departing the cluster.

You can use script-based installer to uninstall VCS on the departing node or perform the following manual steps.

If you have configured VCS as part of the Storage Foundation and High Availability products, you may have to delete other dependent packages before you can delete all of the following ones.

To unconfigure and unload LLT and GAB and remove VCS

1 If you had configured I/O fencing in enabled mode, then stop I/O fencing.

```
# svcadm disable -t vxfen
```

2 Unconfigure GAB and LLT:

```
# /sbin/gabconfig -U
# /sbin/lltconfig -U
```

3 Unload the GAB and LLT modules from the kernel.

- Determine the kernel module IDs:

```
# modinfo | grep gab
# modinfo | grep llt
```

  The module IDs are in the left-hand column of the output.

- Unload the module from the kernel:

```
# modunload -i gab_id
# modunload -i llt_id
```

4 Disable the startup files to prevent LLT, GAB, or VCS from starting up:

```
# /usr/sbin/svcadm disable -t llt
# /usr/sbin/svcadm disable -t gab
# /usr/sbin/svcadm disable -t vcs
```

5 To determine the packages to remove, enter:

```
# pkginfo | grep VRTS
```

6 To permanently remove the VCS packages from the system, use the `pkgrm` command. Start by removing the following packages, which may have been optionally installed, in the order shown below.

On Solaris10:

```
# pkgrm VRTSvcsea
# pkgrm VRTSat
# pkgrm VRTSvcsag
# pkgrm VRTScps
# pkgrm VRTSvcs
# pkgrm VRTSamf
# pkgrm VRTSvxfen
# pkgrm VRTSgab
# pkgrm VRTSllt
# pkgrm VRTSspt
# pkgrm VRTSsfcpi61
# pkgrm VRTSvlic
# pkgrm VRTSperl
```

On Solaris 11:

```
# pkg uninstall VRTSvcsea VRTSsat VRTSvcsag
VRTScps VRTSvcs VRTSamf VRTSvxfen VRTSgab VRTSllt VRTSspt
VRTSsfcpi61 VRTSperl VRTSvlic
```

7 Remove the LLT and GAB configuration files.

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

8 Remove the language packages and patches.

See “Removing VCS packages manually” on page 470.

Section 11

Uninstallation of VCS

- Chapter 32. Uninstalling VCS using the installer
- Chapter 33. Uninstalling VCS using response files
- Chapter 34. Manually uninstalling VCS

Chapter 32

Uninstalling VCS using the installer

This chapter includes the following topics:

- Preparing to uninstall VCS
- Uninstalling VCS using the script-based installer
- Uninstalling VCS with the web-based installer
- Removing language packages using the uninstaller program
- Removing the CP server configuration using the installer program

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.
- 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 81.
- When VRTSvcs package is uninstalled on Solaris 11, the extracted package contents such as VCS configuration files and logs are moved to `/var/pkg/lost+found` directory. Therefore, to access the extracted files, you need to look inside the `/var/pkg/lost+found` directory.

Note: On Solaris 11, if you have VCS packages installed inside non-global zones, perform the steps under *Manually uninstalling VCS packages on non-global zones on Solaris 11* section to uninstall them from non-global zone before attempting to uninstall the packages from global zone.See "Manually uninstalling VCS packages on non-global zones on Solaris 11" on page 475.

Uninstalling VCS using the script-based installer

You must meet the following conditions to use the uninstallvcs 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 on each node in the cluster.

The uninstallvcs removes all VCS packages and VCS language packages.

The following example demonstrates how to uninstall VCS using the uninstallvcs. The uninstallvcs uninstalls VCS on two nodes: sys1 sys2. The example procedure uninstalls VCS from all nodes in the cluster.

Note: If already present on the system, the uninstallation does not remove the VRTSacclib package.

Removing VCS 6.1 packages

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.

   ```
   # cd /opt/VRTS/install
   # ./uninstallvcs<version>
   ```

   Where <version> is the specific release version.

   See "About the script-based installer" on page 49.

   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. The installer lists all the packages that it will remove.
4. Enter **y** to stop all the VCS processes.

   The program stops the VCS processes and proceeds with uninstalling the software.
5. Review the output as the uninstallvcs continues to do the following:
   - Verifies the communication between systems
   - Checks the installations on each system to determine the packages to be uninstalled.
6. Review the output as the uninstaller stops processes, unloads kernel modules, and removes the packages.
7. Note the location of summary, response, and log files that the uninstaller creates after removing all the packages.

Running uninstallvcs from the VCS 6.1 disc

You may need to use the uninstallvcs on the VCS 6.1 disc in one of the following cases:

- You need to uninstall VCS after an incomplete installation.
- The uninstallvcs is not available in /opt/VRTS/install.

If you mounted the installation media to /mnt, access the uninstallvcs by changing directory to:

```
cd /mnt/cluster_server/

./uninstallvcs <version>
```

Uninstalling VCS with the web-based installer

This section describes how to uninstall using the web-based installer.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in VCS 6.1 with a previous version of VCS.

To uninstall VCS

1. Perform the required steps to save any data that you want to preserve. For example, take backups of configuration files.
2. Start the web-based installer.

   See "Starting the web-based installer" on page 189.
3. On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
4. Select **Symantec Cluster Server** from the Product drop-down list, and click **Next**.
5. Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Next**.
6. After the validation completes successfully, click **Next** to uninstall VCS on the selected system.
7. If there are any processes running on the target system, the installer stops the processes. Click **Next**.
8. After the installer stops the processes, the installer removes the products from the specified system.

   Click **Next**.
9. After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.
10. Click **Finish**.

    Most packages have kernel components. To ensure their complete removal, a system restart is recommended after all the packages have been removed.

Note: If already present on the system, the uninstallation does not remove the `VRTSacclib` package.

Removing language packages using the uninstaller program

The uninstallvcs program removes all VCS packages and language packages.

Removing the CP server configuration using the installer program

This section describes how to remove the CP server configuration from a node or a cluster that hosts the CP server.

Warning: Ensure that no VCS cluster (application cluster) uses the CP server that you want to unconfigure. Run the `# cpsadm -s CPS_VIP -p CPS_Port -a list_nodes` to know if any application cluster is using the CP server.

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@cps1.symantecexample.com
# /opt/VRTS/install/installvcsversion  -configcps
```

2 Select option 3 from the menu to unconfigure the CP server.

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

[2] Configure Coordination Point Server on SFHA cluster

[3] Unconfigure Coordination Point Server
```

3 Review the warning message and confirm that you want to unconfigure the CP server.

```
Unconfiguring coordination point server stops the vxcpserv process.
VCS clusters using this server for coordination purpose will have
one less coordination point.
Are you sure you want to take the CP server offline? [y,n,q] (n) y
```

4 Review the screen output as the script performs the following steps to remove the CP server configuration:

- Stops the CP server
- Removes the CP server from VCS configuration
- Removes resource dependencies
- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration
- Successfully unconfigured the Veritas Coordination Point Server

```
The CP server database is not being deleted on the shared storage.
It can be re-used if CP server is reconfigured on the cluster.
The same database location can be specified during CP server
configuration.
```

5 Decide if you want to delete the CP server configuration file.

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files
(in /var/VRTScps)? [y,n,q] (n) y
```

```
Deleting /etc/vxcps.conf and log files on sys1.... Done
Deleting /etc/vxcps.conf and log files on sys2... Done
```

6 Confirm if you want to send information about this installation to Symantec to help improve installation in the future.

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

Upload completed successfully.

Chapter 33

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 an automated uninstallation

1. Make sure that you meet the prerequisites to uninstall VCS.
2. Copy the response file to the system where you want to uninstall VCS.

   See "Sample response file for uninstalling VCS" on page 469.

3 Edit the values of the response file variables as necessary.

See “Response file variables to uninstall VCS” on page 468.

4 Start the uninstallation from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/uninstallvcs<version>
 -responsefile /tmp/response_file
```

Where *<version>* is the specific release version, and /tmp/*response_file* is the response file’s full path name.

See “About the script-based installer” on page 49.

Response file variables to uninstall VCS

Table 33-1 lists the response file variables that you can define to uninstall VCS.

Table 33-1 Response file variables specific to uninstalling VCS

Variable	List or Scalar	Description
CFG{opt}{uninstall}	Scalar	Uninstalls VCS packages. (Required)
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 VCS61 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 *rsh* must be used instead of ssh as the communication method between systems. (Optional)

Table 33-1 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 *logpath* 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 468.

```
#
# Configuration Values:
#
our %CFG;

$CFG{opt}{uninstall}=1;
$CFG{prod}="VCS61";
$CFG{systems}=[ qw(sys1 sys2) ];
1;
```

Chapter 34

Manually uninstalling VCS

This chapter includes the following topics:

- Removing VCS packages manually
- Manually remove the CP server fencing configuration
- Manually deleting cluster details from a CP server
- Manually uninstalling VCS packages on non-global zones on Solaris 11

Removing VCS packages manually

You must remove the VCS packages from each node in the cluster to uninstall VCS.

To manually remove VCS packages on a node

1 Shut down VCS on the local system using the `hastop` command.

```
# hastop -local
```

2 Unconfigure the fencing, GAB, LLT, and AMF modules.

```
# /sbin/vxfenconfig -U
# /sbin/gabconfig -U
# /sbin/lltconfig -U
# /opt/VRTSamf/bin/amfconfig -U
```

3 Determine the GAB kernel module ID:

```
# modinfo | grep gab
```

The module ID is in the left-hand column of the output.

4 Unload the GAB module from the kernel:

```
# modunload -i gab_id
```

5 Determine the LLT kernel module ID:

```
# modinfo | grep llt
```

The module ID is in the left-hand column of the output.

6 Unload the LLT module from the kernel:

```
# modunload -i llt_id
```

7 Determine the AMF kernel module ID:

```
# modinfo | grep amf
```

8 Unload the AMF module from the kernel:

```
# modunload  -i  amf_id
```

9 Remove the VCS 6.1 packages in the following order.

On Solaris 10 systems:

```
# pkgrm VRTSvcswiz
# pkgrm VRTSvbs
# pkgrm VRTSsfmh
# pkgrm VRTSvcsea
# pkgrm VRTSat (if it exists)
# pkgrm VRTSvcsag
# pkgrm VRTScps
# pkgrm VRTSvcs
# pkgrm VRTSamf
# pkgrm VRTSvxfen
# pkgrm VRTSgab
# pkgrm VRTSllt
# pkgrm VRTSspt
# pkgrm VRTSsfcpi6.1
# pkgrm VRTSperl
# pkgrm VRTSvlic
```

On Solaris 11 systems:

```
# pkg uninstall VRTSvcswiz VRTSvbs VRTSsfmh VRTSvcsea VRTScps VRTSvcs
VRTSamf VRTSgab VRTSllt VRTSspt VRTSsfcpi6.1 VRTSperl VRTSvlic
```

Note: The VRTScps package should be removed after manually removing the CP server fencing configuration. See "Manually remove the CP server fencing configuration" on page 473. Moreover, remove the VRTSvcsnr package from logical domains if present using `pkagrm VRTSvcsnr` or `pkg uninstall VRTSvcsnr` on Solaris 10 or Solaris 11 as applicable respectively. On Solaris 11, if you have VCS packages installed inside non-global zones, uninstall them manually from the non-global zone before attempting to uninstall the packages from global zone. See "Manually uninstalling VCS packages on non-global zones on Solaris 11" on page 475.

---

10 Remove the following language packages:

- Remove the Japanese language support packages.
  On Solaris 10:

```
# pkgrm VRTSjacs
# pkgrm VRTSjacse
```

On Solaris 11:

```
# pkg uninstall VRTSjacs VRTSjacse
```

Manually remove the CP server fencing configuration

The following procedure describes how to manually remove the CP server fencing configuration from the CP server. This procedure is performed as part of the process to stop and remove server-based IO fencing.

Note: This procedure must be performed after the VCS cluster has been stopped, but before the VCS cluster software is uninstalled.

This procedure is required so that the CP server database can be reused in the future for configuring server-based fencing on the same VCS cluster(s).

Perform the steps in the following procedure to manually remove the CP server fencing configuration.

Note: The `cpsadm` command is used in the following procedure. For detailed information about the `cpsadm` command, see the *Symantec Cluster Server Administrator's Guide*.

To manually remove the CP server fencing configuration

1. Unregister all VCS cluster nodes from all CP servers using the following command:

   ```
   # cpsadm -s cp_server -a unreg_node -u uuid -n nodeid
   ```

2. Remove the VCS cluster from all CP servers using the following command:

   ```
   # cpsadm -s cp_server -a rm_clus -u uuid
   ```

3 Remove all the VCS cluster users communicating to CP servers from all the CP servers using the following command:

```
# cpsadm -s cp_server -a rm_user -e user_name -g domain_type
```

4 Proceed to uninstall the VCS cluster software.

Manually deleting cluster details from a CP server

You can manually delete the cluster details from a coordination point server (CP server) using the following procedure.

To manually delete cluster details from a CP server

1 List the nodes in the CP server cluster:

```
# cpsadm -s cps1 -a list_nodes

ClusterName      UUID                                 Hostname(Node ID)  Registered
===========   ==================================   ================  ===========
cluster1      {3719a60a-1dd2-11b2-b8dc-197f8305ffc0}  node0(0)          1
```

2 List the CP server users:

```
# cpsadm -s cps1 -a list_users

Username/Domain Type Cluster Name/UUID                                  Role
=================== ==================                                  =======
cpsclient@hostname/vx  cluster1/{3719a60a-1dd2-11b2-b8dc-197f8305ffc0}  Operator
```

3 Remove the privileges for each user of the cluster that is listed in step 2 from the CP server cluster. For example:

```
# cpsadm -s cps1 -a rm_clus_from_user
-c cluster1 -e cpsclient@hostname -g vx -f cps_operator
Cluster successfully deleted from user cpsclient@hostname privileges.
```

4 Remove each user of the cluster that is listed in step 2. For example:

```
# cpsadm -s cps1 -a rm_user -e cpsclient@hostname -g vx
User cpsclient@hostname successfully deleted
```

5 Unregister each node that is registered to the CP server cluster. See the output of step 1 for registered nodes. For example:

```
# cpsadm -s cps1 -a unreg_node -c cluster1 -n 0
Node 0 (node0) successfully unregistered
```

6 Remove each node from the CP server cluster. For example:

```
# cpsadm -s cps1 -a rm_node -c cluster1 -n 0
Node 0 (node0) successfully deleted
```

7 Remove the cluster.

```
# cpsadm -s cps1 -a rm_clus -c cluster1
Cluster cluster1 deleted successfully
```

8 Verify that the cluster details are removed successfully.

```
# cpsadm -s cps1 -a list_nodes

ClusterName      UUID           Hostname(Node ID) Registered
===========   ================  ================  ===========

# cpsadm -s cps1 -a list_users

Username/Domain Type Cluster Name/UUID     Role
==================== ==================    =======
```

Manually uninstalling VCS packages on non-global zones on Solaris 11

1 Log on to the non-global zone as a super user.

2 Uninstall VCS packages from Solaris brand zones.

```
# pkg uninstall VRTSperl VRTSvlic VRTSvcs VRTSvcsag VRTSvcsea
```

3 Uninstall VCS packages from Solaris 10 brand zones.

```
# pkgrm VRTSperl VRTSvlic VRTSvcs VRTSvcsag VRTSvcsea
```

Note: If you have VCS packages installed inside non-global zones, perform the steps mentioned above to uninstall them from non-global zone before attempting to uninstall the packages from global zone.

Section 12

Installation reference

- Appendix A. Services and ports
- Appendix B. VCS installation packages
- Appendix C. Installation command options
- Appendix D. Changes to bundled agents in VCS 6.1
- Appendix E. Configuration files
- Appendix F. Installing VCS on a single node
- Appendix G. Configuring LLT over UDP
- Appendix H. Configuring the secure shell or the remote shell for communications
- Appendix I. Troubleshooting VCS installation
- Appendix J. Sample VCS cluster setup diagrams for CP server-based I/O fencing
- Appendix K. Reconciling major/minor numbers for NFS shared disks
- Appendix L. Compatibility issues when installing Symantec Cluster Server with other products
- Appendix M. Upgrading the Steward process

Appendix A

Services and ports

This appendix includes the following topics:

- About SFHA services and ports

About SFHA services and ports

If you have configured a firewall, ensure that the firewall settings allow access to the services and ports used by SFHA.

Table A-1 lists the services and ports used by SFHA .

Note: The port numbers that appear in bold are mandatory for configuring SFHA.

Table A-1 SFHA services and ports

Port Number	Protocol	Description	Process
4145	TCP/UDP	VVR Connection Server VCS Cluster Heartbeats	vxio.sys
4888	TCP	Veritas Scheduler Service Use to launch the configured schedule.	VxSchedService.exe
5634	HTTPS	Symantec Storage Foundation Messaging Service	xprtld.exe

Table A-1 SFHA services and ports *(continued)*

Port Number	Protocol	Description	Process
7419	TCP	Symantec Plugin Host Service Solutions Configuration Center (SFWConfigPanel.exe) CCF Engine (CEngineDriver.exe	pluginHost.exe
8199	TCP	Volume Replicator Administrative Service	vras.dll
8989	TCP	VVR Resync Utility	vxreserver.exe
14141	TCP	Symantec High Availability Engine Veritas Cluster Manager (Java console) (ClusterManager.exe) VCS Agent driver (VCSAgDriver.exe)	had
14144	TCP/UDP	VCS Notification	Notifier.exe
14149	TCP/UDP	VCS Authentication	vcsauthserver
14150	TCP	Veritas Command Server	CmdServer
14153, 15550 - 15558	TCP/UDP	VCS Cluster Simulator	hasim.exe For more information about the ports used by the VCS Simulator, see the *Symantec Cluster Server Administrator's Guide*.
14155	TCP/UDP	VCS Global Cluster Option (GCO)	wac
14156	TCP/UDP	VCS Steward for GCO	steward
443	TCP	Coordination Point Server	Vxcpserv

Table A-1 SFHA services and ports *(continued)*

Port Number	Protocol	Description	Process
49152-65535	TCP/UDP	Volume Replicator Packets	User configurable ports created at kernel level by `vxio .sys` file

Appendix B

VCS installation packages

This appendix includes the following topics:

- Symantec Cluster Server installation packages

Symantec Cluster Server installation packages

Table B-1 shows the package name and contents for each Symantec Cluster Server package.

Table B-1 Symantec Cluster Server packages

package	Contents	Required/Optional
VRTSamf	Contains the binaries for the Veritas Asynchronous Monitoring Framework kernel driver functionality for all the IMF-aware agents.	Required
VRTScps	Contains the binaries for the Veritas Coordination Point Server.	Optional. Required to Coordination Point Server (CPS).
VRTSgab	Contains the binaries for Symantec Cluster Server group membership and atomic broadcast services.	Required Depends on VRTSllt.
VRTSllt	Contains the binaries for Symantec Cluster Server low-latency transport.	Required
VRTSperl	Contains Perl binaries for Veritas.	Required

Table B-1 Symantec Cluster Server packages *(continued)*

package	Contents	Required/Optional
VRTSsfcpi61	Symantec Storage Foundation Common Product Installer The Storage Foundation Common Product installer package contains the scripts that perform the following: ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ removing nodes ■ etc. You can use this script to simplify the native operating system installations, configurations, and upgrades.	Required
VRTSvcswiz	Contains the wizards for Symantec Cluster Server by Symantec.	Required
VRTSspt	Contains the binaries for Veritas Software Support Tools.	Recommended package, optional
VRTSvcs	VRTSvcs contains the following components: ■ Contains the binaries for Symantec Cluster Server. ■ Contains the binaries for Symantec Cluster Server manual pages. ■ Contains the binaries for Symantec Cluster Server English message catalogs. ■ Contains the binaries for Symantec Cluster Server utilities. These utilities include security services.	Required Depends on VRTSperl and VRTSvlic.
VRTSvcsag	Contains the binaries for Symantec Cluster Server bundled agents.	Required Depends on VRTSvcs.

Table B-1 Symantec Cluster Server packages *(continued)*

package	Contents	Required/Optional
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.
VRTSvlic	Contains the binaries for Symantec License Utilities.	Required
VRTSvxfen	Contains the binaries for Veritas I/O Fencing .	Required to use fencing. Depends on VRTSgab.
VRTSsfmh	Symantec Storage Foundation Managed Host Symantec Storage Foundation Managed Host is now called Veritas Operations Manager (VOM). VOM discovers configuration information on a Storage Foundation managed host. If you want a central server to manage and monitor this managed host, download and install the VRTSsfmcs package on a server, and add this managed host to the Central Server. The VRTSsfmcs package is not part of this release. You can download it separately from: http://www.symantec.com/veritas-operations-manager	Recommended
VRTSvbs	Enables fault management and VBS command line operations on VCS nodes managed by Veritas Operations Manager. For more information, see the *Virtual Business Service–Availability User's Guide*.	Recommended Depends on VRTSsfmh. VRTSsfmh version must be 4.1 or later for VRTSvbs to get installed.

Table B-1 Symantec Cluster Server packages *(continued)*

package	Contents	Required/Optional
VRTSvcsnr	Network reconfiguration service for Oracle VM Server logical domains	Optional You must install VRTSvcsnr manually inside a Oracle VM Server logical domain if the domain is to be configured for disaster recovery.

Table B-2 shows the package name, contents, and type for each Symantec Cluster Server language package.

Table B-2 Symantec Cluster Server language packages

package	Contents	Package type
VRTSmulic	Contains the multi-language Symantec license utilities.	Common L10N package
VRTSatJA		Japanese language package
VRTSjacav	Contains the binaries for Japanese Symantec Cluster Server Agent Extensions for Storage Cluster File System - Manual Pages and Message Catalogs.	Japanese language package
VRTSjacse	Contains Japanese Veritas High Availability Enterprise Agents by Symantec.	Japanese language package
VRTSjacs	Contains the binaries for Symantec Cluster Server Japanese Message Catalogs by Symantec.	Japanese language package
VRTSjacsu	Contains the binaries for Japanese Symantec Cluster Utility Language Pack by Symantec.	Japanese language package
VRTSjadba	Contains the binaries for Japanese RAC support package by Symantec.	Japanese language package
VRTSjadbe	Contains the Japanese Storage Management Software for Databases - Message Catalog.	Japanese language package

Table B-2 Symantec Cluster Server language packages *(continued)*

package	Contents	Package type
VRTSjafs	Contains the binaries for Japanese Language Message Catalog and Manual Pages for VERITAS File System.	Japanese language package
VRTSjaodm	Contains the binaries for Japanese Message Catalog and Man Pages for ODM.	Japanese language package
VRTSjavm	Contains the binaries for Japanese Virtual Disk Subsystem Message Catalogs and Manual Pages.	Japanese language package
VRTSzhvm	Contains the binaries for Chinese Virtual Disk Subsystem Message Catalogs and Manual Pages.	Chinese language package

Appendix C

Installation command options

This appendix includes the following topics:

- Command options for installvcs
- Installation script options
- Command options for uninstallvcs

Command options for installvcs

The `installvcs` command usage takes the following form:

```
installvcs [ system1 system2... ]
        [ -install | -configure | -uninstall | -license
        | -upgrade | -precheck | -requirements | -start | -stop
        | -postcheck ]
        [ -responsefile response_file ]
        [ -logpath log_path ]
        [ -tmppath tmp_path ]
        [ -tunablesfile tunables_file ]
        [ -timeout timeout_value ]
        [ -keyfile ssh_key_file ]
        [ -hostfile hostfile_path ]
        [ -pkgpath pkg_path ]
        [ -rootpath root_path ]
        [ -jumpstart jumpstart_path ]
        [ -flash_archive flash_archive_path ]
        [ -serial | -rsh | -redirect | -installminpkgs
        | -installrecpkgs | -installallpkgs | -minpkgs
```

```
| -recpkgs | -allpkgs | -pkgset | -pkgtable | -pkginfo
| -makeresponsefile | -comcleanup  | -version | -nolic
| -ignorepatchreqs | -settunables | -security | -securityonenode
| -securitytrust | -addnode | -fencing | -upgrade_kernelpkgs
| -upgrade_nonkernelpkgs | -rolling_upgrade
| -rollingupgrade_phase1        | -rollingupgrade_phase2 ]
```

Installation script options

Table C-1 shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Symantec Storage Foundation product scripts, except where otherwise noted.

See "About the script-based installer" on page 49.

Table C-1 Available command line options

Command Line Option	Function
-addnode	Adds a node to a high availability cluster.
-allpkgs	Displays all packages required for the specified product. The packages are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-comsetup	The `-comsetup` option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.
-configcps	The `-configcps` option is used to configure CP server on a running system or cluster.
-configure	Configures the product after installation.
-fencing	Configures I/O fencing in a running cluster.
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases package upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-hotfix_path	Defines the path of a hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-hotfix2_path	Defines the path of a second hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-hotfix3_path	Defines the path of a third hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-hotfix4_path	Defines the path of a fourth hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-hotfix5_path	Defines the path of a fifth hot fix level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-installallpkgs	The `-installallpkgs` option is used to select all packages.
-installrecpkgs	The `-installrecpkgs`option is used to select the recommended packages set.
–installminpkgs	The `-installminpkgs`option is used to select the minimum packages set.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
-ignorepatchreqs	The `-ignorepatchreqs` option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.
–jumpstart *dir_path*	Produces a sample finish file for Solaris JumpStart installation. The *dir_path* indicates the path to the directory in which to create the finish file.
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-I ssh_key_file` to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
-makeresponsefile	Use the `-makeresponsefile` option only to generate response files. No actual software installation occurs when you use this option.
-minpkgs	Displays the minimal packages required for the specified product. The packages are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-noipc	Disables the installer from making outbound networking calls to Symantec Operations Readiness Tool (SORT) in order to automatically obtain hot fix and release information updates.
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
–pkginfo	Displays a list of packages and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS packages.
–pkgpath *package_path*	Designates the path of a directory that contains all packages to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
–pkgset	Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.
-pkgtable	Displays product's packages in correct installation order by group.
–postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
–recpkgs	Displays the recommended packages required for the specified product. The packages are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer hot fix file.
-requirements	The `-requirements` option displays required OS version, required packages and patches, file system space, and other system requirements in order to install the product.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The `-rollingupgrade_phase1` option is used to perform rolling upgrade Phase-I. In the phase, the product kernel packages get upgraded to the latest version.
-rollingupgrade_phase2	The `-rollingupgrade_phase2` option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
–rootpath *root_path*	Specifies an alternative root directory on which to install packages. On Solaris operating systems, `-rootpath` passes `-R path` to `pkgadd` command.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
-securitytrust	The `-securitytrust` option is used to setup trust with another broker.
–serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
-settunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the `-tunablesfile` option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the `-timeout` option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 prevents the script from timing out. The `-timeout` option does not work with the `-serial` option
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file *tunables_file*	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
-version	Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. Lists the installed patches, hot fixes, and available updates for the installed product if an Internet connection is available.

Command options for uninstallvcs

The `uninstallvcs` command usage takes the following form:

On Solaris 10:

```
uninstallvcs [ system1 system2... ]
        [ -uninstall ]
        [ -responsefile response_file ]
        [ -logpath log_path ]
        [ -timeout timeout_value ]
        [ -keyfile ssh_key_file ]
        [ -hostfile hostfile_path ]
        [ -rootpath root_path ]
        [ -serial | -rsh | -redirect | -makeresponsefile
        | -comcleanup | -version | -nolic | -ignorepatchreqs ]
```

On Solaris 11:

```
uninstallvcs [ system1 system2... ]
        [ -responsefile response_file ]
        [ -logpath log_path ]
        [ -tmppath tmp_path ]
        [ -timeout timeout_value ]
        [ -keyfile ssh_key_file ]
        [ -hostfile hostfile_path ]
        [ -rootpath root_path ]
        [ -ai ai_path ]
        [ -serial | -rsh | -redirect | -makeresponsefile
        | -comcleanup | -version | -ignorepatchreqs ]
```

For description of the `uninstallvcs` command options:

See “Installation script options” on page 487.

Appendix D

Changes to bundled agents in VCS 6.1

This appendix includes the following topics:

- Deprecated agents
- New agents
- New and modified attributes for VCS 6.1 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
- CFSQlogckd
- ClusterMonitorConfig
- DiskReservation
- NFSLock
- Service group heartbeat (ServiceGroupHB)
- SANVolume
- VRTSWebApp

Note: No agents were deprecated since the 5.1 SP1 release.

New agents

No new agents were added in the VCS 6.1 release on this platform.

The following new agent is added in the 6.0 release:

- AlternateIO—Monitors VCS storage and network service groups that in turn monitor redundant I/O services exported from the control domain and alternate I/O domain to a guest logical domain.

The following agents were added in the 5.1 SP1 release:

- VolumeSet—Brings Veritas Volume Manager (VxVM) volume sets online and offline, and monitors them.
- Disk—Monitors a physical disk or a partition.
- Project—Adds, deletes, and monitors Solaris projects.

The following agents were added in the 5.1 release:

- CoordPoint—Provides server-based I/O fencing.

The following agents were added in the 5.0 MP3 release:

- DiskGroupSnap—Verifies the configuration and the data integrity in a campus cluster environment.
- LDom—Monitors and manages logical domains on Solaris SPARC.
- Zpool—Monitors ZFS storage pools.
- SambaServer—Monitors the smbd process.
- SambaShare—Use to make a Samba Share highly available or to monitor it.
- NetBios—Use to make the nmbd process highly available or to monitor it.

The following agents were added in the 5.0 release:

- Apache (now bundled on all platforms)—Provides high availability to an Apache Web server.
- NFSRestart—Provides high availability for NFS record locks.
- ProcessOnOnly—Starts and monitors a user-specified process.
- RemoteGroup—Monitors and manages a service group on another system.

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

New and modified attributes for VCS 6.1 agents

Table D-1 lists the attributes that VCS adds or modifies when you upgrade from VCS 6.0.1 to VCS 6.1

Table D-1 Changes to attributes VCS 6.0.1 to 6.1

Agent	New/Modified attributes	Default value
Apache		
	New attributes added:	
	IMF	{ Mode=2, MonitorFreq=5, RegisterRetryLimit=3 }
	IMFRegList	{ httpdDir, ConfigFile }
	Attribute modified: ContainerOpts	{ RunInContainer=1, PassCInfo=1 }
Application		
	New attributes added to IMFRegList: LevelTwoMonitorFreq	{ MonitorProcesses, User, PidFiles, MonitorProgram, StartProgram, LevelTwoMonitorFreq }
	New attribute added: LevelTwoMonitorFreq	1
DiskGroup		
	Attribute removed from ArgList: DiskGroupType	
	New attribute added to ArgList: ClearClone	{ DiskGroup, StartVolumes, StopVolumes, MonitorOnly, MonitorReservation, tempUseFence, PanicSystemO nDGLoss, UmountVolumes, Reservation, ConfidenceLevel, ClearClone }

Table D-1 Changes to attributes VCS 6.0.1 to 6.1 *(continued)*

Agent	New/Modified attributes	Default value
	New attribute added: ClearClone	0
NFS		
	New attribute added to ArgList:	
	MountdOptions	{ Nproc, GracePeriod, NFSSecurity, NFSv4Support, LockFileTimeout, MountdOptions, Protocol }
	Protocol	{ Nproc, GracePeriod, NFSSecurity, NFSv4Support, LockFileTimeout, MountdOptions, Protocol }
	New attributes added:	
	MountdOptions	MountdOptions
	Protocol	all
NSFRestart		
	New attributes added to ArgList:	
	NFSRes:MountdOptions	{ LocksPathName, NFSLockFailover, LockServers, NFSRes, "NFSRes:Nservers", "NFSRes:LockFileTimeout", "NFSRes:UseSMF", Lower, State, "NFSRes:MountdOptions", "NFSRes:Protocol" }

Table D-1 Changes to attributes VCS 6.0.1 to 6.1 *(continued)*

Agent	New/Modified attributes	Default value
	NFSRes:Protocol	{ LocksPathName, NFSLockFailover, LockServers, NFSRes, "NFSRes:Nservers", "NFSRes:LockFileTimeout", "NFSRes:UseSMF", Lower, State, "NFSRes:MountdOptions", "NFSRes:Protocol" }
	Protocol	MountdOptions
	Protocol	all
NotifierMngr		
	New attributes added to ArgList: MessageExpiryInterval	{ EngineListeningPort, MessagesQueue, MessageExpiryInterval, NotifierListeningPort, NotifierSourceI P, SnmpdTrapPort, SnmpCommunity, SnmpConsoles, SmtpServer, SmtpServerVrfyOff, SmtpServerTimeout, SmtpReturnPath, SmtpFromPath, Smt pRecipients }
	New attribute added: MessageExpiryInterval	3600
	New attributes added:	
	MigrateTimeout	600
	MigrateWaitLimit	2
	New attribute added to SupportedOperations: migrate	{ migrate }

Table D-1 Changes to attributes VCS 6.0.1 to 6.1 *(continued)*

Agent	New/Modified attributes	Default value
	Attribute modified: AgentFile	"bin/Script60Agent"
	New attribute added to ArgList:	
	Password	{ State, IState, LDomName, CfgFile, MonitorCPU, NumCPU, ConfigureNetwork, IPAddress, Netmask, Gateway, DNS, Memory, CEInfo, RemoveLDomConfigForMigration, ResyncVMCfg, DomainFailurePolicy, Password, UserName, ReregPGR }
	User	{ State, IState, LDomName, CfgFile, MonitorCPU, NumCPU, ConfigureNetwork, IPAddress, Netmask, Gateway, DNS, Memory, CEInfo, RemoveLDomConfigForMigration, ResyncVMCfg, DomainFailurePolicy, Password, UserName, ReregPGR }
	DomainFailurePolicy	{ State, IState, LDomName, CfgFile, MonitorCPU, NumCPU, ConfigureNetwork, IPAddress, Netmask, Gateway, DNS, Memory, CEInfo, RemoveLDomConfigForMigration, ResyncVMCfg, DomainFailurePolicy, Password, UserName, ReregPGR }

Table D-1 Changes to attributes VCS 6.0.1 to 6.1 *(continued)*

Agent	New/Modified attributes	Default value
	RegregPGR	{ State, IState, LDomName, CfgFile, MonitorCPU, NumCPU, ConfigureNetwork, IPAddress, Netmask, Gatew ay, DNS, Memory, CEInfo, RemoveLDomConfigForMigration, ResyncVMCfg, DomainFailurePolicy, Password, UserName, RegregPGR }
	New attribute added:	
	User	root
	Password	
	MonitorCPU	0
	DomainFailurePolicy	{ primary="stop" }
	RegregPGR	0
Mount		
	New attribute added to RegList: VxFSMountLock	{VxFSMountLock}

- Table D-2 lists the attributes that VCS adds or modifies when you upgrade from VCS 6.0 to VCS 6.0.1
- Table D-3 lists the attributes that VCS adds or modifies when you upgrade from VCS 5.1 SP1 to VCS 6.1.
- Table D-4 lists the attributes that VCS adds or modifies when you upgrade from VCS 5.1 to VCS 5.1 SP1.
- Table D-5 lists the attributes that VCS adds or modifies when you upgrade from VCS 5.0 MP3 to VCS 5.1.
- Table D-6 lists the attributes that VCS adds or modifies when you upgrade from VCS 5.0 to VCS 5.0 MP3.
- Table D-7 lists the attributes that VCS adds or modifies when you upgrade from VCS 4.1 to VCS 5.0.

Table D-2 Changes to attributes VCS 6.0 to 6.0.1

Agent	New/Modified attributes	Default value
Application		
	Attribute modified: IMFRegList	{ MonitorProcesses, User, PidFiles, MonitorProgram, StartProgram }
DiskGroup		
	New attribute added: IMFRegList	{ DiskGroup, Reservation }
	New attribute addded: IMF	{ Mode = 3, MonitorFreq = 5, RegisterRetryLimit = 3 }
	Depricated attribute: DiskGroupType	
IPMultiNICB		
	Attribute modified: ArgList	{ BaseResName, Address, NetMask, DeviceChoice, RouteOptions, DeleteRouteOptions, PrefixLen, IgnoreMultiNICBFailure, "BaseResName:Protocol", Options }
	New attribute added: DeleteRouteOptions	
LDom		
	New attribute added: SupportedActions	{ "vmconfigsync" }
	New attribute added: ResyncVMCfg	0 Do not set ResyncVMCfg attribute manually.
MultiNICB		

Table D-2 Changes to attributes VCS 6.0 to 6.0.1 *(continued)*

Agent	New/Modified attributes	Default value
	Attribute modified: ArgList	{ UseMpathd, MpathdCommand, ConfigCheck, MpathdRestart, Device, NetworkHosts, LinkTestRatio, IgnoreLinkStatus, NetworkTimeout, OnlineTestRepeatCount, OfflineTestRepeatCount, NoBroadcast, DefaultRouter, Failback, GroupName, Protocol, IPMPDevice }
	New attribute added: IPMPDevice	

Table D-3 Changes to attributes from VCS 5.1 SP1 to VCS 6.1

Agent	New and modified attributes	Default value
AlternateIO (New agent)		
	New attributes	
	AgentFile	bin/Script51Agent
	ArgList	{ StorageSG, NetworkSG }
	StorageSG	{}
	NetworkSG	{}
Application		
	Modified attributes	
	ArgList (new attribute added to list)	{ State, IState, User, StartProgram, StopProgram, CleanProgram, MonitorProgram, PidFiles, MonitorProcesses, EnvFile, UseSUDash }

Table D-3 Changes to attributes from VCS 5.1 SP1 to VCS 6.1 *(continued)*

Agent	New and modified attributes	Default value
	IMF	{ Mode = 3, MonitorFreq = 1, RegisterRetryLimit = 3 }
	SupportedActions (new action added to keylist)	{ "program.vfd", "user.vfd", "cksum.vfd", getcksum, propcv }
DNS		
	New attributes	
	UseGSSAPI	0
	RefreshInterval	0
	CleanRRKeys	0
	Modified attribute	
	ArgList (new attribute added to list)	{ Domain, TTL, TSIGKeyFile, StealthMasters, ResRecord, CreatePTR, OffDelRR, UseGSSAPI, RefreshInterval, CleanRRKeys }
DiskGroup		
	Modified attributes	
	PanicSystemOnDGLoss (attribute data type change)	int PanicSystemOnDGLoss = 0
	ArgList (new attribute added to list)	{ DiskGroup, StartVolumes, StopVolumes, MonitorOnly, MonitorReservation, tempUseFence, PanicSystemOnDGLoss, DiskGroupType, UmountVolumes, Reservation, ConfidenceLevel }
DiskGroupSnap		
	New attribute	
	FDType	""
	Modified attribute	

Table D-3 Changes to attributes from VCS 5.1 SP1 to VCS 6.1 *(continued)*

Agent	New and modified attributes	Default value
	ArgList (new attribute added to list)	{ TargetResName, FDSiteName, FDType }
IP		
	Modified attribute	
	RegList	{ NetMask }
IPMultiNIC		
	Modified attribute	
	ToleranceLimit	1
LDom		
	New attributes	
	DNS	""
	Netmask	""
	CEInfo	{ Enabled=0, CESystem=NONE, FaultOnHBLoss=1 }
	Gateway	""
	ConfigureNetwork	0
	IPAddress	""
	Memory	""
	IntentionalOffline	1
	RemoveLDomConfigForMigration	0
	Modified attributes	
	AEPTimeout	1
	ArgList (new attribute added to list)	{ State, IState, LDomName, CfgFile, MonitorCPU, NumCPU, ConfigureNetwork, IPAddress, Netmask, Gateway, DNS, Memory, Memory, RemoveLDomConfigForMigration }

Table D-3 Changes to attributes from VCS 5.1 SP1 to VCS 6.1 *(continued)*

Agent	New and modified attributes	Default value
	AgentFile	bin/Script51Agent
	RegList (new attribute added to list)	{ NumCPU, Memory }
Mount		
	Modified attributes	
	AEPTimeout	1
	IMF	{ Mode = 3, MonitorFreq = 1, RegisterRetryLimit = 3 }
	SecondLevelMonitor (deprecated attribute)	
	SecondLevelTimeout (deprecated attribute)	
	ArgList (list updated for deprecated attributes)	{ MountPoint, BlockDevice, FSType, MountOpt, FsckOpt, SnapUmount, CkptUmount, OptCheck, CreateMntPt, MntPtPermission, MntPtOwner, MntPtGroup, AccessPermissionChk, RecursiveMnt, VxFSMountLock }
Process		
	Modified attribute	
	IMF	{ Mode = 3, MonitorFreq = 5, RegisterRetryLimit = 3 }
RVGPrimary		
	Modified attribute	
	ArgList (new attribute added to list)	{ RvgResourceName, "RvgResourceName:RVG", "RvgResourceName:DiskGroup", AutoTakeover, AutoResync, BunkerSyncTimeOut, BunkerSyncElapsedTime }

Table D-3 Changes to attributes from VCS 5.1 SP1 to VCS 6.1 *(continued)*

Agent	New and modified attributes	Default value
Zone		
	New attributes	
	DROpts	{}
	DeleteVCSZoneUser	0
	Modified attributes	
	IMF (attribute data type change)	static int IMF{} = { Mode = 3, MonitorFreq = 5, RegisterRetryLimit = 3 }
	AEPTimeout	1
	IMFRegList	{ BootState }
	ArgList (new attributes added to list)	{ Pool, BootState, ShutdownGracePeriod, RunFsck, DetachZonePath, ForceAttach, DeleteVCSZoneUser, DROpts }
Zpool		
	New attributes	
	ForceOpt	0
	DeviceDir	[]
	FailMode	continue
	ForceRecoverOpt	0
	Modified attribute	
	ArgList (new attribute added to list)	{ PoolName, AltRootPath, ChkZFSMounts, ZoneResName, "ZoneResName:State", DeviceDir, FailMode, ForceOpt, ForceRecoverOpt }

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1

Agent	New and modified attributes	Default value
Application		

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1 *(continued)*

Agent	New and modified attributes	Default value
	New attributes	
	EnvFile	""
	UseSUDash	0
	RegList	{ MonitorProcesses, User }
	IMFRegList	{ MonitorProcesses, User, PidFiles, MonitorProgram }
	Modified attributes	
	User (change in default value)	"root"
	ArgList (new attribute added to list)	{ User, StartProgram, StopProgram, CleanProgram, MonitorProgram, PidFiles, MonitorProcesses, EnvFile, UseSUDash }
Disk (new agent)		
	New attributes	
	Partition	""
	ArgList	{ Partition }
	OfflineMonitorInterval	60
	Operations	"None"
DiskGroup		
	New attribute	
	Reservation	"ClusterDefault"
	Modified attribute	

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1 *(continued)*

Agent	New and modified attributes	Default value
	ArgList (new attribute added to list)	{ DiskGroup, StartVolumes, StopVolumes, MonitorOnly, MonitorReservation, tempUseFence, PanicSystemOnDGLoss, UmountVolumes, Reservation }
IP		
	New attribute	
	ExclusiveIPZone	0
	Modified attribute	
	ArgList (new attribute added to list)	{ Device, Address, NetMask, Options, ArpDelay, IfconfigTwice, RouteOptions , PrefixLen, ExclusiveIPZone }
IPMultiNICB		
	New attribute	
	Options	""
	Modified attribute	
	ArgList (new attribute added to list)	{ BaseResName, Address, NetMask, DeviceChoice, RouteOptions, PrefixLen, IgnoreMultiNICBFailure, "BaseResName:Protocol", Options }
Mount		
	New attribute	
	IMFRegList	{ MountPoint, BlockDevice, FSType }
MultiNICA		
	Modified attribute	
	Protocol (change in default value)	"IPv4"
MultiNICB		

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1 *(continued)*

Agent	New and modified attributes	Default value
	Modified attribute	
	Protocol (change in default value)	"IPv4"
NetBios		
	New attribute	
	PidFile	""
	Modified attribute	
	ArgList (new attribute added to list)	{ "SambaServerRes:ConfFile", "SambaServerRes:SambaTopDir", "SambaServerRes:LockDir", NetBiosName, NetBiosAliases, Interfaces, WinsSupport, DomainMaster, "SambaServerRes:PidFile", SambaServerRes, PidFile }
NFS		
	New attribute	
	CleanRmtab	0
	Modified attribute	
	ArgList (new attribute added to list)	{ UseSMF, Nservers, LockFileTimeout, CleanRmtab }
NFSRestart		
	New attribute	
	Lower	0
	Modified attribute	
	ArgList (new attribute added to list)	{ LocksPathName, NFSLockFailover, LockServers, NFSRes, "NFSRes:Nservers", "NFSRes:LockFileTimeout", "NFSRes:UseSMF", Lower, State }

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1 *(continued)*

Agent	New and modified attributes	Default value
NIC		
	New attributes	
	ExclusiveIPZone	0
	ContainerOpts	{ RunInContainer=0, PassCInfo=1 }
	Modified attributes	
	Protocol	"IPv4"
	ArgList (new attribute added to list)	{ Device, PingOptimize, NetworkHosts, Protocol, NetworkType, ExclusiveIPZone }
NotifierMngr		
	New attribute	
	NotifierSourceIP	""
	Modified attribute	
	ArgList (new attribute added to list)	{ EngineListeningPort, MessagesQueue, NotifierListeningPort, NotifierSourceIP, SnmpdTrapPort, SnmpCommunity, SnmpConsoles, SmtpServer, SmtpServerVrfyOff, SmtpServerTimeout, SmtpReturnPath, SmtpFromPath, SmtpRecipients }
RemoteGroup		
	New attributes	
	ReturnIntOffline	{ }
	OfflineMonitoringNode	""
	Modified attributes	

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1 *(continued)*

Agent	New and modified attributes	Default value
	IntentionalOffline (change in default value,RemoteGroup agent now supports intentional offline feature.)	1
	ArgList (new attribute added to list)	{ IpAddress, Port, Username, Password, GroupName, VCSSysName, ControlMode, OfflineWaitTime, DomainType, BrokerIp, ReturnIntOffline }
RVGPrimary		
	New attributes	
	BunkerSyncTimeOut	""
	BunkerSyncElapsedTime	0
	Modified attributes	
	ArgList (new attribute added to list)	{ RvgResourceName, AutoTakeover, AutoResync, BunkerSyncTimeOut, BunkerSyncElapsedTime }
	SupportedActions (new action added to keylist)	{ fbsync, ElectPrimary }
RVGSnapshot		
	New attribute	
	VCSResLock	""
SambaServer		
	New attributes	
	PidFile	""
	SocketAddress	""
	Modified attribute	

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1 *(continued)*

Agent	New and modified attributes	Default value
	ArgList (new attribute added to list)	{ ConfFile, SambaTopDir, LockDir, Ports, IndepthMonitorCyclePeriod, ResponseTimeout, PidFile, SocketAddress }
SambaShare		
	Modified attribute	
	ArgList (dependent attributes added to list)	{ "SambaServerRes:ConfFile", "SambaServerRes:SambaTopDir", "SambaServerRes:LockDir", ShareName, ShareOptions, "SambaServerRes:Ports", SambaServerRes, "SambaServerRes:PidFile", "SambaServerRes:SocketAddress" }
VolumeSet (new agent)		
	New attributes	
	DiskGroup	""
	VolumeSet	""
	ArgList	{ DiskGroup, VolumeSet }
Zone		
	New attributes	
	RunFsck	0
	DetachZonePath	1
	ForceAttach	1
	Modified attribute	
	BootState (change in default value)	"multi-user"

Table D-4 Changes to attributes from VCS 5.1 to VCS 5.1 SP1 *(continued)*

Agent	New and modified attributes	Default value
	ArgList (dependent attributes added to list)	{ Pool, BootState, ShutdownGracePeriod, RunFsck, DetachZonePath, ForceAttach }

Table D-5 Changes to attributes from VCS 5.0 MP3 to VCS 5.1

Agent	New and modified attributes	Default value
Apache		
	Modified attribute	
	ContainerName (deleted attribute)	
	ContainerType (deleted attribute)	
	New attribute	
	ContainerOpts	RunInContainer=1, PassCInfo=0
Application		
	Modified attribute	
	ContainerName (deleted attribute)	
	ContainerType (deleted attribute)	
	New attribute	
	ContainerOpts	RunInContainer=1, PassCInfo=0
DNS		
	Modified attributes	
	Alias (deleted attribute)	

Table D-5 Changes to attributes from VCS 5.0 MP3 to VCS 5.1 *(continued)*

Agent	New and modified attributes	Default value
	Hostname (deleted attribute)	
DiskGroup		
	Modified attributes	
	StartVolumes	1
	StopVolumes	1
	PanicSystemOnDGLoss	0
IP		
	New attributes	
	RouteOptions	
	PrefixLen	
	ContainerOpts	RunInContainer=0, PassCInfo=1
	Modified attribute	
	ContainerName (deleted attribute)	
IPMultiNIC		
	New attributes	
	PrefixLen	
	ContainerOpts	RunInContainer=0, PassCInfo=1
	Modified attribute	
	ContainerName (deleted attribute)	
IPMultiNICB		
	New attributes	
	OnlineRetryLimit	1

Table D-5 Changes to attributes from VCS 5.0 MP3 to VCS 5.1 *(continued)*

Agent	New and modified attributes	Default value
	ContainerOpts	RunInContainer=0, PassCInfo=1
	IgnoreMultiNICBFailure	0
	RouteOptions	
	PrefixLen	
	Modified attributes	
	ContainerName (deleted attribute)	
LDOM		
	New attributes	
	MonitorCPU	1
Mount		
	New attributes	
	OptCheck	0
	CreateMountPt	0
	ReuseMntPt	0
	MntPtPermission	
	MntPtOwner	
	MntPtGroup	
	AccessPermissionChk	0
	RecursiveMnt	0
	ContainerOpts	RunInContainer=0, PassCInfo=0
	Modified attributes	
	ContainerName (deleted attribute)	

Table D-5 Changes to attributes from VCS 5.0 MP3 to VCS 5.1 *(continued)*

Agent	New and modified attributes	Default value
	ContainerType (deleted attribute)	Zone
MultiNICA		
	New attributes	
	Protocol	
MultiNICB		
	Modified attribute	
	MpathdCommand	/usr/lib/inet/in.mpathd
	New attribute	
	Protocol	
NFS		
	New attribute	
	UseSMF	0
NFSRestart		
	New attribute	
	LockServers	20
NIC		
	New attribute	
	Protocol	
Phantom		
	Modified attribute	
	Dummy (deleted attribute)	
Process		
	New attribute	

Table D-5 Changes to attributes from VCS 5.0 MP3 to VCS 5.1 *(continued)*

Agent	New and modified attributes	Default value
	ContainerOpts	RunInContainer=1, PassCInfo=0
	Modified attributes	
	ContainerName (deleted attribute)	
	ContainerType (deleted attribute)	Zone
ProcessOnOnly		
	New attribute	
	ContainerOpts	RunInContainer=0, PassCInfo=1
	Modified attributes	
	ContainerName (deleted attribute)	
	ContainerType (deleted attribute)	Zone
Share		
	New attribute	
	NFSRes	
Zone		
	New attributes	
	ContainerOpts	RunInContainer=1, PassCInfo=0
	BootState	
	Pool	
	Modified attribute	
	ZoneName (deleted attribute)	

Table D-6 Changes to attributes from VCS 5.0 to VCS 5.0 MP3

Agent	New and modified attributes	Default value
Apache		
	New attribute	
	SupportedActions	"checkconffile.vfd"
	ContainerType	Zone
	PidFile	
	ContainerName	
	IntentionalOffline	0
DNS		
	New attributes	
	SupportedActions	"dig.vfd", "keyfile.vfd", "master.vfd"
	ResRecord	
	CreatePTR	0
	OffDelRR	0
DiskGroup		
	New attributes	
	SupportedActions	"license.vfd", "disk.vfd", "udid.vfd", "verifyplex.vfd", "checkudid", "campusplex", "numdisks", "joindg", "splitdg", "getvxvminfo", "volinuse"
	UmountVolumes	0
Mount		
	New attribute	
	VxFSMountLock	1
	Modified attribute	
	SupportedActions	"mountpoint.vfd", "mounted.vfd", "vxfslic.vfd", "chgmntlock", "mountentry.vfd"

Table D-6 Changes to attributes from VCS 5.0 to VCS 5.0 MP3 *(continued)*

Agent	New and modified attributes	Default value
NFSRestart		
	New attributes	
	SupportedActions	"lockdir.vfd", "nfsconf.vfd"
Share		
	New attributes	
	SupportedActions	"direxists.vfd"

Table D-7 Changes to attributes from VCS 4.1 to VCS 5.0

Agent	New and modified attributes	Default value
Application		
	New attributes	
	SupportedActions	program.vfd, user.vfd, cksum.vfd, getcksum
DiskGroup		
	New attributes	
	SupportedActions	"license.vfd", "disk.vfd", "numdisks"
	PanicSystemOnDGLoss	1
	DiskGroupType	Private
	Modified attributes	
	tempUseFence	Invalid
IP		
	New attributes	
	SupportedActions	"device.vfd", "route.vfd"
	ContainerName	
	Modified attribute	
	IfconfigTwice	

Table D-7 Changes to attributes from VCS 4.1 to VCS 5.0 *(continued)*

Agent	New and modified attributes	Default value
IPMultiNIC		
	New attributes	
	ContainerName	
	Modified attribute	
	IfconfigTwice	
IPMultiNICB		
	New attributes	
	ToleranceLimit	1
	MonitorInterval	30
	ContainerName	
	Modified attribute	
	DeviceChoice	0
Mount		
	New attributes	
	SupportedActions	"mountpoint.vfd", "mounted.vfd", "vxfslic.vfd"
	VxFSMountLock	1
	ContainerName	
	ContainerType	Zone
	Modified attribute	
	SnapUmount	
MultiNICA		
	Modified attribute	
	IfconfigTwice	
MultiNICB		

Table D-7 Changes to attributes from VCS 4.1 to VCS 5.0 *(continued)*

Agent	New and modified attributes	Default value
	New attribute	
	GrNew attributes oupName	
	Modified attributes	
	NoBroadcast	
	Failback	
NFS		
	New attributes	
	LockFileTimeout	180
NIC		
	New attributes	
	SupportedActions	"device.vfd"
Process		
	New attribute	
	SupportedActions	"program.vfd", getcksum

Manually removing deprecated resource types and modifying attributes

With VCS 6.0, certain resource type definitions are no longer used. Before you start the upgrade process, you must remove the resources of the deprecated resource types from your cluster configuration.

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.
- Modify attribute values that might have changed.
  See Table D-5 on page 514.
  See Table D-6 on page 519.
  See Table D-7 on page 520.
- Save the main.cf.
- Reformat the main.cf file.

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

4 Verify the configuration.

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

5 Start VCS on the local node.

6 Start VCS on other nodes.

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 *Symantec 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, ldap, 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
```

Appendix E

Configuration files

This appendix includes the following topics:

- About the LLT and GAB configuration files
- About the AMF configuration files
- About the VCS configuration files
- About I/O fencing configuration files
- Sample configuration files for CP server
- Packaging related SMF services on Solaris 11

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/lltab files. GAB requires /etc/gabtab file.

Table E-1 lists the LLT configuration files and the information that these files contain.

Table E-1 LLT configuration files

File	Description
/etc/default/llt	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.

/etc/llthosts	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 sys1`
	`1 sys2`

Table E-1 LLT configuration files *(continued)*

File	Description

/etc/llttab	The file `llttab` contains the information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the private network links that correspond to the specific system. For example, the file /etc/llttab contains the entries that resemble the following:
	■ For Solaris 10 SPARC:
	`set-node sys1`
	`set-cluster 2`
	`link net1 /dev/net:0 - ether - -`
	`link net2 /dev/net:1 - ether - -`
	■ For Solaris 11 SPARC :
	`set-node sys1`
	`set-cluster 2`
	`link net1 /dev/net/net1 - ether - -`
	`link net2 /dev/net/net2 - ether - -`
	The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the `link` command. These lines identify the two network cards that the LLT protocol uses.
	If you configured a low priority link under LLT, the file also includes a "link-lowpri" 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.

Table E-2 lists the GAB configuration files and the information that these files contain.

Table E-2 GAB configuration files

File	Description
/etc/default/gab	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.
/etc/gabtab	After you install VCS, the file /etc/gabtab contains a `gabconfig(1)` command that configures the GAB driver for use. The file /etc/gabtab contains a line that resembles: `/sbin/gabconfig -c -nN` 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. Use the `-c` option for `/sbin/gabconfig` to avoid a split-brain condition. **Note:**

About the AMF configuration files

Asynchronous Monitoring Framework (AMF) kernel driver provides asynchronous event notifications to the VCS agents that are enabled for intelligent resource monitoring.

Table E-3 lists the AMF configuration files.

Table E-3 AMF configuration files

File	Description
/etc/default/amf	This file stores the start and stop environment variables for AMF: ■ AMF_START—Defines the startup behavior for the AMF module after a system reboot or when AMF is attempted to start using the init script. Valid values include: 1—Indicates that AMF is enabled to start up. (default) 0—Indicates that AMF is disabled to start up. ■ AMF_STOP—Defines the shutdown behavior for the AMF module during a system shutdown or when AMF is attempted to stop using the init script. Valid values include: 1—Indicates that AMF is enabled to shut down. (default) 0—Indicates that AMF is disabled to shut down.
/etc/amftab	After you install VCS, the file `/etc/amftab` contains a `amfconfig`(1) command that configures the AMF driver for use. The AMF init script uses this `/etc/amftab` file to configure the AMF driver. The `/etc/amftab` file contains the following line by default: `/opt/VRTSamf/bin/amfconfig -c`

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 531.
  See "Sample main.cf file for global clusters" on page 533.
- 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.

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 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 "SecureClus = 1" cluster attribute.
- The installvcs 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 prompts about notification.
  - The installvcs 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 *Symantec Cluster Server Administrator's Guide* for information about managing VCS global clusters.

Refer to the *Symantec Cluster Server Administrator's Guide* to review the configuration concepts, and descriptions of main.cf and types.cf files for Solaris systems.

Sample main.cf file for VCS clusters

The following sample main.cf file is for a three-node cluster in secure mode.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"
include "Db2udbTypes.cf"
include "SybaseTypes.cf"


cluster vcs02 (
    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
    )

NIC csgnic (
    Device = net0
    NetworkHosts = { "10.182.13.1" }
    )

NotifierMngr ntfr (
   SnmpConsoles = { sys4" = SevereError }
   SmtpServer = "smtp.example.com"
   SmtpRecipients =  { "ozzie@example.com" = SevereError }
   )

ntfr requires csgnic

// resource dependency tree
//
//     group ClusterService
//     {
//     NotifierMngr ntfr
//         {
//         NIC csgnic
//         }
// }
```

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.

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 -secure"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac -secure" }
    RestartLimit = 3
    )

IP gcoip (
    Device = net0
    Address = "10.182.13.50"
    NetMask = "255.255.240.0"
    )
```

```
NIC csgnic (
    Device = net0
    NetworkHosts = { "10.182.13.1" }
    )

NotifierMngr ntfr (
   SnmpConsoles = { sys4 = SevereError }
   SmtpServer = "smtp.example.com"
   SmtpRecipients =  { "ozzie@example.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
//             }
//         }
//     }
```

About I/O fencing configuration files

Table E-4 lists the I/O fencing configuration files.

Table E-4 I/O fencing configuration files

File	Description
/etc/default/vxfen	This file stores the start and stop environment variables for I/O fencing: ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: 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: 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing 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, you must make sure to set the values of these environment variables to 1.
/etc/vxfendg	This file includes the coordinator disk group information. This file is not applicable for server-based fencing.

Table E-4 I/O fencing configuration files *(continued)*

File	Description

/etc/vxfenmode	This file contains the following parameters:
	■ vxfen_mode
	  ■ 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
	This parameter is applicable only for server-based fencing. Set the value as cps.
	■ scsi3_disk_policy
	  ■ 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
	Note: You must use the same SCSI-3 disk policy on all the nodes.
	■ security
	This parameter is applicable only for server-based fencing.
	1—Indicates that communication with the CP server is in secure mode. This setting is the default.
	0—Indicates that communication with the CP server is in non-secure mode.
	■ List of coordination points
	This list is required only for server-based fencing configuration.
	Coordination points in server-based fencing can include coordinator disks, CP servers, or both. If you use coordinator disks, you must create a coordinator disk group containing the individual coordinator disks.
	Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points and multiple IP addresses for each CP server.
	■ single_cp
	This parameter is applicable for server-based fencing which uses a single highly available CP server as its coordination point. Also applicable for when you use a coordinator disk group with single disk.
	■ autoseed_gab_timeout
	This parameter enables GAB automatic seeding of the cluster even when some cluster nodes are unavailable. This feature requires that I/O fencing is enabled.
	0—Turns the GAB auto-seed feature on. Any value greater than 0 indicates the number of seconds that GAB must delay before it automatically seeds the cluster.
	-1—Turns the GAB auto-seed feature off. This setting is the default.

Table E-4 I/O fencing configuration files *(continued)*

File	Description

/etc/vxfentab	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/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.
	Note: The /etc/vxfentab file is a generated file; do not modify this file.
	For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk along with its unique disk identifier. A space separates the path and the unique disk identifier. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:
	■ Raw disk:
	`/dev/rdsk/c1t1d0s2 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006804E795D075`
	`/dev/rdsk/c2t1d0s2 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006814E795D076`
	`/dev/rdsk/c3t1d2s2 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006824E795D077`
	■ DMP disk:
	`/dev/vx/rdmp/c1t1d0s2 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006804E795D075`
	`/dev/vx/rdmp/c2t1d0s2 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006814E795D076`
	`/dev/vx/rdmp/c3t1d0s2 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006824E795D077`
	For server-based fencing, the /etc/vxfentab file also includes the security settings information.
	For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.

Sample configuration files for CP server

The `/etc/vxcps.conf` file determines the configuration of the coordination point server (CP server.)

See "Sample CP server configuration (/etc/vxcps.conf) file output" on page 543.

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 "Sample main.cf file for CP server hosted on a single node that runs VCS" on page 538.
- The main.cf file for a CP server that is hosted on an SFHA cluster:
  See "Sample main.cf file for CP server hosted on a two-node SFHA cluster" on page 540.

Note: If you use IPM-based protocol for communication between the CP server and VCS clusters (application clusters), the CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses). If you use HTTPS-based protocol for communication, the CP server only supports Internet Protocol version 4 (IPv4 addresses).

The example main.cf files use IPv4 addresses.

Sample main.cf file for CP server hosted on a single node that runs VCS

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: cps1

```
include "types.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"

// cluster name:  cps1
// CP server: cps1

cluster cps1 (
     UserNames = { admin = bMNfMHmJNiNNlVNhMK, haris = fopKojNvpHouNn,
             "cps1.symantecexample.com@root@vx" = aj,
             "root@cps1.symantecexample.com" = hq }
     Administrators = { admin, haris,
             "cps1.symantecexample.com@root@vx",
             "root@cps1.symantecexample.com" }
     SecureClus = 1
     HacliUserLevel = COMMANDROOT
```

```
    )

system cps1 (
    )

group CPSSG (
      SystemList = { cps1 = 0 }
      AutoStartList = { cps1 }
      )

      IP cpsvip1 (
            Critical = 0
            Device @cps1 = bge0
            Address = "10.209.3.1"
            NetMask = "255.255.252.0"
            )

      IP cpsvip2 (
            Critical = 0
            Device @cps1 = bge1
            Address = "10.209.3.2"
            NetMask = "255.255.252.0"
            )

      NIC cpsnic1 (
           Critical = 0
           Device @cps1 = bge0
           PingOptimize = 0
           NetworkHosts @cps1 = { "10.209.3.10 }
           )

      NIC cpsnic2 (
           Critical = 0
           Device @cps1 = bge1
           PingOptimize = 0
           )

      Process vxcpserv (
           PathName = "/opt/VRTScps/bin/vxcpserv"
           ConfInterval = 30
           RestartLimit = 3
           )
```

```
    Quorum quorum (
          QuorumResources = { cpsvip1, cpsvip2 }
          )

cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
vxcpserv requires quorum


// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1
//     {
//     NIC cpsnic1
//     }
// IP cpsvip2
//     {
//     NIC cpsnic2
//     }
// Process vxcpserv
//     {
//     Quorum quorum
//     }
// }
```

Sample main.cf file for CP server hosted on a two-node SFHA cluster

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: cps1, cps2

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"


// cluster: cps1
```

```
// CP servers:
// cps1
// cps2

cluster cps1 (
     UserNames = { admin = ajkCjeJgkFkkIskEjh,
            "cps1.symantecexample.com@root@vx" = JK,
            "cps2.symantecexample.com@root@vx" = dl }
     Administrators = { admin, "cps1.symantecexample.com@root@vx",
            "cps2.symantecexample.com@root@vx" }
     SecureClus = 1
     )

system cps1 (
     )

system cps2 (
     )

group CPSSG (
      SystemList = { cps1 = 0, cps2 = 1 }
      AutoStartList = { cps1, cps2 } )

      DiskGroup cpsdg (
            DiskGroup = cps_dg
            )

      IP cpsvip1 (
            Critical = 0
            Device @cps1 = bge0
            Device @cps2 = bge0
            Address = "10.209.81.88"
            NetMask = "255.255.252.0"
            )

      IP cpsvip2 (
            Critical = 0
            Device @cps1 = bge1
            Device @cps2 = bge1
            Address = "10.209.81.89"
            NetMask = "255.255.252.0"
            )
```

```
        Mount cpsmount (
              MountPoint = "/etc/VRTScps/db"
              BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
              FSType = vxfs
              FsckOpt = "-y"
              )

        NIC cpsnic1 (
             Critical = 0
             Device @cps1 = bge0
             Device @cps2 = bge0
             PingOptimize = 0
             NetworkHosts @cps1 = { "10.209.81.10 }
             )

        NIC cpsnic2 (
             Critical = 0
             Device @cps1 = bge1
             Device @cps2 = bge1
             PingOptimize = 0
             )

        Process vxcpserv (
              PathName = "/opt/VRTScps/bin/vxcpserv"
              )

        Quorum quorum (
              QuorumResources = { cpsvip1, cpsvip2 }
              )

        Volume cpsvol (
              Volume = cps_volume
              DiskGroup = cps_dg
              )

    cpsmount requires cpsvol
    cpsvip1 requires cpsnic1
    cpsvip2 requires cpsnic2
    cpsvol requires cpsdg
    vxcpserv requires cpsmount
    vxcpserv requires quorum
```

```
// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1
//     {
//     NIC cpsnic1
//     }
// IP cpsvip2
//     {
//     NIC cpsnic2
//     }
// Process vxcpserv
//     {
//     Quorum quorum
//     Mount cpsmount
//         {
//         Volume cpsvol
//             {
//             DiskGroup cpsdg
//             }
//         }
//     }
// }
```

Sample CP server configuration (/etc/vxcps.conf) file output

The following is an example of a coordination point server (CP server) configuration file `/etc/vxcps.conf` output.

```
##  The vxcps.conf file determines the
## configuration for Veritas CP Server.
cps_name=cps1
vip=[10.209.81.88]
vip=[10.209.81.89]:56789
vip_https=[10.209.81.88]:55443
vip_https=[10.209.81.89]
port=14250
port_https=443
security=1
db=/etc/VRTScps/db
ssl_conf_file=/etc/vxcps_ssl.properties
```

Packaging related SMF services on Solaris 11

After installing packages on Solaris 11 system, the following SMF services are present in online state. These SMF services are meant for proper package operation during uninstall operation. Symantec recommends you to not disable these services.

```
svc:/system/gab-preremove:default

svc:/system/llt-preremove:default

svc:/system/vxfen-preremove:default
```

Appendix F

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
- Configuring VCS
- Verifying single-node operation

About installing VCS on a single node

You can install VCS 6.1 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 546.

See “Creating a single-node cluster manually” on page 547.

Creating a single-node cluster using the installer program

Table F-1 specifies the tasks that are involved to install VCS on a single node using the installer program.

Table F-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 546.
Install the VCS software on the system using the installer.	See "Starting the installer for the single node cluster" on page 546.

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, enable 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 25.

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[q,?]
```

Enter a single system name. While you configure, the installer 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.

Selecting **n** disables LLT, GAB, and I/O fencing kernel modules of VCS. So, the kernel programs are not loaded to the node.

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 F-2 specifies the tasks that you need to perform to install VCS on a single node.

Table F-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 547.
Install the VCS software manually and add a license key	See "Installing VCS software manually on a single node" on page 548.
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.	
Start VCS and verify single-node operation.	See "Verifying single-node operation" on page 548.

Setting the path variable for a manual single node installation

Set the path variable.

See "Setting the PATH variable" on page 76.

Installing VCS software manually on a single node

Install the VCS 6.1 packages manually and install the license key.

Refer to the following sections:

- See “Installing VCS software manually” on page 238.
- See “Adding a license key for a manual installation” on page 245.

Configuring VCS

You now need to configure VCS.

See “Configuring VCS manually” on page 263.

Verifying single-node operation

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

To verify single-node cluster

1 Run the SMF command to start VCS as a single-node cluster.

```
# svcadm enable system/vcs-onenode
```

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

```
# ps -ef | grep had
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
```

Appendix G

Configuring LLT over UDP

This appendix includes the following topics:

- Using the UDP layer for LLT
- Manually configuring LLT over UDP using IPv4
- Manually configuring LLT over UDP using IPv6
- LLT over UDP sample /etc/lltab

Using the UDP layer for LLT

VCS 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.

Manually configuring LLT over UDP using IPv4

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on separate subnets. Set the broadcast address in /etc/lltab explicitly depending on the subnet for each link.

See "Broadcast address in the /etc/llttab file" on page 550.

- 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 552.
- Set the broadcast address correctly for direct-attached (non-routed) links.
  See "Sample configuration: direct-attached links" on page 554.
- 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 556.

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 sys1:

```
sys1 # cat /etc/llttab

set-node sys1
set-cluster 1
link link1 /dev/udp - udp  50000  -  192.168.9.1 192.168.9.255
link link2 /dev/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 sys2:

```
sys2 # cat /etc/llttab

set-node sys2
set-cluster 1
link link1 /dev/udp - udp  50000  -  192.168.9.2 192.168.9.255
link link2 /dev/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/lltab file

Review the link command information in this section for the /etc/lltab file. See the following information for sample configurations:

- See "Sample configuration: direct-attached links" on page 554.
- See "Sample configuration: links crossing IP routers" on page 556.

Table G-1 describes the fields of the link command that are shown in the /etc/lltab file examples. Note that some of the fields differ from the command for standard LLT links.

Table G-1 Field description for link command in /etc/lltab

Field	Description
tag-name	A unique string that is used as a tag by LLT; for example link1, link2,....
device	The device path of the UDP protocol; for example /dev/udp.
node-range	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
link-type	Type of link; must be "udp" for LLT over UDP.
udp-port	Unique UDP port in the range of 49152-65535 for the link. See "Selecting UDP ports" on page 552.
MTU	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the `lltstat -l` command to display the current value.
IP address	IP address of the link on the local node.
bcast-address	■ 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/lltab file

The `set-addr` command in the /etc/lltab 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 556.

Table G-2 describes the fields of the set-addr command.

Table G-2 Field description for set-addr command in /etc/lltab

Field	Description
node-id	The node ID of the peer node; for example, 0.
link tag-name	The string that LLT uses to identify the link; for example link1, link2,....
address	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 -a | more
UDP
   Local Address         Remote Address      State
-------------------- -------------------- -------
      *.sunrpc                            Idle
      *.*                                 Unbound
      *.32771                             Idle
      *.32776                             Idle
      *.32777                             Idle
      *.name                              Idle
      *.biff                              Idle
      *.talk                              Idle
      *.32779                             Idle
.
.
.
      *.55098                             Idle
      *.syslog                            Idle
```

```
*.58702                                 Idle
*.*                                     Unbound
```

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:

```
# ifconfig interface_name netmask netmask
```

For example:

- For the first network interface on the node sys1:

  ```
  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 sys2:

  ```
  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 sys1:

  ```
  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 sys2:

  ```
  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/lltab depending on the subnet that the links are on.

An example of a typical /etc/lltab 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 /dev/udp - udp 50000 - 192.168.30.1 192.168.30.255
link link2 /dev/udp - udp 50001 - 192.168.31.1 192.168.31.255
```

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

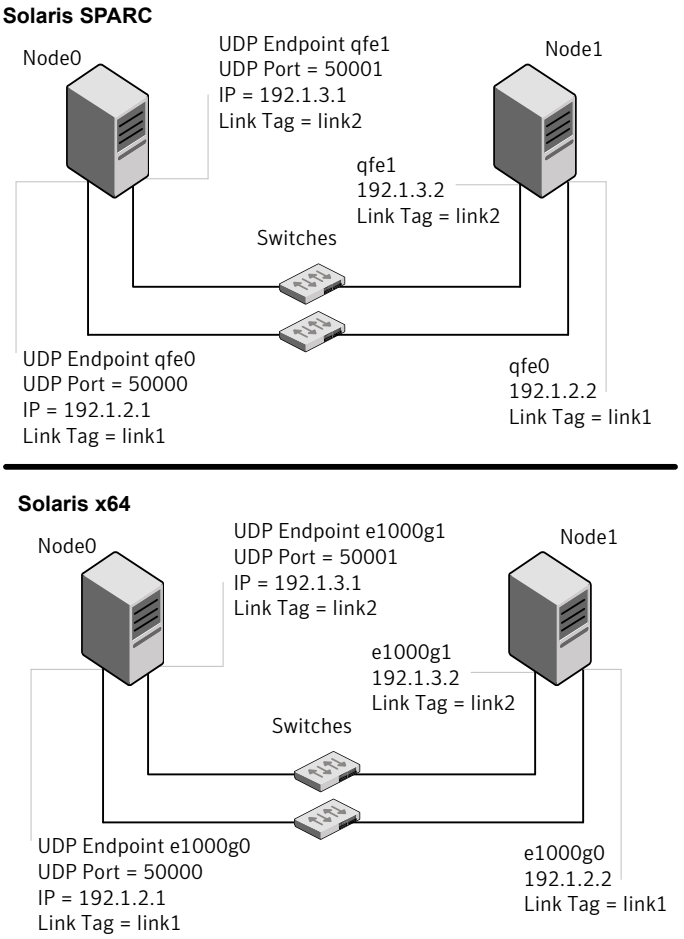

The configuration that the /etc/lltab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT sends broadcast requests to peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the /etc/lltab 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 /dev/udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 /dev/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 /dev/udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 /dev/udp - udp 50001 - 192.1.3.2 192.1.3.255
```

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

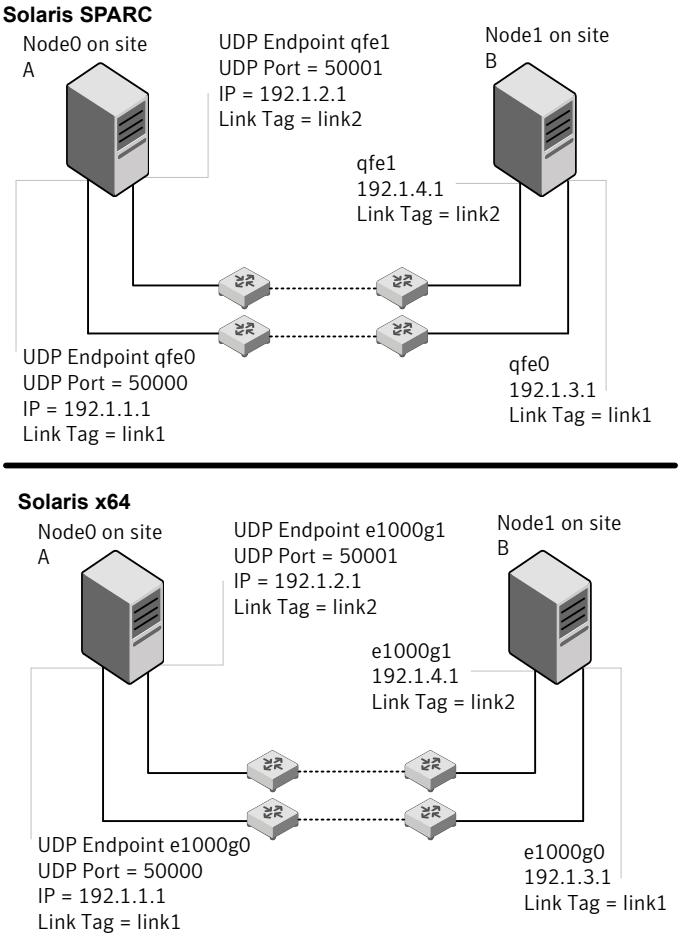

The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1
```

```
link link1 /dev/udp - udp 50000 - 192.1.3.1 -
link link2  /dev/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 /dev/udp - udp 50000 - 192.1.1.1 -
link link2 /dev/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
```

Manually 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/lltab 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 560.
- For the links that cross an IP router, disable multicast features and specify the IPv6 address of each link manually in the /etc/lltab file.
  See "Sample configuration: links crossing IP routers" on page 563.

The link command in the /etc/lltab file

Review the link command information in this section for the /etc/lltab file. See the following information for sample configurations:

- See "Sample configuration: direct-attached links" on page 561.
- See "Sample configuration: links crossing IP routers" on page 563.

Note that some of the fields in Table G-3 differ from the command for standard LLT links.

Table G-3 describes the fields of the link command that are shown in the /etc/lltab file examples.

Table G-3 Field description for link command in /etc/lltab

Field	Description
tag-name	A unique string that is used as a tag by LLT; for example link1, link2,....
device	The device path of the UDP protocol; for example /dev/udp6.
node-range	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
link-type	Type of link; must be "udp6" for LLT over UDP.
udp-port	Unique UDP port in the range of 49152-65535 for the link. See "Selecting UDP ports" on page 560.
MTU	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the `lltstat -l` command to display the current value.
IPv6 address	IPv6 address of the link on the local node.

Table G-3 Field description for link command in /etc/llttab *(continued)*

Field	Description
mcast-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 multicast feature of LLT is disabled, such as when LLT must cross IP routers.

See "Sample configuration: links crossing IP routers" on page 563.

Table G-4 describes the fields of the set-addr command.

Table G-4 Field description for set-addr command in /etc/llttab

Field	Description
node-id	The ID of the peer node; for example, 0.
link tag-name	The string that LLT uses to identify the link; for example link1, link2,....
address	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 -a |  more

UDP: IPv4
   Local Address        Remote Address      State
-------------------- -------------------- ----------
      *.sunrpc                           Idle
```

```
      *.*                                  Unbound
      *.32772                              Idle
      *.*                                  Unbound
      *.32773                              Idle
      *.lockd                              Idle
      *.32777                              Idle
      *.32778                              Idle
      *.32779                              Idle
      *.32780                              Idle
      *.servicetag                         Idle
      *.syslog                             Idle
      *.16161                              Idle
      *.32789                              Idle
      *.177                                Idle
      *.32792                              Idle
      *.32798                              Idle
      *.snmpd                              Idle
      *.32802                              Idle
      *.*                                  Unbound
      *.*                                  Unbound
      *.*                                  Unbound

UDP: IPv6
   Local Address             Remote Address           State      If
------------------------- ------------------------- ---------- -----
      *.servicetag                                  Idle
      *.177                                         Idle
```

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-3 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure G-3 A typical configuration of direct-attached links that use LLT over UDP

Solaris SPARC

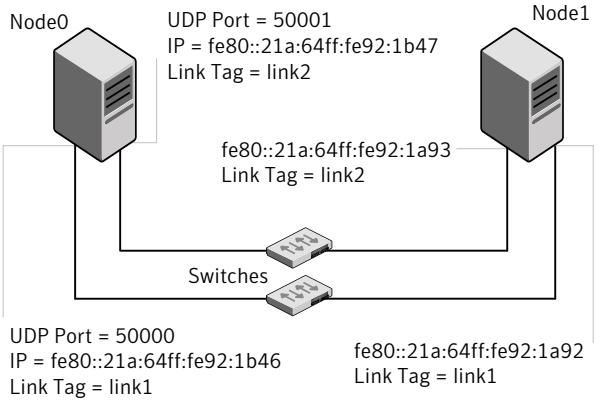

Solaris x64

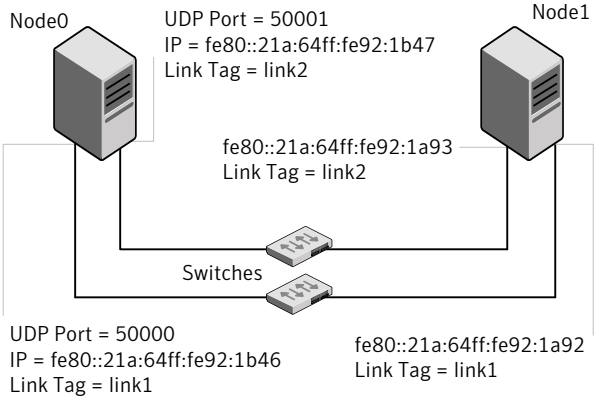

The configuration that the /etc/lltab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

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/lltab 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 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 /dev/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 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

Sample configuration: links crossing IP routers

Figure G-4 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-4 A typical configuration of links crossing an IP router

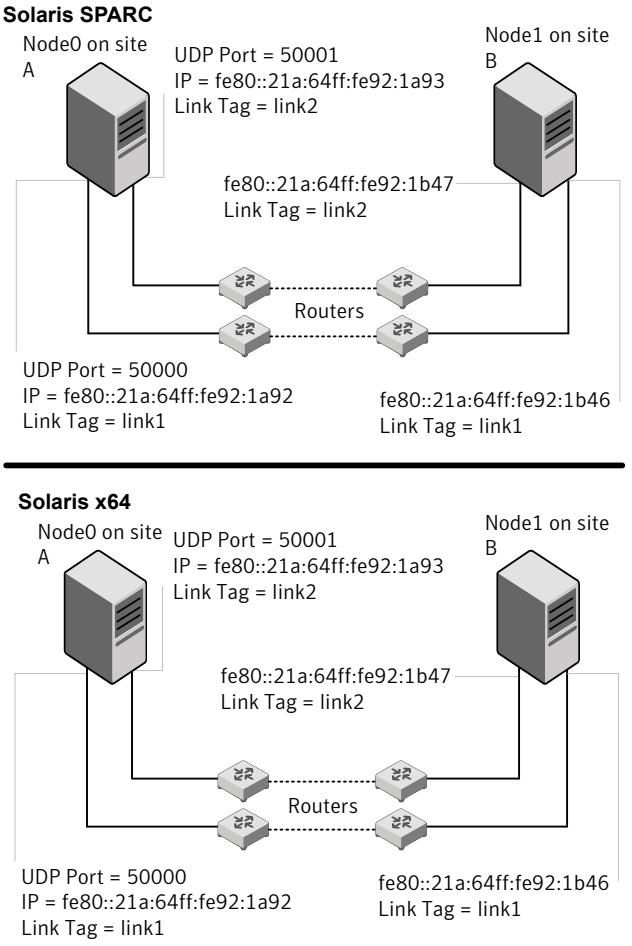

The configuration that the following `/etc/lltab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/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 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 /dev/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
```

LLT over UDP sample /etc/lltab

The following is a sample of LLT over UDP in the etc/lltab file.

```
set-node sys1
set-cluster clus1
link e1000g1 /dev/udp - udp 50000 - 192.168.10.1 -
link e1000g2 /dev/udp - udp 50001 - 192.168.11.1 -
link-lowpri e1000g0 /dev/udp - udp 50004 - 10.200.58.205 -
set-addr 1 e1000g1 192.168.10.2
```

```
set-addr 1 e1000g2 192.168.11.2
set-addr 1 e1000g0 10.200.58.206
set-bcasthb 0
set-arp 0
```

Appendix H

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

- Setting up inter-system communication

Setting up inter-system communication

If you manually need to set up a communication mode, refer to these procedures. You must have root privilege 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, contact Symantec Support. See http://support.symantec.com.

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.

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 in as root on the source system from which you want to install the Symantec product.
2. To generate a DSA key pair on the source system, type the following:

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

   System output similar to the following is displayed:

   ```
   Generating public/private dsa key pair.
   Enter file in which to save the key (//.ssh/id_dsa):
   ```

3. Press **Enter** to accept the default location of /.ssh/id_dsa. System output similar to the following is displayed:

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

4. Do not enter a passphrase. Press **Enter**. Enter same passphrase again. Press **Enter** again.
5. Make sure the /.ssh directory is on all the target installation systems. If that directory is absent, create it on the target system and set the write permission to root only:

   ```
   # mkdir /.ssh
   # chmod go-w /
   # chmod 700 /.ssh
   # chmod go-rwx /.ssh
   ```

6 Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems. To enable SFTP, the /etc/ssh/sshd_config file must contain the following two lines:

```
PermitRootLogin yes
Subsystem sftp /usr/lib/ssh/sftp-server
```

7 If the lines are not there, add them and restart SSH. To restart SSH on Solaris 10 and 11, type the following command:

```
# svcadm restart ssh
```

8 To copy the public DSA key, /.ssh/id_dsa.pub to each target system, type the following commands:

```
# sftp target_sys
```

If you run this step for the first time on a system, output similar to the following appears:

```
Connecting to target_sys...
The authenticity of host 'target_sys (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9e:61:91:9e:44:6b:87:86:ef:68:a6:fd:87:7d.
Are you sure you want to continue connecting (yes/no)?
```

9 Enter **yes**. Output similar to the following is displayed:

```
Warning: Permanently added 'target_sys,10.182.00.00'
(DSA) to the list of known hosts.
root@target_sys password:
```

10 Enter the root password.

11 At the sftp prompt, type the following command:

```
sftp> put /.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /.ssh/id_dsa.pub to /id_dsa.pub
```

12 To quit the SFTP session, type the following command:

```
sftp> quit
```

13 To begin the ssh session on the target system, type the following command:

```
# ssh target_sys
```

14 Enter the root password at the prompt:

```
password:
```

15 After you log in, enter the following command to append the authorization key to the id_dsa.pub file:

```
# cat /id_dsa.pub >> /.ssh/authorized_keys
```

16 Delete the id_dsa.pub public key file. Before you delete this public key file, make sure to complete the following tasks:

- The file is copied to the target (host) system
- The file is added to the authorized keys file

To delete the id_dsa.pub public key file, type the following command:

```
# rm /id_dsa.pub
```

17 To log out of the ssh session, type the following command:

```
# exit
```

18 When you install from a source system that is also an installation target, add the local system id_dsa.pub key to the local /.ssh/authorized_key file. The installation can fail if the installation source system is not authenticated.

19 Run the following commands on the source installation system. These commands bring the private key into the shell environment and makes the key globally available for the user root:

```
# exec /usr/bin/ssh-agent $SHELL
# ssh-add
Identity added: /.ssh/identity
```

This step is shell-specific and is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

20 To verify that you can connect to the target system, type the following command:

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

The commands should execute on the remote system without any requests for a passphrase or password from the system.

Appendix I

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 Symantec products
- Installer cannot create UUID for the cluster
- LLT startup script displays errors
- The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails
- Issues during fencing startup on VCS cluster nodes set up for server-based fencing

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 Symantec Storage
Foundation/Symantec 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:

  ```
  # /opt/VRTS/bin/vxlicrep
  ```

- 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 Symantec products

After the installation and configuration is complete, the Symantec product installer starts the processes that the installed products use. You can use the product installer to stop or start the processes, if required.

To stop the processes

- Use the `-stop` option to stop the product installation script.

  For example, to stop the product's processes, enter the following command:

  ```
  # ./installer -stop
  ```

  or

  ```
  # /opt/VRTS/install/installvcs<version> -stop
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 49.

To start the processes

- Use the `-start` option to start the product installation script.

  For example, to start the product's processes, enter the following command:

  ```
  # ./installer -start
  ```

  or

  ```
  # /opt/VRTS/install/installvcs<version> -start
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 49.

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.

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.

LLT startup script displays errors

If more than one system on the network has the same clusterid-nodeid pair and the same Ethernet sap/UDP port, 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
LLT lltconfig ERROR V-14-2-15245 cluster id 1 is
already being used by nid 0 and has the
address - 00:04:23:AC:24:2D
LLT lltconfig ERROR V-14-2-15664 LLT could not
configure any link
```

Check the log files that get generated in the `/var/svc/log` directory for any errors.

Recommended action: Ensure that all systems on the network have unique clusterid-nodeid pair. You can use the `lltdump -f device -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 fencing startup on VCS cluster nodes set up for server-based fencing

Table I-1 Fencing startup issues on VCS cluster (client cluster) nodes

Issue	Description and resolution
`cpsadm` command on the VCS cluster 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 the /etc/vxfenmode file and ensure that the VCS cluster nodes use the correct CP server virtual IP or virtual hostname and the correct port number. ■ For HTTPS communication, ensure that the virtual IP and ports listed for the server can listen to HTTPS requests.
Authorization failure	Authorization failure occurs when the nodes on the client clusters and or users are not added in the CP server configuration. Therefore, fencing on the VCS cluster (client 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 client cluster node and user in the CP server configuration and restart fencing. See "Preparing the CP servers manually for use by the VCS cluster" on page 278.
Authentication failure	If you had configured secure communication between the CP server and the VCS cluster (client cluster) nodes, authentication failure can occur due to the following causes: ■ The client cluster requires its own private key, a signed certificate, and a Certification Authority's (CA) certificate to establish secure communication with the CP server. If any of the files are missing or corrupt, communication fails. ■ If the client cluster certificate does not correspond to the client's private key, communication fails. ■ If the CP server and client cluster do not have a common CA in their certificate chain of trust, then communication fails.

Appendix J

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 J-1 on page 578.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
- Two node campus cluster that is served be remote CP server and 2 SCSI-3 disks:
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:

Two unique client clusters served by 3 CP servers

Figure J-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 `vxfenmode` file on the client nodes, vxfenmode is set to `customized` with vxfen mechanism set to `cps`.

Figure J-1 Two unique client clusters served by 3 CP servers

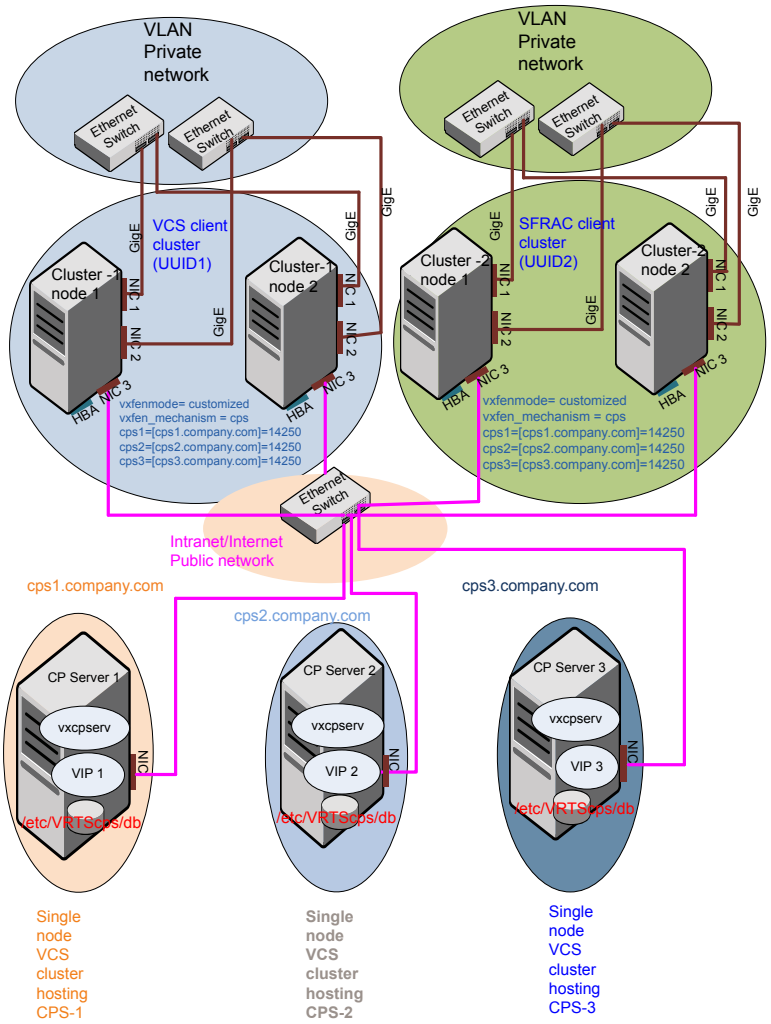

Client cluster served by highly available CPS and 2 SCSI-3 disks

Figure J-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 two SCSI-3 disks are part of the disk group vxfencoorddg. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure J-2 Client cluster served by highly available CP server and 2 SCSI-3 disks

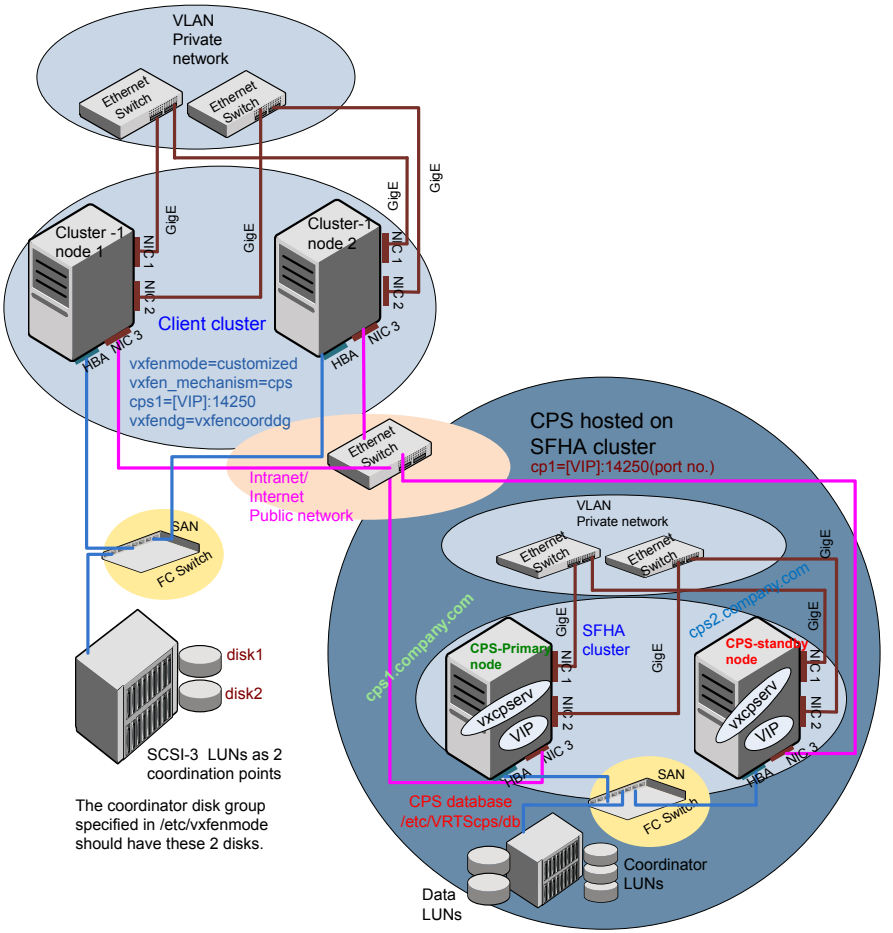

Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure J-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 two SCSI-3 disks (one from each site) are part of disk group vxfencoorddg. The third coordination point is a CP server on a single node VCS cluster.

Figure J-3 Two node campus cluster served by remote CP server and 2 SCSI-3

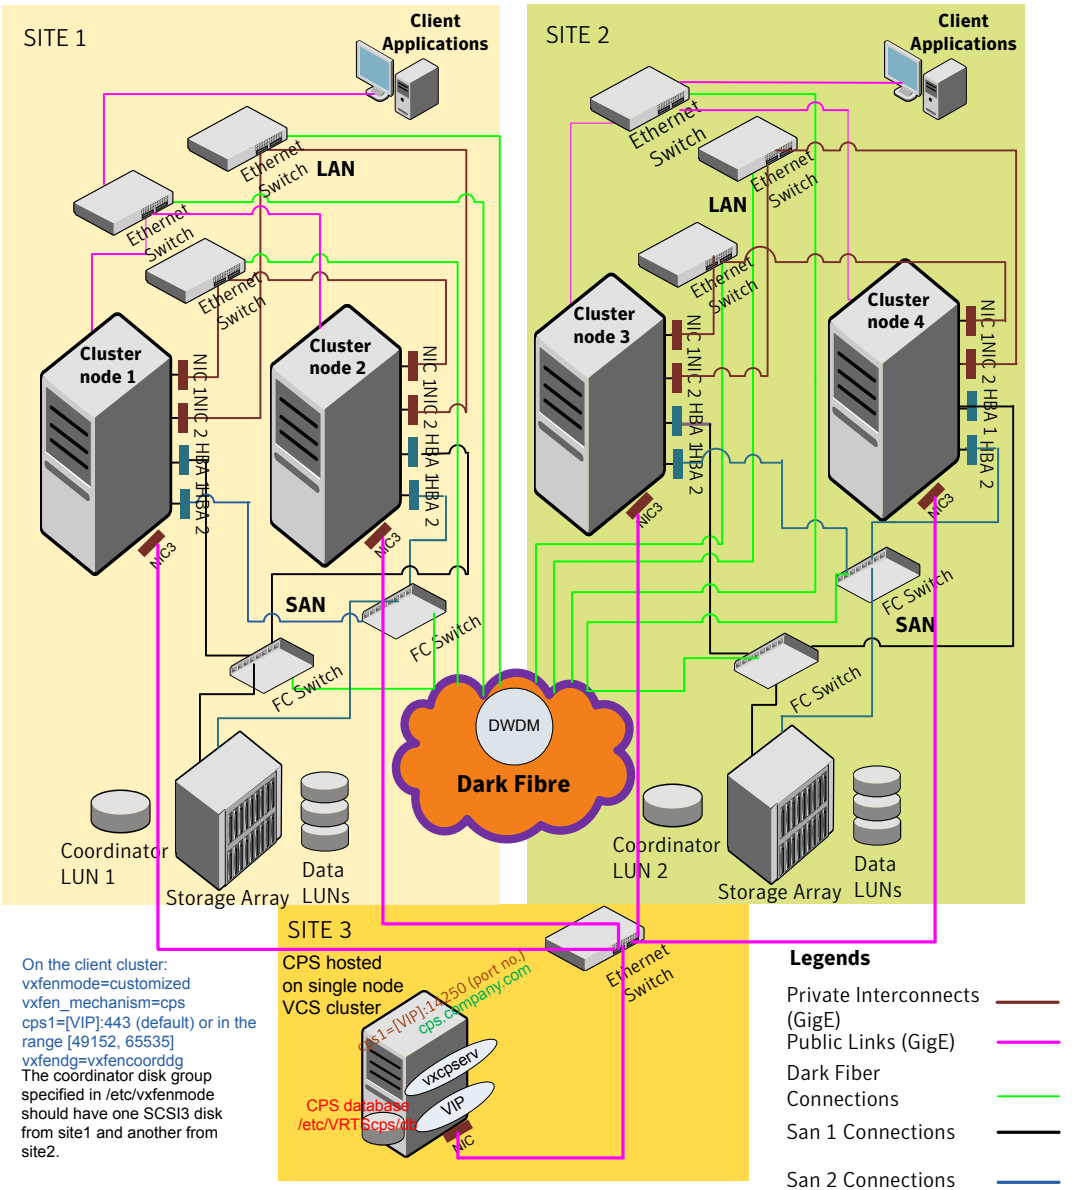

Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

Figure J-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 two SCSI-3 disks are are part of the disk group vxfencoorddg. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure J-4 Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

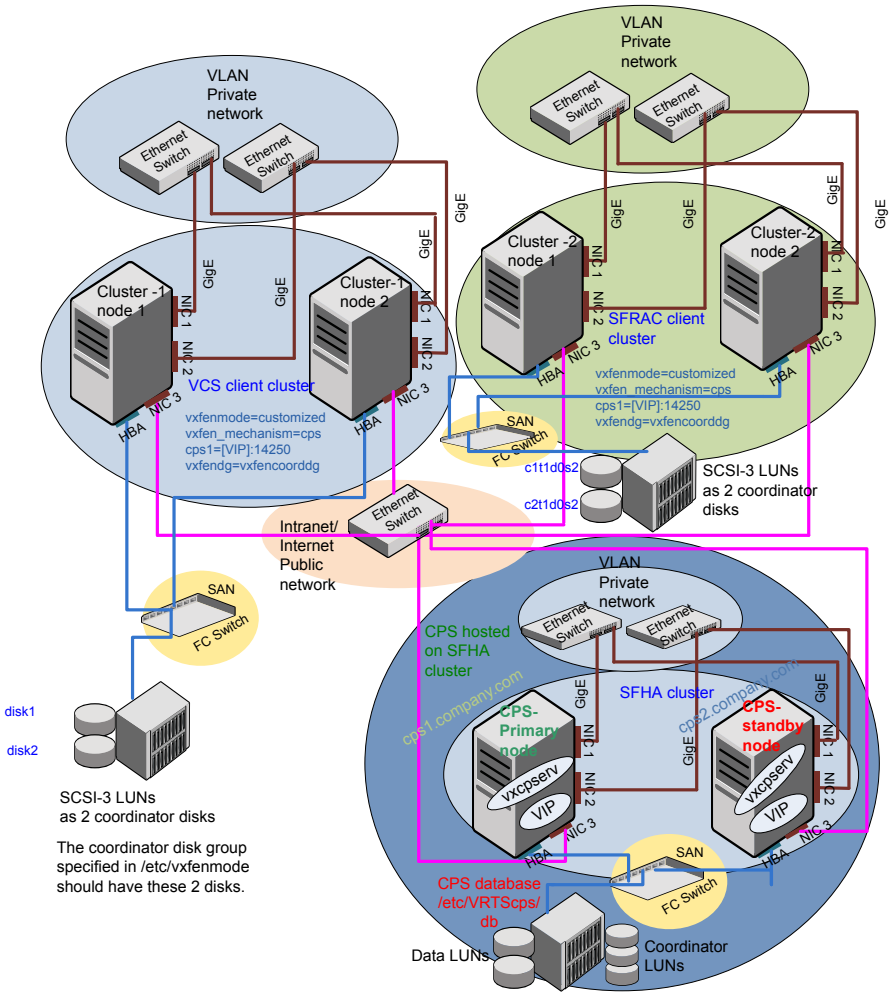

Appendix K

Reconciling major/minor numbers for NFS shared disks

This appendix includes the following topics:

- Reconciling major/minor numbers for NFS shared disks

Reconciling major/minor numbers for NFS shared disks

Your configuration may include disks on the shared bus that support NFS. You can configure the NFS file systems that you export on disk partitions or on Veritas Volume Manager volumes.

An example disk partition name is `/dev/dsk/c1t1d0s2`.

An example volume name is `/dev/vx/dsk/shareddg/vol3`. Each name represents the block device on which the file system is to be mounted.

In a VCS cluster, block devices providing NFS service must have the same major and minor numbers on each cluster node. Major numbers identify required device drivers (such as a Solaris partition or a VxVM volume). Minor numbers identify the specific devices themselves. NFS also uses major and minor numbers to identify the exported file system.

Major and minor numbers must be verified to ensure that the NFS identity for the file system is the same when exported from each node.

Checking major and minor numbers for disk partitions

The following sections describe checking and changing, if necessary, the major and minor numbers for disk partitions used by cluster nodes.

To check major and minor numbers on disk partitions

- Use the following command on all nodes exporting an NFS file system. This command displays the major and minor numbers for the block device.

  ```
  # ls -lL block_device
  ```

  The variable *block_device* refers to a partition where a file system is mounted for export by NFS. Use this command on each NFS file system. For example, type:

  ```
  # ls -lL /dev/dsk/c1t1d0s2
  ```

  Output on Node A resembles:

  ```
  crw-r-----  1 root  sys  32,1 Dec 3 11:50 /dev/dsk/c1t1d0s2
  ```

  Output on Node B resembles:

  ```
  crw-r-----  1 root  sys  32,1 Dec 3 11:55 /dev/dsk/c1t1d0s2
  ```

  Note that the major numbers (32) and the minor numbers (1) match, satisfactorily meeting the requirement for NFS file systems.

To reconcile the major numbers that do not match on disk partitions

1. Reconcile the major and minor numbers, if required. For example, if the output in the previous section resembles the following, perform the instructions beginning step 2:

   Output on Node A:

   ```
   crw-r-----  1 root  sys  32,1 Dec 3 11:50 /dev/dsk/c1t1d0s2
   ```

   Output on Node B:

   ```
   crw-r-----  1 root  sys  36,1 Dec 3 11:55 /dev/dsk/c1t1d0s2
   ```

2. Place the VCS command directory in your path.

   ```
   # export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin
   ```

3 Attempt to change the major number on System B (now 36) to match that of System A (32). Use the command:

```
# haremajor -sd major_number
```

For example, on Node B, enter:

```
# haremajor -sd 32
```

4 If the command succeeds, go to step 8.

5 If the command fails, you may see a message resembling:

```
Error: Preexisting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

6 Notice that the number 36 (the major number on Node A) is not available on Node B. Run the `haremajor` command on Node B and change it to 128,

```
# haremajor -sd 128
```

7 Run the same command on Node A. If the command fails on Node A, the output lists the available numbers. Rerun the command on both nodes, setting the major number to one available to both.

8 Reboot each system on which the command succeeds.

9 Proceed to reconcile the major numbers for your next partition.

To reconcile the minor numbers that do not match on disk partitions

1 In the example, the minor numbers are `1` and `3` and are reconciled by setting to 30 on each node.

2 Type the following command on both nodes using the name of the block device:

```
# ls -1 /dev/dsk/c1t1d0s2
```

Output from this command resembles the following on Node A:

```
lrwxrwxrwx  1 root  root  83 Dec 3 11:50
 /dev/dsk/c1t1d0s2         -> ../../
 devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:d,raw
```

The `device name` (in bold) includes the slash following the word `devices`, and continues to, but does not include, the colon.

3 Type the following command on both nodes to determine the instance numbers that the SCSI driver uses:

```
# grep sd /etc/path_to_inst | sort -n -k 2,2
```

Output from this command resembles the following on Node A:

```
"/sbus@1f,0/QLGC,isp@0,10000/sd@0,0" 0 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@1,0" 1 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@2,0" 2 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@3,0" 3 "sd"
.
.
"/sbus@1f,0/SUNW,fas@e,8800000/sd@d,0" 27 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@e,0" 28 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@f,0" 29 "sd"
```

In the output, the instance numbers are in the second field.

The instance number that is associated with the device name that matches the name for Node A displayed in step 2, is "1."

4 Compare instance numbers for the device in the output on each node.

After you review the instance numbers, perform one of the following tasks:

- If the instance number from one node is unused on the other— it does not appear in the output of step 3—edit `/etc/path_to_inst`.
  You edit this file to make the second node's instance number similar to the number of the first node.
- If the instance numbers in use on both nodes, edit `/etc/path_to_inst` on both nodes. Change the instance number that is associated with the device name to an unused number. The number needs to be greater than the highest number that other devices use. For example, the output of step 3 shows the instance numbers that all devices use (from 0 to 29). You edit the file `/etc/path_to_inst` on each node and reset the instance numbers to 30.

5 Type the following command to reboot each node on which `/etc/path_to_inst` was modified:

```
# reboot -- -rv
```

Checking the major and minor number for VxVM volumes

The following sections describe checking and changing, if necessary, the major and minor numbers for the VxVM volumes that cluster systems use.

To check major and minor numbers on VxVM volumes

1. Place the VCS command directory in your path. For example:

   ```
   # export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin
   ```

2. To list the devices, use the `ls -lL` *block_device* command on each node:

   ```
   # ls -lL /dev/vx/dsk/shareddg/vol3
   ```

   On Node A, the output may resemble:

   ```
   brw-------   1 root  root  32,43000 Mar 22 16:4 1
   /dev/vx/dsk/shareddg/vol3
   ```

   On Node B, the output may resemble:

   ```
   brw-------   1 root  root  36,43000 Mar 22 16:4 1
   /dev/vx/dsk/shareddg/vol3
   ```

3. Import the associated shared disk group on each node.

4 Use the following command on each node exporting an NFS file system. The command displays the major numbers for `vxio` and `vxspec` that Veritas Volume Manager uses . Note that other major numbers are also displayed, but only `vxio` and `vxspec` are of concern for reconciliation:

```
# grep vx /etc/name_to_major
```

Output on Node A:

```
vxdmp 30
vxio 32
vxspec 33
vxfen 87
vxglm 91
```

Output on Node B:

```
vxdmp 30
vxio 36
vxspec 37
vxfen 87
vxglm 91
```

5 To change Node B's major numbers for `vxio` and `vxspec` to match those of Node A, use the command:

```
haremajor -vx major_number_vxio major_number_vxspec
```

For example, enter:

```
# haremajor -vx 32 33
```

If the command succeeds, proceed to step 8. If this command fails, you receive a report similar to the following:

```
Error: Preexisting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

6 If you receive this report, use the `haremajor` command on Node A to change the major number (`32/33`) to match that of Node B (`36/37`). For example, enter:

```
# haremajor -vx 36 37
```

If the command fails again, you receive a report similar to the following:

```
Error: Preexisting major number 36
These are available numbers on this node: 126...
Check /etc/name_to_major on all systems for
available numbers.
```

7 If you receive the second report, choose the larger of the two available numbers (in this example, `128`). Use this number in the `haremajor` command to reconcile the major numbers. Type the following command on both nodes:

```
# haremajor -vx 128 129
```

8 Reboot each node on which `haremajor` was successful.

9 If the minor numbers match, proceed to reconcile the major and minor numbers of your next NFS block device.

10 If the block device on which the minor number does not match is a volume, consult the `vxdg`(1M) manual page. The manual page provides instructions on reconciling the Veritas Volume Manager minor numbers, and gives specific reference to the `reminor` option.

Node where the vxio driver number have been changed require rebooting.

Appendix L

Compatibility issues when installing Symantec Cluster Server with other products

This appendix includes the following topics:

- Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present
- Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present
- Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present

Installing Storage Foundation when other Symantec products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the VOM Central Server and Managed Host packages as is.
- When uninstalling Storage Foundation products where VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpbx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpbx, VRTSicsco, and VRTSat.

Appendix M

Upgrading the Steward process

This appendix includes the following topics:

- Upgrading the Steward process

Upgrading the Steward process

The Steward process can be configured in both secure and non-secure mode. The following procedures provide the steps to upgrade the Steward process.

Upgrading Steward configured in secure mode from 6.0 to 6.1

To upgrade Steward configured on Solaris 10 systems in secure mode:

1. Log on to the Steward system as a root user.
2. Stop the Steward process.

   ```
   # steward -stop -secure
   ```

3. Uninstall the VRTSvcs and VRTSperl packages.
4. Install the VRTSvcs and VRTSperl packages.
5. Start the Steward process.

   ```
   # steward -start -secure
   ```

To upgrade Steward configured on Solaris 11 systems in secure mode:

1. Log on to the Steward system as a root user.
2. Stop the Steward process.

```
# steward -stop -secure
```

3. Upgrade the VRTSvcs and VRTSperl packages.

```
# pkg set-publisher -p <path to p5p package> Symantec
# pkg update VRTSperl VRTSvcs
# pkg unset-publisher Symantec
```

4. Start the Steward process.

```
# steward -start -secure
```

Upgrading Steward configured in non-secure mode from 6.0 to 6.1

To upgrade Steward configured on Solaris 10 systems in non-secure mode:

1. Log on to the Steward system as a root user.
2. Stop the Steward process.

```
# steward -stop
```

3. Uninstall the VRTSvcs and VRTSperl packages.
4. Install the VRTSvcs and VRTSperl packages.
5. Start the Steward process.

```
# steward -start
```

To upgrade Steward configured on Solaris 11 systems in non-secure mode:

1. Log on to the Steward system as a root user.
2. Stop the Steward process.

3 Upgrade the VRTSvcs and VRTSperl package.

```
# pkg set-publisher -p <path to p5p package> Symantec
# pkg update VRTSperl VRTSvcs
# pkg unset-publisher Symantec
```

4 Start the Steward process.

```
# steward -start
```

Refer to *About the Steward process: Split-brain in two-cluster global clusters* in the *Symantec Cluster Server Administrator's Guide* for more information.

Index

Symbols

/etc/llttab
    LLT directives 261

A

abort sequence 76
about
    Deployment Server 301
    global clusters 27
    installation and configuration methods 55
    SORT 30
    Symantec product licensing 61
    Veritas Operations Manager 28
    web-based installer 53
adding
    ClusterService group 271
    users 150
adding node
    to a one-node cluster 424
attributes
    UseFence 275
Automated installer
    about 253
    installing 253
    using 253

B

BE upgrade
    administering BE in Solaris 11 392
    completing Solaris 11 upgrade 391
    creating new Solaris 11 BE 388
    upgrading Solaris 11 using the installer 389
    verifying Solaris 11 upgrade 391
before using
    web-based installer 188
Blade server considerations 78
block device
    partitions
        example file name 584
block device *(continued)*
    volumes
        example file name 584
bundled agents
    types.cf file 263

C

cables
    cross-over Ethernet 436
changing root user 403
checking product versions 44
cluster
    creating a single-node cluster
        installer 546
        manual 547
    four-node configuration 24
    removing a node from 450
    verifying operation 415
Cluster Manager 29
    installing Java Console 404
ClusterService group
    adding manually 271
cold start
    running VCS 26
commands
    format 74
    gabconfig 263, 414
    hastatus 415
    hastop 470
    hasys 416
    lltconfig 526
    lltstat 412
    vxdisksetup (initializing disks) 158
    vxlicinst 156–157, 245
    vxlicrep 156, 247
communication channels 25
communication disk 25
configuration files
    types.cf 263
configuring
    GAB 263

configuring *(continued)*
hardware 37
LLT
manual 259
private network 67
rsh 70
ssh 70, 567
switches 67
configuring VCS
adding users 150
event notification 151–152
global clusters 154
required information 82
script-based installer 136
starting 137
controllers
private Ethernet 67
SCSI 71
coordinator disks
DMP devices 33
for I/O fencing 33
setting up 273
creating
Flash archive 251
post-deployment scripts 252
creating root user 75

D

data disks
for I/O fencing 33
demo key 247
deploying
specific Symantec release 319
Symantec product updates to your environment 316
deployment management
overview 303
deployment preferences
setting 306
Deployment Script
installing 302
Deployment Server
about 301
setting up 304
Deployment Server command line option
for loading and downloading the most recent release information 308
for specifying a non-default repository location 307
directives
LLT 261
disabling
external network connection attempts 46
disk space
directories 37
language pack 37
required 37
disk space requirements 38
disks
adding and initializing 158
coordinator 273
testing with vxfentsthdw 164
verifying node access 165
documentation
accessing 403
downloading maintenance releases and hot fixes 44

E

eeprom
parameters 67
Ethernet controllers 67, 436
existing coordination points
order 207

F

FC-AL controllers 74
fibre channel 37
flarcreate 251
Flash archive 251
post-deployment scripts 252
functions
go 76

G

GAB
description 25
manual configuration 263
port membership information 414
starting 270
verifying 414
gabconfig command 263, 414
-a (verifying GAB) 414
gabtab file
creating 263
verifying after installation 526
global clusters 27
configuration 154

H

hardware
configuration 24
configuring network and storage 37
hastatus -summary command 415
hastop command 470
hasys -display command 416
hubs 67
independent 436

I

I/O fencing
checking disks 164
setting up 272
shared storage 164
I/O fencing requirements
non-SCSI-3 43
Install Bundles
integration options 330
installer
about the script-based installer 49
installer hot fixes
obtaining either manually or automatically 45
installer program
uninstalling language packages 465
Installing
VCS with the web-based installer 190
web-based installer 190
installing
Automated Installer 253
JumpStart 248
language packages 95
manually 244
manual 238
post 155
required disk space 37
simulator 407
Symantec product license keys 63
the Deployment Script 302
using Flash archive 251
using response files 211
installing manually
Japanese language packages 244
installing VCS
required information 82
installvcs
options 51
installvcs prompts
b 52
installvcs prompts *(continued)*
n 52
y 52

J

Japanese language packages 244
Java Console 29
installing 404
installing on UNIX 404
JumpStart
installing 248
Jumpstart
Generating the finish scripts 248
overview 248
Preparing installation resources 249

K

keyless licensing
setting or changing the product level 245

L

language packages 465
disk space 37
Japanese 244
license keys
adding with vxlicinst 156, 245
obtaining 62
replacing demo key 157, 247
licenses
information about 156
showing information 247
licensing
installing Symantec product license keys 63
setting or changing the product level for keyless licensing 245
links
private network 526
Live Upgrade
administering Solaris 10 boot environments 386
completing Solaris 10 upgrade 384
preparing 377
reverting to primary boot environment 386
Solaris 10 systems 376
supported upgrade paths 375
Switching boot environment for Solaris SPARC 386
Symantec Cluster Server exceptions 374
upgrading Solaris 10 on alternate boot disk 379

Live Upgrade *(continued)*
upgrading Solaris 10 using the installer 383
verifying Solaris 10 upgrade 385
LLT
description 25
directives 261
interconnects 78
manual configuration 259
starting 269
verifying 412
LLT directives
link 261
link-lowpri 261
set-cluster 261
set-node 261
lltconfig command 526
llthosts file
verifying after installation 526
lltstat command 412
lltab file
verifying after installation 526
loading and downloading the most recent release information
Deployment Server command line option for 308

M

MAC addresses 67
main.cf file
contents after installation 531
main.cf files 537
major and minor numbers
checking 585, 588
shared devices 584
MANPATH variable
setting 76
media speed 78
optimizing 78
membership information 414
mounting
software disc 80

N

network partition
preexisting 26
protecting against 24
Network partitions
protecting against 25
network switches 67
NFS 23
NFS services
shared storage 584
non-SCSI-3 fencing
manual configuration 292
setting up 292
non-SCSI-3 I/O fencing
requirements 43
non-SCSI3 fencing
setting up 181
using installvcs 181

O

obtaining
installer hot fixes either automatically or manually 45
security exception on Mozilla Firefox 189
optimizing
media speed 78
overview
deployment management 303
VCS 23

P

parameters
eeprom 67
PATH variable
setting 76
VCS commands 411
persistent reservations
SCSI-3 71
phased 337
phased upgrade 337, 339
example 338
port a
membership 414
port h
membership 414
port membership information 414
post-deployment scripts 252
preinstallation check
web-based installer 190
preparing
Live Upgrade 377
prerequisites
uninstalling 461
private network
configuring 67

R

RAM
    installation requirement 37
release images
    viewing or downloading available 309
release information
    updating on systems without Internet access 320
release notes 36
releases
    finding out which releases you have, and which upgrades or updates you may need 317
removing a system from a cluster 450
repository images
    viewing and removing repository images stored in your repository 313
requirements
    Ethernet controllers 37
    fibre channel 37
    hardware 37
    RAM Ethernet controllers 37
    SCSI host bus adapter 37
response files 53
    installation 211
    rolling upgrade 362
    syntax 54
    uninstalling 467
    upgrading 359
rolling upgrade 368
    using response files 362
    using the script-based installer 368
    versions 366
rsh 137, 567
    configuration 70

S

script-based installer
    about 49
    VCS configuration overview 136
SCSI driver
    determining instance numbers 586
SCSI host bus adapter 37
SCSI-3
    persistent reservations 71
SCSI-3 persistent reservations
    verifying 272
seeding 26
    automatic 26
    manual 26
setting
    deployment preferences 306
    MANPATH variable 76
    PATH variable 76
setting up
    Deployment Server 304
shared storage
    Fibre Channel
        setting up 74
    NFS services 584
simulataneous install or upgrade 330
simulator
    installing 407
single-node cluster
    adding a node to 424
single-system cluster
    creating 546–547
SMTP email notification 151
SNMP trap notification 152
specifying
    non-default repository location 307
ssh 137, 567
    configuration 70
    configuring 567
starting
    web-based installer 189
starting configuration
    installvcs program 137
    Symantec product installer 137
starting VCS after manual upgrade 269
storage
    fully shared vs. distributed 24
    setting up shared fibre 74
    shared 24
supported upgrade paths
    Live Upgrade 375
switches 67
Symantec product license keys
    installing 63
Symantec product updates
    deploying to your environment 316
Symantec products
    starting process 573
    stopping process 573
Symantec release
    deploying a specific release 319
system communication using rsh
    ssh 567
system state attribute value 415

T

types.cf 263
    bundled agents 263
types.cf file 263

U

uninstalling
    prerequisites 461
    using response files 467
    using the web-based installer 464
uninstalling language packages 465
updating release information
    on systems without Internet access 320
upgrade
    phased 337, 339
upgrade paths
    SPARC 324
    x64 Platform Edition 324
upgrades or updates
    finding out which releases you have 317
upgrading
    phased 337
    rolling 368
    using response files 359
    using the web-based installer 335
upgrading Steward
    in secure mode 593
    in-non-secure mode 593
using Live Upgrade 373

V

variables
    MANPATH 76
    PATH 76
VCS
    basics 23
    command directory path variable 411
    configuration files
        main.cf 530
    configuring 136
    coordinator disks 273
    documentation 403
    manually installing 238
    notifications 27
    replicated states on each system 24
    starting 269–270
VCS features 27
VCS installation
    preinstallation information 37
    verifying
        cluster operations 411
        GAB operations 411
        LLT operations 411
VCS notifications
    SMTP notification 27
    SNMP notification 27
viewing and removing repository images
    stored in your repository 313
viewing or downloading
    available release images 309
Volume Manager
    Fibre Channel 74
vxdisksetup command 158
vxlicinst command 156, 245
vxlicrep command 156, 247

W

web-based installer 190
    about 53
    before using 188
    installation 190
    preinstallation check 190
    starting 189
    uninstalling 464
    upgrading 335