

Veritas Storage Foundation™ Cluster File System High Availability Installation Guide

Solaris

6.0

Veritas Storage Foundation™ Cluster File System High Availability 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: 6.0

Document version: 6.0.4

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. 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

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 [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

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>

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 About Storage Foundation Cluster File System High Availability	23
About Veritas Storage Foundation Cluster File System High Availability	23
About I/O fencing	24
About Veritas Operations Manager	25
About Veritas Operations Manager	25
Chapter 2 Before you install	26
About planning for SFCFSHA installation	26
About installation and configuration methods	27
Assessing the system for installation readiness	28
Symantec Operations Readiness Tools	29
Prechecking your systems using the Veritas installer	29
Downloading the Veritas Storage Foundation Cluster File System High Availability software	30
Setting environment variables	31
Optimizing LLT media speed settings on private NICs	32
Guidelines for setting the media speed of the LLT interconnects	32
Creating the /opt directory	33
About using ssh or rsh with the Veritas installer	33
Setting up shared storage	33
Setting up shared storage: SCSI disks	34
Setting up shared storage: Fibre Channel	36
Prerequisites for Veritas Storage Foundation Cluster File System High Availability	37
Hardware overview and requirements for Veritas Storage Foundation Cluster File System High Availability	38
Shared storage	39
Fibre Channel switch	39

	Cluster platforms	39
Chapter 3	System requirements	40
	Release notes	40
	Hardware compatibility list (HCL)	41
	Supported operating systems	41
	I/O fencing requirements	41
	Coordinator disk requirements for I/O fencing	41
	CP server requirements	42
	Non-SCSI-3 I/O fencing requirements	45
	Veritas File System requirements	46
	Storage Foundation Cluster File System High Availability memory requirements	46
	Storage Foundation Cluster File System High Availability CPU requirements	47
	Veritas Storage Foundation Cluster File System High Availability node requirements	47
	Database requirements	47
	Disk space requirements	47
	Discovering product versions and various requirement information	48
	Number of nodes supported	48
Chapter 4	Licensing Veritas products	49
	About Veritas SFHA Solutions product licensing	49
	Setting or changing the Veritas SFHA Solutions product level for keyless licensing	50
	Installing Veritas SFHA Solutions product license keys	52
Section 2	Installation of Storage Foundation Cluster File System High Availability	53
Chapter 5	Installing Storage Foundation Cluster File System High Availability using the script-based installer	54
	Installation preparation overview	54
	Mounting the product disc	55
	About the Veritas installer	56
	Installing Storage Foundation Cluster File System High Availability using the product installer	56

	Installing language packages	60
Chapter 6	Installing Storage Foundation Cluster File System High Availability using the web-based installer	61
	About the Web-based installer	61
	Before using the Veritas Web-based installer	62
	Starting the Veritas Web-based installer	63
	Obtaining a security exception on Mozilla Firefox	63
	Performing a pre-installation check with the Veritas Web-based installer	64
	Installing SFCFSHA with the Web-based installer	64
Chapter 7	Installing Storage Foundation Cluster File System High Availability using other methods	67
	Installing using JumpStart	67
	Overview of JumpStart installation tasks	68
	Generating the finish scripts	68
	Preparing installation resources	70
	Adding language pack information to the finish file	71
	Using a Flash archive to install SFCFSHA and the operating system	72
	Creating the Veritas post-deployment scripts	73
	Installing SFCFSHA using the pkgadd command	74
Section 3	Configuration of Veritas Storage Foundation Cluster File System High Availability	77
Chapter 8	Preparing to configure SFCFSHA	78
	About configuring SFCFSHA clusters for data integrity	78
	About I/O fencing for SFCFSHA in virtual machines that do not support SCSI-3 PR	79
	About I/O fencing components	80
	About data disks	80
	About coordination points	80
	About preferred fencing	81
	About I/O fencing configuration files	82
	About planning to configure I/O fencing	84

	Typical SFCFSHA cluster configuration with server-based I/O fencing	88
	Recommended CP server configurations	89
	Setting up the CP server	92
	Planning your CP server setup	92
	Installing the CP server using the installer	93
	Configuring the CP server cluster in secure mode	94
	Setting up shared storage for the CP server database	95
	Configuring the CP server using the configuration utility	96
	Configuring the CP server manually	105
	Verifying the CP server configuration	106
Chapter 9	Configuring Veritas Storage Foundation Cluster File System High Availability	108
	Configuring SFCFSHA using the script-based installer	108
	Overview of tasks to configure SFCFSHA using the script-based installer	108
	Starting the software configuration	109
	Specifying systems for configuration	110
	Configuring the cluster name	111
	Configuring private heartbeat links	111
	Configuring the virtual IP of the cluster	114
	Configuring the cluster in secure mode	116
	Configuring a secure cluster node by node	117
	Adding VCS users	120
	Configuring SMTP email notification	121
	Configuring SNMP trap notification	123
	Configuring global clusters	125
	Completing the SFCFSHA configuration	126
	Verifying and updating licenses on the system	127
	Configuring SFCFSHA using the Web-based installer	129
	Database configuration requirements	134
	Configuring Veritas Storage Foundation Cluster File System High Availability manually	134
	Configuring Veritas Volume Manager	134
	Configuring Veritas File System	141
	Configuring the SFDB repository database after installation	142
Chapter 10	Configuring SFCFSHA for data integrity	144
	Setting up disk-based I/O fencing using installsfcfsha	144
	Initializing disks as VxVM disks	144
	Configuring disk-based I/O fencing using installsfcfsha	145

Checking shared disks for I/O fencing	147
Setting up disk-based I/O fencing manually	152
Identifying disks to use as coordinator disks	152
Setting up coordinator disk groups	153
Creating I/O fencing configuration files	154
Modifying VCS configuration to use I/O fencing	155
Verifying I/O fencing configuration	156
Setting up server-based I/O fencing using <code>installsfcfsha</code>	157
Setting up non-SCSI-3 server-based I/O fencing in virtual environments using <code>installsfcfsha</code>	157
Configuring SFCFSHA for data integrity using the Web-based installer	159
Setting up server-based I/O fencing manually	162
Preparing the CP servers manually for use by the SFCFSHA cluster	163
Configuring server-based fencing on the SFCFSHA cluster manually	166
Configuring CoordPoint agent to monitor coordination points	172
Verifying server-based I/O fencing configuration	173
Setting up non-SCSI-3 fencing in virtual environments manually	174
Sample <code>/etc/vxfenmode</code> file for non-SCSI-3 fencing	176
Enabling or disabling the preferred fencing policy	179

Section 4 Upgrading Storage Foundation Cluster File System High Availability 182

Chapter 11 Preparing to upgrade Veritas Storage Foundation Cluster File System High Availability	183
About upgrading	183
About the different ways that you can upgrade	184
Supported upgrade paths	184
About using the installer to upgrade when the root disk is encapsulated	186
Preparing to upgrade	186
Getting ready for the upgrade	187
Creating backups	189
Determining if the root disk is encapsulated	190
Pre-upgrade planning for Veritas Volume Replicator	190
Preparing to upgrade VVR when VCS agents are configured	192
Verifying that the file systems are clean	196
Upgrading the array support	197

Chapter 12	Performing a typical Storage Foundation Cluster File System High Availability upgrade using the installer	199
	Performing a full upgrade	199
	Ensuring the file systems are clean	199
	Modifying the main.cf file	200
	Performing the upgrade	203
Chapter 13	Performing a rolling upgrade	208
	Performing a rolling upgrade using the installer	208
	About rolling upgrades	208
	Performing a rolling upgrade using the script-based installer	211
	Performing a rolling upgrade of SFCFSHA using the Web-based installer	214
Chapter 14	Performing a phased upgrade	217
	Performing a phased upgrade of SFCFSHA	217
	Prerequisites for a phased upgrade	217
	Planning for a phased upgrade	218
	Phased upgrade limitations	218
	Moving the service groups to the second subcluster	218
	Upgrading the SFCFSHA stack on the first subcluster	221
	Preparing the second subcluster	222
	Activating the first subcluster	224
	Upgrading the operating system on the second subcluster	225
	Upgrading the second subcluster	225
	Completing the phased upgrade	226
Chapter 15	Upgrading the operating system	227
	Upgrading the Solaris operating system	227
Chapter 16	Upgrading Veritas Volume Replicator	231
	Upgrading Veritas Volume Replicator	231
	Upgrading VVR without disrupting replication	231
Chapter 17	Upgrading language packages	233
	Upgrading language packages	233

Chapter 18	Migrating from SFHA to SFCFSHA	234
	Migrating from SFHA to SFCFSHA 6.0	234
Chapter 19	Upgrading with Live Upgrade	237
	About Live Upgrade	237
	About Live Upgrade in a Veritas Volume Replicator (VVR) environment	238
	Supported upgrade paths for Live Upgrade	239
	Performing Live Upgrade in a Solaris zone environment	240
	Before you upgrade SFCFSHA using Solaris Live Upgrade	242
	Upgrading SFCFSHA and Solaris using Live Upgrade	245
	Creating a new boot environment on the alternate boot disk	246
	Upgrading SFCFSHA using the installer for a Live Upgrade	247
	Upgrading SFCFSHA manually	248
	Completing the Live Upgrade	251
	Verifying Live Upgrade of SFCFSHA	252
	Upgrading Solaris using Live Upgrade	253
	Removing and reinstalling SFCFSHA using the installer	253
	Upgrading SFCFSHA using Live Upgrade	254
	Administering boot environments	255
	Reverting to the primary boot environment	255
	Switching the boot environment for Solaris SPARC	255
	Switching the boot environment for Solaris x86-64	257
Chapter 20	Performing post-upgrade tasks	260
	Re-joining the backup boot disk group into the current disk group	260
	Reverting to the backup boot disk group after an unsuccessful upgrade	261
Section 5	Verification of the installation or the upgrade	262
Chapter 21	Verifying the Storage Foundation Cluster File System High Availability installation	263
	Performing a postcheck on a node	264
	Verifying that the products were installed	264
	Installation log files	264
	Using the installation log file	265
	Using the summary file	265

About enabling LDAP authentication for clusters that run in secure mode	265
Enabling LDAP authentication for clusters that run in secure mode	267
Starting and stopping processes for the Veritas products	273
Checking Veritas Volume Manager processes	273
Checking Veritas File System installation	274
Verifying Veritas File System kernel installation	274
Verifying command installation	274
Verifying agent configuration for Storage Foundation Cluster File System High Availability	275
Synchronizing time on Cluster File Systems	275
Configuring VCS for Storage Foundation Cluster File System High Availability	276
main.cf file	276
Storage Foundation Cluster File System HA Only	277
Veritas Cluster Server application failover services	278
Configuring the cluster UUID when creating a cluster manually	278
About the cluster UUID	278
About the LLT and GAB configuration files	278
Verifying the LLT, GAB, and VCS configuration files	281
Verifying LLT, GAB, and cluster operation	282
Verifying LLT	282
Verifying GAB	285
Verifying the cluster	286
Verifying the cluster nodes	287
 Section 6	
Adding and removing nodes	290
 Chapter 22	
Adding a node to a cluster	291
About adding a node to a cluster	291
Before adding a node to a cluster	292
Adding a node to a cluster using the product installer	294
Adding a node using the Web-based installer	298
Adding the node to a cluster manually	298
Starting Volume Manager on the new node	310
Configuring cluster files on the new node	310
Configuring cluster files on the new node	302
Starting fencing on the new node	316
Configuring CVM and CFS on the new node	318
Starting VCS after adding the new node	317

	Configuring the ClusterService group for the new node	319
	Starting Volume Manager on the new node	310
	Configuring cluster files on the new node	310
	Setting up the node to run in secure mode	312
	Starting fencing on the new node	316
	Starting VCS after adding the new node	317
	Configuring CVM and CFS on the new node	318
	Configuring the ClusterService group for the new node	319
	Configuring server-based fencing on the new node	320
	Adding the new node to the vxfen service group	321
	Updating the Storage Foundation for Databases (SFDB) repository after adding a node	322
	Sample configuration file for adding a node to the cluster	322
Chapter 23	Removing a node from Storage Foundation Cluster File System clusters	327
	About removing a node from a cluster	327
	Removing a node from a cluster	328
	Modifying the VCS configuration files on existing nodes	329
	Removing the node configuration from the CP server	331
	Removing security credentials from the leaving node	332
	Updating the Storage Foundation for Databases (SFDB) repository after removing a node	333
	Sample configuration file for removing a node from the cluster	333
Section 7	Setting up and configuring replicated global cluster	336
Chapter 24	Setting up a replicated global cluster	337
	Replication in the SFCFSHA environment	337
	About setting up a global cluster in an SFCFSHA environment	338
	Configuring an SFCFSHA global cluster at the primary site	340
	Configuring an SFCFSHA cluster at the secondary site	342
	Configuring replication for SFCFSHA clusters at both sites	344
	Modifying the ClusterService group for a global SFCFSHA cluster	344
	Modifying the global SFCFSHA cluster configuration using the wizard	345
	Defining the remote SFCFSHA cluster and heartbeat objects	346
	Configuring the VCS service groups for global SFCFSHA clusters	349

Chapter 25	Configuring a global cluster using VVR	350
	About configuring a global SFCFSHA cluster using VVR for replication	350
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the primary site	351
	Creating the data and SRL volumes on the primary site	351
	Setting up the Replicated Volume Group (RVG) on the primary site	353
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the secondary site	354
	Creating the data and SRL volumes on the secondary site	355
	Editing the /etc/vx/vras/.rdg files	355
	Setting up IP addresses for RLINKs on each cluster	356
	Setting up the disk group on secondary site for replication	357
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): starting replication	359
	Starting replication using automatic synchronization	359
	Starting replication using full synchronization with Storage Checkpoint	360
	Verifying replication status	360
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): configuring cluster resources	361
	About modifying the VCS configuration for replication	361
	Modifying the VCS Configuration on the Primary Site	363
	Modifying the VCS Configuration on the Secondary Site	367
	Managing a global SFCFSHA cluster using Veritas Volume Replicator (VVR)	371
	Migrating the role of primary site to the secondary site	373
	Taking over the primary role by the remote cluster	374
Section 8	Uninstallation of Storage Foundation Cluster File System High Availability	378
Chapter 26	Uninstalling Storage Foundation Cluster File System High Availability	379
	About removing Veritas Storage Foundation Cluster File System High Availability	379
	Preparing to uninstall	380
	Shutting down cluster operations	389

	Preparing the tunables file	430
	Setting parameters for the tunables file	430
	Tunables value parameter definitions	431
Appendix D	Configuring I/O fencing using a response file	438
	Configuring I/O fencing using response files	438
	Response file variables to configure disk-based I/O fencing	439
	Sample response file for configuring disk-based I/O fencing	440
	Response file variables to configure server-based I/O fencing	441
	Sample response file for configuring server-based I/O fencing	443
	Response file variables to configure non-SCSI-3 server-based I/O fencing	443
	Sample response file for configuring non-SCSI-3 server-based I/O fencing	444
Appendix E	Configuring the secure shell or the remote shell for communications	446
	About configuring secure shell or remote shell communication modes before installing products	446
	Manually configuring and passwordless ssh	447
	Restarting the ssh session	451
	Enabling and disabling rsh for Solaris	452
Appendix F	Storage Foundation Cluster File System High Availability components	454
	Veritas Storage Foundation Cluster File System High Availability installation packages	454
	Veritas Cluster Server installation packages	457
	Veritas Cluster File System installation packages	458
	Chinese language packages	458
	Japanese language packages	458
	Veritas Storage Foundation obsolete and reorganized installation packages	459
Appendix G	High availability agent information	463
	About agents	463
	VCS agents included within SFCFSHA	464
	Enabling and disabling intelligent resource monitoring for agents manually	464
	Administering the AMF kernel driver	466

CVMCluster agent	467
Entry points for CVMCluster agent	467
Attribute definition for CVMCluster agent	468
CVMCluster agent type definition	469
CVMCluster agent sample configuration	469
CVMVxconfigd agent	469
Entry points for CVMVxconfigd agent	470
Attribute definition for CVMVxconfigd agent	470
CVMVxconfigd agent type definition	472
CVMVxconfigd agent sample configuration	472
CVMVoidg agent	472
Entry points for CVMVoidg agent	472
Attribute definition for CVMVoidg agent	473
CVMVoidg agent type definition	474
CVMVoidg agent sample configuration	475
CFSMount agent	475
Entry points for CFSMount agent	476
Attribute definition for CFSMount agent	476
CFSMount agent type definition	478
CFSMount agent sample configuration	479
CFSfsckd agent	479
Entry points for CFSfsckd agent	479
Attribute definition for CFSfsckd agent	480
CFSfsckd agent type definition	481
CFSfsckd agent sample configuration	482
Appendix H Troubleshooting information	483
Restarting the installer after a failed connection	483
What to do if you see a licensing reminder	483
Troubleshooting information	484
Storage Foundation Cluster File System High Availability installation	
issues	484
Incorrect permissions for root on remote system	485
Inaccessible system	486
Storage Foundation Cluster File System High Availability	
problems	486
Unmount failures	486
Mount failures	486
Command failures	487
Performance issues	488
High availability issues	488

Appendix I	Troubleshooting cluster installation	490
	Installer cannot create UUID for the cluster	490
	The vxfsntsthdw utility fails when SCSI TEST UNIT READY command fails	491
	Troubleshooting CP server	491
	Troubleshooting issues related to the CP server service group	492
	Checking the connectivity of CP server	492
	Troubleshooting server-based fencing on the SFCFSHA cluster nodes	493
	Issues during fencing startup on SFCFSHA cluster nodes set up for server-based fencing	493
	Issues during online migration of coordination points	494
Appendix J	Sample SFCFSHA cluster setup diagrams for CP server-based I/O fencing	495
	Configuration diagrams for setting up server-based I/O fencing	495
	Two unique client clusters served by 3 CP servers	496
	Client cluster served by highly available CPS and 2 SCSI-3 disks	497
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	498
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	500
Appendix K	Reconciling major/minor numbers for NFS shared disks	502
	Reconciling major/minor numbers for NFS shared disks	502
	Checking major and minor numbers for disk partitions	503
	Checking the major and minor number for VxVM volumes	506
Appendix L	Configuring LLT over UDP using IPv6	509
	Using the UDP layer of IPv6 for LLT	509
	When to use LLT over UDP	509
	Manually configuring LLT over UDP using IPv6	509
	The link command in the /etc/lfttab file	510
	The set-addr command in the /etc/lfttab file	511
	Selecting UDP ports	511
	Sample configuration: direct-attached links	512
	Sample configuration: links crossing IP routers	514

Appendix M	Configuring LLT over UDP using IPv4	517
	Using the UDP layer for LLT	517
	When to use LLT over UDP	517
	Manually configuring LLT over UDP using IPv4	517
	Broadcast address in the /etc/lldtab file	518
	The link command in the /etc/lldtab file	518
	The set-addr command in the /etc/lldtab file	519
	Selecting UDP ports	520
	Configuring the netmask for LLT	521
	Configuring the broadcast address for LLT	521
	Sample configuration: direct-attached links	522
	Sample configuration: links crossing IP routers	523
Appendix N	Compatibility issues when installing Storage Foundation Cluster File System High Availability with other products	526
	Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present	526
	Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present	527
	Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present	527
Index		528

Installation overview and planning

- [Chapter 1. About Storage Foundation Cluster File System High Availability](#)
- [Chapter 2. Before you install](#)
- [Chapter 3. System requirements](#)
- [Chapter 4. Licensing Veritas products](#)

About Storage Foundation Cluster File System High Availability

This chapter includes the following topics:

- [About Veritas Storage Foundation Cluster File System High Availability](#)
- [About I/O fencing](#)
- [About Veritas Operations Manager](#)

About Veritas Storage Foundation Cluster File System High Availability

Veritas Storage Foundation Cluster File System High Availability by Symantec extends Veritas Storage Foundation to support shared data in a storage area network (SAN) environment. Using Storage Foundation Cluster File System High Availability, multiple servers can concurrently access shared storage and files transparently to applications.

Veritas Storage Foundation Cluster File System High Availability also provides increased automation and intelligent management of availability and performance.

Veritas Storage Foundation Cluster File System High Availability includes Veritas Cluster Server, which adds high availability functionality to the product.

To install the product, follow the instructions in the *Veritas Storage Foundation Cluster File System High Availability Installation Guide*.

For information on high availability environments, read the Veritas Cluster Server documentation.

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, when you install SFCFSHA. To protect data on shared disks, you must configure I/O fencing after you install and configure SFCFSHA.

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 multiple clusters.

In virtualized environments that do not support SCSI-3 PR, SFCFSHA supports non-SCSI-3 server-based I/O fencing.

See [“About planning to configure I/O fencing”](#) on page 84.

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

See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

About Veritas Operations Manager

Veritas Operations Manager by Symantec gives you a single, centralized management console for the Veritas Storage Foundation and High Availability products. You can use it to monitor, visualize, and manage storage and cluster resources, and generate reports about them. Veritas Operations Manager lets administrators centrally manage diverse datacenter environments.

About Veritas Operations Manager

Symantec recommends use of Veritas Operations Manager to manage Storage Foundation and Cluster Server environments.

Veritas Operations Manager provides a centralized management console for Veritas Storage Foundation and High Availability products. You can use Veritas Operations Manager to monitor, visualize, and manage storage resources and generate reports.

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.

The Veritas Enterprise Administrator (VEA) console is no longer packaged with Storage Foundation products. If you want to continue using VEA, a software version is available for download from <http://www.symantec.com/operations-manager/support>. Veritas Storage Foundation Management Server is deprecated.

If you want to manage a single cluster using Cluster Manager (Java Console), a version is available for download from https://www4.symantec.com/Vrt/offer?a_id=89446. You cannot manage the new features of this release using the Java Console. Veritas Cluster Server Management Console is deprecated.

Before you install

This chapter includes the following topics:

- [About planning for SFCFSHA installation](#)
- [About installation and configuration methods](#)
- [Assessing the system for installation readiness](#)
- [Downloading the Veritas Storage Foundation Cluster File System High Availability software](#)
- [Setting environment variables](#)
- [Optimizing LLT media speed settings on private NICs](#)
- [Guidelines for setting the media speed of the LLT interconnects](#)
- [Creating the /opt directory](#)
- [About using ssh or rsh with the Veritas installer](#)
- [Setting up shared storage](#)
- [Prerequisites for Veritas Storage Foundation Cluster File System High Availability](#)
- [Hardware overview and requirements for Veritas Storage Foundation Cluster File System High Availability](#)

About planning for SFCFSHA installation

Before you continue, make sure that you are using the current version of this guide. The latest documentation is available on the Symantec Symantec Operations Readiness Tools (SORT) website.

<https://sort.symantec.com/documents>

Document version: 6.0.4.

This installation guide is designed for system administrators who already have a knowledge of basic UNIX system and network administration. Basic knowledge includes commands such as `tar`, `mkdir`, and simple shell scripting. Also required is basic familiarity with the specific platform and operating system where SFCFSHA will be installed.

Follow the preinstallation instructions if you are installing Veritas Storage Foundation Cluster File System High Availability.

About installation and configuration methods

You can use one of the following methods to install and configure SFCFSHA.

Table 2-1 Installation and configuration methods

Method	Description
<p>Interactive installation and configuration using the script-based installer</p> <p>Note: If you obtained SFCFSHA from an electronic download site, you must use the <code>installsfcfsha</code> script instead of the <code>installer</code> script.</p>	<p>You can use one of the following script-based installers:</p> <ul style="list-style-type: none"> ■ Common product installer script: <code>installer</code> The common product installer script provides a menu that simplifies the selection of installation and configuration options. ■ Product-specific installation script: <code>installsfcfsha</code> ■ The product-specific installation script provides command-line interface options. Installing and configuring with the <code>installsfcfsha</code> script is identical to specifying SFCFSHA from the <code>installer</code> script. Use this method to install or configure only SFCFSHA.
<p>Silent installation using the response file</p>	<p>The response file automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. You can use the script-based installers with the response file to install silently on one or more systems.</p> <p>See “About response files” on page 409.</p>

Table 2-1 Installation and configuration methods (*continued*)

Method	Description
Web-based installer	<p>The Web-based installer provides an interface to manage the installation and configuration from a remote site using a standard Web browser.</p> <p><code>webinstaller</code></p> <p>See “About the Web-based installer” on page 61.</p>
JumpStart	<p>You can use the Veritas product installer of the product-specific installation script to generate a JumpStart script file. Use the generated script to install Veritas packages from your JumpStart server.</p> <p>See “Installing using JumpStart” on page 67.</p>
Manual installation and configuration	<p>Manual installation uses the Solaris commands to install SFCFSHA. To retrieve a list of all packages and patches required for all products in the correct installation order, enter:</p> <pre># installer -allpkgs</pre> <p>Use the Solaris commands to install SFCFSHA. Then use a manual or an interactive method with <code>installsfcfsha</code> or <code>installer</code> script to configure the SFCFSHA stack.</p>

Assessing the system for installation readiness

Symantec provides the following tools for assessing your system, to ensure that the system meets the requirements for installing Veritas Storage Foundation Cluster File System High Availability 6.0.

Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a Web-based application that is designed to support Symantec enterprise products.

See [“Symantec Operations Readiness Tools”](#) on page 29.

Prechecking your systems using the installer Performs a pre-installation check on the specified systems. The Veritas product installer reports whether the specified systems meet the minimum requirements for installing Veritas Storage Foundation Cluster File System High Availability 6.0.

See “Prechecking your systems using the Veritas installer” on page 29.

Symantec Operations Readiness Tools

[Symantec Operations Readiness Tools \(SORT\)](#) is a Web site that automates and simplifies some of the most time-consuming administrative tasks. SORT helps you manage your datacenter more efficiently and get the most out of your Symantec products.

Among its broad set of features, SORT lets you do the following:

- Generate server-specific reports that describe how to prepare your servers for installation or upgrade of Symantec enterprise products.
- Access a single site with the latest production information, including patches, agents, and documentation.
- Create automatic email notifications for changes in patches, documentation, and array-specific modules.

To access SORT, go to:

<https://sort.symantec.com>

Prechecking your systems using the Veritas installer

The script-based and Web-based installer's precheck option checks for the following:

- Recommended swap space for installation
- Recommended memory sizes on target systems for Veritas programs for best performance
- Required operating system versions

To use the precheck option

- 1 Start the script-based or Web-based installer.
- 2 Select the precheck option:
 - From the Web-based installer, select the **Perform a Pre-Installation Check** from the Task pull-down menu.

- In the script-based installer, from root on the system where you want to perform the check, start the installer.

```
# ./installer
```

In the Task Menu, press the p key to start the precheck.

- 3 Review the output and make the changes that the installer recommends.

Downloading the Veritas Storage Foundation Cluster File System High Availability software

One method of obtaining the Veritas Storage Foundation Cluster File System High Availability software is to download it to your local system from the Symantec Web site.

For a Trialware download, perform the following. Contact your Veritas representative for more information.

To download the trialware version of the software

- 1 Open the following link in your browser:
<http://www.symantec.com/index.jsp>
- 2 On the bottom of the page, click the **Downloads** link.
- 3 In the Business field, click **Trialware**.
- 4 On the next page near the bottom of the page, click **Business Continuity**.
- 5 Under Cluster Server, click **Download Now**.
- 6 In the new window, click **Download Now**.
- 7 You can use existing credentials to log in or create new credentials.
- 8 Review the terms and conditions, and click **I agree**.
- 9 Find the product that you want to download and select it. Continue with the installation.

If you download a standalone Veritas product, the single product download files do not contain the product installer. Use the installation script for the specific product to install the product.

Note: Trialware is the full product version. The enabled licensing places the product in a demo or a trial state.

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 401.

To download the software

- 1 Verify that you have enough space on your filesystem to store the downloaded software.

The estimated space for download, gunzip, and tar extract is:

2 GB for SPARC

1.5 GB for Opteron

See [“Disk space requirements”](#) on page 47.

If you plan to install the software on the same system, make sure that you also have enough space for the installed software.

- 2 To see the space available, you can use the `df` command with the name of the local file system where you intend to download the software.

```
# /usr/bin/df -l filesystem
```

Caution: When you select a location to download files, do not select a directory that contains Veritas products from a previous release or maintenance pack. Make sure that different versions exist in different directories.

- 3 Download the software, specifying the file system with sufficient space for the file.

Setting environment variables

Most of the commands used in the installation are in the `/sbin` or `/usr/sbin` directory. Add these directories to your `PATH` environment variable as necessary.

After installation, SFCFSHA commands are in `/opt/VRTS/bin`. SFCFSHA manual pages are stored in `/opt/VRTS/man`.

Some VCS custom scripts reside in `/opt/VRTSvcs/bin`. If you are installing a high availability product, add `/opt/VRTSvcs/bin` to the `PATH` also.

Add the following directories to your `PATH` and `MANPATH` environment variable:

- If you are using Bourne or Korn shell (`sh` or `ksh`), enter the following:

```
$ PATH=$PATH:/usr/sbin:/opt/VRTS/bin
$ MANPATH=/usr/share/man:/opt/VRTS/man:$MANPATH
$ export PATH MANPATH
```

- If you are using a C shell (`csh` or `tcsh`), enter the following:

```
% set path = ( $path /usr/sbin /opt/VRTS/bin )
% setenv MANPATH /usr/share/man:/opt/VRTS/man:$MANPATH
```

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 for more information.

Creating the /opt directory

The directory `/opt` must exist, be writable and must not be a symbolic link.

If you are upgrading, you cannot have a symbolic link from `/opt` to an unconverted volume. If you do have a symbolic link to an unconverted volume, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

About using ssh or rsh with the Veritas installer

The installer uses passwordless secure shell (ssh) or remote shell (rsh) communications among systems. The installer uses the ssh or rsh daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. You then provide the installer with the superuser passwords for the systems where you plan to install. The ssh 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 or rsh configuration from the systems.

See [“Installation script options”](#) on page 402.

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:

- When you add new nodes to an existing cluster.
- When the nodes are in a subcluster during a phased upgrade.
- When you perform installer sessions using a response file.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 446.

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 84.

See also the *Veritas Storage Foundation Cluster File System High Availability 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 `nvrामrc` 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 `nvrामrc` script. The path information varies from system to system.

- 5 Edit the `nvrामrc` 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 `nvrामrc` 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 `nvrामrc` 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 `nvrामrc` script by entering:

```
{0} ok printenv nvrामrc
```

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 `nvrामrc` script on the node.

```
{0} ok setenv use-nvrामrc? 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###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.

Prerequisites for Veritas Storage Foundation Cluster File System High Availability

Each cluster node must be connected to the public network and each must have a unique host name by which it can be addressed on the public network. The local node from which you install does not have to be part of the cluster.

Provide the following information when installing SFCFSHA:

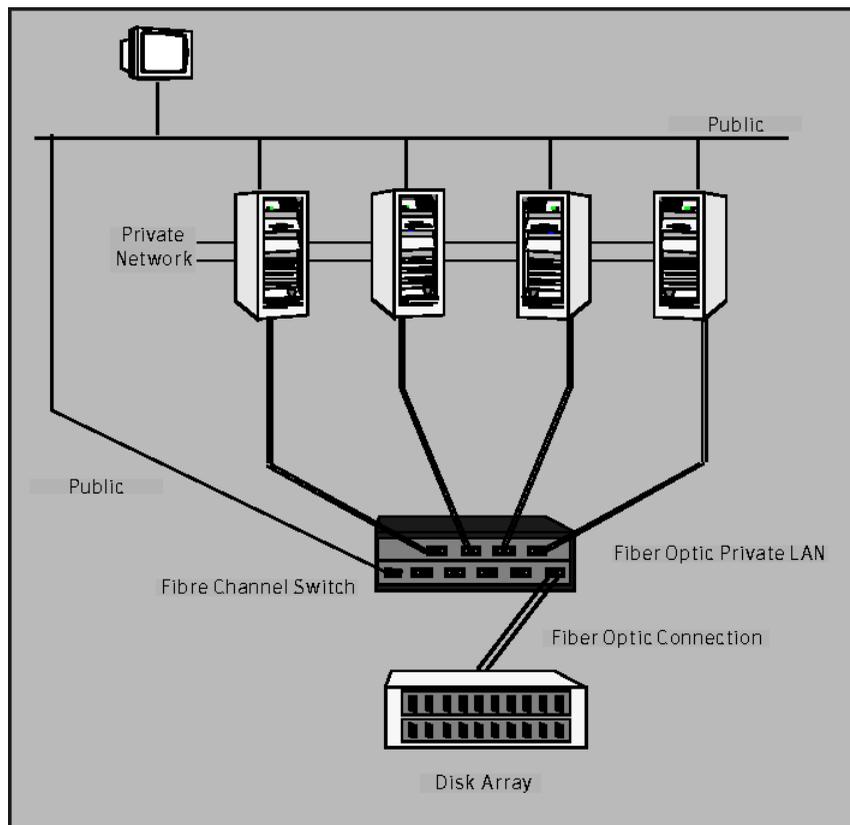
- The cluster name, beginning with a letter (a-z, A-Z).
- A unique ID from 0-65535 for the cluster. Within the public subnet, a new cluster using a duplicate cluster ID can cause existing clusters to fail.
- The host names of the cluster nodes.
- The device names of the network interface cards (NICs) used for the private networks among nodes.
- Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities as root on all cluster nodes or remote systems.
- Symantec recommends configuring the cluster with I/O fencing enabled. I/O fencing requires shared devices to support SCSI-3 Persistent Reservations (PR). Enabling I/O fencing prevents data corruption caused by a split brain scenario.
The Veritas Storage Foundation Cluster File System High Availability is supported without I/O fencing enabled. However, without I/O fencing enabled, split brain scenarios can result in data corruption.
- In a large cluster environment, make sure the first volume of the volume set is large enough to accommodate all of the metadata. A large cluster environment includes more than 14 nodes, and a volume set with more than 40 volumes. The minimum size of the first volume should be more than 900M.

Hardware overview and requirements for Veritas Storage Foundation Cluster File System High Availability

VxFS cluster functionality runs optimally on a Fibre Channel fabric. Fibre Channel technology provides the fastest, most reliable, and highest bandwidth connectivity currently available. By employing Fibre Channel technology, SFCFSHA can be used in conjunction with the latest Veritas Storage Area Network (SAN) applications to provide a complete data storage and retrieval solution.

Figure 2-1 shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

Figure 2-1 Four Node SFCFSHA Cluster Built on Fibre Channel Fabric



Shared storage

Shared storage can be one or more shared disks or a disk array connected either directly to the nodes of the cluster or through a Fibre Channel Switch. Nodes can also have non-shared or local devices on a local I/O channel. It is advisable to have `/`, `/usr`, `/var` and other system partitions on local devices.

Fibre Channel switch

Each node in the cluster must have a Fibre Channel I/O channel to access shared storage devices. The primary component of the Fibre Channel fabric is the Fibre Channel switch.

Cluster platforms

There are several hardware platforms that can function as nodes in a Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) cluster.

See the *Veritas Storage Foundation Cluster File System High Availability Release Notes*.

Note: For a cluster to work correctly, all nodes must have the same time. If you are not running the Network Time Protocol (NTP) daemon, make sure the time on all the systems comprising your cluster is synchronized.

System requirements

This chapter includes the following topics:

- [Release notes](#)
- [Hardware compatibility list \(HCL\)](#)
- [Supported operating systems](#)
- [I/O fencing requirements](#)
- [Veritas File System requirements](#)
- [Storage Foundation Cluster File System High Availability memory requirements](#)
- [Storage Foundation Cluster File System High Availability CPU requirements](#)
- [Veritas Storage Foundation Cluster File System High Availability node requirements](#)
- [Database requirements](#)
- [Disk space requirements](#)
- [Discovering product versions and various requirement information](#)
- [Number of nodes supported](#)

Release notes

The *Release Notes* for each Veritas 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>

Hardware compatibility list (HCL)

The hardware compatibility list contains information about supported hardware and is updated regularly. Before installing or upgrading Storage Foundation and High Availability Solutions products, review the current compatibility list to confirm the compatibility of your hardware and software.

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

<http://www.symantec.com/docs/TECH170013>

For information on specific High Availability setup requirements, see the *Veritas Cluster Server Installation Guide*.

Supported operating systems

For information on supported operating systems, see the *Veritas Storage Foundation Cluster File System High Availability Release Notes*.

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 41.
- CP servers
See “[CP server requirements](#)” on page 42.

If you have installed SFCFSHA 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 45.

Coordinator disk requirements for I/O fencing

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

- For disk-based I/O fencing, you must have three coordinator disks.
- The coordinator disks can be raw devices, DMP devices, or iSCSI devices.
- Each of the coordinator disks must use a physically separate disk or LUN.
Symantec recommends using the smallest possible LUNs for coordinator disks.

- Each of the coordinator disks should exist on a different disk array, if possible.
- The coordinator disks must support SCSI-3 persistent reservations.
- Symantec recommends using hardware-based mirroring for coordinator disks.
- Coordinator disks must not be used to store data or must not be included in disk groups that store user data.
- Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.

CP server requirements

SFCFSHA 6.0 clusters (application clusters) support coordination point servers (CP servers) which are hosted on the following VCS and SFHA versions:

- VCS 6.0, 5.1SP1, or 5.1 single-node cluster
Single-node VCS clusters with VCS 5.1 SP1 RP1 and later or VCS 6.0 and later that hosts CP server does not require LLT and GAB to be configured.
- SFHA 6.0, 5.1SP1, or 5.1 cluster

Warning: Before you upgrade 5.1 CP server nodes to use VCS or SFHA 6.0, you must upgrade all the application clusters that use this CP server to version 6.0. Application clusters at version 5.1 cannot communicate with CP server that runs VCS or SFHA 5.1 SP1 or later.

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

See the *Veritas Cluster Server Installation Guide* or the *Veritas Storage Foundation High Availability Installation Guide*.

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 3-1 lists additional requirements for hosting the CP server.

Table 3-1 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: <ul style="list-style-type: none">■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB)■ 300 MB in /usr■ 20 MB in /var■ 10 MB in /etc (for the CP server database)
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the CP servers.
RAM	Each CP server requires at least 512 MB.
Network	Network hardware capable of providing TCP/IP connection between CP servers and SFCFSHA clusters (application clusters).

Table 3-2 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 3-2 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	<p>CP server supports any of the following operating systems:</p> <ul style="list-style-type: none"> ■ AIX 6.1 and 7.1 ■ HP-UX 11i v3 ■ Linux: <ul style="list-style-type: none"> ■ RHEL 5 ■ RHEL 6 ■ SLES 10 ■ SLES 11 ■ Solaris 10 <p>Review other details such as supported operating system levels and architecture for the supported operating systems.</p> <p>See the <i>Veritas Cluster Server Release Notes</i> or the <i>Veritas Storage Foundation High Availability Release Notes</i> for that platform.</p>

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 14250. This is the default port that can be changed during a CP server configuration. 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 one of the other available virtual IP addresses.
- The CP server supports either Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with the application clusters. If the CP server is configured to use an IPv6 virtual IP address, then the application clusters should also be on the IPv6 network where the CP server is being hosted.
- 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 the number of hops between the nodes.

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

	CP server in secure mode	CP server in non-secure mode
SFCFSHA cluster in secure mode	Yes	Yes
SFCFSHA cluster in non-secure mode	Yes	Yes
CP server cluster in secure mode	Yes	No
CP server cluster in non-secure mode	No	Yes

For secure communications between the SFCFSHA cluster and CP server, 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 *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

Non-SCSI-3 I/O fencing requirements

Supported virtual environment for non-SCSI-3 fencing:

- Solaris 10 Update 7 and later
Oracle VM Server for SPARC 2.0 and 2.1
Guest operating system: Solaris 10

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

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

Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x6000. When you install the Veritas File System package, `VRTSvxfs`, the `VRTSvxfs` packaging scripts check the values of these variables in the kernel. If the values are less than the required values, `VRTSvxfs` increases the values and modifies the `/etc/system` file with the required values. If the `VRTSvxfs` scripts increase the values, the installation proceeds as usual except that you must reboot and restart the installation program. A message displays if a reboot is required.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System. Use the following commands to check the values of the variables:

```
# echo "lwp_default_stksize/X" | mdb -k
lwp_default_stksize:
lwp_default_stksize:          6000

# echo "svc_default_stksize/X" | mdb -k
svc_default_stksize:
svc_default_stksize:          6000
```

If the values shown are less than 6000, you can expect a reboot after installation.

Note: The default value of the `svc_default_stksize` variable is 0 (zero), which indicates that the value is set to the value of the `lwp_default_stksize` variable. In this case, no reboot is required, unless the value of the `lwp_default_stksize` variable is too small.

To avoid a reboot after installation, you can modify the `/etc/system` file with the appropriate values. Reboot the system prior to installing the packages. Add the following lines to the `/etc/system` file:

```
set lwp_default_stksize=0x6000
set rpcmod:svc_default_stksize=0x6000
```

Storage Foundation Cluster File System High Availability memory requirements

2 GB of memory is required.

Storage Foundation Cluster File System High Availability CPU requirements

A minimum of 2 CPUs is required.

Veritas Storage Foundation Cluster File System High Availability node requirements

Veritas Storage Foundation Cluster File System High Availability supports mixed cluster environments with Solaris 10 SPARC operating systems as long as all the nodes in the cluster have the same CPU architecture.

Database requirements

The following TechNote identifies the most current information on supported database and operating system combinations:

<http://www.symantec.com/docs/DOC4039>

Note: SFCFSHA supports running Oracle, DB2, and Sybase on VxFS and VxVM. SFCFSHA does not support running SFDB tools with DB2 and Sybase.

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 or the `-precheck` option of the script-based installer to determine whether there is sufficient space.

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

If you have downloaded SFCFSHA, you must use the following command:

```
# ./installsfcfsha -precheck
```

Discovering product versions and various requirement information

Symantec provides several methods to check the Veritas product you have installed, plus various requirement information.

You can check the existing product versions using the `installer` command with the `-version` option before or after you install. After you have installed the current version of the product, you can use the `showversion` script in the `/opt/VRTS/install` directory to find version information.

Information the `version` option or the `showversion` script discovers on systems includes the following:

- The installed version of all released Storage Foundation and High Availability Suite of products
- The required packages or patches (if applicable) that are missing
- The available updates (including patches or hotfixes) from Symantec Operations Readiness Tools (SORT) for the installed products

To run the version checker

- 1 Mount the media.
- 2 Start the installer with the `-version` option.

```
# ./installer -version system1 system2
```

Number of nodes supported

SFCFSHA supports cluster configurations with up to 64 nodes.

Licensing Veritas products

This chapter includes the following topics:

- [About Veritas SFHA Solutions product licensing](#)
- [Setting or changing the Veritas SFHA Solutions product level for keyless licensing](#)
- [Installing Veritas SFHA Solutions product license keys](#)

About Veritas SFHA Solutions product licensing

You have the option to install Veritas 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.

www.symantec.com/techsupp/

During the installation, you can choose to either:

- Install a license key for the product and features that you want to install.
When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
- Continue to install without a license key.
The installer prompts for the product modes and options that you want to install, and then sets the required product level.
Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server, such as Veritas Operations Manager (VOM). 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 encounter problems while licensing this product, visit the Symantec licensing support website.

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

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

See [About Veritas Storage Foundation and High Availability Solutions 6.0](#)

- 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 Veritas SFHA Solutions product level for keyless licensing”](#) on page 50.

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 Veritas SFHA Solutions product license keys”](#) on page 52.

See the `vxlicinst(1m)` manual page.

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

Note: In order to change from one product group to another, you may need to perform additional steps.

Setting or changing the Veritas SFHA Solutions product level for keyless licensing

The keyless licensing method uses product levels to determine the Veritas products and functionality that are licensed. In order to use keyless licensing, you must set up a Management Server to manage your systems.

For more information and to download the management server, see the following URL:

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

When you set the product license level for the first time, you enable keyless licensing for that system. If you install with the product installer and select the keyless option, you are prompted to select the product and feature level that you want to license.

After you install, you can change product license levels at any time to reflect the products and functionality that you want to license. When you set a product level, you agree that you have the license for that functionality.

To set or change the product level

- 1 Show your current working directory:

```
# pwd
```

Output resembles:

```
/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 Veritas products until you install a new key or set a new product level.

To clear the product license level

- 1 View the current setting for the product license level.

```
# ./vxkeyless [-v] display
```

- 2 If there are keyless licenses installed, remove all keyless licenses:

```
# ./vxkeyless [-q] set NONE
```

For more details on using the `vxkeyless` utility, see the `vxkeyless (1m)` manual page.

Installing Veritas SFHA Solutions 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:

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

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

To install a new license

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

```
# cd /opt/VRTS/bin
```

```
# ./vxlicinst -k xxxx-xxxx-xxxx-xxxx-xxxx-xxx
```

Installation of Storage Foundation Cluster File System High Availability

- [Chapter 5. Installing Storage Foundation Cluster File System High Availability using the script-based installer](#)
- [Chapter 6. Installing Storage Foundation Cluster File System High Availability using the web-based installer](#)
- [Chapter 7. Installing Storage Foundation Cluster File System High Availability using other methods](#)

Installing Storage Foundation Cluster File System High Availability using the script-based installer

This chapter includes the following topics:

- [Installation preparation overview](#)
- [Mounting the product disc](#)
- [About the Veritas installer](#)
- [Installing Storage Foundation Cluster File System High Availability using the product installer](#)
- [Installing language packages](#)

Installation preparation overview

[Table 5-1](#) provides an overview of an installation using the product installer.

Table 5-1 Installation overview

Installation task	Section
Obtain product licenses.	See “About Veritas SFHA Solutions product licensing” on page 49.
Download the software, or insert the product DVD.	See “Downloading the Veritas Storage Foundation Cluster File System High Availability software” on page 30. See “Mounting the product disc” on page 55.
Set environment variables.	See “Setting environment variables” on page 31.
Create the <code>/opt</code> directory, if it does not exist.	See “Creating the /opt directory” on page 33.
Configure the secure shell (ssh) on all nodes.	See “About configuring secure shell or remote shell communication modes before installing products” on page 446.
Verify that hardware, software, and operating system requirements are met.	See “Release notes” on page 40.
Check that sufficient disk space is available.	See “Disk space requirements” on page 47.
Use the installer to install the products.	See “About the Veritas installer” on page 56.

Mounting the product disc

Mounting the product disc

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

To mount the product disc

- 1 Log in as superuser on a system where you want to install SFCFSHA.
 The system from which you install SFCFSHA need not 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.

About the Veritas installer

The installer enables you to install and configure the product, verify preinstallation requirements, and view the product's description.

If you obtained a standalone Veritas product from an electronic download site, the single product download files do not contain the general product installer. Use the product installation script to install the product.

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 401.

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 402.

Installing Storage Foundation Cluster File System High Availability using the product installer

The product installer is the recommended method to license and install Storage Foundation Cluster File System High Availability.

The following sample procedure is based on the installation of a Veritas Storage Foundation Cluster File System High Availability cluster with two nodes: "host1" and "host2". If you are installing on standalone systems only, some steps are unnecessary, and these are indicated.

Default responses are enclosed by parentheses. Press Return to accept defaults.

Note: If you have obtained a Veritas product from an electronic download site, the single product download files do not contain the `installer` installation script, so you must use the product installation script to install the product. For example, if you download Veritas Cluster File System High Availability, use the `installsfcfsha` script instead of the `installer` script.

To install Veritas Storage Foundation Cluster File System High Availability

- 1 To install on multiple systems, set up the systems so that commands between systems execute without prompting for passwords or confirmations.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 446.

- 2 Load and mount the software disc.

- 3 Move to the top-level directory on the disc.

```
# cd /dvd_mount
```

- 4 From this directory, type the following command to install on the local system. Also use this command to install on remote systems provided that the secure shell or remote shell utilities are configured:

```
# ./installer
```

- 5 Enter `I` to install and press Return.

- 6 From the Installation menu, choose the `I` option for Install and enter the number for Storage Foundation Cluster File System High Availability. Press Return.

- 7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement as
specified in the
storage_foundation_cluster_file_system_ha/EULA/lang/EULA_CFSHA_Ux_6.0.pdf
file present
```

```
on the media? [y,n,q,?] y
```

- 8 Select from one of the following install options:
- Minimal packages: installs only the basic functionality for the selected product.
 - Recommended packages: installs the full feature set without optional packages.
 - All packages: installs all available packages.

Each option displays the disk space that is required for installation. Select which option you want to install and press Return.

- 9 You are prompted to enter the system names (in the following example, "host1" and "host2") on which the software is to be installed. Enter the system name or names and then press Return.

```
Enter the system names separated by spaces:[q?] (host1 host2)
```

- 10 During the initial system check, the installer verifies that communication between systems has been set up. The installer prompts you to allow it to set up ssh or rsh. After the installation, the installer cleans up the ssh or rsh as needed.
- 11 After the system checks complete, the installer displays a list of the packages that will be installed. Press Enter to continue with the installation.

12 You are prompted to choose your licensing method.

To comply with the terms of Symantec's End User License Agreement, you have 60 days to either:

- * Enter a valid license key matching the functionality in use on the systems
- * Enable keyless licensing and manage the systems with a Management Server.

For more details visit <http://go.symantec.com/sfhakeyless>. The product is fully functional during these 60 days.

- 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) 2

If you have a valid license key, select 1 and enter the license key at the prompt. Skip to step 16.

To install using keyless licensing, select 2. You are prompted for the product modes and the options that you want to install and license.

Note: The keyless license option enables you to install without entering a key. However, you must still have a valid license to install and use Veritas products.

Keyless licensing requires that you manage the systems with a Management Server. Refer to the following URL for details:

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

13 Select **yes** to enable replication.

14 Select **yes** to enable the Global Cluster Option.

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

- 15 At the prompt, specify whether you want to send your installation information to Symantec.

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

- 16 The product installation completes.

Review the output and summary files. Reboot nodes as requested. Run the following command to configure SFCFSHA.

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

See the Section on "Configuration of Veritas Storage Foundation Cluster File System High Availability" for more information.

Installing language packages

To install SFCFSHA in a language other than English, install the required language packages after installing the English packages.

To install the language packages on the server

- 1 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With Solaris volume management software, the disc is automatically mounted as

```
/cdrom/cdrom0.
```

- 2 Install the language packages using the `install_lp` command.

```
# cd /cdrom/cdrom0
# ./install_lp
```

Installing Storage Foundation Cluster File System High Availability using the web-based installer

This chapter includes the following topics:

- [About the Web-based installer](#)
- [Before using the Veritas Web-based installer](#)
- [Starting the Veritas Web-based installer](#)
- [Obtaining a security exception on Mozilla Firefox](#)
- [Performing a pre-installation check with the Veritas Web-based installer](#)
- [Installing SFCFSHA with the Web-based installer](#)

About the Web-based installer

Use the Web-based installer interface to install Veritas 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 `xprt1wid` process. The `webinstaller` script can also be used to check the status of the XPortal Server.

When the `webinstaller` script starts the `xprt1wid` process, the script displays a URL. Use this URL to access the Web-based installer from a Web browser such as Internet Explorer or FireFox.

The Web installer creates log files whenever the Web installer is operating. While the installation processes are operating, the log files are located in a session-based directory under the `/var/tmp` directory. After the install process completes, the log files are located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep these files for auditing, debugging, and future use.

The location of the Veritas XPortal Server configuration file is `/var/opt/webinstaller/xprt1wid.conf`.

See [“Before using the Veritas Web-based installer”](#) on page 62.

See [“Starting the Veritas Web-based installer”](#) on page 63.

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 6-1 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Veritas products.	Must be a supported platform for Veritas Storage Foundation Cluster File System High Availability 6.0.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must use the same operating system as the target systems and 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: <ul style="list-style-type: none"> ■ Internet Explorer 6, 7, and 8 ■ Firefox 3.x and later

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprt1wid`, on the installation server:

```
# ./webinstaller start
```

The webinstaller script displays a URL. Note this URL.

Note: If you do not see the URL, run the command again.

The default listening port is 14172. If you have a firewall that blocks port 14172, use the `-port` option to use a free port instead.

- 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 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 **Add Exception** button.
- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.

- 5 Click **Confirm Security Exception** button.
- 6 Enter root in User Name field and root password of the web server in the Password field.

Performing a pre-installation check with the Veritas Web-based installer

This section describes performing a pre-installation check with the Veritas Web-based installer.

To perform a pre-installation check

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 63.
- 2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.
- 3 Select the Veritas Storage Foundation Cluster File System High Availability 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 Click **Finish**. The installer prompts you for another task.

Installing SFCFSHA with the Web-based installer

This section describes installing SFCFSHA with the Veritas Web-based installer.

To install SFCFSHA using the Web-based installer

- 1 Perform preliminary steps.
See [“Performing a pre-installation check with the Veritas Web-based installer”](#) on page 64.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 63.
- 3 Select **Install a Product** from the **Task** drop-down list.

- 4 Select **Veritas Storage Foundation Cluster File System HA** 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 SFCFSHA on the selected system.
- 10 After the installation completes, you must choose your licensing method.

On the license page, select one of the following tabs:

- Keyless licensing

Note: The keyless license option enables you to install without entering a key. However, in order to ensure compliance you must manage the systems with a management server.

For more information, go to the following website:

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

Complete the following information:

- Choose whether you want to enable Veritas Replicator.
- Choose whether you want to enable Global Cluster option.

Click **Register**.

- Enter license key

If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

11 The product installation completes.

Review the output. Reboot nodes as requested. The installer may prompt you to perform other tasks.

12 If prompted, select the checkbox to specify whether you want to send your installation information to Symantec.

Would you like to send the information about this installation to Symantec to help improve installation in the future?

Click **Finish**. The installer prompts you for another task.

Installing Storage Foundation Cluster File System High Availability using other methods

This chapter includes the following topics:

- [Installing using JumpStart](#)
- [Installing SFCFSHA using the pkgadd command](#)

Installing using JumpStart

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 stand-alone 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 SFCFSHA and the operating system in conjunction with JumpStart.

See [“Using a Flash archive to install SFCFSHA and the operating system”](#) on page 72.

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 68.
- 4 Prepare shared storage installation resources.
See [“Preparing installation resources”](#) on page 70.
- 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 Veritas software.

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

Generating the finish scripts

Perform these steps to generate the finish scripts to install SFCFSHA.

To generate the script

- 1 Run the product installer program to generate the scripts for all products.

```
./installer -jumpstart directory_to_generate_scripts
```

Or

```
./installprod -jumpstart directory_to_generate_script
```

Where *prod* is the product's installation command, and *directory_to_generate_scripts* is where you want to put the product's script.

For example:

```
# ./installsfcfs -jumpstart /js_scripts
```

- 2 When you are prompted to encapsulate the root disk automatically, choose **yes** to do so. If you do not want to encapsulate it automatically, choose **no** and go to step 6.

- 3 Specify a disk group name for the root disk.

Specify the disk group name of the root disk to be encapsulated:
rootdg

- 4 Specify private region length.

Specify the private region length of the root disk to be encapsulated: **(65536)**

- 5 Specify the disk's media name of the root disk to encapsulate.

Specify the disk media name of the root disk to be encapsulated:
(**rootdg_01**)

- 6 JumpStart finish scripts and encapsulation scripts are generated in the directory you specified in step 1.

Output resembles:

```
The finish scripts for VM is generated at /js_scripts/  
jumpstart_sfcfsha.fin  
The encapsulation boot disk script for SFCFSHA is generated at  
/js_scripts/encap_bootdisk_vm.fin
```

List the `js_scripts` directory.

```
# ls /js_scripts
```

Output resembles:

```
encap_bootdisk_vm.fin jumpstart_sfcfsha.fin
```

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
```

- 4 If you want to encapsulate the root disk automatically when you perform the JumpStart installation, copy the scripts `encap_bootdisk_vm.fin` generated previously to *ENCAPSRC*.

See [“Generating the finish scripts”](#) on page 68.

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/pkg  
# cp -r * BUILDSRC/pkg
```

If you downloaded the language pack:

```
# cd /path_to_language_pack_installation_media/pkg  
# cp -r * BUILDSRC/pkg
```

- 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 SFCFSHA 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 Veritas 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. Reboot the cloned system.
4. Install the Veritas 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 Veritas post-deployment scripts”](#) on page 73.
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 Veritas product on all nodes in the cluster. Start configuration with the following command:

```
# /opt/VRTS/install/installsfcfsha -configure
```
10. Perform post-installation and configuration tasks.

Creating the Veritas 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 Veritas 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_postdeployment.sh` file and the `vrts_postdeployment.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 Veritas products.

Installing SFCFSHA using the pkgadd command

On Solaris 10, the packages must be installed while in the global zone.

To install SFCFSHA using the pkgadd command

- 1 Mount the software disc.
See [“Mounting the product disc”](#) on page 55.
- 2 Copy the supplied VRTS* files from the installation media to a temporary location. Modify them if needed.

```
# cp /cdrom/cdrom0/pkgs/VRTS* \  
    /tmp/pkgs
```

- 3 Create the admin file in the current directory. Specify the `-a adminfile` option when you use the `pkgadd` command:

```
mail=  
instance=overwrite  
partial=nocheck  
  runlevel=quit  
idepend=quit  
rdepend=nocheck  
space=quit  
setuid=nocheck  
conflict=nocheck  
action=nocheck  
basedir=default
```

- 4 Use the product-specific install command with one of the following options to get a list of packages in the order to be installed:

- minpkgs
- recpkgs
- allpkgs

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 401.

See [“Installation script options”](#) on page 402.

- 5 Install the packages listed in step 4.

```
# pkgadd -a adminfile -d /tmp/pkgs pkgname.pkg
```

On Solaris 10, these packages must be installed while in the global zone. If a package's `pkginfo` file contains the variable `SUNW_PKG_ALLZONES` set not equal to `true`, the `-G` option should additionally be specified to the `pkgadd` command.

- 6 Use the product-specific install command with one of the following options to get a list of packages in the order to be installed:

- minpkgs
- recpkgs
- allpkgs

7 Verify that the packages are installed:

```
# pkginfo -l  
   packagename
```

8 Start the processes.

Configuration of Veritas Storage Foundation Cluster File System High Availability

- [Chapter 8. Preparing to configure SFCFSHA](#)
- [Chapter 9. Configuring Veritas Storage Foundation Cluster File System High Availability](#)
- [Chapter 10. Configuring SFCFSHA for data integrity](#)

Preparing to configure SFCFSHA

This chapter includes the following topics:

- [About configuring SFCFSHA clusters for data integrity](#)
- [About I/O fencing for SFCFSHA in virtual machines that do not support SCSI-3 PR](#)
- [About I/O fencing components](#)
- [About I/O fencing configuration files](#)
- [About planning to configure I/O fencing](#)
- [Setting up the CP server](#)

About configuring SFCFSHA clusters for data integrity

When a node fails, SFCFSHA 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. SFCFSHA 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 SFCFSHA, you must configure I/O fencing in SFCFSHA to ensure data integrity.

See [“About planning to configure I/O fencing”](#) on page 84.

About I/O fencing for SFCFSHA 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, Veritas Clustered Volume Manager 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 *Veritas Cluster Server Administrator's Guide* for more information on how I/O fencing works.

In virtualized environments that do not support SCSI-3 PR, SFCFSHA attempts to provide reasonable safety for the data disks. SFCFSHA 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 installsfcfsha”](#) on page 157.

See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 174.

About I/O fencing components

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

- Data disks—Store shared data
See [“About data disks”](#) on page 80.
- Coordination points—Act as a global lock during membership changes
See [“About coordination points”](#) on page 80.

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

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. SFCFSHA 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 SFCFSHA configuration.
You can configure coordinator disks to use Veritas Volume Manager 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.

See the *Veritas 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 SFCFSHA cluster nodes to perform the following tasks:

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

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

- 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.

- Disable preferred fencing policy to use the default node count-based race policy.

See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide* for more details.

See [“Enabling or disabling the preferred fencing policy”](#) on page 179.

About I/O fencing configuration files

[Table 8-1](#) lists the I/O fencing configuration files.

Table 8-1 I/O fencing configuration files

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

Table 8-1 I/O fencing configuration files (*continued*)

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

Table 8-1 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	<p>When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.</p> <p>Note: The /etc/vxfentab file is a generated file; do not modify this file.</p> <p>For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:</p> <ul style="list-style-type: none"> ■ Raw disk: <ul style="list-style-type: none"> /dev/rdisk/c1t1d0s2 /dev/rdisk/c2t1d0s2 /dev/rdisk/c3t1d2s2 ■ DMP disk: <ul style="list-style-type: none"> /dev/vx/rdmp/c1t1d0s2 /dev/vx/rdmp/c2t1d0s2 /dev/vx/rdmp/c3t1d0s2 <p>For server-based fencing, the /etc/vxfentab file also includes the security settings information.</p> <p>For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.</p>

About planning to configure I/O fencing

After you configure SFCFSHA with the installer, you must configure I/O fencing in the cluster for data integrity.

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

See [Figure 8-2](#) on page 87.

[Figure 8-1](#) illustrates a high-level flowchart to configure I/O fencing for the SFCFSHA cluster.

Figure 8-1 Workflow to configure I/O fencing

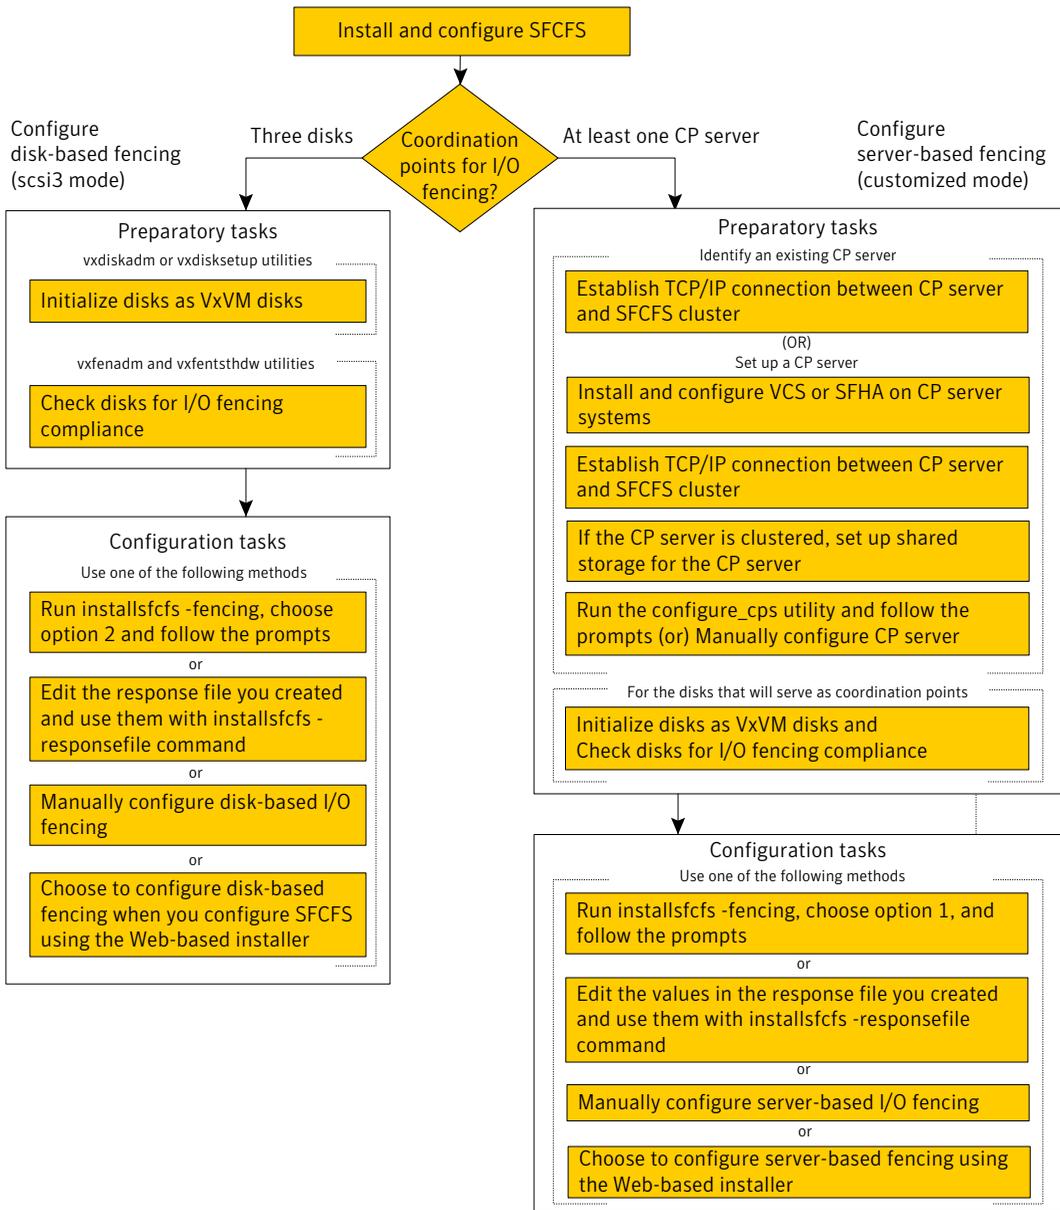
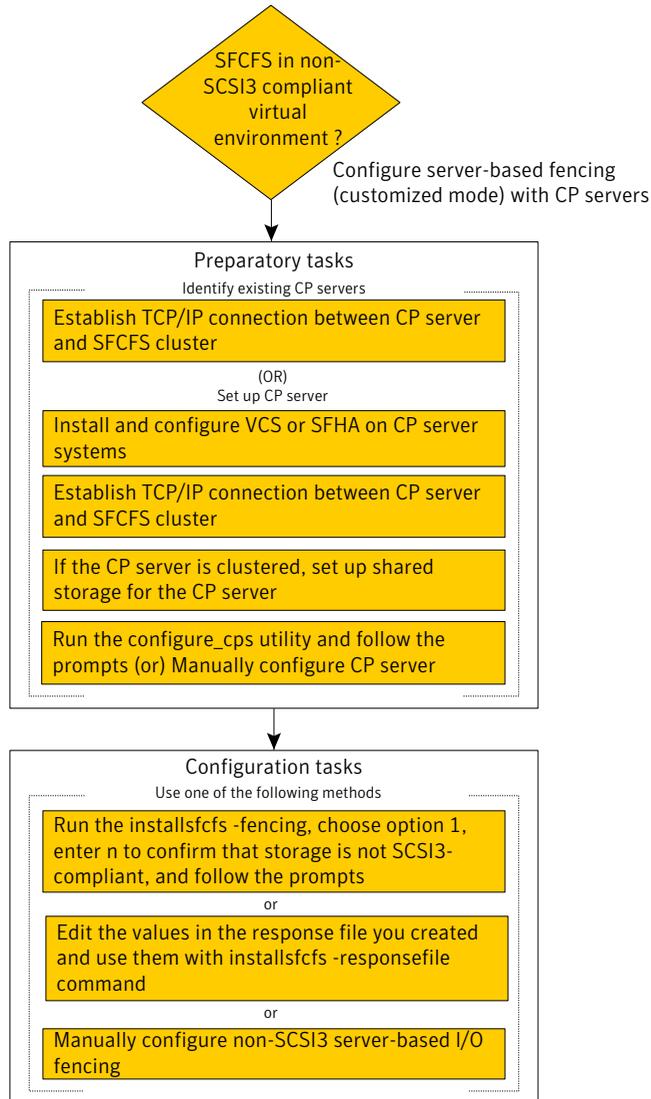


Figure 8-2 illustrates a high-level flowchart to configure non-SCSI-3 server-based I/O fencing for the SFCFSHA cluster in virtual environments that do not support SCSI-3 PR.

Figure 8-2 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 `installsfcfs` See [“Setting up disk-based I/O fencing using `installsfcfs`”](#) on page 144.
 See [“Setting up server-based I/O fencing using `installsfcfs`”](#) on page 157.
 See [“Setting up non-SCSI-3 server-based I/O fencing in virtual environments using `installsfcfs`”](#) on page 157.
- Using the Web-based installer See [“Configuring SFCFSHA for data integrity using the Web-based installer”](#) on page 159.
- Using response files See [“Response file variables to configure disk-based I/O fencing”](#) on page 439.
 See [“Response file variables to configure server-based I/O fencing”](#) on page 441.
 See [“Response file variables to configure non-SCSI-3 server-based I/O fencing”](#) on page 443.
 See [“Configuring I/O fencing using response files”](#) on page 438.
- Manually editing configuration files See [“Setting up disk-based I/O fencing manually”](#) on page 152.
 See [“Setting up server-based I/O fencing manually”](#) on page 162.
 See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 174.

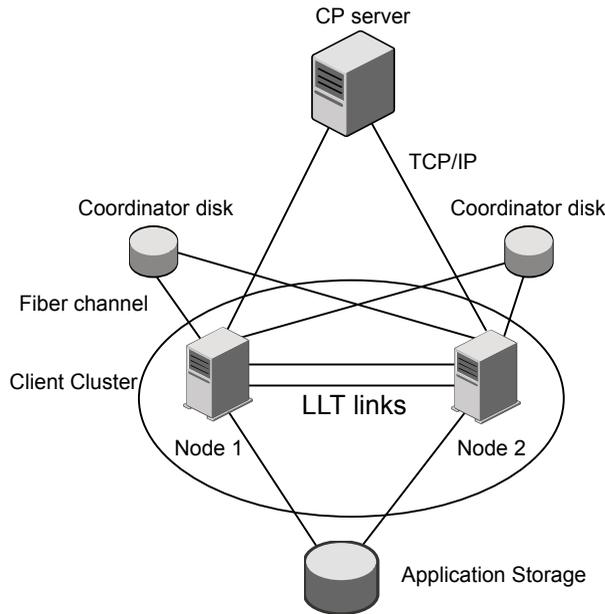
You can also migrate from one I/O fencing configuration to another.

See the *Veritas Storage Foundation Cluster File System High Availability Administrator’s Guide* for more details.

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

[Figure 8-3](#) displays a configuration using a SFCFSHA cluster (with two nodes), a single CP server, and two coordinator disks. The nodes within the SFCFSHA cluster are connected to and communicate with each other using LLT links.

Figure 8-3 CP server, SFCFSHA 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 8-4](#) on page 90.
- Multiple application clusters use a single CP server and multiple pairs of coordinator disks (two) as their coordination points
 See [Figure 8-5](#) on page 91.
- 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 8-6](#) on page 91.

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 CP servers 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 8-4 displays a configuration using three CP servers that are connected to multiple application clusters.

Figure 8-4 Three CP servers connecting to multiple application clusters

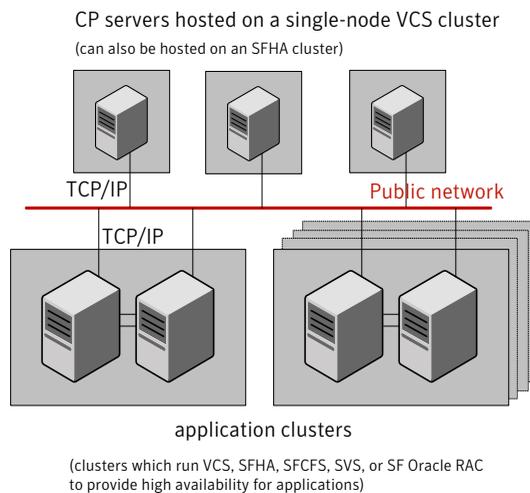


Figure 8-5 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.

Figure 8-5 Single CP server with two coordinator disks for each application cluster

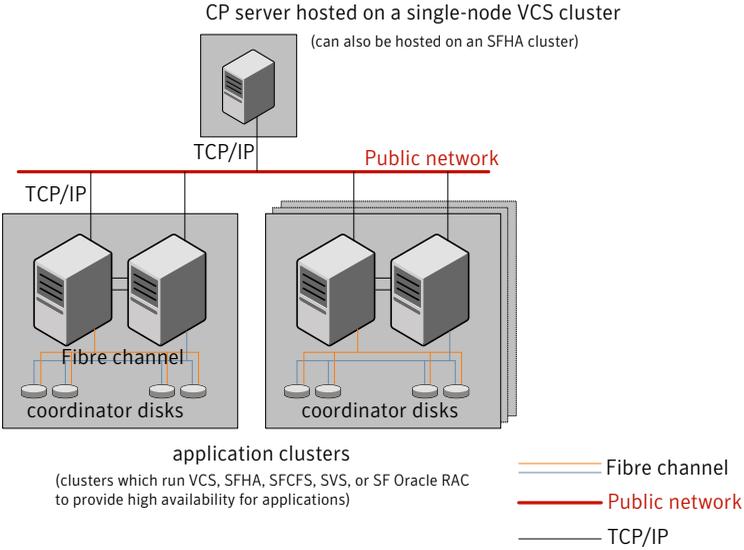
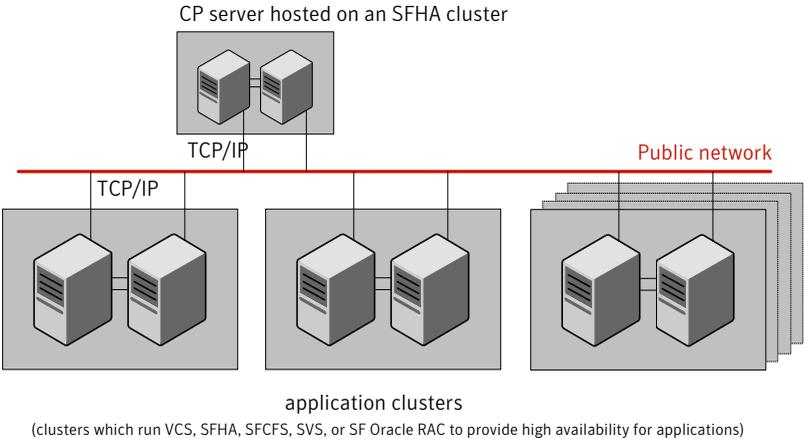


Figure 8-6 displays a configuration using a single CP server that is connected to multiple application clusters.

Figure 8-6 Single CP server connecting to multiple application clusters



See "Configuration diagrams for setting up server-based I/O fencing" on page 495.

Setting up the CP server

Table 8-2 lists the tasks to set up the CP server for server-based I/O fencing.

Table 8-2 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 92.
Install the CP server	See “Installing the CP server using the installer” on page 93.
Configure the CP server cluster in secure mode	See “Configuring the CP server cluster in secure mode” on page 94.
Set up shared storage for the CP server database	See “Setting up shared storage for the CP server database” on page 95.
Configure the CP server	See “Configuring the CP server using the configuration utility” on page 96. See “Configuring the CP server manually” on page 105.
Verify the CP server configuration	See “Verifying the CP server configuration” on page 106.

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 configure disk-based fencing during the SFHA configuration.
 - You must set up shared storage for the CP server database during your CP server setup.

CP server setup uses a single system	<p>Install and configure VCS to create a single-node VCS cluster.</p> <p>During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.</p> <p>See the <i>Veritas Cluster Server Installation Guide</i> for instructions on installing and configuring VCS.</p> <p>Proceed to configure the CP server.</p> <p>See “Configuring the CP server using the configuration utility” on page 96.</p> <p>See “Configuring the CP server manually” on page 105.</p>
CP server setup uses multiple systems	<p>Install and configure SFHA to create an SFHA cluster. This makes the CP server highly available.</p> <p>Meet the following requirements for CP server:</p> <ul style="list-style-type: none"> ■ During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages. ■ During configuration, configure disk-based fencing (scsi3 mode). <p>See the <i>Veritas Storage Foundation and High Availability Installation Guide</i> for instructions on installing and configuring SFHA.</p> <p>Proceed to set up shared storage for the CP server database.</p>

Configuring the CP server cluster in secure mode

You must configure security on the CP server only if you want to secure the communication between the CP server and the SFCFSHA cluster (CP client).

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:

```
# installvcs -security
```

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

```
# installsfha -security
```

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.

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 Import the disk group if it is not already imported.

For example:

```
# vxdg import cps_dg
```

- 3 Create a mirrored volume over the disk group.

For example:

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

- 4 Create a file system over the volume.

The CP server configuration utility only supports vxfs file system type. If you use an alternate file system, then 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/rdisk/cps_dg/cps_volume
```

```
HP-UX # mkfs -F vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

```
Linux # mkfs -t vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

```
Solaris # mkfs -F vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

Configuring the CP server using the configuration utility

The CP server configuration utility (`configure_cps.pl`) is part of the VRTScps package.

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 96.

For CP servers on an SFHA cluster: See [“To configure the CP server on an SFHA cluster”](#) on page 100.

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 CP server configuration script on the node where you want to configure the CP server:

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

- 3 Enter **1** at the prompt to configure CP server on a single-node VCS cluster.

The configuration utility 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.
- Checks to see if the CP server is already configured on the system.
If the CP server is already configured, then the configuration utility informs the user and requests that the user unconfigure the CP server before trying to configure it.

- 4 Enter the name of the CP server.

```
Enter the name of the CP Server: mycps1
```

- 5 Enter valid virtual IP addresses on which the CP server process should depend on:

- Enter the number of virtual IP addresses you want to configure:

```
Enter the number of virtual IP(s) to configure : 2
```

- Enter valid virtual IP addresses:

```
Enter a valid IP address for Virtual IP - 1 which the CP Server  
process should depend on : 10.209.83.85  
Enter a valid IP address for Virtual IP - 2 which the CP Server  
process should depend on : 10.209.83.87
```

You can also use IPv6 address.

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

```
Enter a port number for virtual IP 10.209.83.85 in range [49152,  
65535], or press enter for default port (14250) :
```

```
Using default port: 14250
```

```
Enter a port number for virtual IP 10.209.83.87 in range  
[49152, 65535], or press enter for default port (14250) :
```

```
Using default port: 14250
```

7 Choose whether the communication between the CP server and the SFCFSHA clusters has to be made secure.

If you have not configured the CP server cluster in secure mode, enter **n** at the prompt.

Warning: If the CP server cluster is not configured in secure mode, and if you enter **y**, then the script immediately exits. You must configure the CP server cluster in secure mode and rerun the CP server configuration script.

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

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

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

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

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

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

- 9** Verify and confirm the CP server configuration information.

Following is the CP Server configuration information:

```
-----
(a)CP Server Name: mycps1
(b)CP Server Virtual IP(s): 10.209.83.85 10.209.83.87
(c)CP Server Port(s): 14250 14250
(d)CP Server Security : 1
(e)CP Server Database Dir: /etc/VRTScps/db
-----
```

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

- 10** The configuration utility 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.

Configuring CP Server Service Group (CPSSG) for this cluster

- 11** Enter the number of NIC resources that you want to configure. You must use a public NIC.

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.

- 12** Enter a valid network interface for the virtual IP address for the CP server process.

```
Enter a valid network interface for virtual IP 10.209.83.85
on mycps1.symantecexample.com: bge0
Enter a valid network interface for virtual IP 10.209.83.87
on mycps1.symantecexample.com: bge0
```

- 13** 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 10.209.83.85 [1 to 2] : 1
Enter the NIC resource you want to associate with the
virtual IP 10.209.83.87 [1 to 2] : 2
```

- 14** Enter networkhosts information for each NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure
NIC resource to be online always.
Do you want to add NetworkHosts attribute for the NIC device
bge0 on system mycps1? [y/n] : y
Enter a valid IP address to configure NetworkHosts for
NIC bge0 on system mycps1 : 10.209.83.86
Do you want to add another Network Host ?[y/n] : n
```

- 15** Enter the netmask for each virtual IP address. For example:

```
Enter the netmask for virtual IP 10.209.83.85 :
255.255.252.0
Enter the netmask for virtual IP 10.209.83.87 :
255.255.252.0
```

If you entered an IPv6 address, enter the prefix details at the prompt.

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

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

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

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

For example:

```
# hagrps -state CPSSG

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

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

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

For information about the CPSSG, refer to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

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

To configure the CP server on an SFHA cluster

- 1 Verify that the `VRTScps` package is installed on each node.
- 2 Make sure that you have configured passwordless `ssh` or `rsh` on the CP server cluster nodes.
- 3 Run the CP server configuration script on any node in the cluster:

```
# /opt/VRTScps/bin/configure_cps.pl [-n]
```

The CP server configuration utility uses `ssh` by default to communicate between systems. Use the `-n` option for `rsh` communication.

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

The configuration utility 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.

- Checks to see if the CP server is already configured on the system. If the CP server is already configured, then the configuration utility informs the user and requests that the user unconfigure the CP server before trying to configure it.

5 Enter the name of the CP server.

Enter the name of the CP Server: **mycps1**

6 Enter valid virtual IP addresses on which the CP server process should depend on:

- Enter the number of virtual IP addresses you want to configure:

Enter the number of virtual IP(s) to configure : 2

- Enter valid virtual IP addresses:

Enter a valid IP address for Virtual IP - 1 which the CP Server process should depend on : 10.209.83.85

Enter a valid IP address for Virtual IP - 2 which the CP Server process should depend on : 10.209.83.87

You can also use IPv6 address.

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

Enter a port number for virtual IP 10.209.83.85 in range [49152, 65535], or press enter for default port (14250) :

Using default port: 14250

Enter a port number for virtual IP 10.209.83.87 in range [49152, 65535], or press enter for default port (14250) :

Using default port: 14250

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

If you have not configured the CP server cluster in secure mode, enter **n** at the prompt.

Warning: If the CP server cluster is not configured in secure mode, and if you enter **y**, then the script immediately exits. You must configure the CP server cluster in secure mode and rerun the CP server configuration script.

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

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

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

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

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

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

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

- 10 Verify and confirm the CP server configuration information.

Following is the CP Server configuration information:

```

-----
(a)CP Server Name: mycps1
(b)CP Server Virtual IP(s): 10.209.83.85 10.209.83.87
(c)CP Server Port(s): 14250 14250
(d)CP Server Security : 1
(e)CP Server Database Dir: /etc/VRTScps/db
-----
  
```

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

- 11** The configuration utility proceeds with the configuration process, and creates a `vxcps.conf` configuration file on each node.

The following output is for one node:

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

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

- 12** Enter the number of NIC resources that you want to configure. You must use a public NIC.

```
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.

- 13** Confirm whether you use the same NIC name for the virtual IP on all the systems in the cluster.

```
Is the name of network interfaces for NIC resource - 1
same on all the systems?[y/n] : y
```

- 14** Enter a valid network interface for the virtual IP address for the CP server process.

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

- 15** 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 10.209.83.85 [1 to 2] : 1
Enter the NIC resource you want to associate with the
virtual IP 10.209.83.87 [1 to 2] : 2
```

16 Enter networkhosts information for each NIC resource.

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

Do you want to add NetworkHosts attribute for the NIC device bge0 on system mycps1? [y/n] : y

Enter a valid IP address to configure NetworkHosts for NIC bge0 on system mycps1 : 10.209.83.86

Do you want to add another Network Host ?[y/n] : n

17 Enter the netmask for each virtual IP address.

Enter the netmask for virtual IP 10.209.83.85 :
255.255.252.0

If you entered an IPv6 address, enter the prefix details at the prompt.

18 Enter the name of the disk group for the CP server database.

Enter the name of diskgroup for cps database :
cps_dg

19 Enter the name of the volume that is created on the above disk group.

Enter the name of volume created on diskgroup cps_dg :
cps_volume

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

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

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

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

For example:

```
# hagrp -state CPSSG
```

#Group	Attribute	System	Value
CPSSG	State	mycps1	ONLINE
CPSSG	State	mycps2	OFFLINE

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

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

For information about the CPSSG, refer to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

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

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

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 `main.cf` as an example:

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 have configured the CP server cluster in secure mode or not, do 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`.

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

- 5 Start VCS on all the cluster nodes.

```
# hstart
```

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

```
# hagrps -state CPSSG
```

Output similar to the following appears:

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

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)

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

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

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

Configuring Veritas Storage Foundation Cluster File System High Availability

This chapter includes the following topics:

- [Configuring SFCFSHA using the script-based installer](#)
- [Configuring SFCFSHA using the Web-based installer](#)
- [Database configuration requirements](#)
- [Configuring Veritas Storage Foundation Cluster File System High Availability manually](#)
- [Configuring the SFDB repository database after installation](#)

Configuring SFCFSHA using the script-based installer

Overview of tasks to configure SFCFSHA using the script-based installer

[Table 9-1](#) lists the tasks that are involved in configuring SFCFSHA using the script-based installer.

Table 9-1 Tasks to configure SFCFSHA using the script-based installer

Task	Reference
Start the software configuration	See “Starting the software configuration” on page 109.
Specify the systems where you want to configure SFCFSHA	See “Specifying systems for configuration” on page 110.
Configure the basic cluster	See “Configuring the cluster name” on page 111. See “Configuring private heartbeat links” on page 111.
Configure virtual IP address of the cluster (optional)	See “Configuring the virtual IP of the cluster” on page 114.
Configure the cluster in secure mode (optional)	See “Configuring the cluster in secure mode” on page 116.
Add VCS users (required if you did not configure the cluster in secure mode)	See “Adding VCS users” on page 120.
Configure SMTP email notification (optional)	See “Configuring SMTP email notification” on page 121.
Configure SNMP email notification (optional)	See “Configuring SNMP trap notification” on page 123.
Configure global clusters (optional) Note: You must have enabled Global Cluster Option when you installed SFCFSHA.	See “Configuring global clusters” on page 125.
Complete the software configuration	See “Completing the SFCFSHA configuration” on page 126.

Starting the software configuration

You can configure SFCFSHA using the Veritas product installer or the `installscfsha` command.

Note: If you want to reconfigure SFCFSHA, 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 SFCFSHA 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:

Storage Foundation Cluster File System High Availability

To configure SFCFSHA using the `installsfcfsha` program

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

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

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

Specifying systems for configuration

The installer prompts for the system names on which you want to configure SFCFSHA. 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 SFCFSHA.

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

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

The installer does the following tasks:

- Checks that the local node running the installer can communicate with remote nodes

If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases.

- Makes sure that the systems are running with the supported operating system
 - Makes sure the installer started from the global zone
 - Checks whether SFCFSHA is installed
 - Exits if Veritas Storage Foundation Cluster File System High Availability 6.0 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 84.

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

See [“Using the UDP layer for LLT”](#) on page 517.

The following procedure helps you configure LLT over Ethernet.

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 UDP.
 - Option 1: 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: 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.
- 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 `bge0`.)

```
Enter the NIC for the first private heartbeat link on galaxy:
[b, q, ?] bge0
Would you like to configure a second private heartbeat link?
[y, n, q, b, ?] (y)
Enter the NIC for the second private heartbeat link on galaxy:
[b, q, ?] bge1
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)
```

- For Solaris x64:

You must not enter the network interface card that is used for the public network (typically `bge0`.)

```
Enter the NIC for the first private heartbeat link on galaxy:
[b, q, ?] e1000g1
Would you like to configure a second private heartbeat link?
[y, n, q, b, ?] (y)
```

```
Enter the NIC for the second private heartbeat link on galaxy:
[b,q,?] e1000g2
Would you like to configure a third private 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 galaxy: [b,q,?] private_NIC1
Do you want to use address 192.168.0.1 for the
first private heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the first private heartbeat
link on galaxy: [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 galaxy: [b,q,?] private_NIC2
Do you want to use address 192.168.1.1 for the
second private heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the second private heartbeat
link on galaxy: [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 galaxy: [b,q,?] (private_NIC0)
Do you want to use address 192.168.3.1 for
the low priority heartbeat link on galaxy: [y,n,q,b,?] (y)
Enter the UDP port for the low priority heartbeat
link on galaxy: [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 galaxy, 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.

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

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

To configure the virtual IP of the cluster

- 1 Review the required information to configure the virtual IP of the cluster.
- 2 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 galaxy: bge0
```

```
Enter the NIC for Virtual IP of the Cluster to use on galaxy:  
[b,q,?] (bge0)
```

- 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 bge0 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: bge0
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: bge0
IP: 2001:454e:205a:110:203:baff:feee:10
Prefix: 64
```

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

Configuring the cluster in secure mode

The installer prompts whether you want to configure a secure cluster.

Would you like to configure the SFCFSHA cluster in secure mode?
 [y,n,q,?] (n)

To configure a secure cluster, enter `y`.

If you want to confirm that the configured cluster is in secure mode, verify that the output of the following command is 1.

```
# haclus -value SecureClus
```

```
1
```

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 `-securityonnode` option to configure a secure cluster node by node.

[Table 9-2](#) lists the tasks that you must perform to configure a secure cluster.

Table 9-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See “Configuring the first node” on page 117.
Configure security on the remaining nodes	See “Configuring the remaining nodes” on page 118.
Complete the manual configuration steps	See “Completing the secure cluster configuration” on page 119.

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/installsfcfsha -securityonnode
```

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 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 `/var/VRTSvcs/vcsauth/bkup` directory 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/installsfcfsha -securityonnode
```

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/CmdServer/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 On all nodes, create the `/etc/VRTSvcs/conf/config/.secure` file.

```
# touch /etc/VRTSvcs/conf/config/.secure
```

- 6 On the first node, start VCS. Then start VCS on the remaining nodes.

```
# /opt/VRTSvcs/bin/hastart
```

- 7 On all nodes, start CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer
```

- 8 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 *Veritas Cluster Server Administrator's Guide* for more information.

To configure SMTP email notification

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

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

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

See [“Configuring SNMP trap notification”](#) on page 123.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: bge0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (bge0)
Is bge0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SMTP server's host name.

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

- Enter the email address of each recipient.

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

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

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

- 4 Add more SMTP recipients, if necessary.

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

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

```
Enter the full email address of the SMTP recipient
```

```
(example: user@yourcompany.com): [b,q,?] harriet@example.com
```

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

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

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

5 Verify and confirm the SMTP notification information.

```
NIC: bge0
```

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

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

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

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

Configuring SNMP trap notification

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

Note that you can also configure the notification after installation.

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

To configure the SNMP trap notification

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

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

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

See [“Configuring global clusters”](#) on page 125.

3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: bge0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (bge0)
Is bge0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SNMP trap daemon port.

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

- Enter the SNMP console system name.

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

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

```
Enter the minimum severity of events for which SNMP traps
should be sent to saturn [I=Information, W=Warning, E=Error,
S=SevereError]: [b,q,?] E
```

4 Add more SNMP consoles, if necessary.

- If you want to add another SNMP console, enter `y` and provide the required information at the prompt.

```
Would you like to add another SNMP console? [y,n,q,b] (n) y
Enter the SNMP console system name: [b,q,?] jupiter
Enter the minimum severity of events for which SNMP traps
should be sent to jupiter [I=Information, W=Warning,
E=Error, S=SevereError]: [b,q,?] S
```

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

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

5 Verify and confirm the SNMP notification information.

NIC: bge0

SNMP Port: 162

Console: saturn receives SNMP traps for Error or higher events

Console: jupiter receives SNMP traps for SevereError or higher events

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

Configuring global clusters

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster.

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

See the *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for instructions to set up SFCFSHA global clusters.

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

To configure the global cluster option

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

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

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

3 Provide information to configure this cluster as global cluster.

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

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

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

4 Verify and confirm the configuration of the global cluster. For example:

For IPv4: Global Cluster Option configuration verification:

```
NIC: bge0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

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

On Solaris x64, an example for the NIC's port is bge0.

For IPv6 Global Cluster Option configuration verification:

```
NIC: bge0
IP: 2001:454e:205a:110:203:baff:feee:10
Prefix: 64
```

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

On Solaris x64, an example for the NIC's port is bge0.

Completing the SFCFSHA configuration

After you enter the SFCFSHA configuration information, the installer prompts to stop the SFCFSHA 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 SFCFSHA, it restarts SFCFSHA and its related processes.

To complete the SFCFSHA configuration

- 1 If prompted, press Enter at the following prompt.

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

- 2 Review the output as the installer stops various processes and performs the configuration. The installer then restarts SFCFSHA 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 SFCFSHA 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 SFCFSHA using response files”](#) on page 410.

Verifying and updating licenses on the system

After you install SFCFSHA, 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 127.

See [“Updating product licenses using vxlicinst”](#) on page 128.

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.

Updating product licenses using `vxlicinst`

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

See [“Replacing a SFCFSHA demo license with a permanent license”](#) on page 128.

To update product licenses

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

```
# vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```

Replacing a SFCFSHA demo license with a permanent license

When a SFCFSHA 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 SFCFSHA on all nodes in the cluster:

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

This command does not shut down any running applications.

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

```
# vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```

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

```
# vxlicrep
```

- 5 Start SFCFSHA on each node:

```
# hastart
```

Configuring SFCFSHA using the Web-based installer

Before you begin to configure SFCFSHA using the Web-based installer, review the configuration requirements.

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 SFCFSHA on a cluster

- 1 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 63.

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

Task	Configure a Product
Product	Storage Foundation for Cluster File System High Availability

Click **Next**.

- 3 On the Select Systems page, enter the system names where you want to configure SFCFSHA, and click **Next**.

Example: **galaxy nebula**

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 I/O fencing.

To configure I/O fencing, click **Yes**.

To configure I/O fencing later, click **No**. You can configure I/O fencing later using the Web-based installer.

See [“Configuring SFCFSHA for data integrity using the Web-based installer”](#) on page 159.

You can also configure I/O fencing later using the `installsfcfsha -fencing` command, the response files, or manually configure.

- 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 over Ethernet. If you choose Auto detect over Ethernet , the installer auto-detects the LLT links over Ethernet. Verify the links and click Yes in the Confirmation dialog box. Skip to step 7. If you click No , you must manually enter the details to configure LLT over Ethernet.
Number of Heartbeats	Choose the number of heartbeat links you want to configure.
NICs	Choose the NICs that you want to configure.
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.
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:

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 *Veritas Storage Foundation Cluster File System High Availability Installation Guide* for instructions to set up SFCFSHA 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.

Security

To configure a secure SFCFSHA 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 `installsfcfsha`.

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 SFCFSHA for data integrity using the Web-based installer”](#) on page 159.

- 11 Click **Next** to complete the process of configuring SFCFSHA.

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.

Database configuration requirements

Most relational database management system (RDBMS) software requires operating system parameters to be set prior to operation. In Solaris 10, system parameters are managed through the Resource Controls facility. The most critical settings are normally located in the Shared Memory and Semaphore settings on Solaris. For precise settings, consult your current database installation and configuration documentation.

Configuring Veritas Storage Foundation Cluster File System High Availability manually

You can manually configure different products within Veritas Storage Foundation Cluster File System High Availability.

Configuring Veritas Volume Manager

Use the following procedures to configure Veritas Volume Manager. If you have installed and configured VxVM using the product installer, you do not need to complete the procedures in this section.

For information on setting up VxVM disk groups and volumes after installation, see "Configuring Veritas Volume Manager" in the *Veritas Storage Foundation Administrator's Guide*.

In releases of VxVM (Volume Manager) before 4.0, a system that was installed with VxVM was configured with a default disk group, `rootdg`. The `rootdg` disk group had to contain at least one disk. By default, operations were directed to the `rootdg` disk group. From release 4.0 onward, VxVM can function without any disk group having been configured.

Starting and enabling the configuration daemon

The VxVM configuration daemon (`vxconfigd`) maintains VxVM disk and disk group configurations. The `vxconfigd` communicates configuration changes to the kernel and modifies configuration information stored on disk.

Startup scripts usually invoke `vxconfigd` at system boot time. The `vxconfigd` daemon must be running for VxVM to operate properly.

The following procedures describe how to check that `vxconfigd` is started, whether it is enabled or disabled, how to start it manually, or how to enable it as required.

To determine whether `vxconfigd` is enabled, use the following command:

```
# vxctl mode
```

The following message indicates that the `vxconfigd` daemon is running and enabled:

```
mode: enabled
```

This message indicates that `vxconfigd` is not running:

```
mode: not-running
```

This message indicates that `vxconfigd` is running, but not enabled:

```
mode: disabled
```

To start the `vxconfigd` daemon, enter the following command:

```
# vxconfigd
```

To enable the volume daemon, enter the following command:

```
# vxctl enable
```

Once started, `vxconfigd` automatically becomes a background process.

By default, `vxconfigd` writes error messages to the console. However, you can configure it to write errors to a log file. For more information, see the `vxconfigd(1M)` and `vxdtl(1M)` manual pages.

Starting the volume I/O daemon

The volume I/O daemon (`vxiod`) provides extended I/O operations without blocking calling processes. Several `vxiod` daemons are usually started at system boot time after initial installation, and they should be running at all times. The procedure below describes how to verify that the `vxiod` daemons are running, and how to start them if necessary.

To verify that `vxiod` daemons are running, enter the following command:

```
# vxiod
```

The `vxiod` daemon is a kernel thread and is not visible using the `ps` command.

If, for example, 16 `vxiod` daemons are running, the following message displays:

```
16 volume I/O daemons running
```

where 16 is the number of `vxiod` daemons currently running. If no `vxiod` daemons are currently running, start some by entering this command:

```
# vxiod set no_of_daemons
```

where the number of daemons ranges from 1 to 16. Symantec recommends that at least one `vxiod` daemon should be run for each CPU in the system.

For more information, see the `vxiod(1M)` manual page.

Using `vxinstall` to configure Veritas Volume Manager

If you used the Veritas Installation Menu or the `installvm` script, you do not need to carry out the instructions in this section. Licensing, configuration of enclosure based naming and creation of a default disk group are managed by the menu installer and the `installvm` script.

Because you are no longer required to configure VxVM disks immediately, the `vxinstall` command no longer invokes the `vxdiskadm` program, making it much simpler than in previous releases.

The utility provides the following functions:

- Licensing VxVM.
- Setting up a system-wide default disk group.

- Starting VxVM daemons in case installation of SF has been done manually.

To run the command, enter

```
# vxinstall
```

which will prompt you to enter a license key:

```
Are you prepared to enter a license key [y,n,q,?] (default: y) y
```

The `vxinstall` program then asks if you want to set up a system-wide default disk group, which is optional:

```
Do you want to setup a system wide default disk group ?
[y,n,q,?] (default: y)
```

VxVM will continue with the question:

```
Which disk group [<group>,list,q,?] ?
```

If you know the name of the disk group that you want to use as the default disk group, enter it at the prompt, or use the `list` option and make a selection.

In releases prior to VxVM 4.0, the default disk group was `rootdg` (the root disk group). For VxVM to function, the `rootdg` disk group had to exist and it had to contain at least one disk. This requirement no longer exists, however you may find it convenient to create a system-wide default disk group. For operations that require a disk group, the system-wide default disk group will be used if the VxVM command is not specified with the `-g` option. The main benefit of creating a default disk group is that VxVM commands default to the default disk group and you will not need to use the `-g` option. To verify the default disk group after it has been created, enter the command:

```
# vxdg defaultdg
```

VxVM does not allow you to use the following names for the default disk group because they are reserved words: `bootdg`, `defaultdg` and `nodg`.

At this stage, the installation of VxVM is complete. To carry out further tasks such as disk encapsulation or initialization, see the *Veritas Storage Foundation Administrator's Guide*.

Excluding a device that VxVM controls

This section describes how to exclude a device that is under VxVM control. The option to prevent paths from being multi-pathed by the Dynamic Multi-Pathing (DMP) driver, `vxdmpp`, is deprecated.

To suppress devices from being seen by VxVM

- 1 Enter the command

```
# vxdiskadm
```

- 2 Select menu item `VolumeManager/Disk/ExcludeDevices` from the `vxdiskadm` main menu.

The following message displays:

```
VxVM INFO V-5-2-5950 This operation might lead to
some devices being suppressed from VxVM's view. (This operation
can be reversed using the vxdiskadm command).
```

```
Do you want to continue? [y,n,q,?] (default: n) y
```

- 3 Enter `y`.

- 4 Select one of the following operations:

- Suppress all paths through a controller from VxVM's view:

Select Option 1.

Enter a controller name when prompted:

```
Enter a controller name:[ctrl_name,all,list,list-exclude,q,?]
```

- Suppress a path from VxVM's view:

Select Option 2.

Enter a path when prompted.

```
Enter a pathname or pattern:[<Pattern>,all,list,list-exclude,q?]
```

- Suppress disks from VxVM's view by specifying a VID:PID combination:

Select Option 3 and read the messages displayed on the screen.

Enter a VID:PID combination when prompted.

```
Enter a VID:PID combination:[<Pattern>,all,list,exclude,q,?]
```

The disks that match the VID:PID combination are excluded from VxVM. Obtain the Vendor ID and Product ID from the Standard SCSI inquiry data returned by the disk.

If you selected any of the options, reboot the system for device exclusion to take effect.

Enabling optional cluster support in VxVM

An optional cluster feature enables you to use VxVM in a cluster environment. The cluster functionality in VxVM allows multiple hosts to simultaneously access and manage a set of disks under VxVM control. A cluster is a set of hosts sharing a set of disks; each host is referred to as a node in the cluster.

Converting existing VxVM disk groups to shared disk groups

If you want to convert existing private disk groups to shared disk groups, use the following procedure. Use these steps if you are moving from a single node to a cluster, or if you are already in a cluster and have existing private disk groups.

To convert existing disk groups to shared disk groups

- 1 Ensure that all systems that are running are part of the same cluster.
- 2 Start the cluster on all of the nodes on which you are converting the disk groups.

3 Configure the disk groups using the following procedure.

To list all disk groups, use the following command:

```
# vxdg list
```

To deport disk groups to be shared, use the following command:

```
# vxdg deport disk_group_name
```

Make sure that CVM is started. To check the master node:

```
# vxdctl -c mode
```

To import disk groups to be shared, use the following command on the master node:

```
# vxdg -s import disk_group_name
```

This procedure marks the disks in the shared disk groups as shared and stamps them with the ID of the cluster, enabling other nodes to recognize the shared disks.

If dirty region logs exist, ensure they are active. If not, replace them with larger ones.

To display the shared flag for all the shared disk groups, use the following command:

```
# vxdg list disk_group_name
```

The disk groups are now ready to be shared.

- ### 4 If the cluster is only running with one node, bring up the other cluster nodes.
- Enter the `vxdg list` command on each node to display the shared disk groups. This command displays the same list of shared disk groups displayed earlier.

Configuring shared disks

This section describes how to configure shared disks. If you are installing VxVM for the first time or adding disks to an existing cluster, you need to configure new shared disks. If you are upgrading VxVM, verify that your shared disks still exist.

The shared disks should be configured from one node only. Since the VxVM software cannot tell whether a disk is shared or not, you must specify which are the shared disks.

Make sure that the shared disks are not being accessed from another node while you are performing the configuration. If you start the cluster on the node where you

perform the configuration only, you can prevent disk accesses from other nodes because the quorum control reserves the disks for the single node.

Also, hot-relocation can be configured.

Verifying existing shared disks

If you are upgrading from a previous release of VxVM, verify that your shared disk groups still exist.

To verify that your shared disk groups exist

- 1 Start the cluster on all nodes.
- 2 Enter the following command on all nodes:

```
# vxdg -s list
```

This displays the existing shared disk groups.

Configuring Veritas File System

After installing Veritas File System, you can create a file system on a disk slice or Veritas Volume Manager volume with the `mkfs` command. Before you can use this file system, you must mount it with the `mount` command. You can unmount the file system later with the `umount` command. A file system can be automatically mounted at system boot time if you add an entry for it in the following file:

```
/etc/vfstab
```

The Veritas-specific commands are described in the Veritas File System guides and online manual pages.

See the *Veritas File System Administrator's Guide*.

Loading and unloading the file system module

The `vxfs` file system module automatically loads on the first reference to a VxFS file system. This occurs when a user tries to mount a VxFS disk layout. In some instances, you may want to load the file system module manually. To do this, first load `vxfs`, then `vxportal`. `vxportal` is a pseudo device driver that enables VxFS commands to issue ioctls to the VxFS modules even when there are no file systems mounted on the system.

```
# modload /kernel/fs/vxfs  
# modload /kernel/drv/vxportal
```

If you have a license for the Veritas Quick I/O feature, you can load its kernel modules:

```
# modload /usr/kernel/drv/sparcv9/fdd
```

To determine if the modules successfully loaded, enter:

```
# modinfo | grep vxportal
# modinfo | grep vxfs
```

The above commands provide information about the modules. The first field in the output is the module ID.

You can unload the module by entering:

```
# modunload -i portal_module_id
# modunload -i vxfs_module_id
```

The `modunload` command fails if any mounted VxFS file systems exist. To determine if any VxFS file systems are mounted, enter:

```
# df -F vxfs
```

vxtunefs command permissions and Cached Quick I/O

By default, you must have superuser (`root`) privileges to use the `/opt/VRTS/bin/vxtunefs` command. The `vxtunefs` command is a tool that lets you change caching policies to enable Cached Quick I/O and change other file system options. Database administrators can be granted permission to change default file system behavior in order to enable and disable Cached Quick I/O. The system administrator must change the `vxtunefs` executable permissions as follows:

```
# chown root /opt/VRTS/bin/vxtunefs
# chgrp dba /opt/VRTS/bin/vxtunefs
# chmod 4550 /opt/VRTS/bin/vxtunefs
```

Setting the permissions for `/opt/VRTS/bin/vxtunefs` to 4550 allows all users in the `dba` group to use the `vxtunefs` command to modify caching behavior for Quick I/O files.

For more information, see the *Veritas File System Administrator's Guide*.

Configuring the SFDB repository database after installation

If you want to use the Storage Foundation for Databases (SFDB) tools, you must set up the SFDB repository after installing and configuring SFCFSHA and Oracle. For SFDB repository set up procedures:

See Veritas Storage Foundation: Storage and Availability Management for Oracle Databases

Configuring SFCFSHA for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing using `installsfcfsha`](#)
- [Setting up disk-based I/O fencing manually](#)
- [Setting up server-based I/O fencing using `installsfcfsha`](#)
- [Setting up non-SCSI-3 server-based I/O fencing in virtual environments using `installsfcfsha`](#)
- [Configuring SFCFSHA for data integrity using the Web-based installer](#)
- [Setting up server-based I/O fencing manually](#)
- [Setting up non-SCSI-3 fencing in virtual environments manually](#)
- [Enabling or disabling the preferred fencing policy](#)

Setting up disk-based I/O fencing using `installsfcfsha`

You can configure I/O fencing using the `-fencing` option of the `installsfcfsha`.

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:

```
# devfsadm
```

- 2 To initialize the disks as VxVM disks, use one of the following methods:
 - Use the interactive `vxdiskadm` utility to initialize the disks as VxVM disks. For more information see the *Veritas 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 `installsfcfsha`

Note: The installer stops and starts SFCFSHA to complete I/O fencing configuration. Make sure to unfreeze any frozen VCS service groups in the cluster for the installer to successfully stop SFCFSHA.

To set up disk-based I/O fencing using the `installsfcfsha`

- 1 Start the `installsfcfsha` with `-fencing` option.

```
# /opt/VRTS/install/installsfcfsha -fencing
```

The `installsfcfsha` 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 SFCFSHA 6.0 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 [1-4,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.
 - Enter the numbers corresponding to the disks that you want to use as coordinator disks.
 - Enter the disk group name.

- 6 Verify that the coordinator disks you chose meet the I/O fencing requirements.

You must verify that the disks are SCSI-3 PR compatible using the `vxfcntlsthdw` utility and then return to this configuration program.

See [“Checking shared disks for I/O fencing”](#) on page 147.

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,?] raw
```

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
 - Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information
- 9 Verify and confirm the I/O fencing configuration information that the installer summarizes.
- 10 Review the output as the configuration program does the following:
- Stops VCS and I/O fencing on each node.
 - Configures disk-based I/O fencing and starts the I/O fencing process.
 - Updates the VCS configuration file `main.cf` if necessary.
 - Copies the `/etc/vxfenmode` file to a date and time suffixed file `/etc/vxfenmode-date-time`. This backup file is useful if any future fencing configuration fails.
 - Starts VCS on each node to make sure that the SFCFSHA is cleanly configured to use the I/O fencing feature.
- 11 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.
- 12 Configure the Coordination Point agent to monitor the coordinator disks.

See [“Configuring CoordPoint agent to monitor coordination points”](#) on page 172.

Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure SFCFSHA meets the I/O fencing requirements. You can test the shared disks using the `vxfentsth` 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 `vxfsentsthdw` utility has additional options suitable for testing many disks. Review the options for testing the disk groups (`-g`) and the disks that are listed in a file (`-f`). You can also test disks without destroying data using the `-r` option.

See the *Veritas Storage Foundation Cluster File System High Availability 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 148.
- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 149.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfsentsthdw utility”](#) on page 150.

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	SUN	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 vxfcntl utility, you must verify that the systems see the same disk.

To verify that the nodes have access to the same disk

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed SFCFSHA.
- 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/rdisk/c1t1d0s2` path on node A and the `/dev/rdisk/c2t1d0s2` path on node B.

From node A, enter:

```
# vxfenadm -i /dev/rdisk/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/rdisk/c2t1d0s2` path.

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

```
# vxfenadm -i /dev/rdisk/c3t1d2s2
```

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

Testing the disks using `vxfcntl` utility

This procedure uses the `/dev/rdisk/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 `vxfcntl` utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/rdisk/c1t1d0s2 is ready to be configured for I/O Fencing on  
node galaxy
```

For more information on how to replace coordinator disks, refer to the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

To test the disks using `vxfcntl` utility

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

- 2 From one node, start the utility.

Run the utility with the `-n` option if you use `rsh` for communication.

```
# vxfcntlsthdw [-n]
```

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

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

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

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

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

```
Enter the first node of the cluster: galaxy
```

```
Enter the second node of the cluster: nebula
```

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

```
Enter the disk name to be checked for SCSI-3 PGR on node
```

```
IP_adrs_of_galaxy in the format:
```

```
for dmp: /dev/vx/rdmp/cxtxdxsx
```

```
for raw: /dev/rdisk/cxtxdxsx
```

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

```
IP_adrs_ofgalaxy and IP_adrs_of_nebula
```

```
/dev/rdsk/c2t13d0s2
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
```

```
IP_adrs_of_nebula in the format:
```

```
for dmp: /dev/vx/rdmp/cxtxdxsx
```

```
for raw: /dev/rdisk/cxtxdxsx
```

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

```
IP_adrs_ofgalaxy and IP_adrs_of_nebula
```

```
/dev/rdsk/c2t13d0s2
```

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

- 5 Review the output as the utility performs the checks and reports its activities.

- 6 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 galaxy.

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

```
ALL tests on the disk /dev/rdisk/c1t1d0s2 have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

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

Setting up disk-based I/O fencing manually

Table 10-1 lists the tasks that are involved in setting up I/O fencing.

Table 10-1 Tasks to set up I/O fencing manually

Task	Reference
Initializing disks as VxVM disks	See "Initializing disks as VxVM disks" on page 144.
Identifying disks to use as coordinator disks	See "Identifying disks to use as coordinator disks" on page 152.
Checking shared disks for I/O fencing	See "Checking shared disks for I/O fencing" on page 147.
Setting up coordinator disk groups	See "Setting up coordinator disk groups" on page 153.
Creating I/O fencing configuration files	See "Creating I/O fencing configuration files" on page 154.
Modifying SFCFSHA configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 155.
Configuring CoordPoint agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 172.
Verifying I/O fencing configuration	See "Verifying I/O fencing configuration" on page 156.

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 144.

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 147.

Setting up coordinator disk groups

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

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

Refer to the *Veritas 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 `vxencoorddg` disk group

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

```
# vxdg init vxencoorddg c1t1d0s2 c2t1d0s2 c3t1d0s2
```

- 2 Set the coordinator attribute value as "on" for the coordinator disk group.

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

- 1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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

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

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

- 4 Edit the following file on each node in the cluster to change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1:

```
/etc/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 If the I/O fencing driver `vxfen` is already running, stop the I/O fencing driver.

```
# svcadm disable -t vxfen
```

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

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

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

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

Regardless of whether the fencing configuration is disk-based or server-based, the value of the cluster-level attribute `UseFence` is set to `SCSI3`.

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

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

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

For example, on each remaining node, enter:

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

- 9 Start the I/O fencing driver and VCS. Perform the following steps on each node:

- Start the I/O fencing driver.
The `vxfen` startup script also invokes the `vxfenconfig` command, which configures the `vxfen` driver to start and use the coordination points that are listed in `/etc/vxfentab`.

```
# svcadm enable vxfen
```

- Start VCS.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

- 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 (galaxy)  
1 (nebula)
```

```
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 using `installsfcfsha`

See [“Setting up server-based I/O fencing using `installsfcfsha`”](#) on page 157.

Setting up non-SCSI-3 server-based I/O fencing in virtual environments using `installsfcfsha`

If you have installed VCS in virtual environments that do not support SCSI-3 PR-compliant storage, you can configure non-SCSI-3 fencing.

Setting up non-SCSI-3 server-based I/O fencing in virtual environments using `installsfcfsha`**To configure I/O fencing using the `installsfcfsha` in a non-SCSI-3 PR-compliant setup**

- 1 Start the `installsfcfsha` with `-fencing` option.

```
# /opt/VRTS/install/installsfcfsha -fencing
```

The `installsfcfsha` 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 SFCFSHA 6.0 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
[1-4,b,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 14250. 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 SFCFSHA 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 SFCFSHA 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 SFCFSHA cluster with the CP server.
 - Adds CP server user to the CP server.
 - Adds SFCFSHA cluster to the CP server user.
 - Updates the following configuration files on each node of the SFCFSHA cluster
 - `/etc/vxfenmode` file
 - `/etc/vxenviron` file
 - `/etc/llttab` file
- 10 Review the output as the installer stops SFCFSHA on each node, starts I/O fencing on each node, updates the VCS configuration file `main.cf`, and restarts SFCFSHA with non-SCSI-3 server-based fencing.
- Confirm to configure the CP agent on the SFCFSHA 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.

Configuring SFCFSHA for data integrity using the Web-based installer

After you configure SFCFSHA, you must configure the cluster for data integrity. Review the configuration requirements.

See [“Configuring SFCFSHA using the Web-based installer”](#) on page 129.

See [“About planning to configure I/O fencing”](#) on page 84.

To configure SFCFSHA for data integrity

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 63.
- 2 On the Select a task and a product page, select the task and the product as follows:

Task	I/O fencing configuration
Product	Storage Foundation for Cluster File System/HA

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, choose whether to configure disk-based fencing or server-based fencing.

If you chose to configure disk-based fencing, go to step [7](#).
If you chose to configure server-based fencing, go to step [10](#).
- 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.
Go to step [16](#).
 - 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**. Go to step [9](#).

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.

Go to step [16](#).

10 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**.

11 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**.

12 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.
- 13 In the Confirmation dialog box that appears, confirm whether the coordination points information you provided is correct, and click **Yes**.
 - 14 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.
 - 15 Configure the CP agent on the SFCFSHA (application cluster), and click **Next**.
 - 16 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.
 - 17 Select the checkbox to specify whether you want to send your installation information to Symantec.
 Click **Finish**. The installer prompts you for another task.

Setting up server-based I/O fencing manually

Tasks that are involved in setting up server-based I/O fencing manually include:

Table 10-2 Tasks to set up server-based I/O fencing manually

Task	Reference
Preparing the CP servers for use by the SFCFSHA cluster	See “Preparing the CP servers manually for use by the SFCFSHA cluster” on page 163.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See “Configuring server-based fencing on the SFCFSHA cluster manually” on page 166.
Modifying SFCFSHA configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 155.
Configuring Coordination Point agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 172.
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 173.

Preparing the CP servers manually for use by the SFCFSHA cluster

Use this procedure to manually prepare the CP server for use by the SFCFSHA cluster or clusters.

[Table 10-3](#) displays the sample values used in this procedure.

Table 10-3 Sample values in procedure

CP server configuration component	Sample name
CP server	mycps1
Node #1 - SFCFSHA cluster	galaxy
Node #2 - SFCFSHA cluster	nebula
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the SFCFSHA cluster

- 1 Determine the cluster name and uuid on the SFCFSHA cluster.

For example, issue the following commands on one of the SFCFSHA cluster nodes (*galaxy*):

```
# 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 SFCFSHA cluster and nodes are present in the CP server.

For example:

```
# cpsadm -s mycps1.symantecexample.com -a list_nodes
```

ClusName	UUID	Hostname (Node ID)	Registered
clus1	{f0735332-1dd1-11b2-bb31-00306eea460a}	galaxy(0)	0
clus1	{f0735332-1dd1-11b2-bb31-00306eea460a}	nebula(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 *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

3 Add the SFCFSHA cluster and nodes to each CP server.

For example, issue the following command on the CP server (mycps1.symantecexample.com) to add the cluster:

```
# cpsadm -s mycps1.symantecexample.com -a add_clus\  
-c clus1 -u {f0735332-1dd1-11b2}
```

```
Cluster clus1 added successfully
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the first node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c clus1 -u {f0735332-1dd1-11b2} -h galaxy -n0
```

```
Node 0 (galaxy) successfully added
```

Issue the following command on the CP server (mycps1.symantecexample.com) to add the second node:

```
# cpsadm -s mycps1.symantecexample.com -a add_node\  
-c clus1 -u {f0735332-1dd1-11b2} -h nebula -n1
```

```
Node 1 (nebula) successfully added
```

4 If security is to be enabled, check whether the CPSADM@VCS_SERVICES@cluster_uuid users are created in the CP server.

If the output below does not show the users, then add them as described in the next step.

```
# cpsadm -s mycps1.symantecexample.com -a list_users
```

```
Username/Domain Type Cluster Name / UUID Role  
  
CPSADM@VCS_SERVICES@f0735332-1dd1-11b2/vx  
clus1/{f0735332-1dd1-11b2} Operator
```

If security is to be disabled, then add the user name "cpsclient@hostname" to the server instead of the CPSADM@VCS_SERVICES@cluster_uuid (for example, cpsclient@galaxy).

The CP server can only run in either secure mode or non-secure mode, both connections are not accepted at the same time.

5 Add the users to the CP server.

Issue the following commands on the CP server (mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
CPSADM@VCS_SERVICES@cluster_uid\  
-f cps_operator -g vx
```

```
User CPSADM@VCS_SERVICES@cluster_uid  
successfully added
```

6 Authorize the CP server user to administer the SFCFSHA cluster. You must perform this task for the CP server users corresponding to each node in the SFCFSHA cluster.

For example, issue the following command on the CP server (mycps1.symantecexample.com) for SFCFSHA cluster clus1 with two nodes galaxy and nebula:

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c clus1\  
-u {f0735332-1dd1-11b2}\  
-e CPSADM@VCS_SERVICES@cluster_uid\  
-f cps_operator -g vx
```

```
Cluster successfully added to user  
CPSADM@VCS_SERVICES@cluster_uid privileges.
```

Configuring server-based fencing on the SFCFSHA cluster manually

The configuration process for the client or SFCFSHA cluster to use CP server as a coordination point requires editing the `/etc/vxfenmode` file.

You need to edit this file to specify the following information for your configuration:

- Fencing mode
- Fencing mechanism
- Fencing disk policy (if applicable to your I/O fencing configuration)
- Appropriate value for the security configuration
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)

Note: Whenever coordinator disks are used as coordination points in your I/O fencing configuration, you must create a disk group (vxencoorddg). You must specify this disk group in the `/etc/vxfenmode` file.

See [“Setting up coordinator disk groups”](#) on page 153.

The customized fencing framework also generates the `/etc/vxfentab` file which has security setting and the coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

To configure server-based fencing on the SFCFSHA 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.

The following sample file output displays what the `/etc/vxfenmode` file contains:

See [“Sample vxfenmode file output for server-based fencing”](#) on page 167.

- 3 After editing the `/etc/vxfenmode` file, run the `vxfen` init script to start fencing.

For example:

```
# svcadm enable vxfen
```

- 4 For CP servers in secure mode, make sure that the security is enabled on the cluster and the credentials for the CPSADM are present in the `/var/VRTSvcs/vcsauth/data/CPSADM` directory.

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 when enabled uses secure communication to the cp server
# using VxAT (Veritas Authentication Service)
# available options:
# 0 - don't use Veritas Authentication Service for cp server
# communication
# 1 - use Veritas Authentication Service for cp server
# communication
security=1

#
# Specify 3 or more odd number of coordination points in this file,
# one in 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:
```

```
#
# 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,[mycps1.company.com]
# cps2=[192.168.0.25]
# cps3=[mycps2.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 [mycps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#   [192.168.0.23]
#   [mycps1.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]
#   [mycps1.company.com]
#   [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoordg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
```

[Table 10-4](#) defines the vxfenmode parameters that must be edited.

Table 10-4 vxfenmode file parameters

vxfenmode File Parameter	Description
vxfen_mode	Fencing mode of operation. This parameter must be set to "customized".
vxfen_mechanism	Fencing mechanism. This parameter defines the mechanism that is used for fencing. If one of the three coordination points is a CP server, then this parameter must be set to "cps".
scsi3_disk_policy	Configure the vxfen module to use either DMP devices, "dmp" or the underlying raw character devices, "raw". Note: The configured disk policy is applied on all the nodes.
security	Security parameter 1 indicates that secure mode is used for CP server communications. Security parameter 0 indicates that communication with the CP server is made in non-secure mode. The default security value is 1.
cps1, cps2, or vxfendg	Coordination point parameters. Enter either the virtual IP address or the FQHN (whichever is accessible) of the CP server. <i>cps<number>=[virtual_ip_address/virtual_host_name]:port</i> Where <i>port</i> is optional. The default port value is 14250. If you have configured multiple virtual IP addresses or host names over different subnets, you can specify these as comma-separated values. For example: <i>cps1=[192.168.0.23], [192.168.0.24]:58888, [mycps1.company.com]</i> Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxencoordg) 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).

Table 10-4 vxfenmode file parameters (*continued*)

vxfenmode File Parameter	Description
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 14250. 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.

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

To configure CoordPoint agent to monitor coordination points

- 1 Ensure that your SFCFSHA 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
# hagr -add vxfen
# hagr -modify vxfen SystemList galaxy 0 nebula 1
# hagr -modify vxfen AutoFailOver 0
# hagr -modify vxfen Parallel 1
# hagr -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 1
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

- 3 Verify the status of the agent on the SFCFSHA 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     galaxy  ONLINE
coordpoint     State     nebula  ONLINE
```

- 4 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed 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 all such information 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 *Veritas Storage Foundation Cluster File System High Availability 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 162.

- 2 Make sure that the SFCFSHA 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_vxind_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 vxmfend process wait for a customized fencing
# script to complete. Only used with vxmfen_mode=customized
vxmfen_script_timeout=25

#
# security when enabled uses secure communication to the cp server
# using VxAT (Veritas Authentication Service)
# available options:
# 0 - don't use Veritas Authentication Service for cp server
#   communication
# 1 - use Veritas Authentication Service for cp server
#   communication
security=1

#
# Specify 3 or more odd number of coordination points in this file,
# one in 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:
#
# 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,[mycps1.company.com]
# cps2=[192.168.0.25]
# cps3=[mycps2.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 [mycps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#   [192.168.0.23]

```

```
# [mycps1.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]
# [mycps1.company.com]
# [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoordg
#
# 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=[mycps1.company.com]
cps2=[mycps2.company.com]
cps3=[mycps3.company.com]
port=14250
=====
```

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 81.

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 galaxy five times more weight compared to nebula, run the following commands:

```
# hasys -modify galaxy FencingWeight 50  
# hasys -modify nebula FencingWeight 10
```

- Save the VCS configuration.

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

- 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:

```
# hagrps -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 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
```

Upgrading Storage Foundation Cluster File System High Availability

- [Chapter 11. Preparing to upgrade Veritas Storage Foundation Cluster File System High Availability](#)
- [Chapter 12. Performing a typical Storage Foundation Cluster File System High Availability upgrade using the installer](#)
- [Chapter 13. Performing a rolling upgrade](#)
- [Chapter 14. Performing a phased upgrade](#)
- [Chapter 15. Upgrading the operating system](#)
- [Chapter 16. Upgrading Veritas Volume Replicator](#)
- [Chapter 17. Upgrading language packages](#)
- [Chapter 18. Migrating from SFHA to SFCFSHA](#)
- [Chapter 19. Upgrading with Live Upgrade](#)
- [Chapter 20. Performing post-upgrade tasks](#)

Preparing to upgrade Veritas Storage Foundation Cluster File System High Availability

This chapter includes the following topics:

- [About upgrading](#)
- [About the different ways that you can upgrade](#)
- [Supported upgrade paths](#)
- [About using the installer to upgrade when the root disk is encapsulated](#)
- [Preparing to upgrade](#)

About upgrading

There are many types of upgrades available. Before you start to upgrade, review the types of upgrades for the Veritas products.

See [“About the different ways that you can upgrade”](#) on page 184.

Review the supported upgrade paths that are available for the different methods of upgrading.

See [“Supported upgrade paths”](#) on page 184.

After you determine the type of upgrade that you want to perform and its upgrade paths, review the steps to prepare for the upgrade.

If you want to upgrade Coordination Point (CP) server systems that use Veritas Cluster Server (VCS) or Storage Foundation High Availability (SFHA) to 6.0, make sure you upgraded all application clusters to 6.0. Then, upgrade VCS or SFHA on the CP server systems.

About the different ways that you can upgrade

Symantec offers you several different ways to upgrade. You need to decide which upgrade method best suits your environment, your expertise, and the downtime required.

Table 11-1 Review this table to determine how you want to perform the upgrade

Upgrade types and considerations	Methods available for upgrade
Typical upgrades—use a Veritas provided tool or you can perform the upgrade manually. Requires some server downtime.	Script-based—you can use this to upgrade for the supported upgrade paths Web-based—you can use this to upgrade for the supported upgrade paths Manual—you can use this to upgrade from the previous release Response file—you can use this to upgrade from the previous release
Native operating system upgrade—use the upgrade software that comes with the operating system. Note that not all operating systems support native upgrades.	Operating system specific methods Operating system upgrades

Supported upgrade paths

[Table 11-2](#) describes upgrading to 6.0 on Solaris SPARC.

Table 11-2 Solaris SPARC upgrades using the script- or Web-based installer

Veritas software versions	Solaris 8 or older	Solaris 9	Solaris 10
3.5 3.5 MP4 4.0 4.0 MP1 4.0 MP2	Upgrade the operating system to at least Solaris 10, then upgrade product to 5.0 MP3. Upgrade to 6.0 using the installer script.	Upgrade the operating system to at least Solaris 10, then upgrade product to 5.0 MP3. Upgrade to 6.0 using the installer script.	N/A
4.1 4.1 MP1 4.1 MP2 5.0 5.0 MP1	Upgrade the operating system to at least Solaris 10, then upgrade product to 5.0 MP3. Upgrade to 6.0 using the installer script.	Upgrade the operating system to at least Solaris 10, then upgrade product to 5.0 MP3. Upgrade to 6.0 using the installer script.	Upgrade the product to 5.0 MP3. Upgrade to 6.0 using the installer script.
5.0 MP3 5.0 MP3 RPx	Upgrade the operating system to at least Solaris 10. Upgrade to 6.0 using the installer script.	Upgrade operating system to at least Solaris 10. Upgrade to 6.0 using the installer script.	Upgrade directly to 6.0 using the installer script.
5.1 5.1 RPx 5.1 SP1 5.1 SP1 RPx	N/A	Upgrade the operating system to at least Solaris 10. Upgrade to 6.0 using the installer script.	Upgrade directly to 6.0 using the installer script.

[Table 11-3](#) describes upgrading to 6.0 on Solaris x64.

Table 11-3 Solaris x64 upgrades using the script- or Web-based installer

Veritas software versions	Solaris 10
4.1 4.1 Phase 2	Upgrade to 5.0 MP3. Upgrade to 6.0 using the installer script.
5.0 5.0 MP3 5.0 MP3 RPx	Upgrade to 5.0 MP3. Upgrade to 6.0 using the installer script.

Table 11-3 Solaris x64 upgrades using the script- or Web-based installer
(continued)

Veritas software versions	Solaris 10
5.1	Use the installer to upgrade to 6.0.
5.1 RPx	
5.1 SP1*	
5.1 SP1 RPx	

* When you upgrade to 6.0 from 5.1 SP1 using the Web-based installer, you must first upgrade to 5.1 SP1 RP1 if you want the installer to create a backup of the boot disk. You can upgrade directly to 6.0 from 5.1 SP1 if you do not want the installer to create a backup of the boot disk.

About using the installer to upgrade when the root disk is encapsulated

In prior versions of SFCFSHA, when upgrading a system with an encapsulated root disk, you first had to unencapsulate. When upgrading to SFCFSHA 6.0, that is no longer necessary, as shown in the table below.

Table 11-4 Upgrading using installer when the root disk is encapsulated

Starting version	Ending version	Action required
5.0 MP3 RPx	6.0	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 or 5.1 RPx	6.0	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 SP1 or 5.1 SP1 RPx	6.0	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Preparing to upgrade

Before you upgrade, you need to prepare the systems and storage. Review the following procedures and perform the appropriate tasks.

Getting ready for the upgrade

Complete the following tasks before you perform the upgrade:

- Review the *Veritas Storage Foundation Cluster File System High Availability Release Notes* for any late-breaking information on upgrading your system.
- Review the Symantec Technical Support website for additional information: <http://www.symantec.com/techsupp/>
- Perform the following system-level settings:
 - Set `diag-level` to `min` to perform the minimum amount of diagnostics when the system boots. Depending on the configuration of your systems you may want to re-enable this after you perform the upgrade.

```
{1} ok setenv diag-level min  
  
diag-level=min
```

- Set `auto-boot?` to `false`. For tight control when systems reboot, set this variable to `false`. Re-enable this variable after the upgrade.

```
{1} ok setenv auto-boot? false  
  
auto-boot?=false
```

- Deactivate cron to make sure that extraneous jobs are not performed while you upgrade the systems. Do one of the following:
Solaris 9:

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

Solaris 10:

```
# svcadm disable -t system/cron:default
```

- For Solaris 10, make sure that all non-global zones are booted and in the running state before you use the Veritas product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must upgrade each package in each non-global zone manually.

For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you reboot the alternative root, you can install `VRTSodm`.

- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the packages, for example `/packages/Veritas` when the root file system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.
 Do not put the files under `/tmp`, which is erased during a system reboot. Do not put the files on a file system that is inaccessible prior to running the upgrade script.
 You can use a Veritas-supplied disc for the upgrade as long as modifications to the upgrade script are not required. If `/usr/local` was originally created as a slice, modifications are required.
- Unmount all the file systems not on the `root` disk. Comment out their entries in `/etc/vfstab`. Stop the associated volumes and deport the associated disk groups. Any file systems that the Solaris operating system or Storage Foundation assumes should be in `rootdg` but are not, must be unmounted and the associated entry in `/etc/vfstab` commented out.
- For any startup scripts in `/etc/rcS.d`, comment out any application commands or processes that are known to hang if their file systems are not present.
- Make sure that the current operating system supports version 6.0 of the product. If the operating system does not support it, plan for a staged upgrade.
- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Veritas products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/vfstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/vfstab` and not mounted during the upgrade. Active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure the file systems are clean before upgrading.
 See [“Verifying that the file systems are clean”](#) on page 196.
- Symantec recommends that you upgrade VxFS disk layouts to a supported version prior to installing VxFS 6.0. Unsupported disk layout versions 4, 5, and 6 can be mounted for the purpose of online upgrading in VxFS 6.0. You can upgrade unsupported layout versions online before installing VxFS 6.0.

- Upgrade arrays (if required).
- If replication using VVR is configured, Symantec recommends that the disk group version is at least 110 prior to upgrading.

```
# vxvg list diskgroup
```

- If replication using VVR is configured, make sure the size of the SRL volume is greater than 110 MB.
Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.
- If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date on all the hosts.

```
# vxrlink -g diskgroup status rlink_name
```

Note: Do not continue until the primary RLINKs are up-to-date.

- If VCS is used to manage VVR replication, follow the preparation steps to upgrade VVR and VCS agents.
- Make sure that you have worked out all terminal emulation issues. Make sure that the terminal you use is fully functional for OpenBoot prompts and single-user and multi-user run levels.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. Removing the disk in this manner offers a fallback point.
- Determine if the root disk is encapsulated.
See [“Determining if the root disk is encapsulated”](#) on page 190.

Creating backups

Save relevant system information before the upgrade.

To create backups

- 1 Log in as superuser.
- 2 Before the upgrade, ensure that you have made backups of all data that you want to preserve.

Back up the `/etc/system` file.

- 3 Copy the `vfstab` file to `vfstab.orig`:

```
# cp /etc/vfstab /etc/vfstab.orig
```
- 4 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 5 If you are installing the high availability version of the Veritas Storage Foundation 6.0 software, follow the guidelines given in the *Veritas Cluster Server Installation Guide* and *Veritas Cluster Server Release Notes* for information on preserving your VCS configuration across the installation procedure.

Determining if the root disk is encapsulated

Before you upgrade, you need to determine if the root disk is encapsulated by running the following command:

```
# mount | grep "/" on"
```

If the output from this command includes a path name that contains `vx` and `rootvol` as in `/dev/vx/dsk/bootdg/rootvol`, then the root disk is encapsulated.

If the root disk is encapsulated, follow the appropriate upgrade procedures.

Pre-upgrade planning for Veritas Volume Replicator

Before installing or upgrading Veritas Volume Replicator (VVR):

- Confirm that your system has enough free disk space to install VVR.
- Make sure you have root permissions. You must have root permissions to perform the install and upgrade procedures.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for more information.

See the *Getting Started Guide* for more information on the documentation.

Planning an upgrade from the previous VVR version

If you plan to upgrade VVR from the previous VVR version, you can upgrade VVR with reduced application downtime by upgrading the hosts at separate times. While the Primary is being upgraded, the application can be migrated to the Secondary, thus reducing downtime. The replication between the (upgraded) Primary and the Secondary, which have different versions of VVR, will still continue. This feature facilitates high availability even when the VVR upgrade is not complete on both the

sites. Symantec recommends that the Secondary hosts be upgraded before the Primary host in the Replicated Data Set (RDS).

See the *Veritas Storage Foundation Cluster File System High Availability Release Notes* for information regarding VVR support for replicating across Storage Foundation versions

Replicating between versions is intended to remove the restriction of upgrading the Primary and Secondary at the same time. VVR can continue to replicate an existing RDS with Replicated Volume Groups (RVGs) on the systems that you want to upgrade. When the Primary and Secondary are at different versions, VVR does not support changing the configuration with the `vradmin` command or creating a new RDS.

Also, if you specify TCP as the network protocol, the VVR versions on the Primary and Secondary determine whether the checksum is calculated. As shown in [Table 11-5](#), if either the Primary or Secondary are running a version of VVR prior to 6.0, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 6.0, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 11-5 VVR versions and checksum calculations

VVR prior to 6.0 (DG version <= 140)	VVR 6.0 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes
Primary and Secondary		Yes
	Primary and Secondary	No

Note: When replicating between versions of VVR, avoid using commands associated with new features. The earlier version may not support new features and problems could occur.

If you do not need to upgrade all the hosts in the RDS simultaneously, you can use replication between versions after you upgrade one host. You can then upgrade the other hosts in the RDS later at your convenience.

Note: If you have a cluster setup, you must upgrade all the nodes in the cluster at the same time.

Planning and upgrading VVR to use IPv6 as connection protocol

Veritas Storage Foundation Cluster File System High Availability supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

- VVR continues to support replication between IPv4-only nodes with IPv4 as the internet protocol
- VVR supports replication between IPv4-only nodes and IPv4/IPv6 dual-stack nodes with IPv4 as the internet protocol
- VVR supports replication between IPv6-only nodes and IPv4/IPv6 dual-stack nodes with IPv6 as the internet protocol
- VVR supports replication between IPv6 only nodes
- VVR supports replication to one or more IPv6 only nodes and one or more IPv4 only nodes from a IPv4/IPv6 dual-stack node
- VVR supports replication of a shared disk group only when all the nodes in the cluster that share the disk group are at IPv4 or IPv6

Additional settings for using VVR in a localized environment

If the language packages for VVR are installed, VVR displays localized messages, if the client locale is a supported non-English locale. The client locale is the locale from which you are accessing the VVR command line or GUI. For example, if the Japanese version of VVR is installed, then the messages are displayed in the Japanese locale, if the client locale is Japanese.

Make sure that the appropriate locale has been installed on all the hosts that are intended to be a part of the VVR RDS setup. Otherwise, some VVR error messages will be displayed in English, because it is the default locale. Make sure the following settings are done on all hosts that are intended to be part of the RDS:

- Install the required client locale from the Operating System disc.
- Install the required Volume Manager and VVR localized packages.
- Set the client locale, before using any of the VVR interfaces:
 - For the VVR command line, set the locale using the appropriate method for your operating system.
 - For VRW, select the locale from the VRW login page.

Preparing to upgrade VVR when VCS agents are configured

To prepare to upgrade VVR when VCS agents for VVR are configured, perform the following tasks in the order presented:

- [Freezing the service groups and stopping all the applications](#)
- [Preparing for the upgrade when VCS agents are configured](#)

Freezing the service groups and stopping all the applications

This section describes how to freeze the service groups and stop all applications.

To freeze the service groups and stop applications

Perform the following steps for the Primary and Secondary clusters:

- 1 Log in as the superuser.
- 2 Make sure that `/opt/VRTS/bin` is in your PATH so that you can execute all the product commands.
- 3 Before the upgrade, cleanly shut down all applications.

In a shared disk group environment:

- OFFLINE all application service groups that do not contain RVGShared resources. Do not OFFLINE the ClusterService, cvm and RVGLogowner groups.
- If the application resources are part of the same service group as an RVGShared resource, then OFFLINE only the application resources.

In a private disk group environment:

- OFFLINE all application service groups that do not contain RVG resources. Do not OFFLINE the service groups containing RVG resources.
- If the application resources are part of the same service group as an RVG resource, then OFFLINE only the application resources. In other words, ensure that the RVG resource remains ONLINE so that the private disk groups containing these RVG objects do not get deported.

Note: You must also stop any remaining applications not managed by VCS.

- 4 On any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 5 On any node in the cluster, list the groups in your configuration:

```
# hagrps -list
```

- 6 On any node in the cluster, freeze all service groups except the ClusterService group by typing the following command for each group name displayed in the output from step 5.

```
# hagrps -freeze group_name -persistent
```

Note: Make a note of the list of frozen service groups for future use.

- 7 On any node in the cluster, save the configuration file (`main.cf`) with the groups frozen:

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

Note: Continue only after you have performed steps 3 to step 7 for each cluster.

- 8 Display the list of service groups that have RVG resources and the nodes on which each service group is online by typing the following command on any node in the cluster:

```
# hares -display -type RVG -attribute State
Resource      Attribute      System      Value
VVRGrp        State          system02    ONLINE
ORAGrp        State          system02    ONLINE
```

Note: For the resources that are ONLINE, write down the nodes displayed in the System column of the output.

- 9 Repeat step 8 for each cluster.
- 10 For private disk groups, determine and note down the hosts on which the disk groups are imported.
 See [“Determining the nodes on which disk groups are online”](#) on page 195.
- 11 For shared disk groups, run the following command on any node in the CVM cluster:

```
# vxdctl -c mode
```

Note the master and record it for future use.

Determining the nodes on which disk groups are online

For private disk groups, determine and note down the hosts on which the disk groups containing RVG resources are imported. This information is required for restoring the configuration after the upgrade.

To determine the online disk groups

- 1 On any node in the cluster, list the disk groups in your configuration, and note down the disk group names listed in the output for future use:

```
# hares -display -type RVG -attribute DiskGroup
```

Note: Write down the list of the disk groups that are under VCS control.

- 2 For each disk group listed in the output in step 1, list its corresponding disk group resource name:

```
# hares -list DiskGroup=diskgroup Type=DiskGroup
```

- 3 For each disk group resource name listed in the output in step 2, get and note down the node on which the disk group is imported by typing the following command:

```
# hares -display dg_resname -attribute State
```

The output displays the disk groups that are under VCS control and nodes on which the disk groups are imported.

Preparing for the upgrade when VCS agents are configured

If you have configured the VCS agents, it is recommended that you take backups of the configuration files, such as `main.cf` and `types.cf`, which are present in the `/etc/VRTSvcs/conf/config` directory.

To prepare a configuration with VCS agents for an upgrade

- 1 List the disk groups on each of the nodes by typing the following command on each node:

```
# vxdisk -o alldgs list
```

The output displays a list of the disk groups that are under VCS control and the disk groups that are not under VCS control.

Note: The disk groups that are not locally imported are displayed in parentheses.

- 2 If any of the disk groups have not been imported on any node, import them. For disk groups in your VCS configuration, you can import them on any node. For disk groups that are not under VCS control, choose an appropriate node on which to import the disk group. Enter the following command on the appropriate node:

```
# vxdg -t import diskgroup
```

- 3 If a disk group is already imported, then recover the disk group by typing the following command on the node on which it is imported:

```
# vxrecover -bs
```

- 4 Verify that all the Primary RLINKs are up to date.

```
# vxrlink -g diskgroup status rlink_name
```

Note: Do not continue until the Primary RLINKs are up-to-date.

Verifying that the file systems are clean

Verify that all file systems have been cleanly unmounted.

To make sure the file systems are clean

- 1 Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | /opt/VRTSvxfs/sbin/fsdb filesystem | \  
    grep clean  
    flags 0 mod 0 clean clean_value
```

A *clean_value* value of `0x5a` indicates the file system is clean. A value of `0x3c` indicates the file system is dirty. A value of `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.

- 2 If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem  
# mount -F vxfs Block_Device  
    mountpoint  
# umount mountpoint
```

These commands should complete any extended operations on the file system and unmount the file system cleanly.

A pending large package clone removal extended operation might be in progress if the `umount` command fails with the following error:

```
file system device busy
```

An extended operation is in progress if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system  
file system still in progress.
```

- 3 If an extended operation is in progress, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large package clone can take several hours.
- 4 Repeat step 1 to verify that the unclean file system is now clean.

Upgrading the array support

The Storage Foundation 6.0 release includes all array support in a single package, VRTSaslapm. The array support package includes the array support previously included in the VRTSvxvm package. The array support package also includes support previously packaged as external array support libraries (ASLs) and array policy modules (APMs).

See the 6.0 Hardware Compatibility List for information about supported arrays.

See [“Hardware compatibility list \(HCL\)”](#) on page 41.

When you upgrade Storage Foundation products with the product installer, the installer automatically upgrades the array support. If you upgrade Storage Foundation products with manual steps, you should remove any external ASLs or APMs that were installed previously on your system. Installing the VRTSvxvm package exits with an error if external ASLs or APMs are detected.

After you have installed Storage Foundation 6.0, Symantec provides support for new disk arrays through updates to the `VRTSaslapm` package.

For more information about array support, see the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

Performing a typical Storage Foundation Cluster File System High Availability upgrade using the installer

This chapter includes the following topics:

- [Performing a full upgrade](#)

Performing a full upgrade

Performing a full upgrade involves the following tasks:

- Ensuring that the file systems are clean
- Updating the main.cf file
- Performing the upgrade
- Updating the configuration and confirming startup
- Upgrading the remaining nodes

Ensuring the file systems are clean

Before upgrading to SFCFSHA 6.0, ensure that the file systems are clean. To ensure that the logs have been replayed and the file systems are marked clean:

To ensure the file systems are clean

- 1 Log in as superuser onto any node in the cluster.
- 2 Take the service group offline on each node of the cluster, which contains VxFS and CFS resources:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFSSMount resource.

Repeat this step for each SFCFSHA service group.

Note: This unmounts the CFS file systems.

- 3 Unmount all VxFS file systems not under VCS control:

```
# umount mount_point
```

- 4 Check and repair each VxFS file system:

```
# fsck -F vxfs /dev/vx/rdisk/diskgroup/volume
```

The `fsck` command in `/opt/VRTS/bin` accepts either the block or character device (`/dev/vx/dsk/dg/vol`) or (`/dev/vx/rdisk/dg/vol`). The operating system version of `fsck` may limit the device types it accepts.

For more information, see the `fsck` and `fsck_vxfs` man pages.

Repeat this step for each file system.

Modifying the main.cf file

Save a copy of the `main.cf` file and modify the configuration information in the `main.cf` file.

To modify the main.cf file

- 1 On any node, make a copy of the current `main.cf` file. For example:

```
# cp /etc/VRTSvcs/conf/config/main.cf \
    /etc/VRTSvcs/conf/main.save
```

- 2 Choose one node from the cluster to execute step 3 through step 9.
- 3 On the node you selected in step 2, enter:

```
# haconf -makerw
# hares -unlink vxfsckd qlogckd
# hares -unlink qlogckd cvm_clus
# hares -link vxfsckd cvm_clus
# hares -delete qlogckd
# haconf -dump -makero
```

- 4 On all the nodes in the cluster, enter:

```
# ps -ef | grep qlogckd
# kill -9 pid_of_qlogckd
# modinfo | grep -i qlog
# modunload -i module_id_of_qlog
```

- 5 On the node you selected in step 2, stop VCS on all nodes:

```
# /opt/VRTS/bin/hastop -all -force
```

- 6 On the node you selected in step 2 and if you have configured the VCS Cluster Manager (Web Console), complete the following to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

- Remove VRTSweb:

```
Process VRTSweb (
    PathName = "/opt/VRTSvcs/bin/haweb"
    Arguments = "10.129.96.64 8181"
)
```

- Replace it with:

```
VRTSWebApp VCSweb (
    Critical =0
    AppName = vcs
    InstallDir = "/opt/VRTSweb/VERITAS"
```

```
TimeForOnline = 5
)
```

- Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named `csgnic` and the public NIC device is `hme0`, add:

```
NIC csgnic (
    Device = hme0
```

- Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMgr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
ntfr requires csgnic
webip requires csgnic
```

- 7 On the node you selected in step 2, remove `qlogckd` from the `/etc/VRTSvcs/conf/config/main.cf` file. For example:

```
CFSQlogckd qlogckd (
    Critical = 0
)
```

Make sure you remove all dependencies on `qlogckd` from the `main.cf` file.

- 8 On the node you selected in step 2, verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

- 9 On the node you selected in step 2, start VCS:

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

- 10 On the remaining nodes in the cluster, start VCS:

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

- 11 If VVR is configured, freeze the service groups and stop the applications.

See [“Freezing the service groups and stopping all the applications”](#) on page 193.

Performing the upgrade

To perform the upgrade

- 1 Log in as superuser.
- 2 Insert the appropriate media disc into your system's DVD-ROM drive.
- 3 If volume management software is running on your system, the software disc automatically mounts as `/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

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

where `c##d#` is the location of the CD drive.

- 4 Change to the top-level directory on the disc:

```
# cd /cdrom
```

- 5 Verify there are no VxFS file systems mounted on the nodes being upgraded:

```
# mount -p | grep vxfs
```

If any VxFS file systems are mounted, offline the group on each node of the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where `group` is the VCS service group that has the CVMVolDg and CFSSMount resource.

If VxFS are not managed by VCS then unmount them manually:

```
# umount mount_point
```

Repeat this step for each SFCFSHA service group.

- 6 Start the upgrade from any node in the cluster. Enter the following command, and then press `y` to upgrade the cluster configuration.

```
# ./installsfcfsha -upgrade
```

- 7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement
as specified in the storage_foundation_cluster_file_system_ha/EULA/
en/EULA_CFSHA_Ux_6.0.pdf file present on media? [y,n,q,?] y
```

- 8 The installer discovers if any of the systems that you are upgrading have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the system's boot disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer **y**.
- 9 The installer then prompts you to name the backup boot disk group. Enter the name for it or press **Enter** to accept the default.
- 10 If you are prompted to start the split operation. Press **y** to continue.

Note: The split operation can take some time to complete.

- 11 You are prompted to enter the system names (in the following example, "host1" and "host2") on which the software is to be upgraded. Enter the system name or names and then press Return.

```
Enter the system names separated by spaces on which to
install SFCFSHA: host1 host2
```

- 12 During the initial system check, the installer verifies that communication between systems has been set up.

If the installer hangs or asks for a login password, stop the installer and set up ssh or rsh. Then run the installer again.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 446.
- 13 After the system checks complete, the installer displays a list of the packages that will be upgraded. Press Enter to continue with the upgrade.
- 14 Output shows information that SFCFSHA must be stopped on a running system. Enter **y** to continue.
- 15 Press **Return** to begin removing the previous packages and installing the new.
- 16 Press **Return** again for summary information about logs and reboots.

Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.

- 17 Update the configuration.
- 18 Only perform this step if you have split the mirrored root disk to back it up. After a successful reboot, verify the upgrade and re-join the backup disk group. If the upgrade fails, revert to the backup disk group.

See [“Re-joining the backup boot disk group into the current disk group”](#) on page 260.

See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 261.

Updating the configuration and confirming startup

Perform the following steps on each upgraded node.

To update the configuration and confirm startup

- 1 Remove the `/etc/VRTSvcs/conf/config/.stale` file, if it exists.

```
# rm -f /etc/VRTSvcs/conf/config/.stale
```

- 2 Reboot the upgraded nodes, if root disk is encapsulated.

```
# /usr/sbin/shutdown -i6 -g0 -y
```

- 3 After the nodes reboot, verify that LLT is running:

```
# lltconfig  
LLT is running
```

- 4 Verify GAB is configured:

```
# gabconfig -l | grep 'Driver.state' | \  
grep Configured  
Driver state : Configured
```

- 5 Verify VxVM daemon is started and enabled:

```
# /opt/VRTS/bin/vxdctl mode  
mode: enabled
```

- 6 Confirm all upgraded nodes are in a running state.

```
# gabconfig -a
```

- 7 Log in as superuser.

- 8 Insert the appropriate media disc into your system's CD-ROM drive.
- 9 If volume management software is running on your system, the software disc automatically mounts as `/cdrom`.

If volume management software is not available to mount the CD-ROM, you must mount it manually, enter:

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

where `c##d#` is the location of the CD drive.

- 10 Change to the top-level directory on the disc:

```
# cd /cdrom
```

- 11 Run `installsfcfsha` from any node in the cluster:

```
# ./installsfcfsha -start system01
                        system02
```

- 12 After the configuration is complete, the CVM and SFCFSHA groups may come up frozen. To find out the frozen CVM and SFCFSHA groups, enter the following command:

```
# /opt/VRTS/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFSHA groups using the following commands for each group:

- Make the configuration read/write.

```
# /opt/VRTS/bin/haconf -makerw
```

- Unfreeze the group.

```
# /opt/VRTS/bin/hagrp -unfreeze group -persistent
```

- Save the configuration.

```
# /opt/VRTS/bin/haconf -dump -makero
```

- 13** If VVR is configured, and the CVM and SFCFSHA groups are offline, bring the groups online in the following order:

Bring online the CVM groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group -sys system01
# /opt/VRTS/bin/hagrp -online group -sys system02
```

where *group* is the VCS service group that has the CVMVolDg resource.

Bring online the RVGShared groups and the virtual IP on the master node using the following commands:

```
# hagrp -online RVGShared -sys masterhost
# hares -online ip_name
    masterhost
```

Bring online the SFCFSHA groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group -sys system01
# /opt/VRTS/bin/hagrp -online group -sys system02
```

where *group* is the VCS service group that has the CFMount resource.

If the SFCFSHA service groups do not come online then your file system could be dirty.

- 14** Upgrade the remaining nodes.

Note: If you upgrade to SFCFSHA 6.0 and the file systems are dirty, you have to deport the shared disk group and import it as non-shared. After the import, run `fsck`. `fsck` should succeed. Then deport the disk group and import it back as shared.

- 15** Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

- 16** On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

Performing a rolling upgrade

This chapter includes the following topics:

- [Performing a rolling upgrade using the installer](#)

Performing a rolling upgrade using the installer

Use a rolling upgrade to upgrade Veritas Storage Foundation Cluster File System High Availability to the latest release with minimal application downtime.

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 and VCS agent packages in phase 2.

Note: You need to perform a rolling upgrade on a completely configured cluster.

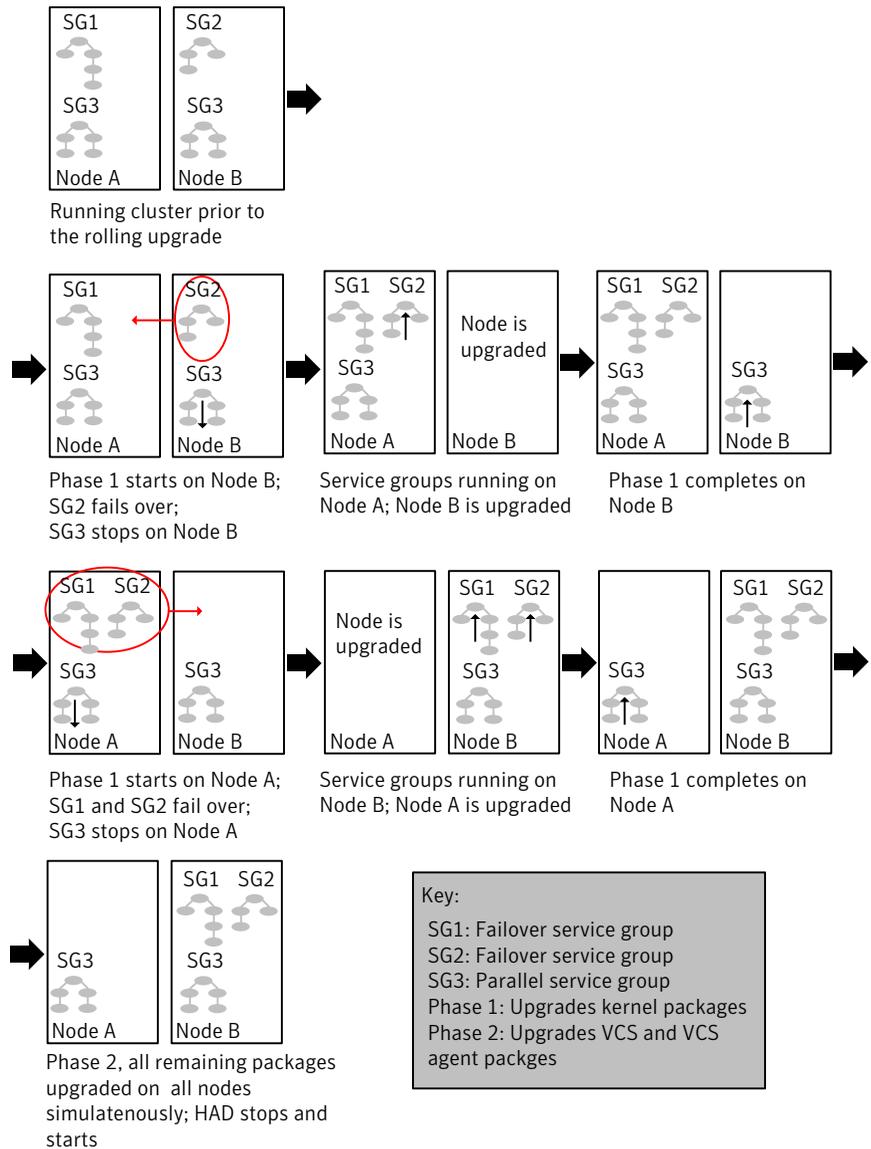
The following is an overview of the flow for a rolling upgrade:

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.

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 Veritas Cluster Server (VCS) engine HAD, but does not include application downtime.

[Figure 13-1](#) illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 13-1 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 script-based installer

Before you start the rolling upgrade, make sure that Veritas Cluster Server (VCS) is running.

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.0, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in [Patch Lookup](#) on the [SORT](#) website.

To perform a rolling upgrade

- 1 Complete the preparatory steps on the first sub-cluster.
- 2 Complete updates to the operating system, if required. For instructions, see the operating system documentation.

The nodes are restarted after the operating system update.

If you are performing a major OS upgrade (such as Solaris 9 to Solaris 10), do the following:

- Mount the software disc for the previous version of SFCFSHA.
- From the top level disc directory, start the installer.
- Uninstall the SFCFSHA packages.
- To finish the uninstallation, start the installer again, and install the packages for the previous version of SFCFSHA on the upgraded OS.
When the installation completes, the CPI prompts you to reboot the nodes again.

Restart the nodes again manually. Failing to perform this additional reboot prevents the upgrade from proceeding further.

```
# shutdown -g0 -y -i6
```

- 3 Log in as superuser and mount the SFCFSHA Veritas Storage Foundation Cluster File System High Availability 6.0 installation media.
- 4 Complete the preparatory steps again for the first sub-cluster.
- 5 From root, start the installer.

```
# ./installer
```

- 6 From the menu select **G) Upgrade a product**; then select **2)Rolling Upgrade**;
- 7 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 8 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.
- 9 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.
- 10 Review the EULA, and type **y** if you agree to its terms.
- 11 The installer prompts you to stop the applicable processes. Type **y** to continue.
 The installer fails over failover service groups to the node or nodes that are not upgraded at this time. The downtime is the time that it normally takes for the service group's failover. 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.
 It performs the configuration for the upgrade and re-starts processes.
 In case of failure in the startup of some of the processes, you may need to reboot the nodes and manually check the cluster's status.
- 13 Complete the preparatory steps on the nodes that you have not yet upgraded.
- 14 Complete updates to the operating system, if required, on the nodes that you have not yet upgraded. For instructions, see the operating system documentation.
 The nodes are restarted after the operating system update.
 If you are performing a major OS upgrade (such as Solaris 9 to Solaris 10), do the following:
 - Mount the software disc for the previous version of SFCFSHA.
 - From the top level disc directory, start the installer.
 - Uninstall the SFCFSHA packages.
 - To finish the uninstallation, start the installer again, and install the packages for the previous version of SFCFSHA on the upgraded OS.

When the installation completes, the CPI prompts you to reboot the nodes again.

Restart the nodes again manually. Failing to perform this additional reboot prevents the upgrade from proceeding further.

```
# shutdown -g0 -y -i6
```

- 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 reboots nodes, restart the installer.

The installer repeats step 8 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.

If the boot disk is encapsulated, you need to reboot nodes when phase 1 completes.

- 17 The installer determines the remaining packages to upgrade. Press **Enter** to continue.

- 18 The installer stops Veritas Cluster Server (VCS) processes. Type **y** to continue.

The installer performs prechecks, uninstalls old packages, and installs the new packages. It performs post-installation tasks, and the configuration for the upgrade.

- 19 Type **y** or **n** to help Symantec improve the automated installation.

- 20 If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

- 21 Upgrade application.

- 22 To upgrade VCS or Storage Foundation High Availability (SFHA) on the Coordination Point (CP) server systems to version 6.0, upgrade all the application clusters to 6.0. You then upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, refer to the appropriate installation guide.

Performing a rolling upgrade of SFCFSHA using the Web-based installer

This section describes using the Veritas 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 208.

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.0, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in [Patch Lookup](#) on the [SORT](#) website.

To start the rolling upgrade—phase 1

- 1 Complete the preparatory steps on the first sub-cluster.
- 2 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 3 Complete updates to the operating system, if required. For instructions, see the operating system documentation.

The nodes are restarted after the operating system update.

If you are performing a major OS upgrade (such as Solaris 9 to Solaris 10), do the following:

- Mount the software disc for the previous version of SFCFSHA.
- From the top level disc directory, start the installer.
- Uninstall the SFCFSHA packages.
- To finish the uninstallation, start the installer again, and install the packages for the previous version of SFCFSHA on the upgraded OS.
When the installation completes, the CPI prompts you to reboot the nodes again.

Restart the nodes again manually. Failing to perform this additional reboot prevents the upgrade from proceeding further.

```
# shutdown -g0 -y -i6
```

- 4 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 63.

- 5 In the Task pull-down menu, select G) Upgrade a product; then select 2) Rolling Upgrade;.

Click the **Next** button to proceed.
- 6 Review the systems that the installer has chosen to start the rolling upgrade. These systems are chosen to minimize downtime during the upgrade.

Click **Yes** to proceed.

The installer validates systems. If it throws an error, address the error and return to the installer.
- 7 Review the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.
- 8 The installer stops all processes. Click **Next** to proceed.
- 9 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. The installer starts all the relevant processes and brings all the service groups online.

If the installer reboots nodes, restart the installer.
- 10 Complete the preparatory steps on the nodes that you have not yet upgraded.
- 11 Complete updates to the operating system on the nodes that you have not yet upgraded. For instructions, see the operating system documentation.

The nodes are restarted after the operating system update.

If you are performing a major OS upgrade (such as Solaris 9 to Solaris 10), do the following:

 - Mount the software disc for the previous version of SFCFSHA.
 - From the top level disc directory, start the installer.
 - Uninstall the SFCFSHA packages.
 - To finish the uninstallation, start the installer again, and install the packages for the previous version of SFCFSHA on the upgraded OS.

When the installation completes, the CPI prompts you to reboot the nodes again.

Restart the nodes again manually. Failing to perform this additional reboot prevents the upgrade from proceeding further.

```
# shutdown -g0 -y -i6
```

- 12 When prompted, perform step 5 through step 9 on the nodes that you have not yet upgraded.
- 13 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 Veritas Web-based installer”](#) on page 63.

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 and stops processes. 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 validates systems. If it throws an error, address the error and return to the installer. Click **Next** to proceed.
- 5 The installer stops all processes. Click **Next** to proceed.
- 6 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. The installer starts all the relevant processes and brings all the service groups online.
- 7 If you have network connection to the Internet, the installer checks for updates. If updates are discovered, you can apply them now.
- 8 Upgrade application.

The upgrade is complete.

Performing a phased upgrade

This chapter includes the following topics:

- [Performing a phased upgrade of SFCFSHA](#)

Performing a phased upgrade of SFCFSHA

Performing a phased upgrade involves the following tasks:

- Moving the service groups to the second subcluster
- Upgrading the SFCFSHA stack on the first subcluster
- Preparing the second subcluster
- Activating the first subcluster
- Upgrading the operating system on the second subcluster
- Upgrading the second subcluster
- Finishing the phased upgrade

Before you start the upgrade on the first half of the cluster, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade. Note that your applications have downtime during this procedure.

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 basic guidelines follow:

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

Phased upgrade limitations

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

- While you perform the upgrades, do not start any modules. Also, do not add or remove service groups to any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.
- You can perform a phased upgrade when the root disk is encapsulated.

Moving the service groups to the second subcluster

To move the service groups to the second subcluster

- 1 Switch failover groups from the first half of the cluster to one of the nodes in the second half of the cluster. In this procedure, `galaxy` is a node in the first half of the cluster and `jupiter` is a node in the second half of the cluster. Enter the following:

```
# hagrps -switch failover_group -to jupiter
```

- 2 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the applications.

- 3 On the first half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS.

```
# mount -p | grep vxfs
```

Verify that no processes use the VxFS or CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS or CFS mount point with the mechanism provided by the application.

Unmount the VxFS or CFS file system. Enter the following:

```
# umount /mount_point
```

- 4 On the first half of the cluster, bring all the VCS service groups offline including CVM group. Enter the following:

```
# hagrps -offline group_name -sys galaxy
```

When the CVM group becomes OFFLINE, all the parallel service groups such as the CFS file system will also become OFFLINE on the first half of the cluster nodes.

- 5 Verify that the VCS service groups are offline on all the nodes in first half of the cluster. Enter the following:

```
# hagrps -state group_name
```

- 6 Freeze the nodes in the first half of the cluster. Enter the following:

```
# haconf -makerw  
# hasys -freeze -persistent galaxy  
# haconf -dump -makero
```

- 7 If the cluster-wide attribute UseFence is set to SCSI3, then reset the value to NONE in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

- 8 Verify that only GAB ports a, b, d and h are open. Enter the following:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port h gen ada40f membership 01
```

Do not stop VCS. Port h should be up and running.

- 9 In the first half of the cluster, stop all VxVM and CVM volumes. Enter the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open. Enter the following:

```
# vxprint -Aht -e v_open
```

- 10 On first half of the cluster, upgrade the operating system on all the nodes, if applicable. For instructions, see the upgrade paths for the operating system. See [“Supported upgrade paths”](#) on page 184.

Upgrading the SFCFSHA stack on the first subcluster

To upgrade the SFCFSHA stack on the first subcluster

- ◆ **Note:** This procedure is based on an "in-place" upgrade path; that is, if the operating system is upgraded, the release will be the same, but only the path level will change. If you are moving from major operating system release to another, you must uninstall the SFCFSHA stack before you upgrade the operating system. After the operating system is upgraded, you must reinstall SFCFSHA.
-

On the first half of the cluster, upgrade SFCFSHA by using the `installsfcfs` script. For example use the `installsfcfs` script as shown below:

```
# ./installsfcfs galaxy
```

After the upgrade for first half of the cluster is complete, no GAB ports will be shown in `gabconfig -a` output.

To upgrade your operating system, follow the normal procedures for your platform.

Note: After the installation completes, you can safely ignore any instructions that the installer displays.

Preparing the second subcluster

To prepare the second subcluster

- 1 On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application. [Downtime starts now.]

- 2 On the second half of the cluster, unmount the VxFS and CFS file systems that are not managed by VCS. Enter the following:

```
# mount -p | grep vxfs
```

Verify that no processes use the VxFS and CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS and CFS mount point with the mechanism provided by the application.

Unmount the VxFS and CFS file system. Enter the following:

```
# umount /mount_point
```

- 3 On the second half of the cluster, unfreeze all the VCS service groups on all the nodes using the following commands:

```
# haconf -makerw  
# hagr -unfreeze group_name -persistent  
# haconf -dump -makero
```

- 4 On the second half of the cluster, bring all the VCS service groups offline, including CVM group. Enter the following:

```
# hagr -offline group_name -sys jupiter
```

- 5 On the second half of the cluster, verify that the VCS service groups are offline. Enter the following:

```
# hagr -state group_name
```

- 6 Stop VCS on the second half of the cluster. Enter the following:

```
# hastop -local
```

- 7 On each node of the second half of the cluster, change the contents of the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cp /etc/vxfen.d/vxfenmode_disabled /etc/vxfenmode
# cat /etc/vxfenmode#
# 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=disabled
```

- 8 If the cluster-wide attribute `UseFence` is set to `SCSI3`, reset the value to `NONE` in the `/etc/VRTSvcs/conf/config/main.cf` file, in second half of the cluster.
- 9 On the second half on cluster, stop the following SFCFSHA modules: `VCS`, `VxFEN`, `ODM`, `GAB`, and `LLT`. Enter the following:

```
# /etc/init.d/vxglm stop
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- 10 If the cluster-wide attribute `UseFence` is set to `NONE`, reset the value to `SCSI3` in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

Activating the first subcluster

To activate the first subcluster

- 1 Restart the upgraded nodes in the first half of the cluster:

```
# /usr/sbin/shutdown -i6 -g0 -y
```

When the first half of the cluster nodes come up, no GAB ports are OPEN. The following command does not show any GAB ports:

```
# /sbin/gabconfig -a
GAB Port Memberships
```

```
=====
```

- 2 If required, force gab to form a cluster after the upgraded nodes are rebooted in first half of the cluster.

```
# /sbin/gabconfig -x
```

GAB ports a, b, d and h appear in `gabconfig -a` command output.

Note: If port b and h are not up, you need to bring fencing and VCS manually online.

- 3 For nodes that use Solaris 10, start VCS in first half of the cluster:

```
# svcadm enable system/vcs
```

- 4 On the first half of the cluster, bring the VCS service groups online. Enter the following:

```
# hagrps -online group_name -sys node_name
```

After you bring the CVM service group ONLINE, all the GAB ports u, v, w and f come ONLINE and all the CFS mounts service groups also come ONLINE automatically. Only failover service groups need to be brought ONLINE manually.

- 5 Manually mount the VxFS and CFS file systems that are not managed by VCS. [Downtime ends now.]

Upgrading the operating system on the second subcluster

To upgrade the operating system on the second subcluster

- ◆ Enter the following.

```
# /etc/init.d/vxglm stop
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

On the second half of the cluster, upgrade the operating system, if applicable. For instructions, see the upgrade paths for the operating system.

Upgrading the second subcluster

To upgrade the second subcluster

- ◆ Enter the following:

```
# ./installsfcfsha node_name
```

Completing the phased upgrade

To complete the phased upgrade

- 1 Verify that the cluster UUID on the nodes in the second subcluster is the same as the cluster UUID on the nodes in the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -display nodename
```

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 Restart the upgraded nodes in the second half of the cluster:

```
# /usr/sbin/shutdown -i6 -g0 -y
```

When second half of the nodes come up, all the GAB ports a, b, d, h, u, v, w and f are ONLINE. Also all the CFS mounts service groups come online automatically.

- 3 For nodes that use Solaris 10, start VCS in first half of the cluster:

```
# svcadm enable system/vcs
```

- 4 Manually mount the VxFS and CFS file systems that are not managed by VCS in the second half of the cluster.

- 5 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

- 6 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

Upgrading the operating system

This chapter includes the following topics:

- [Upgrading the Solaris operating system](#)

Upgrading the Solaris operating system

If you are running Veritas Storage Foundation Cluster File System High Availability 6.0 with an earlier release of the Solaris operating system, you can upgrade the Solaris operating system using the following procedure.

Warning: You should only use this procedure to upgrade the Solaris operating system if you are running Veritas Storage Foundation Cluster File System High Availability 6.0.

The directory `/opt` must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by the `upgrade_start` are unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

To upgrade the Solaris operating system only

- 1 Bring the system down to single-user mode using the following command:

```
# init s
```

You must mount `/opt` manually if `/opt` is on its own partition.

- 2 Load and mount the software disc from the currently installed version of Veritas Storage Foundation Cluster File System High Availability.

See [“Mounting the product disc”](#) on page 55.

- 3 Change directory:

```
# cd /mount_point/scripts
```

- 4 Run the `upgrade_start` with the `-check` argument to detect any problems that exist which could prevent a successful upgrade. Use the `upgrade_start` script that was supplied with the currently installed SF release. If this command reports success, you can proceed with running the `upgrade_start` script, but if it reports errors, correct the problem(s) and rerun `upgrade_start -check`.

```
# ./upgrade_start -check
```

- 5 Run the `upgrade_start` script so that the system can come up with partitions. The `upgrade_start` script searches for volumes containing file systems, and if any are found, converts them to partitions:

```
# ./upgrade_start
```

- 6 Bring the system down to run level 0.

```
# init 0
```

- 7 Upgrade the operating system to a supported version of Solaris.

You should boot up the system from run level 0 depending on the Solaris upgrade procedure that you want to follow. Refer to the Solaris installation documentation for instructions on how to upgrade the Solaris operating system.

- 8 After installing the Solaris operating system, install any Solaris patches required by Veritas Storage Foundation Cluster File System High Availability 6.0.

See the *Veritas Storage Foundation Cluster File System High Availability Release Notes*.

- 9 After the system is up with the upgraded Solaris operating system, bring the system down to single-user mode by entering:

```
# init s
```

- 10 Ensure that `/opt` is mounted.
- 11 Load and mount the software disc from the currently installed version of Veritas Storage Foundation Cluster File System High Availability.
- 12 If you upgraded to Solaris 10, you must reinstall certain Veritas Storage Foundation Cluster File System High Availability packages in order to support Solaris 10 functionality.

To reinstall the required packages, follow the steps below:

- Remove the existing packages in the reverse order of their installation. For example, if you chose the installation of all packages and patches then uninstall those in the following order.

```
# pkgrm VRTSsfmh VRTSodm VRTSgms  
VRTScavf VRTSglm VRTSdbed VRTSvcsea VRTSvcscag  
VRTScps VRTSvcsc VRTSamf VRTSvxfen VRTSgab  
VRTSllt VRTSat VRTSfssdk VRTSvxfs VRTSob  
VRTSaslapm VRTSvxvm VRTSspt VRTSperl VRTSvlic
```

- Run the following commands.
To obtain a list of recommended packages to install:

```
# ./installsfcfsha -recpkgs
```

Or

To obtain a list of all packages to install:

```
# ./installsfcfsha -allpkgs
```

- Change to the directory containing the appropriate packages.

```
# cd /mount_point/pkg
```

- Use the `pkgadd` command to install the packages from the list you generated.

- Reboot the system.
- 13 Complete the upgrade from the software disc from the currently installed version of Storage Foundation by entering:

```
# devlinks  
# ./upgrade_finish
```

Upgrading Veritas Volume Replicator

This chapter includes the following topics:

- [Upgrading Veritas Volume Replicator](#)

Upgrading Veritas Volume Replicator

If a previous version of Veritas Volume Replicator (VVR) is configured, the product installer upgrades VVR automatically when you upgrade the Storage Foundation products.

When upgrading from 4.1 MP1 or later, you have the option to upgrade without disrupting replication.

See [“Upgrading VVR without disrupting replication”](#) on page 231.

Upgrading VVR without disrupting replication

This section describes the upgrade procedure from an earlier version of VVR to the current version of VVR when replication is in progress, assuming that you do not need to upgrade all the hosts in the RDS simultaneously.

You may also need to set up replication between versions.

See [“Planning an upgrade from the previous VVR version”](#) on page 190.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade can be performed either on the Primary or on the Secondary. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS. This section includes separate sets of steps, for the Primary upgrade and for the Secondary upgrade.

Note: If you have a cluster setup, you must upgrade all the nodes in the cluster at the same time.

Upgrading VVR on the Secondary

Follow these instructions to upgrade the Secondary hosts.

To upgrade the Secondary

- 1 Stop replication to the Secondary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvgname
```

- 2 Upgrade from VVR after 4.1 MP1 and prior to 6.0 on the Secondary.

- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname sec_hostname
```

Upgrading VVR on the Primary

After you upgrade the Secondary, use the Veritas product installer to upgrade the Primary.

To upgrade the Primary

- 1 Stop replication to the Primary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvgname
```

- 2 Upgrade from VVR after 4.1 MP1 and prior to 6.0 on the Secondary.

- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname  
sec_hostname
```

See [“Planning an upgrade from the previous VVR version”](#) on page 190.

Upgrading language packages

This chapter includes the following topics:

- [Upgrading language packages](#)

Upgrading language packages

If you are upgrading Veritas products in a language other than English, you must install the required language packages after installing the English packages. Verify that the English installation is correct before proceeding.

Install the language packages as for an initial installation.

See [“Installing language packages”](#) on page 60.

Migrating from SFHA to SFCFSHA

This chapter includes the following topics:

- [Migrating from SFHA to SFCFSHA 6.0](#)

Migrating from SFHA to SFCFSHA 6.0

This section describes how to migrate Storage Foundation High Availability (SFHA) 6.0 to Storage Foundation Cluster File System High Availability (SFCFSHA) 6.0.

The product installer does not support direct upgrades from a previous version of SFHA or SFCFSHA6.0. Ensure that you upgrade the existing SFHA to 6.0 before beginning this procedure.

To migrate from SFHA 6.0 to SFCFSHA 6.0

- 1 Back up the `main.cf` file before beginning the upgrade.
- 2 Confirm that the storage disks are visible on all the nodes in the 6.0 SFHA cluster.
- 3 Bring all the failover service groups offline, using the following command:

```
# hagrps -offline group_name -any
```

The above command brings the service group offline on the node where the service group is currently online.

- 4 Unmount all the VxFS file systems which are not under VCS control. If the local file systems are under VCS control, then VCS unmounts the file systems when the failover service group is brought offline in step 3.

On the nodes that have any mounted VxFS local file systems that are not under VCS control:

```
# umount -F vxfs -a
```

- 5 Stop all of the activity on the volumes and deport the local disk groups. If the local disk groups are part of VCS failover service groups, then VCS deports the disk groups when the failover service group is brought offline in step 3.

```
# vxvol -g diskgroup_name stopall  
# vxdg deport diskgroup_name
```

- 6 Upgrade the existing SFHA to SFCFSHA 6.0:

```
# ./installsfcfsha
```

- 7 After installation is completed, the install script asks you to install licenses. Enter the correct license key to register the key.
- 8 The installer prompts to reconfigure the VCS. Provide the same cluster name, cluster ID, and LLT link interfaces details that were used during configuration of the SFHA cluster.
- 9 Find out which node is the CVM master, using the following command:

```
# vxdctl -c mode
```

- 10 On the CVM Master node, re-import all the required disk groups which must be in shared mode:

```
# vxdg -s import diskgroup_name
```

- 11 Start all the volumes whose disk groups have been imported as shared in step 10. Use the following command:

```
# vxdg -g diskgroup_name startall
```

- 12 Run the following command for each of the file systems you want to mount as CFS:

```
# cfsmntadm add diskgroup_name volume_name mount_point \  
all=cluster_mount_options
```

- 13 Run the following command to mount CFS file systems on all the nodes:

```
# cfsmount mount_point
```

- 14 Import all other local disk groups which have not been imported in shared mode in step 10.

```
# vxdg import diskgroup_name
```

Start all the volumes of these disk groups using:

```
# vxvol -g diskgroup_name startall
```

Mount these volumes.

- 15 For any of the file systems which VCS needs to monitor through failover service groups, create these failover service groups by adding the Mount, Diskgroup & Volume resources for VxFS file systems under VCS control.

Upgrading with Live Upgrade

This chapter includes the following topics:

- [About Live Upgrade](#)
- [Supported upgrade paths for Live Upgrade](#)
- [Performing Live Upgrade in a Solaris zone environment](#)
- [Before you upgrade SFCFSHA using Solaris Live Upgrade](#)
- [Upgrading SFCFSHA and Solaris using Live Upgrade](#)
- [Upgrading Solaris using Live Upgrade](#)
- [Upgrading SFCFSHA using Live Upgrade](#)
- [Administering boot environments](#)

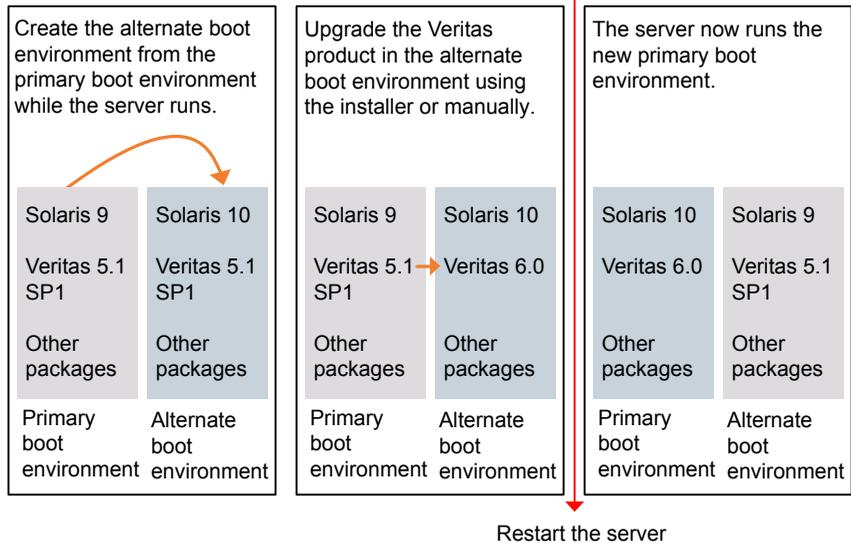
About Live Upgrade

You can use Live Upgrade to perform the following types of upgrade:

- Upgrade the operating system and SFCFSHA.
See [“Upgrading SFCFSHA and Solaris using Live Upgrade”](#) on page 245.
- Upgrade the operating system.
See [“Upgrading Solaris using Live Upgrade”](#) on page 253.
- Upgrade SFCFSHA.
See [“Upgrading SFCFSHA using Live Upgrade”](#) on page 254.

Figure 19-1 illustrates an example of an upgrade of Veritas products from 5.1 SP1 to 6.0, and the operating system from Solaris 9 to Solaris 10.

Figure 19-1 Live Upgrade process



Some service groups (failover and parallel) may be online in this cluster and they are not affected by the Live Upgrade process. The only downtime experienced is when the server is rebooted to boot into the alternate boot disk.

About Live Upgrade in a Veritas Volume Replicator (VVR) environment

In a SFCFSHA environment that uses Veritas Volume Replicator, the following scripts provide the means to upgrade the VVR configuration:

- `vvr_upgrade_lu_start`
- `vvr_upgrade_lu_finish`

This section provides an overview of the VVR upgrade process. See the Live Upgrade procedures for SFCFSHA for the complete procedure.

See [“Upgrading SFCFSHA and Solaris using Live Upgrade”](#) on page 245.

- Use the `vxlustart` script to perform upgrade steps for SFCFSHA.
- Immediately before rebooting the system to switch over to the alternate boot environment, run the `vvr_upgrade_lu_start` script.

Note: Use the `vvr_upgrade_lu_start` script only when the applications are stopped and the next step is to switch over to the alternate boot environment.

- After the `vvr_upgrade_lu_start` script completes successfully, reboot the system. This reboot results in the system booting from the alternate boot environment.
- After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

Supported upgrade paths for Live Upgrade

The systems where you plan to use Live Upgrade must run Solaris 9 or Solaris 10. You can upgrade from systems that run Solaris 9, but SFCFSHA 6.0 is not supported on Solaris 9.

For Solaris 10, make sure that all non-global zones are booted and in the running state before you use the Symantec product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must upgrade each package in each non-global zone manually.

For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you reboot the alternative root, you can install `VRTSodm`.

SFCFSHA version must be at least 5.0 MP3.

Symantec requires that both global and non-global zones run the same version of Veritas products.

Note: If you use Live Upgrade on a system where non-global zones are configured, make sure that all the zones are in the `installed` state before you start Live Upgrade.

You can use Live Upgrade in the following virtualized environments:

Table 19-1 Live Upgrade support in virtualized environments

Environment	Procedure
Solaris native zones	<p>Perform Live Upgrade to upgrade both global and local zones.</p> <p>If you have a zone root that resides on a VxVM volume, use the following procedure.</p> <p>See “Performing Live Upgrade in a Solaris zone environment” on page 240.</p> <p>Use the standard procedure for the other standby nodes.</p> <p>See “Upgrading SFCFSHA and Solaris using Live Upgrade” on page 245.</p>
Solaris branded zones (BrandZ)	<p>Perform Live Upgrade to upgrade the global zone.</p> <p>See “Upgrading SFCFSHA and Solaris using Live Upgrade” on page 245.</p> <p>Manually upgrade the branded zone separately.</p> <p>Note that while you can perform a Live Upgrade in the presence of branded zones, the branded zones are not upgraded.</p>
Oracle VM Server for SPARC	<p>Perform Live Upgrade on the Control domain only.</p> <p>Perform Live Upgrade on the Guest domain only.</p> <p>Use the standard Live Upgrade procedure for both types of logical domains.</p> <p>See “Upgrading SFCFSHA and Solaris using Live Upgrade” on page 245.</p>

Performing Live Upgrade in a Solaris zone environment

If you have a zone root that resides on a VxVM volume, then you must use the following procedure to perform a Live Upgrade on the nodes where zones are online.

Use the standard procedure for the other standby nodes.

See [“Upgrading SFCFSHA and Solaris using Live Upgrade”](#) on page 245.

To perform a Live Upgrade on a node that has a zone root on a VxVM volume

- 1 Unmount all file systems that do not contain local zone root on shared storage.
- 2 Shut down any application that runs on local zone. Take its resources offline and leave only the zone running.

By default, Zone agent `BootState` is set to "multi-user." After you complete the upgrade, you may need to adjust this attribute to the appropriate value before you start your zone through VCS.

Note: Symantec recommends that you set `BootState` to "multi-user-server" to run applications inside non-global zones.

- 3 Freeze the service group that contains the local zone. Make sure that the boot environment disk has enough space for local zone root being copied over during the Live Upgrade.
- 4 Follow the instruction to upgrade using Live Upgrade (which includes `vxlustart`, the product upgrade, and `vxlufinish`).

Before rebooting the systems to complete the Live Upgrade, perform the following steps.

- 5 On the system that houses the local zone, copy all files and directories before the upgrade on the local zone root on shared storage to another location.

```
# zoneadm list -cv
  ID NAME           STATUS    PATH                BRAND  IP
   0 global           running   /                   native  shared
   6 ora-lzone       running   /oralzones          native  shared

# zoneadm -z ora-lzone halt
# cd /oralzones
# ls
dev  lost+found root SUNWattached.xml
# mv dev dev.41
# mv root root.41
# mv SUNWattached.xml SUNWattached.xml.41
```

- 6 Migrate all files and directories after the upgrade on the local zone root on BE to the shared storage using the tar utility:

```
# cd /altroot.5.10/oralzones
# ls
dev lost+found lu root SUNWattached.xml
# tar cf - . | (cd /oralzones; tar xfbp -)
# cd /oralzones
# ls
dev .41 lost+found root.41 SUNWattached.xml.41
dev lost+found lu root SUNWattached.xml
```

- 7 Unfreeze the service group that contains the local zone.
- 8 Shut down all systems.

Before you upgrade SFCFSHA using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

To prepare for the Live Upgrade

- 1 Make sure that the SFCFSHA 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 could potentially cause a root file system to run out of space.
- 4 On the primary boot disk, patch the operating system for Live Upgrade. Patch 137477-01 is required. Verify that this patch is 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 are upgrading 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.

Note: While you can perform Live Upgrade in the presence of branded zones, they must be halted, and the branded zones themselves are not upgraded.

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
```

- 6 Symantec provides the `vxlustart` script that runs a series of commands to create the alternate boot disk for the upgrade.

To preview the commands, specify the `vxlustart` script with the `-v` option.

Symantec recommends that you preview the commands to ensure there are no problems before beginning the Live Upgrade process.

The `vxlustart` script is located on the distribution media, in the `scripts` directory.

```
# cd /cdrom/scripts
```

```
# ./vxlustart -V -u targetos_version -s osimage_path -d diskname
```

- 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 being upgraded, the user will be prompted to compare the patches that are installed on the image with the patches installed on the primary boot disk to determine if any critical patches are missing from the new operating system image.
- u Specifies the operating system version for the upgrade on the alternate boot disk. For example, use `5.10` for Solaris 10.
- U Specifies that only the Storage Foundation products are upgraded. The operating system is cloned from the primary boot disk.
- s Indicates the path of the operating system image to be installed on the alternate boot disk. If this option is omitted, you are prompted to insert the discs that contain the operating system image.

If the `-U` option is specified, you can omit the `-s` option. The operating system is cloned from the primary boot disk.
- d Indicates the name of the alternate boot disk on which you intend to upgrade. If you do not specify this option with the script, you are prompted for the disk information.
- v Indicates verbose, the executing commands display before they run.
- Y Indicates a default yes with no questions asked.
- D Prints with debug option on, and is for debugging.
- F Specifies the rootdisk's file system, where the default is `ufs`.
- t Specifies the number of CDs involved in upgrade.
- r Specifies that if the machine crashes or reboots before the `vxlufinish` command is run, the alternate disk is remounted using this option.

For example, to preview the commands to upgrade only the Veritas 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 c0t1d0
```

Note: This command prompts you to compare the patches that are installed on the image with the patches installed on the primary boot disk. If any patches are missing from the new operating system's image, note the patch numbers. To ensure the alternate boot disk is the same as the primary boot disk, you will need to install these patches on the alternate boot disk.

- 7 If the specified image is 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.

In the procedure examples, the primary or current boot environment resides on Disk0 (c0t0d0) and the alternate or inactive boot environment resides on Disk1 (c0t1d0).

Upgrading SFCFSHA and Solaris using Live Upgrade

Upgrading SFCFSHA using Live Upgrade involves the following steps:

- Prepare to upgrade using Solaris Live Upgrade.
See [“Before you upgrade SFCFSHA using Solaris Live Upgrade”](#) on page 242.
- Create a new boot environment on the alternate boot disk.
See [“Creating a new boot environment on the alternate boot disk”](#) on page 246.
- Upgrade to Veritas Storage Foundation Cluster File System High Availability 6.0 on the alternate boot environment manually or using the installer.
To upgrade SFCFSHA manually, refer to the following procedure:
 - See [“Upgrading SFCFSHA manually”](#) on page 248.To upgrade SFCFSHA using the installer, refer to the following procedure:
 - See [“Upgrading SFCFSHA using the installer for a Live Upgrade”](#) on page 247.
- Switch the alternate boot environment to be the new primary.
See [“Completing the Live Upgrade”](#) on page 251.

- Verify Live Upgrade of SFCFSHA.
See [“Verifying Live Upgrade of SFCFSHA”](#) on page 252.

Creating a new boot environment on the alternate boot disk

Run the `vxlustart` command on each node in the cluster to create a new boot environment on the alternate boot disk.

Note: This step can take several hours to complete. Do not interrupt the session as it may leave the boot environment unstable.

At the end of the process:

- The Solaris operating system on the alternate boot disk is upgraded, if you have chosen to upgrade the operating system.
- A new boot environment is created on the alternate boot disk by cloning the primary boot environment.

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 View the list of VxVM disks on which you want to create the new boot environment.

```
# vxdisk list
```

- 3 Run one of the following commands to perform the upgrade:

To upgrade the operating system, by itself or together with upgrading the Veritas products:

```
# ./vxlustart -v -u targetos_version \  
-s osimage_path -d disk_name
```

Where *targetos_version* is the version of the operating system

osimage_path is the full path to the operating system image

disk_name is the name of the disk as displayed in the output of step 2.

To upgrade the Veritas product only:

```
# ./vxlustart -v -u 5.10 -U -d disk_name
```

The options to the `vxlustart` command are listed in the preupgrade section.

See “[Before you upgrade SFCFSHA using Solaris Live Upgrade](#)” on page 242.

For example, to upgrade to Solaris 10 update 6:

```
# ./vxlustart -v -u 5.10 -s /mnt/Solaris_10u6
```

- 4 Review the output and note the new mount points. If the system is rebooted 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
```

- 5 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 Veritas product installation:

```
# pkgadd -R /altroot.5.10 -d pkg_dir
```

Upgrading SFCFSHA using the installer for a Live Upgrade

You can use the Veritas product installer to upgrade SFCFSHA as part of the Live Upgrade.

On a node in the cluster, run the installer on the alternate boot disk to upgrade SFCFSHA on all the nodes in the cluster. The program uninstalls the existing version of SFCFSHA on the alternate boot disk during the process.

At the end of the process the following occurs:

- Veritas Storage Foundation Cluster File System High Availability 6.0 is installed on the alternate boot disk.

To perform Live Upgrade of SFCFSHA using the installer

- 1 Insert the product disc with Veritas Storage Foundation Cluster File System High Availability 6.0 or access your copy of the software on the network.
- 2 Run the installer script specifying the root path as the alternate boot disk, enter the following:

```
# ./installsfcfsha -upgrade -rootpath /altroot.5.10
```

See [“Removing and reinstalling SFCFSHA using the installer”](#) on page 253.

- 3 Enter the names of the nodes that you want to upgrade to Veritas Storage Foundation Cluster File System High Availability 6.0.

Note: Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the SFCFSHA installation.

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 will not update the VCS configurations for Oracle, Netlsnr, and Sybase resources. If cluster configurations include these resources, you will be 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 6.0.

```
# pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSvxvm
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/logs`.

Upgrading SFCFSHA manually

You can perform a manual upgrade of SFCFSHA using Live Upgrade. On each node, remove and install the appropriate SFCFSHA packages.

Note: #cp /mnt/etc/VRTSvcs/conf/config/PrivNIC.cf /tmp/PrivNIC.cf.save

At the end of the process the following occurs:

- Veritas Storage Foundation Cluster File System High Availability 6.0 is installed on the alternate boot disk.

To perform Live Upgrade of SFCFSHA manually

- 1 Remove the SFCFSHA packages on the alternate boot disk in the following order:

```
# pkgrm -R /altroot.5.10 \  
VRTSodm VRTSgms VRTSvxmsa VRTScavf VRTSglm VRTSfsmnd \  
VRTSfssdk VRTSfsman VRTSvrv VRTSvcsvr VRTSvrpro \  
VRTSddlpr VRTSvdid VRTSalloc VRTSdcli VRTSvmpro \  
VRTSvmman VRTSfspro VRTSdsa VRTSvxvm VRTScmccc \  
VRTScmcs VRTSacclib VRTScssim VRTScscm VRTSweb \  
VRTScscw VRTScutil VRTSjrel5 VRTSvcsmn VRTSvcsag \  
VRTSvcsmg VRTSvcs VRTSvxfen VRTSgab VRTSllt \  
VRTSvxfs VRTSspt VRTSaa VRTSmh VRTSccg VRTSobgui \  
VRTSob VRTSobc33 VRTSat VRTSpbx VRTSicsco VRTSvlic \  
VRTSperl
```

The `-R` option removes the packages from the root path `/altroot.5.10` on the alternate boot disk.

Note that this package list is an example. Full package lists vary from release to release and by product option.

- 2 Install SFCFSHA 6.0 packages in the following order one at a time to the alternate boot disk using the `pkgadd` command:

```
VRTSvlic VRTSperl VRTSspt VRTSvxvm \  
VRTSaslapm VRTSob VRTSsfmh VRTSvxfs \  
VRTSfsadv VRTSfssdk VRTSllt VRTSgab \  
VRTSvxfen VRTSarf VRTSvcs VRTScps \  
VRTSvcsag VRTSvcssea VRTSdbed VRTSglm \  
VRTScavf VRTSgms VRTSodm VRTSsfcp160
```

For example:

```
# pkgadd -R /altroot.5.10 -d package_name.pkg
```

Where you replace `package_name.pkg` with a package's name, for example `VRTSvxvm.pkg`.

```
# pkgadd -R /altroot.5.10 -d VRTSvxvm.pkg
```

- 3 Verify that the version of the Veritas packages on the alternate boot disk is 6.0.

```
# pkginfo -R /altrootpath -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSvxvm
```

- 4 Set the `INSTALL_ROOT_PATH` environment variable to the root path, and then configure a VCS cluster UUID on the alternative root path. Enter the following commands:

```
# export INSTALL_ROOT_PATH=/altroot.5.10
# /altroot.5.10/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure \
-use_llthost
```

- 5 Confirm that you have created the Universal Unique Identifier for the cluster:

```
# /altroot.5.10/opt/VRTSvcs/bin/uuidconfig.pl -clus -display \
-use_llthost
```

- 6 In a zones or branded zones environment, perform the following steps to ensure that all non-global zones contain a universally unique identifier (UUID):

```
# zoneadm -z zone1 detach
# zoneadm -z zone1 attach
# zoneadm -z zone1 boot
# zoneadm list -p
0:global:running:/::native:shared
3:zone1:running:/zone1:3770b7b9-f96a-ef34-f4c5-bc125d56ec27:
native:shared
```

For a Solaris environment without zones, run the following command on the alternate root path of any one node in the cluster to configure a unique VCS cluster ID:

```
# /mnt/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure -use_llthost
```

The `-use_llthost` option indicates that the `/etc/llthost` file is used to determine the names of the nodes in the cluster. Alternatively, you can specify the node names instead of the file name.

Completing the Live Upgrade

At the end of the process:

- If the original primary boot disk was encapsulated, the alternate boot disk is encapsulated.
- 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 using one of the following commands. You must enter the command on all nodes in the cluster.

If the primary root disk is not encapsulated, run the following command:

```
# ./vxlufinish -u target_os_version  
Live Upgrade finish on the Solaris release <5.10>
```

If the primary root disk is encapsulated by VxVM, run the following command:

```
# ./vxlufinish -u target_os_version -g diskgroup  
Live Upgrade finish on the Solaris release <5.10>
```

The Live Upgrade process encapsulates the alternate root disk if the primary root disk was encapsulated.

- 2 If the system crashes or reboots before Live Upgrade completes successfully, you can remount the alternate disk using the following command:

```
# ./vxlustart -r -u target_os_version
```

Then, rerun the vxlufinish command from step 1

```
# ./vxlufinish -u target_os_version
```

- 3 If you are upgrading VVR, run the `vvr_upgrade_lu_start` command.

Note: Only run the `vvr_upgrade_lu_start` command when you are ready to reboot the nodes and switch over to the alternate boot environment.

-
- 4 **Note:** Do not use the `reboot`, `halt`, or `uadmin` commands to reboot 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
```

- 5 After the alternate boot environment is activated, you can switch boot environments. If the root disk is encapsulated, refer to the procedure to switch the boot environments manually.
- See [“Administering boot environments”](#) on page 255.
- 6 After the upgrade, perform any required post-upgrade tasks such as upgrading the disk group.
- 7 After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.
- 8 If you want to upgrade CP server systems that use VCS or SFHA to this version, make sure that you 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 VCS or SFHA Installation Guide.

Verifying Live Upgrade of SFCFSHA

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 is not active, you can revert to the primary boot environment.

See [“Reverting to the primary boot environment”](#) on page 255.

- 2 In a cluster environment, make sure that all the GAB ports are up. Note different ports appear for different products.

- 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.

Upgrading Solaris using Live Upgrade

If you are upgrading Solaris only, you must remove and reinstall SFCFSHA from the alternate boot environment prior to completing the Live Upgrade. You must remove and reinstall because SFCFSHA has kernel components that are specific to Solaris operating system versions. The correct version of the SFCFSHA packages must be installed.

Upgrading Solaris using Live Upgrade involves the following steps:

- Preparing to upgrade using Solaris Live Upgrade.
See [“Before you upgrade SFCFSHA using Solaris Live Upgrade”](#) on page 242.
- Creating a new boot environment on the alternate boot disk
See [“Creating a new boot environment on the alternate boot disk”](#) on page 246.
- Removing and reinstalling Veritas Storage Foundation Cluster File System High Availability 6.0 on the alternate boot environment:
Using manual steps:
See [“Upgrading SFCFSHA manually”](#) on page 248.
Using the installer:
See [“Removing and reinstalling SFCFSHA using the installer”](#) on page 253.

Note: Do NOT configure the Veritas Storage Foundation Cluster File System High Availability 6.0

- Switching the alternate boot environment to be the new primary
See [“Completing the Live Upgrade ”](#) on page 251.
- Verifying Live Upgrade of SFCFSHA.
See [“Verifying Live Upgrade of SFCFSHA”](#) on page 252.

Removing and reinstalling SFCFSHA using the installer

SFCFSHA has kernel components that are specific for Solaris operating system versions. When you use Solaris Live Upgrade to upgrade the Solaris operating system, you must complete these steps to ensure the correct version of SFCFSHA components are installed.

On a node in the cluster, run the installer on the alternate boot disk to remove and reinstall Veritas Storage Foundation Cluster File System High Availability 6.0 on all the nodes in the cluster.

At the end of the process the following occurs:

- Veritas Storage Foundation Cluster File System High Availability 6.0 is installed on the alternate boot disk, with the correct binaries for the new operating system version

To remove and reinstall SFCFSHA using the installer

- 1 Enter the names of the nodes that you want to uninstall.

The installer displays the list of packages that will be uninstalled.

- 2 Press **Return** to continue.

- 3 Install using the installer script, specifying the root path as the alternate boot disk as follows:

```
# /cdrom/storage_foundation_cluster_file_system_ha/installsfcfsha \  
-install -rootpath /altrootpath
```

- 4 Press **Return** to continue.

- 5 Verify that the version of the Veritas packages on the alternate boot disk is 6.0.

```
# pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
# pkginfo -R /altroot.5.10 -l VRTSvxvm
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/log`.

Upgrading SFCFSHA using Live Upgrade

Perform the Live Upgrade manually or use the installer. The nodes will not form a cluster until all of the nodes are upgraded to Veritas Storage Foundation Cluster File System High Availability 6.0. 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.

Upgrading SFCFSHA using Live Upgrade involves the following steps:

- Prepare to upgrade using Solaris Live Upgrade.
See [“Before you upgrade SFCFSHA using Solaris Live Upgrade”](#) on page 242.
- Create a new boot environment on the alternate boot disk.
See [“Creating a new boot environment on the alternate boot disk”](#) on page 246.

- Upgrade to Veritas Storage Foundation Cluster File System High Availability 6.0 on the alternate boot environment manually or using the installer. Refer to one of the following:
 - To upgrade SFCFSHA manually:
 - See [“Upgrading SFCFSHA manually”](#) on page 248.
 - To upgrade SFCFSHA using the installer:
 - See [“Upgrading SFCFSHA using the installer for a Live Upgrade”](#) on page 247.
- Switch the alternate boot environment to be the new primary.
See [“Completing the Live Upgrade ”](#) on page 251.
- Verify Live Upgrade of SFCFSHA.
See [“Verifying Live Upgrade of SFCFSHA”](#) on page 252.

Administering boot environments

Use the following procedures to perform relevant administrative tasks for boot environments.

Reverting to the primary boot environment

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.

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 `vxlufinish` scripts to process Live Upgrade. You must perform the following procedures when you perform a manual Live Upgrade.

Two different procedures exist to switch the boot environment, choose one of the following procedures based on the encapsulation of the root disk:

- See [“To switch the boot environment if the root disk is not encapsulated”](#) on page 256.
- See [“To switch the boot environment if the root disk is encapsulated”](#) on page 257.

The switching procedures for Solaris SPARC vary, depending on whether VxVM encapsulates the root disk.

To switch the boot environment if the root disk is not encapsulated

- 1 Display the status of Live Upgrade boot environments.

```
# lustatus
```

Boot Environment Name	Is Complete	Active Now	Active On Reboot	Can Delete	Copy Status
source.2657	yes	yes	yes	no	-
dest.2657	yes	no	no	yes	-

In this example, the primary boot disk is currently (source.2657). You want to activate the alternate boot disk (dest.2657)

- 2 Unmount any file systems that are mounted on the alternate root disk (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	/vxf	-	-

```
# luumount dest.2657
```

- 3 Activate the Live Upgrade boot environment.

```
# luactivate dest.2657
```

- 4 Reboot the system.

```
# shutdown -g0 -i6 -y
```

The system automatically selects the boot environment entry that was activated.

To switch the boot environment if the root disk is encapsulated**1** Display the current boot disk device and device aliases

```
# eeprom
boot-device=vx-rootdg vx-int_disk
use-nvramrc?=true
nvramrc=devalias vx-int_disk /pci@1c,600000/scsi@2/disk@0,0:a
devalias vx-rootdg01 /pci@1c,600000/scsi@2/disk@1,0:a
```

2 Set the device from which to boot using the eeprom command. This example shows booting from the primary root disk.

```
# eeprom boot-device=vx-rootdg01
```

3 Reboot the system.

```
# shutdown -g0 -i6 -y
```

Switching the boot environment for Solaris x86-64

You do not have to perform the following procedures to switch the boot environment when you use the `vxlufinish` scripts to process Live Upgrade. You must perform the following procedures when you perform a manual Live Upgrade.

Two different procedures exist to switch the boot environment, choose one of the following procedures based on the encapsulation of the root disk:

- See [“To switch the boot environment if root disk is not encapsulated”](#) on page 258.
- See [“To switch the boot environment if root disk is encapsulated”](#) on page 259.

To switch the boot environment if root disk is not encapsulated

- 1 Display the status of Live Upgrade boot environments.

```
# lustatus
```

Boot Environment Name	Is Complete	Active Now	Active On Reboot	Can Delete	Copy Status
source.2657	yes	yes	yes	no	-
dest.2657	yes	no	no	yes	-

In this example, the primary boot disk is currently (source.2657). You want to activate the alternate boot disk (dest.2657)

- 2 Unmount any file systems that are mounted on the alternate root disk (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	/vxf	-	-

```
# luumount dest.2657
```

- 3 Activate the Live Upgrade boot environment.

```
# luactivate dest.2657
```

- 4 Reboot the system.

```
# shutdown -g0 -i6 -y
```

When the system boots up, the GRUB menu displays the following entries for the Live Upgrade boot environments:

```
source.2657
dest.2657
```

The system automatically selects the boot environment entry that was activated.

To switch the boot environment if root disk is encapsulated

- ◆ If the root disk is encapsulated, for releases before Solaris 10 update 6 (2.10u6), you can use the `luactivate` method. For Solaris 10 update 6 and subsequent Solaris 10 updates, do one of the following:
 - Select the GRUB entry for the source boot environment or destination boot environment when the system is booted. You can also use the following procedure to manually set the default GRUB menu.lst entry to the source (PBE) or destination (ABE) grub entry:
 - If the system is booted from the alternate boot environment, perform the following steps to switch to the primary boot environment:

```
# mkdir /priroot
# mount rootpath /priroot
# bootadm list-menu -R /priroot
# bootadm list-menu
# bootadm set-menu -R /priroot default=PBE_menu_entry
# bootadm set-menu default=PBE_menu_entry
# shutdown -g0 -i6 -y
```

Where:

rootpath is the path to the root device, such as

```
/dev/vx/dsk/rootdg/rootvol
```

priroot is the primary root device

PBE_menu_entry is the number of the primary boot environment in the GRUB menu.

- If the system is booted from the primary boot environment, perform the following steps to switch to the alternate boot environment:

```
# bootadm list-menu
# bootadm set-menu default=ABE_menu_entry
ABE booting
```

Performing post-upgrade tasks

This chapter includes the following topics:

- [Re-joining the backup boot disk group into the current disk group](#)
- [Reverting to the backup boot disk group after an unsuccessful upgrade](#)

Re-joining the backup boot disk group into the current disk group

Perform this procedure to rejoin the backup boot disk if you split the mirrored boot disk during upgrade. After a successful upgrade and reboot, you no longer need to keep the boot disk group backup.

See [“Performing a rolling upgrade using the installer”](#) on page 208.

To re-join the backup boot disk group

- ◆ Re-join the *backup_bootdg* disk group to the boot disk group.

```
# /etc/vx/bin/vxrootadm -Y join backup_bootdg
```

where the `-Y` option indicates a silent operation, and *backup_bootdg* is the name of the backup boot disk group that you created during the upgrade.

Reverting to the backup boot disk group after an unsuccessful upgrade

Perform this procedure if your upgrade was unsuccessful and you split the mirrored boot disk to back it up during upgrade. You can revert to the backup that you created when you upgraded.

See [“Performing a rolling upgrade using the installer”](#) on page 208.

To revert the backup boot disk group after an unsuccessful upgrade

- 1 To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

```
# vxprint
```

- 2 Use the `vx dg` command to find the boot disk group where you are currently booted.

```
# vx dg bootdg
```

- 3 Boot the operating system from the backup boot disk group.
- 4 Join the original boot disk group to the backup disk group.

```
# /etc/vx/bin/vxrootadm -Y join original_bootdg
```

where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

Verification of the installation or the upgrade

- [Chapter 21. Verifying the Storage Foundation Cluster File System High Availability installation](#)

Verifying the Storage Foundation Cluster File System High Availability installation

This chapter includes the following topics:

- [Performing a postcheck on a node](#)
- [Verifying that the products were installed](#)
- [Installation log files](#)
- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [Starting and stopping processes for the Veritas products](#)
- [Checking Veritas Volume Manager processes](#)
- [Checking Veritas File System installation](#)
- [Verifying agent configuration for Storage Foundation Cluster File System High Availability](#)
- [Synchronizing time on Cluster File Systems](#)
- [Configuring VCS for Storage Foundation Cluster File System High Availability](#)
- [About the cluster UUID](#)
- [About the LLT and GAB configuration files](#)

- [Verifying the LLT, GAB, and VCS configuration files](#)
- [Verifying LLT, GAB, and cluster operation](#)

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 407.

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.

Verifying that the products were installed

Verify that the SFCFSHA products are installed.

Use the `pkginfo` command to check which packages have been installed.

```
# pkginfo -l VRTSvlic package_name package_name ...
```

You can verify the version of the installed product. Use the following command:

```
# /opt/VRTS/install/installsfcfsha -version
```

Use the following sections to further verify the product installation.

Installation log files

After every product installation, the installer creates three text files:

- Installation log file
- Response file
- Summary file

The name and location of each file is displayed at the end of a product installation, and are always located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

Using the installation log file

The installation log file contains all commands executed during the procedure, their output, and errors generated by the commands. This file is for debugging installation problems and can be used for analysis by Veritas Support.

Using the summary file

The summary file contains the results of the installation by the installer or product installation scripts. The summary includes the list of the packages, and the status (success or failure) of each package. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

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 Oracle Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

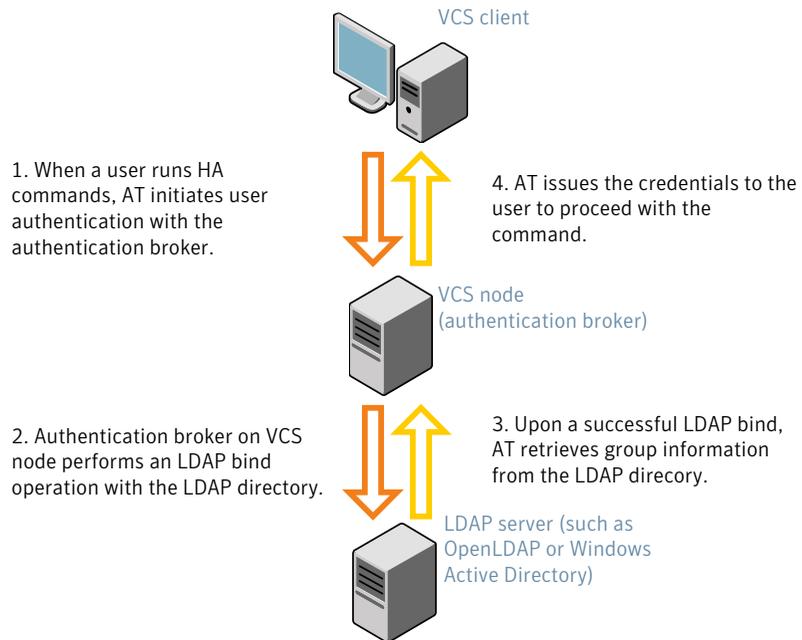
See [“Enabling LDAP authentication for clusters that run in secure mode”](#) on page 267.

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

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

[Figure 21-1](#) depicts the SFCFSHA cluster communication with the LDAP servers when clusters run in secure mode.

Figure 21-1 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
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

To enable OpenLDAP authentication for clusters that run in secure mode

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat addldapdomain \
--domainname "MYENTERPRISE.symantecdomain.com"\
--server_url "ldap://my_openldap_host.symantecexample.com"\
--user_base_dn "ou=people,dc=symantecdomain,dc=myenterprise,dc=com"\
--user_attribute "cn" --user_object_class "account"\
--user_gid_attribute "gidNumber"\
--group_base_dn "ou=group,dc=symantecdomain,dc=myenterprise,dc=com"\
--group_attribute "cn" --group_object_class "posixGroup"\
--group_gid_attribute "member"\
--admin_user "cn=manager,dc=symantecdomain,dc=myenterprise,dc=com"\
--admin_user_password "password"--auth_type "FLAT"
```

- 2 Verify that you can successfully authenticate an LDAP user on the SFCFSHA nodes.

You must have a valid LDAP user ID and password to run the command. In the following example, authentication is verified for the MYENTERPRISE domain for the LDAP user, `vcsadmin1`.

```
galaxy# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat authenticate
--domain ldap:MYENTERPRISE.symantecdomain.com
--prplname vcsadmin1 --broker galaxy:14149
```

Enter password for vcsadmin1: #####

```
authenticate
-----
-----
```

```
Authenticated User vcsadmin1
-----
```

3 Add the LDAP user to the main.cf file.

```
# haconf makerw
# hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
# haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
# hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

4 Verify that the main.cf file has the following lines:

```
# cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster clus1 (
  SecureClus = 1
  Administrators = {
    "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" }
  AdministratorGroups = {
    "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com " }
  )
...
...
```

5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=myenterprise.symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh) or the Korn shell (ksh), run the following commands:

```
# export VCS_DOMAIN=myenterprise.symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute    Value
galaxy       Attribute    RUNNING
nebula       Attribute    RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFCFSHA node using the VCS Cluster Manager (Java Console).

- 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

To enable Windows Active Directory authentication for clusters that run in secure mode

- 1 Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -d \  
-s domain_controller_name_or_ipaddress \  
-u domain_user -g domain_group
```

For example:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-d -s 192.168.20.32 -u Administrator -g "Domain Admins"  
Search User provided is invalid or Authentication is required to  
proceed further.  
Please provide authentication information for LDAP server.
```

```
Username/Common Name: symantecdomain\administrator  
Password:
```

Attribute file created.

- 2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-c -d windows_domain_name
```

For example:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-c -d symantecdomain.com  
Attribute list file not provided, using default AttributeList.txt.  
CLI file name not provided, using default CLI.txt.
```

CLI for addldapdomain generated.

- 3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -x
```

- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat listldapdomains
```

```
Domain Name :          symantecdomain.com
Server URL :          ldap://192.168.20.32:389
SSL Enabled :         No
User Base DN :        CN=people,DC=symantecdomain,DC=com
User Object Class :   account
User Attribute :      cn
User GID Attribute :  gidNumber
Group Base DN :       CN=group,DC=symantecdomain,DC=com
Group Object Class :  group
Group Attribute :     cn
Group GID Attribute : cn
Auth Type :           FLAT
Admin User :
Admin User Password :
Search Scope :        SUB
```

- 5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh) or the Korn shell (ksh), run the following commands:

```
# export VCS_DOMAIN=symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute    Value
galaxy       Attribute    RUNNING
nebula       Attribute    RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFCFSHA node using the VCS Cluster Manager (Java Console).

- 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

To stop the processes

- ◆ Use the `-stop` option to stop the product installation script.

For example, to stop the product's processes, enter the following command:

```
# ./installer -stop
```

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
```

Checking Veritas Volume Manager processes

Use the following procedure to verify that Volume Manager processes are running.

To confirm that key Volume Manager processes are running

- ◆ Type the following command:

```
# ps -ef | grep vx
```

Entries for the `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxcached`, and `vxconfigbackupd` processes should appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

Checking Veritas File System installation

The Veritas File System package consists of a kernel component and administrative commands.

Verifying Veritas File System kernel installation

To ensure that the file system driver is loaded, enter:

```
# modinfo | grep vxfs
```

The `modinfo` command displays information about all modules loaded on the system. If the `vxfs` module is loaded, you will see an entry corresponding to `vxfs`. If not, follow the instructions load and then unload the file system module to complete the process.

See [“Loading and unloading the file system module”](#) on page 141.

Verifying command installation

[Table 21-1](#) lists the directories with Veritas File System commands.

Table 21-1 VxFS command locations

Location	Contents
<code>/etc/fs/vxfs</code>	Contains the Veritas <code>mount</code> command and QuickLog commands required to mount file systems.
<code>/usr/lib/fs/vxfs/bin</code>	Contains the VxFS type-specific switch-out commands.
<code>/opt/VRTSvxfs/sbin</code>	Contains the Veritas-specific commands.
<code>/opt/VRTS/bin</code>	Contains symbolic links to all Veritas-specific commands installed in the directories listed above.

Determine whether these subdirectories are present:

```
# ls /etc/fs/vxfs
# ls /usr/lib/fs/vxfs/bin
# ls /opt/VRTSvxfs/sbin
# ls /opt/VRTS/bin
```

Make sure you have adjusted the environment variables accordingly.

See [“Setting environment variables”](#) on page 31.

Verifying agent configuration for Storage Foundation Cluster File System High Availability

This section describes how to verify the agent configuration.

To verify the agent configuration

- ◆ Enter the cluster status command from any node in the cluster:

```
# cfscluster status
```

Output resembles:

```
Node           : system01
Cluster Manager : running
CVM state      : running
No mount point registered with cluster configuration

Node           : system02
Cluster Manager : running
CVM state      : running
No mount point registered with cluster configuration
```

Synchronizing time on Cluster File Systems

SFCFSHA requires that the system clocks on all nodes are synchronized using some external component such as the Network Time Protocol (NTP) daemon. If the nodes are not in sync, timestamps for change (`ctime`) and modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

Configuring VCS for Storage Foundation Cluster File System High Availability

Configuring VCS means conveying to the VCS engine the definitions of the cluster, service groups, resources, and resource dependencies. VCS uses two configuration files in a default configuration:

- The `main.cf` file defines the entire cluster.
- The `types.cf` file defines the resource types.

By default, both files reside in the directory `/etc/VRTSvcs/conf/config`. Additional files similar to `types.cf` may be present if agents have been added, such as `OracleTypes.cf`.

In a VCS cluster, the first system to be brought online reads the configuration file and creates an internal (in-memory) representation of the configuration. Systems brought online after the first system derive their information from systems running in the cluster. You must stop the cluster while you are modifying the files from the command line. Changes made by editing the configuration files take effect when the cluster is restarted. The node on which the changes were made should be the first node to be brought back online.

main.cf file

The VCS configuration file `main.cf` is created during the installation procedure. After installation, the `main.cf` file contains the base definitions of the cluster and its nodes. Additionally, the file `types.cf` listed in the include statement defines the bundled agents for VCS resources.

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

A typical VCS configuration file for SFCFSHA file resembles:

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTTypes.cf"
include "OracleASMTTypes.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.cf"

cluster sfcfs_1 (
    HacliUserLevel = COMMANDROOT
)
```

```
system thor150 (
)

system thor151 (
)

group cvm (
    SystemList = { thor150 = 0, thor151 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { thor150, thor151 }
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = sfcfs_1
    CVMNodeId = { thor150 = 0, thor151 = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

// resource dependency tree
//
//     group cvm
//     {
//         CVMCluster cvm_clus
//         {
//             CVMVxconfigd cvm_vxconfigd
//         }
//     }
// }
```

Storage Foundation Cluster File System HA Only

If you configured VCS Cluster Manager (Web Console), a service group, "ClusterService," was created that includes IP, Process, and Notifier resources.

These resources were configured according to information you provided during the installation procedure. A resource dependency was also created.

Veritas Cluster Server application failover services

If you installed SFCFS HA, you can begin implementing the application monitoring failover services provided by the Veritas Cluster Server. Information about setting up VCS services is beyond the scope of this document.

See the *Veritas Cluster Server* documentation.

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.

About the cluster UUID

You can verify the existence of the cluster UUID.

To verify the cluster UUID exists

- ◆ From the prompt, run a cat command.

```
cat /etc/vx/.uuids/clusuuid
```

About the LLT and GAB configuration files

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

[Table 21-2](#) lists the LLT configuration files and the information that these files contain.

Table 21-2 LLT configuration files

File	Description
/etc/default/llt	<p>This file stores the start and stop environment variables for LLT:</p> <ul style="list-style-type: none"> ■ LLT_START—Defines the startup behavior for the LLT module after a system reboot. Valid values include: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 1—Indicates that LLT is enabled to shut down. 0—Indicates that LLT is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of SFCFSHA configuration.</p>
/etc/llthosts	<p>The file <code>llthosts</code> 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.</p> <p>For example, the file <code>/etc/llthosts</code> contains the entries that resemble:</p> <pre style="margin-left: 40px;">0 galaxy 1 nebula</pre>

Table 21-2 LLT configuration files (*continued*)

File	Description
/etc/llttab	<p>The file <code>llttab</code> contains the information that is derived during installation and used by the utility <code>lltconfig(1M)</code>. After installation, this file lists the private network links that correspond to the specific system. For example, the file <code>/etc/llttab</code> contains the entries that resemble the following:</p> <ul style="list-style-type: none"> ■ For Solaris SPARC: <pre> set-node galaxy set-cluster 2 link bge0 /dev/bge0 - ether - - link bge1 /dev/bge1 - ether - - </pre> ■ For Solaris x64: <pre> set-node galaxy set-cluster 2 link e1000g1 /dev/e1000g:1 - ether - - link e1000g2 /dev/e1000g:2 - ether - - </pre> <p>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 <code>link</code> command. These lines identify the two network cards that the LLT protocol uses.</p> <p>If you configured a low priority link under LLT, the file also includes a "link-lowpri" line.</p> <p>Refer to the <code>llttab(4)</code> manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the <code>llttab</code> file.</p>

[Table 21-3](#) lists the GAB configuration files and the information that these files contain.

Table 21-3 GAB configuration files

File	Description
/etc/default/gab	<p>This file stores the start and stop environment variables for GAB:</p> <ul style="list-style-type: none"> ■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 1—Indicates that GAB is enabled to shut down. 0—Indicates that GAB is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of SFCFSHA configuration.</p>
/etc/gabtab	<p>After you install SFCFSHA, the file <code>/etc/gabtab</code> contains a <code>gabconfig(1)</code> command that configures the GAB driver for use.</p> <p>The file <code>/etc/gabtab</code> contains a line that resembles:</p> <pre style="text-align: center;">/sbin/gabconfig -c -nN</pre> <p>The <code>-c</code> option configures the driver for use. The <code>-nN</code> specifies that the cluster is not formed until at least <i>N</i> nodes are ready to form the cluster. Symantec recommends that you set <i>N</i> to be the total number of nodes in the cluster.</p> <p>Note: Symantec does not recommend the use of the <code>-c -x</code> option for <code>/sbin/gabconfig</code>. Using <code>-c -x</code> can lead to a split-brain condition.</p>

Verifying the LLT, GAB, and VCS configuration files

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

To verify the LLT, GAB, and VCS configuration files

- 1 Navigate to the location of the configuration files:
 - LLT
/etc/llthosts
/etc/llttab
 - GAB
/etc/gabtab
 - VCS

```
/etc/VRTSvcs/conf/config/main.cf
```

- 2 Verify the content of the configuration files.

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

Verifying LLT, GAB, and cluster operation

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

To verify LLT, GAB, and cluster operation

- 1 Log in to any node in the cluster as superuser.
- 2 Make sure that the PATH environment variable is set to run the VCS commands.
- 3 Verify LLT operation.
See [“Verifying LLT”](#) on page 282.
- 4 Verify GAB operation.
See [“Verifying GAB”](#) on page 285.
- 5 Verify the cluster operation.
See [“Verifying the cluster”](#) on page 286.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

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

```
lltstat -n
```

The output on galaxy resembles:

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

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

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

LLT node information:

Node	State	Links
* 0 galaxy	OPEN	2
1 nebula	OPEN	2
2 saturn	OPEN	1

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

```
lltstat -n
```

The output on nebula resembles:

LLT node information:

Node	State	Links
0 galaxy	OPEN	2
*1 nebula	OPEN	2

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

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

```
lltstat -nvv active
```

The output on galaxy resembles the following:

- For Solaris SPARC:

Node	State	Link	Status	Address
*0 galaxy	OPEN	<i>bge1</i>	UP	08:00:20:93:0E:34
		<i>bge2</i>	UP	08:00:20:93:0E:38
1 nebula	OPEN	<i>bge1</i>	UP	08:00:20:8F:D1:F2
		<i>bge2</i>	DOWN	

- For Solaris x64:

Node	State	Link	Status	Address
*0 galaxy	OPEN	e1000g:1	UP	08:00:20:93:0E:34
		e1000g:2	UP	08:00:20:93:0E:38
1 nebula	OPEN	e1000g:1	UP	08:00:20:8F:D1:F2
		e1000g:2	DOWN	

The command reports the status on the two active nodes in the cluster, galaxy and nebula.

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

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

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

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

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

```
lltstat -p
```

The output resembles:

```
LLT port information:
Port  Usage      Cookie
0     gab        0x0
      opens:    0 2 3 4 5 6 7 8 9 10 11 ... 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:

a	GAB
b	I/O fencing
d	Oracle Disk Manager (ODM)
f	Cluster File System (CFS)
h	Veritas Cluster Server (VCS: High Availability Daemon)
u	Cluster Volume Manager (CVM) (to ship commands from slave node to master node) Port u in the <code>gabconfig</code> output is visible with CVM protocol version ≥ 100 .
v	Cluster Volume Manager (CVM)
w	vxconfigd (module for CVM)
y	Cluster Volume Manager (CVM) I/O shipping

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

To verify GAB

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

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

- 2 Review the output of the command:

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

```
GAB Port Memberships
=====
Port a gen  ada401 membership 01
Port b gen  ada40d membership 01
Port d gen  ada409 membership 01
Port f gen  ada41c membership 01
Port h gen  ada40f membership 01
Port u gen  ada41a membership 01
```

```
Port v gen ada416 membership 01
Port w gen ada418 membership 01
Port y gen ada42a membership 01
```

Note that port b in the `gabconfig` command output may not indicate that I/O fencing feature is configured. After you configure SFCFSHA using the installer, the installer starts I/O fencing in disabled mode. You can use the `vxfenadm -d` command to verify the I/O fencing configuration.

- If GAB does not operate, the command does not return any GAB port membership information:

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

Verifying the cluster

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

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

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

To verify the cluster

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

```
# hastatus -summary
```

The output resembles:

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

A galaxy          RUNNING       0
A nebula          RUNNING       0

-- GROUP STATE
-- Group           System        Probed  AutoDisabled  State

B cvm             galaxy        Y      N              ONLINE
B cvm             nebula        Y      N              ONLINE
```

- 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.

Verifying the cluster nodes

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

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

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

To verify the cluster nodes

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

```
hasys -display
```

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

Note: The following example is for SPARC. x64 clusters have different command output.

#System	Attribute	Value
galaxy	AgentsStopped	0
galaxy	AvailableCapacity	100
galaxy	CPUBinding	BindTo None CPUNumber 0
galaxy	CPUUsage	0
galaxy	CPUUsageMonitoring	Enabled 0 ActionThreshold 0 ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
galaxy	Capacity	100
galaxy	ConfigBlockCount	130
galaxy	ConfigChecksum	46688

#System	Attribute	Value
galaxy	ConfigDiskState	CURRENT
galaxy	ConfigFile	/etc/VRTSvcs/conf/config
galaxy	ConfigInfoCnt	0
galaxy	ConfigModDate	Wed 14 Oct 2009 17:22:48
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	
galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	5.1.00.0
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	HostUtilization	CPU 0 Swap 0
galaxy	LLTNodeId	0
galaxy	LicenseType	DEMO
galaxy	Limits	
galaxy	LinkHbStatus	
galaxy	LoadTimeCounter	0
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	ShutdownTimeout	
galaxy	SourceFile	./main.cf

#System	Attribute	Value
galaxy	SysInfo	Solaris:galaxy,Generic_118558-11,5.9,sun4u
galaxy	SysName	galaxy
galaxy	SysState	RUNNING
galaxy	SystemLocation	
galaxy	SystemOwner	
galaxy	TFrozen	0
galaxy	TRSE	0
galaxy	UpDownState	Up
galaxy	UserInt	0
galaxy	UserStr	
galaxy	VCSFeatures	DR
galaxy	VCSMode	VCS_CFS_VRTS

Adding and removing nodes

- [Chapter 22. Adding a node to a cluster](#)
- [Chapter 23. Removing a node from Storage Foundation Cluster File System clusters](#)

Adding a node to a cluster

This chapter includes the following topics:

- [About adding a node to a cluster](#)
- [Before adding a node to a cluster](#)
- [Adding a node to a cluster using the product installer](#)
- [Adding a node using the Web-based installer](#)
- [Adding the node to a cluster manually](#)
- [Configuring server-based fencing on the new node](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node](#)
- [Sample configuration file for adding a node to the cluster](#)

About adding a node to a cluster

After you install SFCFSHA and create a cluster, you can add and remove nodes from the cluster. You can create clusters of up to 64 nodes.

The following procedure provides a summary of the tasks required to add a node to an existing SFCFSHA cluster.

Table 22-1 Tasks for adding a node to a cluster

Step	Description
Complete the prerequisites and preparatory tasks before adding a node to the cluster.	See “Before adding a node to a cluster” on page 292.

Table 22-1 Tasks for adding a node to a cluster (*continued*)

Step	Description
Add a new node to the cluster: <ul style="list-style-type: none"> ■ Using the product installer ■ Using the Web installer ■ Manually 	See “Adding a node to a cluster using the product installer” on page 294. See “Adding a node using the Web-based installer” on page 298. See “Adding the node to a cluster manually” on page 298.
Complete the configuration of the new node after adding it to the cluster.	
If you are using the Storage Foundation for Databases (SFDB) tools, you must update the repository database.	See “Updating the Storage Foundation for Databases (SFDB) repository after adding a node” on page 322.

The example procedures describe how to add a node to an existing cluster with two nodes.

Before adding a node to a cluster

Before preparing to add the node to an existing cluster, perform the required preparations.

- Verify hardware and software requirements are met.
- Set up the hardware.
- Prepare the new node.

To verify hardware and software requirements are met

- 1 Review the hardware and software requirements.
 See [“Assessing the system for installation readiness”](#) on page 28.
- 2 Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
- 3 Verify the existing cluster is running SFCFSHA 6.0 on all nodes.

- 4 If the cluster is upgraded from a previous version, you must check the cluster protocol version to make sure it has the same version as the node to be added. If there is a protocol mismatch, the node is unable to join the existing cluster.

Check the cluster protocol version using:

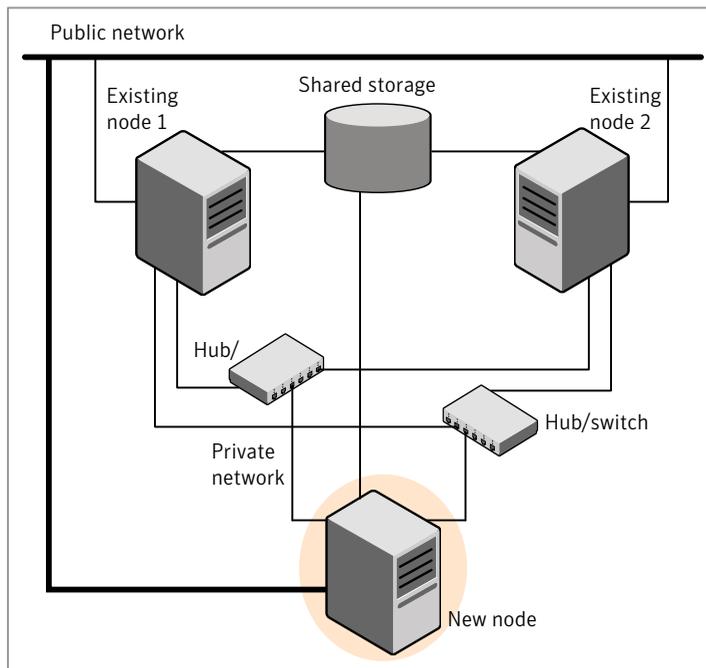
```
# vxctl protocolversion
Cluster running at protocol 110
```

- 5 If the cluster protocol is below 110, upgrade it on the master node using:

```
# vxctl upgrade [version]
```

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in [Figure 22-1](#).

Figure 22-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the SFCFSA private Ethernet controllers.
 - When you add nodes to a 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 22-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.

- 2 Connect the system to the shared storage as required. Make sure that you meet the following requirements:
 - The new node must be connected to the same shared storage devices as the existing nodes.
 - The node must have private network connections to two independent switches for the cluster.
 - The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

Prepare the new node before you add it to an existing cluster.

To prepare the new node

- 1 Verify that the new node meets installation requirements.

```
# ./installsfcfsha -precheck
```

You can also use the Web-based installer for the precheck.

- 2 Install SFCFSHA on the new system using the `-install` option to install SFCFSHA. Do not configure SFCFSHA when prompted.
- 3 You can restart the new node after installation is complete. Configure the new node using the configuration from the existing cluster nodes.

See [“About installation and configuration methods”](#) on page 27.

Adding a node to a cluster using the product installer

You can add a node using the `-addnode` option with the product installer.

The product 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`
- Copies the following files on the new node:
 - `/etc/llthosts`
 - `/etc/gabtab`
 - `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files on the new node:
 - `/etc/llthosts`
 - `/etc/gabtab`
 - `/etc/VRTSvcs/conf/config/main.cf`
 - `/etc/vxfenmode`
 - `/etc/vxfendg`
 - `/etc/vx.uuids/clusuuid`
 - `/etc/default/llt`
 - `/etc/default/gab`
 - `/etc/default/vxfen`
 - `/etc/vcsmmtab`
- Generates security credentials on the new node if the CPS server of existing cluster is secure
- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.
Configures disk-based fencing.
- Adds the new node to the CVM, ClusterService, and VxSS service groups in the VCS configuration.

Note: For other service groups configured under VCS, update the configuration for the new node manually.

Note:

- Starts product processes and configures CVM and CFS on the new node.

Note:

At the end of the process, the new node joins the existing 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 cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the product installer with the `-addnode` option.

```
# cd /opt/VRTS/install
# ./installsfcfsha -addnode
```

The installer displays the copyright message and the location where it stores the temporary installation logs.

- 3 Enter the name of a node in the existing cluster.

The installer uses the node information to identify the existing cluster.

```
Enter one node of the SFCFSHA cluster to which
you would like to add one or more new nodes: galaxy
```

- 4 Review and confirm the cluster information.
- 5 Enter the names of the systems that you want to add as new nodes to the cluster.

```
Enter the system names separated by spaces
to add to the cluster: saturn
```

The installer checks the installed products and packages on the nodes and discovers the network interfaces.

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

Note: The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

```
Enter the NIC for the first private heartbeat
link on saturn: [b,q,?] bge1
```

```
Enter the NIC for the second private heartbeat
link on saturn: [b,q,?] bge2
```

Note: At least two private heartbeat links must be configured for high availability of the cluster.

- 7 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.
The product installer verifies the network interface settings and displays the information.
- 8 Review and confirm the information.
- 9 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: bge3
```

- 10 The installer prompts you with an option to mount the shared volumes on the new node. Select **y** to mount them.
When completed, the installer confirms the volumes are mounted and indicates the location of the log file with details of the actions performed.
- 11 If the existing cluster uses server-based fencing in secure mode, the installer will configure server-based fencing in secure mode on the new nodes.
The installer then starts all the required Veritas processes and joins the new node to cluster.
- 12 Confirm that the new node has joined the existing cluster using `lltstat -n` and `gabconfig -a` commands.

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.

Adding the node to a cluster manually

Perform the following tasks in order after you install to add the node to the cluster manually.

Table 22-2 Tasks for adding a cluster manually

Task	Description
Start the Volume Manager.	See “Starting Volume Manager on the new node” on page 310.

Table 22-2 Tasks for adding a cluster manually (*continued*)

Task	Description
Add the cluster files.	<p>Configure LLT, GAB, and VCSMM.</p> <p>Configure LLT and GAB.</p> <p>See “Configuring cluster files on the new node” on page 310.</p>
Generate security credentials.	<p>If the CP server of the existing cluster is secure, generate security credentials on the new node.</p> <p>See “Setting up the node to run in secure mode” on page 312.</p>
Set up fencing.	<p>Configure fencing for the new node to match the fencing configuration on the existing cluster.</p> <p>If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node.</p> <p>See “Starting fencing on the new node” on page 316.</p>
Start VCS.	<p>See “Starting VCS after adding the new node” on page 317.</p>
Configure Cluster Volume Manager (CVM) and Cluster File System (CFS).	<p>If CVM and CFS are configured on the existing cluster, configure them on the new node.</p> <p>See “Configuring CVM and CFS on the new node” on page 318.</p>
Configure the ClusterService group.	<p>If the ClusterService group is configured on the existing cluster, add the node to the group.</p> <p>See “Configuring the ClusterService group for the new node” on page 319.</p>

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfcfsha` program.

Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfcfsha` program.

To start Volume Manager on the new node

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:

```
# vxinstall
```

- 2 Enter `n` when prompted to set up a system wide disk group for the system. The installation completes.

- 3 Verify that the daemons are up and running. Enter the command:

```
# vxdisk list
```

Make sure the output displays the shared disks without errors.

Configuring cluster files on the new node

Perform the steps in the following procedure to configure LLT and GAB on the new node.

To configure LLT and GAB on the new node

- 1 Edit the `/etc/llthosts` file on the existing nodes. Using `vi` or another text editor, add the line for the new node to the file. The file resembles:

```
0 galaxy
1 nebula
2 saturn
```

- 2 Copy the `/etc/llthosts` file from one of the existing systems over to the new system. The `/etc/llthosts` file must be identical on all nodes in the cluster.
- 3 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link bge1 /dev/bge:1 - ether - -
link bge2 /dev/bge:2 - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 4 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

- 5 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 6 For the following files on the new node:

```
/etc/default/llt  
/etc/default/gab  
/etc/default/vxfen  
/etc/default/vcs
```

Verify if the attributes in each file are set as follows before using smf on Solaris 10 to start the related processes and to load drivers:

```
LLT_START/LLT_STOP=1  
GAB_START/GAB_STOP=1  
VXFEN_START/VXFEN_STOP=1  
VCS_START/VCS_STOP=1
```

- 7 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

- 9 Start the LLT, GAB, and ODM drivers on the new node:

For Solaris 10:

```
# svcadm enable llt
# svcadm enable gab
# svcadm restart vxodm
```

- 10 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen df204 membership 012
Port d gen df207 membership 012
```

Configuring cluster files on the new node

Perform the steps in the following procedure to configure cluster files on the new node.

To configure cluster files on the new node

- 1 Edit the `/etc/llthosts` file on the existing nodes. Using `vi` or another text editor, add the line for the new node to the file. The file resembles:

```
0 galaxy
1 nebula
2 saturn
```

- 2 Copy the `/etc/llthosts` file from one of the existing systems over to the new system. The `/etc/llthosts` file must be identical on all nodes in the cluster.
- 3 Create an `/etc/llttab` file on the new system.

For example:

```
set-node saturn
set-cluster 101
```

For AIX:

For HP-UX:

For Linux:

4 For Solaris:

```
link bge1 /dev/bge:1 - ether - -  
link bge2 /dev/bge:2 - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

5 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

- 6** Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 7** Copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/default/llt  
/etc/default/gab  
/etc/default/vcs
```

Verify if the attributes in each file are set as follows before using `smf` on Solaris 10 to start the related processes and to load drivers:

```
LLT_START/LLT_STOP=1  
GAB_START/GAB_STOP=1  
VXFEN_START/VXFEN_STOP=1  
VCS_START/VCS_STOP=1
```

- 8** Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 9** Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

Start the cluster components

1 Start LLT.

For AIX:

For HP-UX:

For Linux:

For Solaris:

```
# svcadm enable llt
```

2 Start GAB.

For AIX:

For HP-UX:

For Linux:

For Solaris:

```
# svcadm enable gab
```

3 Start disk-based fencing.

For AIX:

For HP-UX:

For Linux:

For Solaris:

4 Start VCSMM.

For AIX:

For HP-UX:

For Linux:

For Solaris:

5 Start LMX.

For AIX:

For HP-UX:

For Solaris:

6 Start ODM.

For AIX:

For HP-UX:

For Linux:

For Solaris:

```
# svcadm enable vxodm
```

7 Start the LLT, GAB, and ODM drivers on the new node:

For AIX:

For HP-UX:

For Linux:

For Solaris:

```
# svcadm enable llt
```

```
# svcadm enable gab
```

```
# svcadm restart vxodm
```

8 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
```

```
Port a gen df204 membership 012
```

```
Port d gen df207 membership 012
```

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

For AIX:

For HP-UX:

For Linux:

For Solaris:

```
/etc/default/vxfen
/etc/vxfendg
/etc/vxfenmode
```

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:

```
# svcadm enable vxfen
```

- 3 On the new node, verify that the GAB port memberships are a, b, and d:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen      df204 membership 012
Port b gen      df20d membership 012
Port d gen      df20a membership 012
```

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

For AIX:

For HP-UX:

For Linux:

For Solaris:

```
/etc/default/vxfen
/etc/vxfendg
/etc/vxfenmode
```

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:

```
# svcadm enable vxfen
```

- 3 On the new node, verify that the GAB port memberships are a, b, and d:

```
# gabconfig -a
```

```
GAB Port Memberships
=====
Port a gen      df204 membership 012
Port b gen      df20d membership 012
Port d gen      df20a membership 012
```

Configuring CVM and CFS on the new node

Modify the existing cluster configuration to configure CVM and CFS for the new node.

To configure CVM and CFS on the new node

- 1 Make a backup copy of the main.cf file on the existing node, if not backed up in previous procedures. For example:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.2node
```

- 2 On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Add the new node to the VCS configuration, if not already added:

```
# hasys -add saturn
```

- 4 To enable the existing cluster to recognize the new node, run the following commands on one of the existing nodes:

```
# hagrps -modify cvm SystemList -add saturn 2
# hagrps -modify cvm AutoStartList -add saturn
# hares -modify cvm_clus CVMNodeId -add saturn 2
# haconf -dump -makero
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 5 On the remaining nodes of the existing cluster, run the following commands:

```
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 6 Copy the configuration files from one of the nodes in the existing cluster to the new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \  
saturn:/etc/VRTSvcs/conf/config/main.cf  
# rcp /etc/VRTSvcs/conf/config/CFSTypes.cf \  
saturn:/etc/VRTSvcs/conf/config/CFSTypes.cf  
# rcp /etc/VRTSvcs/conf/config/CVMTypes.cf \  
saturn:/etc/VRTSvcs/conf/config/CVMTypes.cf
```

- 7 The `/etc/vx/tunefstab` file sets non-default tunables for local-mounted and cluster-mounted file systems.

If you have configured a `/etc/vx/tunefstab` file to tune cluster-mounted file systems on any of the existing cluster nodes, you may want the new node to adopt some or all of the same tunables.

To adopt some or all tunables, review the contents of the file, and copy either the file, or the portions desired, into the `/etc/vx/tunefstab` file on the new cluster node.

Starting VCS after adding the new node

Start VCS on the new node.

To start VCS on the new node

- 1 Start VCS on the new node:

```
# hastart
```

VCS brings the CVM group online.

- 2 Verify that the CVM group is online:

```
# hagrps -state
```

Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.

To configure the ClusterService group for the new node

- 1 On an existing node, for example galaxy, write-enable the configuration:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing ClusterService group.

```
# hagrps -modify ClusterService SystemList -add saturn 2
```

```
# hagrps -modify ClusterService AutoStartList -add saturn
```

- 3 Modify the IP address and NIC resource in the existing group for the new node.

```
# hares -modify gcoip Device bge0 -sys saturn
```

```
# hares -modify gconic Device bge0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer `n` to prompts about licensing. You installed the appropriate license when you ran the `installsfscfsha` program.

To start Volume Manager on the new node

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:

```
# vxinstall
```

- 2 Enter `n` when prompted to set up a system wide disk group for the system.
The installation completes.

- 3 Verify that the daemons are up and running. Enter the command:

```
# vxdisk list
```

Make sure the output displays the shared disks without errors.

Configuring cluster files on the new node

Perform the steps in the following procedure to configure LLT and GAB on the new node.

To configure LLT and GAB on the new node

- 1 Edit the `/etc/llthosts` file on the existing nodes. Using `vi` or another text editor, add the line for the new node to the file. The file resembles:

```
0 galaxy
1 nebula
2 saturn
```

- 2 Copy the `/etc/llthosts` file from one of the existing systems over to the new system. The `/etc/llthosts` file must be identical on all nodes in the cluster.
- 3 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link bge1 /dev/bge:1 - ether - -
link bge2 /dev/bge:2 - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 4 Use `vi` or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

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

Where `N` represents the number of systems in the cluster including the new node. For a three-system cluster, `N` would equal 3.

- 5 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 6 For the following files on the new node:

```
/etc/default/llt
/etc/default/gab
/etc/default/vxfen
/etc/default/vcs
```

Verify if the attributes in each file are set as follows before using `smf` on Solaris 10 to start the related processes and to load drivers:

```
LLT_START/LLT_STOP=1
GAB_START/GAB_STOP=1
VXFEN_START/VXFEN_STOP=1
VCS_START/VCS_STOP=1
```

- 7 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

- 9 Start the LLT, GAB, and ODM drivers on the new node:

For Solaris 10:

```
# svcadm enable llt  
  
# svcadm enable gab  
  
# svcadm restart vxodm
```

- 10 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a  
GAB Port Memberships  
=====
```

Port a gen df204 membership 012	Port d gen df207 membership 012
---------------------------------	---------------------------------

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.

[Table 22-3](#) uses the following information for the following command examples.

Table 22-3 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
saturn	saturn.nodes.example.com	The new node that you are adding to the cluster.

Configuring the authentication broker on node saturn

To configure the authentication broker on node saturn

- 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 saturn by copying the entire `bkup` directory.

The `bkup` directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/

# ls

CMDSERVER  CPSADM  CPSEVER  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`, `CPSADM`, `CPSEVER`, 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 saturn.

```
# /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
```

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

For AIX:

For HP-UX:

For Linux:

For Solaris:

```
/etc/default/vxfen
/etc/vxfendg
/etc/vxfenmode
```

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:

```
# svcadm enable vxfen
```

- 3 On the new node, verify that the GAB port memberships are a, b, and d:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen      df204 membership 012
Port b gen      df20d membership 012
Port d gen      df20a membership 012
```

Starting VCS after adding the new node

Start VCS on the new node.

To start VCS on the new node

- 1 Start VCS on the new node:

```
# hastart
```

VCS brings the CVM group online.

- 2 Verify that the CVM group is online:

```
# hagr -state
```

Configuring CVM and CFS on the new node

Modify the existing cluster configuration to configure CVM and CFS for the new node.

To configure CVM and CFS on the new node

- 1 Make a backup copy of the main.cf file on the existing node, if not backed up in previous procedures. For example:

```
# cd /etc/VRTSvcs/conf/config  
# cp main.cf main.cf.2node
```

- 2 On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Add the new node to the VCS configuration, if not already added:

```
# hasys -add saturn
```

- 4 To enable the existing cluster to recognize the new node, run the following commands on one of the existing nodes:

```
# hagr -modify cvm SystemList -add saturn 2  
# hagr -modify cvm AutoStartList -add saturn  
# hares -modify cvm_clus CVMNodeId -add saturn 2  
# haconf -dump -makero  
# /etc/vx/bin/vxclustadm -m vcs reinit  
# /etc/vx/bin/vxclustadm nidmap
```

- 5 On the remaining nodes of the existing cluster, run the following commands:

```
# /etc/vx/bin/vxclustadm -m vcs reinit
# /etc/vx/bin/vxclustadm nidmap
```

- 6 Copy the configuration files from one of the nodes in the existing cluster to the new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \
saturn:/etc/VRTSvcs/conf/config/main.cf
# rcp /etc/VRTSvcs/conf/config/CFSTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CFSTypes.cf
# rcp /etc/VRTSvcs/conf/config/CVMTypes.cf \
saturn:/etc/VRTSvcs/conf/config/CVMTypes.cf
```

- 7 The `/etc/vx/tunefstab` file sets non-default tunables for local-mounted and cluster-mounted file systems.

If you have configured a `/etc/vx/tunefstab` file to tune cluster-mounted file systems on any of the existing cluster nodes, you may want the new node to adopt some or all of the same tunables.

To adopt some or all tunables, review the contents of the file, and copy either the file, or the portions desired, into the `/etc/vx/tunefstab` file on the new cluster node.

Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.

To configure the ClusterService group for the new node

- 1 On an existing node, for example galaxy, write-enable the configuration:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing ClusterService group.

```
# hagrps -modify ClusterService SystemList -add saturn 2
# hagrps -modify ClusterService AutoStartList -add saturn
```

- 3 Modify the IP address and NIC resource in the existing group for the new node.

```
# hares -modify gcoip Device bge0 -sys saturn
# hares -modify gconic Device bge0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

Configuring server-based fencing on the new node

This section describes the procedures to configure server-based fencing on a new node. Depending on whether server-based fencing is configured in secure or non-secure mode on the existing cluster, perform the tasks in one of the following procedures:

- Server-based fencing in non-secure mode:
[To configure server-based fencing in non-secure mode on the new node](#)
- Server-based fencing in secure mode:
[To configure server-based fencing with security on the new node](#)

To configure server-based fencing in non-secure mode on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s mycps1.symantecexample.com \  
-a list_nodes
```

The new node must be listed in the command output.

- 4 Add the VCS user cpsclient@saturn to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e cpsclient@saturn \  
-f cps_operator -g vx
```

```
User cpsclient@saturn successfully added
```

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 mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s mycps1.symantecexample.com -a list_nodes
```

The new node must be listed in the 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.

Updating the Storage Foundation for Databases (SFDB) repository after adding a node

To add the new node to the vxfen group using the CLI

- 1 On one of the nodes in the existing SFCFSHA cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing vxfen group.

```
# hagrps -modify vxfen SystemList -add saturn 2
```

- 3 Save the configuration by running the following command from any node in the SFCFSHA cluster:

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

Updating the Storage Foundation for Databases (SFDB) repository after adding a node

If you are using Database Storage Checkpoints, Database FlashSnap, or SmartTier for Oracle in your configuration, update the SFDB repository to enable access for the new node after it is added to the cluster.

To update the SFDB repository after adding a node

- 1 Copy the `/var/vx/vxdba/rep_loc` file from one of the nodes in the cluster to the new node.
- 2 If the `/var/vx/vxdba/auth/user-authorizations` file exists on the existing cluster nodes, copy it to the new node.

If the `/var/vx/vxdba/auth/user-authorizations` file does not exist on any of the existing cluster nodes, no action is required.

This completes the addition of the new node to the SFDB repository.

For information on using SFDB tools features, refer to *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases*.

Sample configuration file for adding a node to the cluster

You may use this sample file as reference information to understand the configuration changes that take place when you add a node to a cluster.

The existing sample configuration before adding the node `saturn` is as follows:

- The existing cluster `rac_cluster101` comprises two nodes `galaxy` and `nebula` and hosts a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle. The agent starts, stops, and monitors the database.
- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

The following sample configuration file shows the changes (in **bold**) effected in the configuration after adding a node "saturn" to the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)
system nebula (
)
system saturn (
)
```

Note: In the following group `oradb1_grp`, the `saturn` node has been added.

```
group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
    AutoFailOver = 0
    Parallel = 1
```

```
AutoStartList = { galaxy, nebula, saturn }
)
```

Note: In the following Oracle resource, the saturn node information has been added.

```
Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Sid @saturn = vrts3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
oral requires oradata_mnt
oradata_mnt requires oradata_voldg
```

Note: In the following CVM and CVMCluster resources, the saturn node information has been added.

```
group cvm (
    SystemList = { galaxy = 0, nebula = 1, saturn =2}
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
```

```

)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
    OnlineRetryLimit = 20
)

CFSSMount ocrvote_mnt (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
    MountOpt= "mincache=direct"
)

CVMVolDg ocrvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1, saturn =2 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

```

Note: In the following PrivNIC resource, the saturn node information has been added.

```
PrivNIC ora_priv (  
    Critical = 0  
    Device@galaxy = { bge1 = 0, bge2 = 1}  
    Device@nebula = { bge1 = 0, bge2 = 1}  
    Device@saturn = { bge1 = 0, bge2 = 1}  
    Address@galaxy = "192.168.12.1"  
    Address@nebula = "192.168.12.2"  
    Address@saturn = "192.168.12.5"  
    NetMask = "255.255.255.0"  
)  
  
cssd requires ocrvote_mnt  
cssd requires ora_priv  
ocrvote_mnt requires ocrvote_voldg  
ocrvote_mnt requires vxfsckd  
ocrvote_voldg requires cvm_clus  
vxfsckd requires cvm_clus  
cvm_clus requires cvm_vxconfigd
```

Removing a node from Storage Foundation Cluster File System clusters

This chapter includes the following topics:

- [About removing a node from a cluster](#)
- [Removing a node from a cluster](#)
- [Modifying the VCS configuration files on existing nodes](#)
- [Removing the node configuration from the CP server](#)
- [Removing security credentials from the leaving node](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after removing a node](#)
- [Sample configuration file for removing a node from the cluster](#)

About removing a node from a cluster

You can remove one or more nodes from an SFCFSHA cluster.

Overview of tasks for removing a node from a cluster:

- Prepare the node to be removed.
 - Stop applications that use File System or Cluster File System mount points not configured under VCS.
See [“Removing a node from a cluster”](#) on page 328.
- Remove the node from the cluster.

- Stop VCS on the node to be removed.
- Unmount the File System and Cluster File System file systems not configured under VCS.
- Uninstall SFCFSHA from the node.
See [“Removing a node from a cluster”](#) on page 328.
- Modify the VCS configuration files on the existing nodes.
See [“Modifying the VCS configuration files on existing nodes”](#) on page 329.
- Remove the node configuration from the CP server if it is configured.
- Remove the security credentials from the node if it is part of a secure cluster.
- Update the SFDB repository if you use SFDB tools.
See [“Updating the Storage Foundation for Databases \(SFDB\) repository after removing a node”](#) on page 333.

The Veritas product installer does not support removing a node. You must remove a node manually. The example procedures describe how to remove a node from a cluster with three nodes.

Removing a node from a cluster

Perform the following steps to remove a node from a cluster. The procedure can be done from any node remaining in the cluster or from a remote host.

To prepare to remove a node from a cluster

- 1 Take your application service groups offline (if under VCS control) on the node you want to remove.

```
# hagrps -offline app_group -sys saturn
```

- 2 Stop the applications that use VxFS/CFS mount points and are not configured under VCS. Use native application commands to stop the applications.

To remove a node from a cluster

- 1 Stop VCS on the node:

```
# hastop -local
```

- 2 Unmount the VxFS/CFS file systems that are not configured under VCS.

```
# umount mount_point
```

- 3 Uninstall SFCFSHA from the node using the SFCFSHA installer.

```
# cd /opt/VRTS/install  
  
# ./uninstallsfcfsha saturn
```

The installer stops all SFCFSHA processes and uninstalls the SFCFSHA packages.

- 4 Modify the VCS configuration files on the existing nodes to remove references to the deleted node.

See [“Modifying the VCS configuration files on existing nodes”](#) on page 329.

Modifying the VCS configuration files on existing nodes

Modify the configuration files on the remaining nodes of the cluster to remove references to the deleted nodes.

The process involves:

- Editing the `/etc/llhosts` file
- Editing the `/etc/gabtab` file
- Modifying the VCS configuration to remove the node

For an example `main.cf`:

See [“Sample configuration file for removing a node from the cluster”](#) on page 333.

To edit the `/etc/llhosts` file

- ◆ On each of the existing nodes, edit the `/etc/llhosts` file to remove lines that contain references to the removed nodes.

For example, if `saturn` is the node removed from the cluster, remove the line `"2 saturn"` from the file:

```
0 galaxy  
1 nebula  
2 saturn
```

Change to:

```
0 galaxy  
1 nebula
```

To edit the `/etc/gabtab` file

- ◆ Modify the following command in the `/etc/gabtab` file to reflect the number of systems after the node is removed:

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

where N is the number of remaining nodes in the cluster.

For example, with two nodes remaining, the file resembles:

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

Modify the VCS configuration file `main.cf` to remove all references to the deleted node.

Use one of the following methods to modify the configuration:

- Edit the `/etc/VRTSvcs/conf/config/main.cf` file
 This method requires application down time.
- Use the command line interface
 This method allows the applications to remain online on all remaining nodes.

The following procedure uses the command line interface and modifies the sample VCS configuration to remove references to the deleted node. Run the steps in the procedure from one of the existing nodes in the cluster. The procedure allows you to change the VCS configuration while applications remain online on the remaining nodes.

To modify the VCS configuration using the CLI

- 1 Back up the `/etc/VRTSvcs/conf/config/main.cf` file.

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.cf.3node.bak
```

- 2 Change the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 3 Remove the node from the `AutoStartList` attribute of the service group by specifying the remaining nodes in the desired order:

```
# hagrps -modify cvm AutoStartList galaxy nebula
```

- 4 Remove the node from the `SystemList` attribute of the service group:

```
# hagrps -modify cvm SystemList -delete saturn
```

- 5 Remove the node from the CVMNodeId attribute of the service group:

```
# hares -modify cvm_clus CVMNodeId -delete saturn
```

- 6 If you have the other service groups (such as the database service group or the ClusterService group) that have the removed node in their configuration, perform step 4 and step 5 for each of them.
- 7 Remove the deleted node from the NodeList attribute of all CFS mount resources:

```
# hares -modify CFSMount NodeList -delete saturn
```

- 8 Remove the deleted node from the system list of any other service groups that exist on the cluster. For example, to delete the node saturn:

```
# hagrps -modify appgrp SystemList -delete saturn
```

- 9 Remove the deleted node from the cluster system list:

```
# hasys -delete saturn
```

- 10 Save the new configuration to disk:

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

- 11 Verify that the node is removed from the VCS configuration.

```
# grep -i saturn /etc/VRTSvcs/conf/config/main.cf
```

If the node is not removed, use the VCS commands as described in this procedure to remove the node.

Removing the node configuration from the CP server

After removing a node from a SFCFSHA 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 *Veritas Storage Foundation Cluster File System High Availability 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, using the following command:

```
# cpsadm -s cp_server -a list_users
```

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

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
# cpsadm -s cp_server -a rm_node -h saturn -c clus1 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
# cpsadm -s cp_server -a list_nodes
```

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node saturn. Perform the following steps.

To remove the security credentials

- 1 Stop the AT process.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh \  
stop
```

- 2 Remove the credentials.

```
# rm -rf /var/VRTSvcs/vcsauth/data/
```

Updating the Storage Foundation for Databases (SFDB) repository after removing a node

After removing a node from a cluster, you do not need to perform any steps to update the SFDB repository.

For information on updating the SFDB repository after adding a node to the cluster:

See [“Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node”](#) on page 322.

For information on removing the SFDB repository after removing the product:

See [“Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product”](#) on page 397.

Sample configuration file for removing a node from the cluster

You may use this sample file as reference information to understand the configuration changes involved when you remove a node from a cluster.

The existing sample configuration before removing the node *system3* is as follows:

- The existing cluster *cluster1* comprises three nodes *system1*, *system2*, and *system3* and hosts a single database.
- The database is stored on CFS.
- The database is managed by a VCS database agent.
The agent starts, stops, and monitors the database.

Note: The following sample file shows in **bold** the configuration information that is removed when the node *system3* is removed from the cluster.

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"

cluster cluster1 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)
```

```
cluster cluster1 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system system1 (
)
system system2 (
)
system system3 (
)
```

Note: In the following group *app_grp*, the *system3* node must be removed.

```
group app_grp (
    SystemList = { system1 = 0, system2 = 1, system3 = 2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system1, system2, system3 }
)
```

Note: In the following application resource, the *system3* node information must be removed.

```
App app1 (
    Critical = 0
    Sid @system1 = vrts1
    Sid @system2 = vrts2
    Sid @system3 = vrts3
)

CFSMount appdata_mnt (
    Critical = 0
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/appdatadg/appdatavol"
)

CVMVolDg appdata_voldg (
    Critical = 0
    CVMDiskGroup = appdatadg
    CVMVolume = { appdatavol }
```

```

        CVMActivation = sw
    )

requires group cvm online local firm
appl requires appdata_mnt
appdata_mnt requires appdata_voldg

```

Note: In the following CVM and CVMCluster resources, the *system3* node information must be removed.

```

group cvm (
    SystemList = { system1 = 0, system2 = 1, system3 =2}
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system1, system2, system3 }
)

CFSfsckd vxfsckd (

)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { system1 = 0, system2 = 1, system3 =2 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)

vxfsckd requires cvm_clus
cvm_clus requires cvm_vxconfigd

```

7

Section

Setting up and configuring replicated global cluster

- [Chapter 24. Setting up a replicated global cluster](#)
- [Chapter 25. Configuring a global cluster using VVR](#)

Setting up a replicated global cluster

This chapter includes the following topics:

- [Replication in the SFCFSHA environment](#)
- [About setting up a global cluster in an SFCFSHA environment](#)
- [Configuring an SFCFSHA global cluster at the primary site](#)
- [Configuring an SFCFSHA cluster at the secondary site](#)
- [Configuring replication for SFCFSHA clusters at both sites](#)
- [Modifying the ClusterService group for a global SFCFSHA cluster](#)
- [Defining the remote SFCFSHA cluster and heartbeat objects](#)
- [Configuring the VCS service groups for global SFCFSHA clusters](#)

Replication in the SFCFSHA environment

You can set up a primary SFCFSHA cluster for replication to a secondary SFCFSHA cluster by configuring global VCS service groups and using a replication technology. The application cluster at the secondary site can be a single node cluster. For example, you can have a two-node cluster on the primary site and a two-node or single-node cluster on the secondary site.

You can use one of the following replication technologies:

- Veritas Volume Replicator (VVR), which provides host-based volume replication. Using VVR you can replicate data volumes on a shared disk group in SFCFSHA.

- Supported hardware-based replication technologies. Using hardware-based replication you can replicate data from a primary array to a secondary array.
- Using SFCFSHA with VVR you can run a fire drill to verify the disaster recovery capability of your configuration.
See the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

About setting up a global cluster in an SFCFSHA environment

Configuring a global SFCFSHA cluster requires the coordination of many component setup tasks. The following procedures provide guidelines.

Tasks required to set up a global cluster:

- Configure a cluster at the primary site.
See [“Configuring an SFCFSHA global cluster at the primary site”](#) on page 340.
- Configure a cluster at the secondary site.
See [“Configuring an SFCFSHA cluster at the secondary site”](#) on page 342.
- Configure replication on clusters at both sites.
See [“Configuring replication for SFCFSHA clusters at both sites”](#) on page 344.
- Configure VCS service groups for replication.
See [“Modifying the ClusterService group for a global SFCFSHA cluster”](#) on page 344.
See [“Defining the remote SFCFSHA cluster and heartbeat objects”](#) on page 346.
See [“Configuring the VCS service groups for global SFCFSHA clusters”](#) on page 349.
- Test the HA/DR configuration.
- Upon successful testing, bring the environment into production.

SFCFSHA HA/DR configuration tasks may require adjustments depending upon your particular starting point, environment, and configuration, as the details of your configuration may differ from the examples given in the procedures. Review the installation requirements and sample cluster configuration files for primary and secondary clusters.

Review the requirements information to make sure your configuration is supported for SFCFSHA.

- For product licensing information:
See [“About Veritas SFHA Solutions product licensing”](#) on page 49.

- For supported hardware and software:
 See “[Hardware overview and requirements for Veritas Storage Foundation Cluster File System High Availability](#)” on page 38.
- To confirm the compatibility of your hardware, see the current compatibility list in the Symantec Technical Support website:
<http://www.symantec.com/docs/TECH170013>

SFCFSHA supports the following replication technologies through the use of Veritas replication agents:

Table 24-1 Supported replication options for SFCFSHA global clusters

Replication technology	Supported modes	Supported software
Veritas Volume Replicator (VVR) Supporting agents <ul style="list-style-type: none"> ■ RVGShared ■ RVGSharedPri ■ RVGLogOwner 	<ul style="list-style-type: none"> ■ Asynchronous replication ■ Synchronous replication 	Host-based replication
EMC SRDF Supporting agent: SRDF	<ul style="list-style-type: none"> ■ Asynchronous replication ■ Synchronous replication 	All versions of Solutions Enabler
Hitachi True Copy Supporting agent: HTC	<ul style="list-style-type: none"> ■ Asynchronous replication ■ Synchronous replication 	All versions of the Hitachi CCI
IBM Metro Mirror Supporting agent: MetroMirror	Synchronous replication	All versions of IBM DSCLI. The MetroMirror agent is supported for DS6000 and DS8000 arrays
IBM SVC SVC CopyServices	<ul style="list-style-type: none"> ■ Asynchronous replication ■ Synchronous replication 	SSH access to the SVC

Table 24-1 Supported replication options for SFCFSHA global clusters
 (continued)

Replication technology	Supported modes	Supported software
EMC Mirror View Supporting agent: MirrorView	<ul style="list-style-type: none"> ■ Asynchronous replication ■ Synchronous replication: only individual LUNs may be replicated 	All versions of NaviCLI

Note: Check your vendor's compatibility list for the supported software versions. The support listed above only exists if the host, HBA, and array combination is in your vendor's hardware compatibility list. Check your array documentation.

Note: All arrays must support SCSI-3 persistent reservations for SFCFSHA.

You can use the Veritas replication agents listed in the table above for global clusters that run SFCFSHA. The Veritas replication agents provide application failover and recovery support to your replication configuration. The agents provide this support for environments where data is replicated between clusters.

VCS agents control the direction of replication. They do not monitor the progress or status of replication. The replication agents manage the state of replicated devices that are attached to SFCFSHA nodes. The agents make sure that the system which has the resource online also has safe and exclusive access to the configured devices.

For instructions for configuring AT in a global cluster:

See the *Veritas Cluster Server Administrator's Guide*

Configuring an SFCFSHA global cluster at the primary site

You can use an existing SFCFSHA cluster or you can install a new SFCFSHA cluster for your primary site.

For planning information:

See [“About planning for SFCFSHA installation”](#) on page 26.

If you are using an existing cluster as the primary and you want to set up a global cluster, skip the steps below and proceed to configure your secondary cluster.

See [“Configuring an SFCFSHA cluster at the secondary site”](#) on page 342.

Note: You must have a GCO license enabled for a global cluster. If you are using VVR for replication, you must have a VVR license enabled.

If you do not have an existing cluster and you are setting up two new sites for an SFCFSHA global cluster, follow the steps below.

To set up the cluster and database at the primary site

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.
- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option enabled for a global cluster. If you are using VVR for replication, you must have it enabled.

- 4 Install and configure SFCFSHA.

For preparation:

See [“Prerequisites for Veritas Storage Foundation Cluster File System High Availability”](#) on page 37.

For installation:

See [“About the Web-based installer”](#) on page 61.

For configuration:

For a multi-node cluster, configure I/O fencing.

- 5 Verify the CVM group is online on all nodes in the primary cluster:

```
# hagrpd -state cvm
```

- 6 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing your database software.

You will need to set up:

- Local storage for database software
- Shared storage for resources which are not replicated as part of the hardware-based or host-based replication

- Replicated storage for database files
- 7 Install and configure the database binaries. Consult your database documentation.

Note: Resources which will not be replicated must be on non-replicated shared storage.

After successful database installation and configuration, verify that database resources are up on all nodes.

- 8 Identify the disks that will be replicated, create the required CVM disk group, volume, and file system.
- 9 Create the database on the file system you created in the previous step.
- 10 Configure the VCS service groups for the database.
- 11 Verify that all VCS service groups are online.

Configuring an SFCFSHA cluster at the secondary site

The setup requirements for the secondary site parallel the requirements for the primary site with a few additions or exceptions as noted below.

Important requirements for global clustering:

- Cluster names on the primary and secondary sites must be unique.
- You must use the same OS user and group IDs for your database for installation and configuration on both the primary and secondary clusters.

To set up the cluster on secondary site

- 1 Install and configure servers and storage.
- 2 If you are using hardware-based replication, install the software for managing your array.
- 3 Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster. If you are using VVR for replication, you must have it enabled.

- 4 Install and configure SFCFSHA.
For preparation:
See [“Prerequisites for Veritas Storage Foundation Cluster File System High Availability”](#) on page 37.
For installation:
See [“About the Web-based installer”](#) on page 61.
For configuration:
For a multi-node cluster, configure I/O fencing.
- 5 For a single-node cluster, do not enable I/O fencing. Fencing will run in disabled mode.
- 6 Prepare systems and storage for a global cluster. Identify the hardware and storage requirements before installing your database software.
You will need to set up:
 - Local storage for database software
 - Shared storage for resources which are not replicated
 - Replicated storage for database files
- 7 Install and configure the database binaries. Consult your database documentation.

Note: Resources which will not be replicated must be on non-replicated shared storage.

After successful database installation and configuration, verify that database resources are up on all nodes.

To set up the database for the secondary site

- 1 Do not create the database. The database will be replicated from the primary site.
 - If you are using hardware-based replication, the database, disk group, and volumes will be replicated from the primary site.
Create the directory for the CFS mount point which will host the database data and control files.
 - If you are using VVR for replication, create an identical disk group and volumes for the replicated content with the same names and size as listed on the primary site.

Create the directories for the CFS mount points as they are on the primary site. These will be used to host the database and control files when the failover occurs and the secondary is promoted to become the primary site.

- 2 Copy the `init$ORACLE_SID.ora` file from `$ORACLE_HOME/dbs` at the primary to `$ORACLE_HOME/dbs` at the secondary.
- 3 Create subdirectories for the database as you did on the primary site.

Configuring replication for SFCFSHA clusters at both sites

You must configure replication for the database files. Once replication is configured, make sure it is functioning correctly by testing before proceeding.

To configure replication at both sites

- 1 At both sites, identify the disks on which the database resides at the primary site and associate them with the corresponding disks at the secondary site.
[See “Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): setting up the primary site”](#) on page 351.
- 2 Start replication between the sites.
[See “Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): starting replication”](#) on page 359.

Modifying the ClusterService group for a global SFCFSHA cluster

You have configured VCS service groups for the database on each cluster. Each cluster requires an additional virtual IP address associated with the cluster for cross-cluster communication. The VCS installation and creation of the ClusterService group typically involves defining this IP address.

Configure a global cluster by setting:

- Heartbeat
- Wide area cluster (wac)
- GCO IP (gcoip)
- remote cluster resources

See the *Veritas Cluster Server Administrator's Guide* for complete details on global clustering.

Modifying the global SFCFSHA cluster configuration using the wizard

The global clustering wizard completes the following tasks:

- Validates the ability of the current configuration to support a global cluster environment.
- Creates the components that enable the separate clusters, each of which contains a different set of GAB memberships, to connect and operate as a single unit.
- Creates the ClusterService group, or updates an existing ClusterService group.

Run the global clustering configuration wizard on each of the clusters; you must have the global clustering license in place on each node in the cluster.

To modify the ClusterService group for global clusters using the global clustering wizard

- 1 On the primary cluster, start the GCO Configuration wizard:

```
# /opt/VRTSvcs/bin/gcoconfig
```

- 2 The wizard discovers the NIC devices on the local system and prompts you to enter the device to be used for the global cluster. Specify the name of the device and press Enter.
- 3 If you do not have NIC resources in your configuration, the wizard asks you whether the specified NIC will be the public NIC used by all the systems. Enter **y** if it is the public NIC; otherwise enter **n**. If you entered **n**, the wizard prompts you to enter the names of NICs on all systems.
- 4 Enter the virtual IP address for the local cluster.
- 5 If you do not have IP resources in your configuration, the wizard prompts you for the netmask associated with the virtual IP. The wizard detects the netmask; you can accept the suggested value or enter another one.

The wizard starts running commands to create or update the ClusterService group. Various messages indicate the status of these commands. After running these commands, the wizard brings the ClusterService failover group online on any one of the nodes in the cluster.

Defining the remote SFCFSHA cluster and heartbeat objects

After configuring global clustering, add the remote cluster object to define the IP address of the cluster on the secondary site, and the heartbeat object to define the cluster-to-cluster heartbeat.

Heartbeats monitor the health of remote clusters. VCS can communicate with the remote cluster only after you set up the heartbeat resource on both clusters.

To define the remote cluster and heartbeat

- 1 On the primary site, enable write access to the configuration:

```
# haconf -makerw
```

- 2 On the primary site, define the remote cluster and its virtual IP address.

In this example, the remote cluster is `clus2` and its IP address is `10.11.10.102`:

```
# haclus -add clus2 10.11.10.102
```

- 3 Complete step 1 and step 2 on the secondary site using the name and IP address of the primary cluster.

In this example, the primary cluster is `clus1` and its IP address is `10.10.10.101`:

```
# haclus -add clus1 10.10.10.101
```

- 4 On the primary site, add the heartbeat object for the cluster. In this example, the heartbeat method is ICMP ping.

```
# hahb -add Icmp
```

- 5 Define the following attributes for the heartbeat resource:

- `ClusterList` lists the remote cluster.
- Arguments enable you to define the virtual IP address for the remote cluster.

For example:

```
# hahb -modify Icmp ClusterList clus2  
# hahb -modify Icmp Arguments 10.11.10.102 -clus clus2
```

- 6 Save the configuration and change the access to read-only on the local cluster:

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

- 7 Complete step 4-6 on the secondary site using appropriate values to define the cluster on the primary site and its IP as the remote cluster for the secondary cluster.

- 8 Verify cluster status with the `hastatus -sum` command on both clusters.

```
# hastatus -sum
```

9 Display the global setup by executing `haclus -list` command.

```
# haclus -list
    clus1
    clus2
```

Example of heartbeat additions to the `main.cf` file on the primary site:

```
.
.
remotecluster clus2 (
Cluster Address = "10.11.10.102"
)
heartbeat Icmp (
    ClusterList = { clus2 }
    Arguments @clus2 = { "10.11.10.102" }
)

system galaxy (
)

.
.
```

Example heartbeat additions to the `main.cf` file on the secondary site:

```
.
.
remotecluster clus1 (
    Cluster Address = "10.10.10.101"
)

heartbeat Icmp (
    ClusterList = { clus1 }
    Arguments @clus1 = { "10.10.10.101" }
)

system mercury (
)

.
.
```

See the *Veritas Cluster Server Administrator's Guide* for details for configuring the required and optional attributes of the heartbeat object.

Configuring the VCS service groups for global SFCFSHA clusters

To configure VCS service groups for global clusters

- 1 Configure and enable global groups for databases and resources.
 - Configure VCS service groups at both sites.
 - Configure the replication agent at both sites.
 - For example:
See [“Modifying the VCS Configuration on the Primary Site”](#) on page 363.
- 2 To test real data in an environment where HA/DR has been configured, schedule a planned migration to the secondary site for testing purposes.
For example:
See [“Migrating the role of primary site to the secondary site”](#) on page 373.

- 3 Upon successful testing, bring the environment into production.

For more information about VCS replication agents:

See the *Veritas Cluster Server Bundled Agents Guide*

For complete details on using VVR in a shared disk environment:

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator’s Guide*.

Configuring a global cluster using VVR

This chapter includes the following topics:

- [About configuring a global SFCFSHA cluster using VVR for replication](#)
- [Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): setting up the primary site](#)
- [Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): setting up the secondary site](#)
- [Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): starting replication](#)
- [Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): configuring cluster resources](#)
- [Managing a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\)](#)

About configuring a global SFCFSHA cluster using VVR for replication

Before configuring clusters for global clustering, make sure both clusters have product and database software installed and configured.

Verify that you have the correct installation options enabled, whether you are using keyless licensing or installing keys manually. You must have the GCO option for a global cluster and VVR enabled.

See [“About Veritas SFHA Solutions product licensing”](#) on page 49.

After setting up two clusters running SFCFSHA, you can configure a global cluster environment with VVR. You must modify both cluster configurations to support replication in the global cluster environment.

Configuring SFCFSHA for global clusters requires:

- Setting up both clusters as part of a global cluster environment.
See [“About setting up a global cluster in an SFCFSHA environment”](#) on page 338.
- Setting up replication for clusters at both sites.
See [“Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): setting up the primary site”](#) on page 351.
See [“Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): setting up the secondary site”](#) on page 354.
- Starting replication of the database.
See [“Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): starting replication”](#) on page 359.
- Configuring VCS for replication on clusters at both sites.
See [“Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): configuring cluster resources”](#) on page 361.

Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the primary site

Setting up replication with VVR in a global cluster environment involves the following tasks:

- If you have not already done so, create a disk group to hold data volume, SRL, and RVG on the storage on the primary site. For example:
- Creating the Storage Replicator Log (SRL) in the disk group for the database.
See [“Creating the data and SRL volumes on the primary site”](#) on page 351.
- Creating the Replicated Volume Group (RVG) on the primary site.
See [“Setting up the Replicated Volume Group \(RVG\) on the primary site”](#) on page 353.
- Editing the `/etc/vx/vras/.rdg` files on the primary site after creating the data and SRL volumes on the secondary site.
See [“Editing the `/etc/vx/vras/.rdg` files”](#) on page 355.

Creating the data and SRL volumes on the primary site

Create the data volume if you do not have one already.

Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the primary site

- The data volume on the secondary site has the same name and the same size as the data volume on the primary site.
- The data volume and SRL volume should exist in the same disk group.
- Mirror the data volume in the absence of hardware-based mirroring.

To create the data volume on the primary site

- ◆ In the disk group created for the application database, create a data volume of same size as that in primary for data; in this case, the *rac_vol1* volume on the primary site is 6.6 GB:

```
# vxassist -g oradatadg make rac_vol1 6600M nmirror=2 disk1 disk2
```

Create the SRL. The SRL is a volume in the RVG. The RVG also holds the data volumes for replication.

- The SRL on the secondary site has the same name and the same size as the SRL on the primary site.
- If possible, create SRLs on disks without other volumes.
- Mirror SRLs and in the absence of hardware-based mirroring.

To create the SRL volume on the primary site

- 1 On the primary site, determine the size of the SRL volume based on the configuration and amount of use.

See the Veritas Volume Replicator documentation for details.

- 2 Using the following command, determine whether a node is the CVM master or the slave:

```
# vxdctl -c mode
```

Configuring a global SFCSHA cluster using Veritas Volume Replicator (VVR): setting up the primary site

- 3 On the CVM master node, issue the following command:

```
# vxassist -g oradatadg make rac1_srl 6800M nmirror=2 disk4 disk5
```

Note: Assume that for the example setup that *disk4* and *disk5* are already added and are part of the same disk group. They are used in this step for mirroring and creation of the SRL.

- 4 If the SRL volume is not already started, start the SRL volume by starting all volumes in the disk group:

```
# vxvol -g oradatadg startall
```

Setting up the Replicated Volume Group (RVG) on the primary site

Before creating the RVG on the primary site, make sure the volumes and CVM group are active and online.

To review the status of replication objects on the primary site

- 1 Verify the volumes you intend to include in the group are active.
- 2 Review the output of the `hagrp -state cvm` command to verify that the CVM group is online.
- 3 On each site, verify `vradmin` is running:

```
# ps -ef |grep vradmin
  root   536594   598036   0 12:31:25      0 0:00 grep vradmin
```

If `vradmin` is not running start it:

```
# vxstart_vvr
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradmin: 8199
vxrsyncd: 8989
data: Anonymous-Ports
To change, see vrport(1M) command
# ps -ef |grep vradmin
  root   536782     1    0 12:32:47      - 0:00 /usr/sbin/vradmin
  root  1048622   598036   0 12:32:55      0 0:00 grep vradmin
# netstat -an |grep 4145
tcp4      0      0 *.4145          *.*             LISTEN
udp4      0      0 *.4145          *.*
```

The command to create the primary RVG takes the form:

```
vradmin -g disk_group createpri rvg_name data_volume srl_volume
```

where:

- `disk_group` is the name of the disk group containing the database
- `rvg_name` is the name for the RVG
- `data_volume` is the volume that VVR replicates
- `srl_volume` is the volume for the SRL

To create the primary RVG

- 1 Determine which node is the CVM master node by entering:

```
# vxdctl -c mode
```

- 2 To create the `rac1_rvg` RVG, you must run the following on the master node:

```
# vradmin -g oradatadg createpri rac1_rvg rac1_vol rac1_srl
```

The command creates the RVG on the primary site and adds a Data Change Map (DCM) for each data volume. In this case, a DCM exists for `rac1_vol`.

Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the secondary site

To create objects for replication on the secondary site, use the `vradmin` command with the `addsec` option. To set up replication on the secondary site, perform the following tasks:

- Create a disk group to hold the data volume, SRL, and RVG on the storage on the secondary site. You must match the names and sizes of these volumes with the names and sizes of the volumes on the primary site.
See [“Creating the data and SRL volumes on the secondary site”](#) on page 355.
- Edit the `/etc/vx/vras/.rdg` file on the secondary site.
See [“Editing the /etc/vx/vras/.rdg files”](#) on page 355.
- Use resolvable virtual IP addresses that set network RLINK connections as host names of the primary and secondary sites.
See [“Setting up IP addresses for RLINKs on each cluster”](#) on page 356.

- Create the replication objects on the secondary site.
See [“Setting up the disk group on secondary site for replication”](#) on page 357.

Creating the data and SRL volumes on the secondary site

Note the following when creating volumes for the data and SRL:

- The sizes and names of the volumes must match the sizes and names of the corresponding volumes in the primary site before you create the disk group.
- The disk group must match the size and name of the disk group at the primary site.
- Create the data and SRL volumes on different disks in the disk group. Use the `vxdisk -g diskgroup list` command to list the disks in the disk group.
- Mirror the volumes.

To create the data and SRL volumes on the secondary site

- 1 In the disk group created for the application database, create a data volume of same size as that in primary for data; in this case, the `rac_vol1` volume on the primary site is 6.6 GB:

```
# vxassist -g oradatadg make rac_vol1 6800M nmirror=2 disk1 disk2
```

- 2 Create the volume for the SRL, using the same name and size of the equivalent volume on the primary site. Create the volume on different disks from the disks for the database volume, but on the same disk group that has the data volume:

```
# vxassist -g oradatadg make rac1_srl 1500M nmirror=2 disk4 disk6
```

Editing the `/etc/vx/vras/.rdg` files

Editing the `/etc/vx/vras/.rdg` file on the secondary site enables VVR to replicate the disk group from the primary site to the secondary site. On each node, VVR uses the `/etc/vx/vras/.rdg` file to check the authorization to replicate the RVG on the primary site to the secondary site. The file on each node in the secondary site must contain the primary disk group ID, and likewise, the file on each primary system must contain the secondary disk group ID.

To edit the /etc/vx/vras/.rdg files

- 1 On a node in the primary site, display the primary disk group ID:

```
# vxprint -l diskgroup
.....
```

- 2 On each node in the secondary site, edit the /etc/vx/vras/.rdg file and enter the primary disk group ID on a single line.
- 3 On each cluster node of the primary cluster, edit the /etc/vx/vras/.rdg file and enter the secondary disk group ID on a single line.

Setting up IP addresses for RLINKs on each cluster

Creating objects with the vradm command requires resolvable virtual IP addresses that set network RLINK connections as host names of the primary and secondary sites.

To set up IP addresses for RLINKs on each cluster

- 1 Using the following command, determine whether a node is the CVM master or the slave:

```
# vxdctl -c mode
```

You must configure rlinks on the CVM master node.

- 2 For each RVG running on each cluster, set up a virtual IP address on one of the nodes of the cluster. These IP addresses are part of the RLINK.

The example assumes for the cluster on the primary site:

- The public network interface is bge0:1
- The virtual IP address is 10.10.9.101
- The net mask is 255.255.255.0
- ```
ifconfig bge0:1 plumb
ifconfig bge0:1 inet 10.10.9.101 netmask 255.255.255.0
ifconfig bge0:1 up
```

- 3 Use the same commands with appropriate values for the interface, IP address, and net mask on the secondary site.

The example assumes for the secondary site:

- The public network interface is bge0:1

**Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the secondary site**

- virtual IP address is 10.11.9.102
  - net mask is 255.255.255.0
- 4 Define the virtual IP addresses to correspond to a host name in the virtual cluster on the primary site and a host name in the virtual cluster on the secondary site.

Update the `/etc/hosts` file on all the nodes on both the primary and secondary sites.

The examples assume:

- `clus1` has IP address 10.10.9.101
  - `clus2` has IP address 10.11.9.102
- 5 Use the `ping` command to verify the links are functional.

## Setting up the disk group on secondary site for replication

Create the replication objects on the secondary site from the master node of the primary site, using the `vradmin` command.

### To set up the disk group on the secondary site for replication

- 1 Issue the command in the following format from the cluster on the primary site:

```
vradmin -g dg_pri addsec rvg_pri pri_host sec_host
```

where:

- `dg_pri` is the disk group on the primary site that VVR will replicate. For example: `rac1_vol`
- `rvg_pri` is the RVG on the primary site. For example: `rac1_rvg`
- `pri_host` is the virtual IP address or resolvable virtual host name of the cluster on the primary site.  
For example: `clus1_1`
- `sec_host` is the virtual IP address or resolvable virtual host name of the cluster on the secondary site.  
For example: `clus2_1`

For example, the command to add the cluster on the primary site to the Replicated Data Set (RDS) is:

```
vradmin -g oradatadg addsec rac1_rvg \
```

## Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the secondary site

```
clus1_1
```

```
clus2_1
```

On the secondary site, the above command performs the following tasks:

- Creates an RVG within the specified disk group using the same name as the one for the primary site
- Associates the data and SRL volumes that have the same names as the ones on the primary site with the specified RVG
- Adds a data change map (DCM) for the data volume
- Creates cluster RLINKS for the primary and secondary sites with the default names

2 Verify the list of RVGs in the RDS by executing the following command.

```
vradmin -g oradatadg -l printrvg
```

For example:

```
Replicated Data Set: rac1_rvg
Primary:
HostName: 10.180.88.187 <localhost>
RvgName: rac1_rvg
DgName: rac1_vol
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_clus2_1_rac1_rvg, detached=on,
synchronous=off
Secondary:
HostName: 10.190.99.197
RvgName: rac1_rvg
DgName: oradatadg
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_clus1_1_rac1_rvg, detached=on,
synchronous=off
```

---

**Note:** Once the replication is started the value of the detached flag will change the status from ON to OFF.

---

## Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): starting replication

When you have both the primary and secondary sites set up for replication, you can start replication from the primary site to the secondary site.

Start with the default replication settings:

- Mode of replication: `synchronous=off`
- Latency Protection: `latencyprot=off`
- SRL overflow protection: `srprot_autodcm`
- Packet size: `packet_size=8400`
- Network protocol: `protocol=UDP`

Method of initial synchronization:

- Automatic synchronization
- Full synchronization with Storage Checkpoint

For guidelines on modifying these settings and information on choosing the method of replication for the initial synchronization:

See the *Veritas Volume Replicator Administrator's Guide*

### Starting replication using automatic synchronization

Use the `vradmin` command to start replication or the transfer of data from the primary site to the secondary site over the network. Because the cluster on the secondary site uses only one host name, the command does not require the `sec_host` argument.

**To start replication using automatic synchronization**

- ◆ From the primary site, use the following command to automatically synchronize the RVG on the secondary site:

```
vradmin -g disk_group -a startrep pri_rvg sec_host
```

where:

- `disk_group` is the disk group on the primary site that VVR will replicate

- `pri_rvg` is the name of the RVG on the primary site
- `sec_host` is the virtual host name for the secondary site

For example:

```
vradmin -g oradatadg -a startrep rac1_rvg
clus2
```

## Starting replication using full synchronization with Storage Checkpoint

Use the `vradmin` command with the Storage Checkpoint option to start replication using full synchronization with Storage Checkpoint.

### To start replication using full synchronization with Storage Checkpoint

- 1 From the primary site, synchronize the RVG on the secondary site with full synchronization (using the `-c checkpoint` option):

```
vradmin -g disk_group -full -c ckpt_name syncrvg pri_rvg sec_host
```

where:

- `disk_group` is the disk group on the primary site that VVR will replicate
- `ckpt_name` is the name of the Storage Checkpoint on the primary site
- `pri_rvg` is the name of the RVG on the primary site
- `sec_host` is the virtual host name for the secondary site

For example:

```
vradmin -g oradatadg -c rac1_ckpt syncrvg rac1_rvg
clus2
```

- 2 To start replication after full synchronization, enter the following command:

```
vradmin -g oradatadg -c rac1_ckpt startrep rac1_rvg
clus2
```

## Verifying replication status

Verify that replication is properly functioning.

**To verify replication status**

- 1 Check the status of VVR replication:

```
vradmin -g disk_group_name repstatus rvg_name
```

- 2 Review the `flags` output for the status. The output may appear as `connected` and `consistent`. For example:

```
vxprint -g oradatadg -l rlk_clus2_1_rac1_rvg
Rlink: rlk_clus2_1_rac1_rvg
info: timeout=500 packet_size=8400 rid=0.1078
 latency_high_mark=10000 latency_low_mark=9950
 bandwidth_limit=none
state: state=ACTIVE
 synchronous=off latencyprot=off srlprot=autodcm
.
.
protocol: UDP/IP
checkpoint: rac1_ckpt
flags: write enabled attached consistent connected
asynchronous
```

## Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): configuring cluster resources

After configuring both clusters for global clustering and setting up the application database for replication, configure VCS to provide high availability for the database. Specifically, configure VCS agents to control the cluster resources, including the replication resources.

---

**Note:** The example procedures illustrate the configuration process using a manual file editing method. If you are using the Java Console, some steps do not apply in the same order.

---

### About modifying the VCS configuration for replication

The following resources must be configured or modified for replication:

**Table 25-1** Cluster resources required for replication

| Cluster resources                  | Configuration required                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Log owner group                    | <p>Create a log owner group including the RVGLogowner resources. The RVGLogowner resources are used by:</p> <ul style="list-style-type: none"> <li>■ RLINKs for the RVG</li> <li>■ RVGLogowner resource. The RVG and its associated disk group are defined as attributes for the RVGLogowner resource.</li> </ul> <p>The RVG log owner service group has an online local firm dependency on the service group containing the RVG.</p> <p>The VCS uses the following agents to control the following resources:</p> <ul style="list-style-type: none"> <li>■ RVGLogowner agent to control the RVGLogowner resource</li> <li>■ RVGShared agent to control the RVGShared resource</li> </ul> |
| RVG group                          | <p>Create an RVG group that includes the RVGShared resource replication objects. Define the RVGShared resource and CMMVolDg resource together within a parallel service group. The group is defined as parallel because it may be online at the same time on all cluster nodes.</p>                                                                                                                                                                                                                                                                                                                                                                                                       |
| CMMVolDg resource                  | <p>The CMMVolDg resource does not have volumes specified for the CMMVolume attribute; the volumes are contained in the RVG resource. The CMMVolume attribute for the CMMVolDg resource is empty because all volumes in the RVG are defined by the RVG attribute of the RVGShared resource. The RVG service group has an online local firm dependency on the CMM service group.</p> <p>For a detailed description of the CMMVolDg agent in this guide: See "<a href="#">CMMVolDg agent</a>" on page 472.</p>                                                                                                                                                                               |
| RVGSharedPri resource              | <p>Add the RVGSharedPri resource to the existing application database service group. The CMMVolDg resource must be removed from the existing application database service group.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| application database service group | <p>The existing application database service group is a parallel group consisting of the application database resource, CMMVolDg resource, and CFMount resource (if the database resides in a cluster file system). Define the application service group as a global group by specifying the clusters on the primary and secondary sites as values for the ClusterList group attribute.</p>                                                                                                                                                                                                                                                                                               |

For more information on service replication resources:

See the *Veritas™ Cluster Server Bundled Agents Reference Guide*.

## Modifying the VCS Configuration on the Primary Site

The following are the procedural highlights required to modify the existing VCS configuration on the primary site:

- Configure two service groups:
  - A log owner group including the RVGLogowner resource.
  - An RVG group including the RVGShared resource replication objects.
- Add the RVGSharedPri resource to the existing application database service group and define this group as a global group by setting the ClusterList and ClusterFailOverPolicy attributes.
- Move the CVMVoIdg resource from the existing application database service group to the newly created RVG group.

---

**Note:** The example procedure illustrates the configuration process using a manual file editing method. If you are using the Java Console, some steps do not apply in the same order.

---

### To modify VCS on the primary site

- 1 Log into one of the nodes on the primary cluster.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while you make changes:

```
haconf -dump -makero
```

- 3 Use the following command to make a backup copy of the main.cf file:

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

- 4 Use vi or another text editor to edit the main.cf file. Review the sample configuration file after the SFCFSHA installation.

Add a failover service group using the appropriate values for your cluster and nodes. Include the following resources:

- RVGLogowner resource. The node on which the group is online functions as the log owner (node connected to the second cluster for the purpose of replicating data).

**Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): configuring cluster resources**

- IP resource
- NIC resources

The following are examples of RVGLogowner service group for the different platforms.

```
group rlogowner (
 SystemList = { galaxy = 0, nebula = 1 }
 AutoStartList = { galaxy,nebula }
)

IP logowner_ip (
 Device = bge0
 Address = "10.10.9.101"
 NetMask = "255.255.255.0"
)

NIC nic (
 Device = bge0
 NetworkType = ether
 NetworkHosts = "10.10.8.1"
)

RVGLogowner logowner (
 RVG = rac1_rvg
 DiskGroup = oradatadg
)
requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic
```

- 5 Add the RVG service group using the appropriate values for your cluster and nodes.

Example `RVGgroup` service group:

```
group RVGgroup (
 SystemList = { galaxy = 0, nebula = 1 }
 Parallel = 1
 AutoStartList = { galaxy,nebula }
)

RVGShared racdata_rvg (
 RVG = rac1_rvg
 DiskGroup = oradatadg
)
 CVMVolDg racdata_voldg (
 CVMDiskGroup = oradatadg
 CVMActivation = sw
)
requires group cvm online local firm
racdata_rvg requires racdata_voldg
```

- 6 Modify the application service group using the appropriate values for your cluster and nodes:
- Define the application service group as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute. See the attribute in bold in the example that follows.

---

**Note:** This action must be performed on the primary or secondary site, but not on both.

---



---

**Note:** If you are using the Java Console, the secondary cluster must be configured before you can configure a service group as a global group.

---

- Add the `ClusterFailOverPolicy` cluster attribute. Symantec recommends using the `Manual` value. See the attribute in bold in the example.
- Add the `RVGSharedPri` resource to the group configuration.
- Remove the `CVMVolDg` resource, if it has been configured in your previous configuration. This resource is now part of the RVG service group.
- Specify the service group (`online`, `local`, `firm`) to depend on the RVG service group.

## Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): configuring cluster resources

- Remove the existing dependency of the Database service group on the CVM service group. Remove the line:

```
requires group CVM online local firm
```

- Remove the existing dependency between the CFSSMount for the database and the CVMVoldg for the application database. Remove the line:

```
oradata_mnt requires oradata_voldg
```

The following is an example of an application database service group configured for replication:

```
group database_grp (
 SystemList = { galaxy = 0, nebula = 1 }
 ClusterList = { clus1 = 0, clus2 = 1 }
 Parallel = 1
 ClusterFailOverPolicy = Manual
 Authority = 1
 AutoStartList = { galaxy,nebula }
 OnlineRetryLimit = 3
 TriggerResStateChange = 1
 OnlineRetryInterval = 120
)

CFSSMount oradata_mnt (
 MountPoint = "/oradata"
 BlockDevice = "/dev/vx/dsk/oradatadg/rac1_vol"
)

Process vxfend (
 PathName = "/sbin/vxfend"
 Arguments = "-m sybase -k /tmp/vcmp_socket"
)

RVGSharedPri ora_vvr_shpri (
 RvgResourceName = racdata_rvg
 OnlineRetryLimit = 0
)

requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri
```

### 7 Save and close the main.cf file.

- 8 Use the following command to verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file:

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

- 9 Stop and restart VCS.

```
hastop -all -force
```

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
hastart
```

## Modifying the VCS Configuration on the Secondary Site

The following are highlights of the procedure to modify the existing VCS configuration on the secondary site:

- Add the log owner and RVG service groups.
- Add a service group to manage the application database and the supporting resources.
- Define the replication objects and agents, such that the cluster at the secondary site can function as a companion to the primary cluster.

The following steps are similar to those performed on the primary site.

---

**Note:** The example procedure illustrates the configuration process using a manual file editing method. If you are using the Java Console, some steps do not apply in the same order.

---

### To modify VCS on the secondary site

- 1 Log into one of the nodes on the secondary site as root.
- 2 Use the following command to save the existing configuration to disk, and make the configuration read-only while making changes:

```
haconf -dump -makero
```

- 3 Use the following command to make a backup copy of the `main.cf` file:

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

- 4 Use vi or another text editor to edit the main.cf file. Edit the CVM group on the secondary site.

Review the sample configuration file after the SFCFSHA installation to see the CVM configuration.

In our example, the secondary site has `clus2` consisting of the nodes `mercury` and `jupiter`. To modify the CVM service group on the secondary site, use the CVM group on the primary site as your guide.

- 5 Add a failover service group using the appropriate values for your cluster and nodes. Include the following resources:
  - `RVGLogowner` resource. The node on which the group is online functions as the log owner (node connected to the second cluster for the purpose of replicating data).
  - `IP` resource
  - `NIC` resources

**Example `RVGLogowner` service group:**

```
group rlogowner (
 SystemList = { mercury = 0, jupiter = 1 }
 AutoStartList = { mercury, jupiter }
)

IP logowner_ip (
 Device = bge0
 Address = "10.11.9.102"
 NetMask = "255.255.255.0"
)

NIC nic (
 Device = bge0
 NetworkHosts = { "10.10.8.1" }
 NetworkType = ether
)

RVGLogowner logowner (
 RVG = rac1_rvg
 DiskGroup = oradatadg
)

requires group RVGgroup online local firm
logowner requires logowner_ip
logowner_ip requires nic
```

- 6 Add the RVG service group using the appropriate values for your cluster and nodes.

The following is an example `RVGgroup` service group:

```
group RVGgroup (
 SystemList = { mercury = 0, jupiter = 1 }
 Parallel = 1
 AutoStartList = { mercury, jupiter }
)

RVGShared racdata_rvg (
 RVG = rac1_rvg
 DiskGroup = oradatadg
)

CVMVolDg racdata_voldg (
 CVMDiskGroup = oradatadg
 CVMActivation = sw
)

requires group cvm online local firm
racdata_rvg requires racdata_voldg
```

- 7 Add an application service group. Use the application service group on the primary site as a model for the application service group on the secondary site.
- Define the application service group as a global group by specifying the clusters on the primary and secondary sites as values for the `ClusterList` group attribute.

---

**Note:** This action must be performed on the primary or secondary site, but not on both.

---

- Assign this global group the same name as the group on the primary site; for example, *database\_grp*.
- Include the `ClusterList` and `ClusterFailOverPolicy` cluster attributes. Symantec recommends using the `Manual` value.
- Add the `RVGSharedPri` resource to the group configuration.
- Remove the `CVMVolDg` resource, if it has been configured in your previous configuration. This resource is now part of the RVG service group.

- Specify the service group to depend (online, local, firm) on the RVG service group.

Example of the application group on the secondary site:

```
group database_grp (
 SystemList = { mercury = 0, jupiter = 1 }
 ClusterList = { clus2 = 0, clus1 = 1 }
 Parallel = 1
 OnlineRetryInterval = 300
 ClusterFailOverPolicy = Manual
 Authority = 1
 AutoStartList = { mercury, jupiter }
)

RVGSharedPri ora_vvr_shpri (
 RvgResourceName = racdata_rvg
 OnlineRetryLimit = 0
)

CFSMount oradata_mnt (
 MountPoint = "/oradata"
 BlockDevice = "/dev/vx/dsk/oradatadg/racdb_vol"
 Critical = 0
)

Process vxfend (
 PathName = "/sbin/vxfend"
 Arguments = "-m sybase -k /tmp/vcmp_socket"
)

 RVGSharedPri ora_vvr_shpri (
 RvgResourceName = racdata_rvg
 OnlineRetryLimit = 0
)

requires group RVGgroup online local firm
oradata_mnt requires ora_vvr_shpri
```

- 8 Save and close the `main.cf` file.

- 9 Use the following command to verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file:

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

- 10 Stop and restart VCS.

```
hastop -all -force
```

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
hastart
```

- 11 Verify that VCS brings all resources online. On one node, enter the following command:

```
hagrps -display
```

The application, RVG, and CVM groups are online on both nodes of the primary site. The RVGLogOwner and ClusterService groups are online on one node of the cluster. If either the RVG group or the RVGLogOwner group is partially online, manually bring the groups online using the `hagrps -online` command. This information applies to the secondary site, except for the application group which must be offline.

- 12 Verify the service groups and their resources that are brought online. On one node, enter the following command:

```
hagrps -display
```

The application service group is offline on the secondary site, but the ClusterService, CVM, RVG log owner, and RVG groups are online.

This completes the setup for an SFCFSHA global cluster using VVR for replication. Symantec recommends testing a global cluster before putting it into production.

## Managing a global SFCFSHA cluster using Veritas Volume Replicator (VVR)

For information on using the VCS commands for global clusters:

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

If you have two SFCFSHA clusters configured to use VVR for replication, the following administrative functions are available:

**Table 25-2**

| Migration of the role of the primary site to the remote site | Migration of the role of the primary site to the remote site is a planned transfer of the role of primary replication host from one cluster to a remote cluster. This transfer enables the application on the remote cluster to actively use the replicated data. The former primary cluster becomes free for maintenance or other activity. |
|--------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Takeover of the primary site role by the secondary site      | Takeover of the primary site role by the secondary site occurs when an unplanned event (such as a disaster) causes a failure, making it necessary for the applications using the replicated data to be brought online on the remote cluster.                                                                                                 |

VCS agents manage external objects that are part of wide-area failover. These objects include replication, DNS updates, and so on. These agents provide a robust framework for specifying attributes and restarts, and can be brought online upon fail over.

**Table 25-3**

| VCS replication agents bundled with VVR | Description                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DNS agent                               | The DNS agent updates the canonical name-mapping in the domain name server after a wide-area failover. See the Veritas Cluster Server Bundled Agents Reference Guide for more information about the agent.                                                                                                                                                                                                                         |
| RVG agent                               | The RVG agent manages the Replicated Volume Group (RVG). Specifically, it brings the RVG online, monitors read-write access to the RVG, and takes the RVG offline. Use this agent when using VVR for replication.<br>RVGPrimary agent The RVGPrimary agent attempts to migrate or take over a Secondary to a Primary following an application failover. The agent has no actions associated with the offline and monitor routines. |

VCS provides agents for other array-based or application-based solutions. For more information about the RVG and RVGPrimary agents.

See the *Veritas™ Cluster Server Bundled Agents Reference Guide*.

---

**Note:** The RVGSnapshot agent is not supported for SFCFSHA.

---

## Migrating the role of primary site to the secondary site

After configuring the replication objects within VCS, you can use VCS commands to migrate the role of the cluster on the primary site to the remote cluster. In the procedure below, VCS takes the replicated database service group, *database\_grp*, offline on the primary site and brings it online on the secondary site; the secondary site now assumes the role of the primary site.

---

**Note:** The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

---

### To migrate the role of primary site to the remote site

- 1 From the primary site, use the following command to take the Oracle service group offline on all nodes.

```
hagrp -offline database_grp -any
```

Wait for VCS to take all Oracle service groups offline on the primary site.

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the `vxrlink -g` command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the primary cluster.

For example:

```
#
```

Where `rlk_clus1_priv_rac1_rvg` is the RLINK.

- 3 On the secondary site, which is now the new primary site, bring the Oracle service group online on all nodes:

```
hagrp -online database_grp -any
```

After migrating the role of the primary site to the secondary site, you can use VCS commands to migrate the role of the cluster on the new primary site to the original primary site. In the procedure below, VCS takes the replicated database service

group, *database\_grp*, offline on the new primary (former secondary) site and brings it online on the original primary site; the original primary site now resumes the role of the primary site.

---

**Note:** The `hagrp -switch` command cannot migrate a parallel group within a cluster or between clusters in a global cluster environment.

---

#### To migrate the role of new primary site back to the original primary site

- 1 Make sure that all CRS resources are online, and switch back the group *database\_grp* to the original primary site.

Issue the following command on the remote site:

```
hagrp -offline database_grp -any
```

- 2 Verify that the RLINK between the primary and secondary is up to date. Use the `vxrlink -g` command with the status option and specify the RLINK for the primary cluster. You can use the command from any node on the current primary cluster.

For example:

```
vxrlink -g data_disk_group status rlk_clus1_priv_rac1_rvg
```

Where *rlk\_clus1\_priv\_rac1\_rvg* is the RLINK.

- 3 Make sure that *database\_grp* is offline on the new primary site. Then, execute the following command on the original primary site to bring the *database\_grp* online:

```
hagrp -online database_grp -any
```

## Taking over the primary role by the remote cluster

Takeover occurs when the remote cluster on the secondary site starts the application that uses replicated data. This situation may occur if the secondary site perceives the primary site as dead, or when the primary site becomes inaccessible (perhaps for a known reason). For a more detailed description of concepts of taking over the primary role:

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

Before enabling the secondary site to take over the primary role, the administrator on the secondary site must "declare" the type of failure at the remote (primary, in

this case) site and designate the failure type using one of the options for the `haclus` command.

**Table 25-4** Options for the remote cluster to take over the primary role

| Takeover options | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disaster         | When the cluster on the primary site is inaccessible and appears dead, the administrator declares the failure type as "disaster." For example, fire may destroy a data center, including the primary site and all data in the volumes. After making this declaration, the administrator can bring the service group online on the secondary site, which now has the role as "primary" site.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| Outage           | <p>When the administrator of a secondary site knows the primary site is inaccessible for a known reason, such as a temporary power outage, the administrator may declare the failure as an "outage." Typically, an administrator expects the primary site to return to its original state.</p> <p>After the declaration for an outage occurs, the <code>RVGSharedPri</code> agent enables DCM logging while the secondary site maintains the primary replication role. After the original primary site becomes alive and returns to its original state, DCM logging makes it possible to use fast fail back resynchronization when data is resynchronized to the original cluster.</p> <p>Before attempting to resynchronize the data using the fast fail back option from the current primary site to the original primary site, take the precaution at the original primary site of making a snapshot of the original data. This action provides a valid copy of data at the original primary site for use in the case the current primary site fails before the resynchronization is complete.</p> |
| Disconnect       | When both clusters are functioning properly and the heartbeat link between the clusters fails, a split-brain condition exists. In this case, the administrator can declare the failure as "disconnect," which means no attempt will occur to take over the role of the primary site at the secondary site. This declaration is merely advisory, generating a message in the VCS log indicating the failure results from a network outage rather than a server outage.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Replica          | In the rare case where the current primary site becomes inaccessible while data is resynchronized from that site to the original primary site using the fast fail back method, the administrator at the original primary site may resort to using a data snapshot (if it exists) taken before the start of the fast fail back operation. In this case, the failure type is designated as "replica".                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

The examples illustrate the steps required for an outage takeover and resynchronization.

#### To take over after an outage

- 1 From any node of the secondary site, issue the `haclus` command:

```
haclus -declare outage -clus clus1
```

- 2 After declaring the state of the remote cluster, bring the `database_grp` service group online on the secondary site. For example:

```
hagrps -online -force database_grp -any
```

#### To resynchronize after an outage

- 1 On the original primary site, create a snapshot of the RVG before resynchronizing it in case the current primary site fails during the resynchronization. Assuming the disk group is `data_disk_group` and the RVG is `rac1_rvg`, type:

```
vxrvrg -g data_disk_group -F snapshot rac1_rvg
```

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for details on RVG snapshots.

- 2 Resynchronize the RVG. From any node of the current primary site, issue the `hares` command and the `-action` option with the `fbsync` action token to resynchronize the `RVGSharedPri` resource. For example:

```
hares -action ora_vvr_shpri fbsync -sys mercury
```

- 3 Perform one of the following commands, depending on whether the resynchronization of data from the current primary site to the original primary site is successful:
  - If the resynchronization of data is successful, use the `vxrvrg` command with the `snapback` option to reattach the snapshot volumes on the original primary site to the original volumes in the specified RVG:

```
vxrvrg -g data_disk_group snapback rac1_rvg
```

- A failed attempt at the resynchronization of data (for example, a disaster hits the primary RVG when resynchronization is in progress) could generate inconsistent data.

You can restore the contents of the RVG data volumes from the snapshot taken in step 1:

```
vxrvg -g data_disk_group snaprestore rac1_rvg
```

## Troubleshooting Veritas Volume Replicator (VVR) components of SFCFSHA

If the rlink is not up to date, use the `hares -action` command with the `resync` action token to synchronize the RVG.

### To update the rlink

- ◆ Enter the following command example on any node in the primary cluster, specifying the RVGSharedPri resource:

```
hares -action ora_vvr_shpri resync -sys galaxy
```

# Uninstallation of Storage Foundation Cluster File System High Availability

- [Chapter 26. Uninstalling Storage Foundation Cluster File System High Availability](#)

# Uninstalling Storage Foundation Cluster File System High Availability

This chapter includes the following topics:

- [About removing Veritas Storage Foundation Cluster File System High Availability](#)
- [Shutting down cluster operations](#)
- [Disabling the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling SFCFSHA packages using the script-based installer](#)
- [Uninstalling SFCFSHA with the Veritas Web-based installer](#)
- [Uninstalling Storage Foundation using the pkgrm command](#)
- [Removing the CP server configuration using the removal script](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product](#)

## About removing Veritas Storage Foundation Cluster File System High Availability

This section covers uninstallation requirements and steps to uninstall the Veritas software.

Only users with superuser privileges can uninstall Veritas Storage Foundation Cluster File System High Availability.

---

**Warning:** Failure to follow the instructions in the following sections may result in unexpected behavior.

---

## Preparing to uninstall

Review the following removing the Veritas software.

### Remote uninstallation

You must configure remote communication to uninstall SFCFSHA on remote systems. In a High Availability environment, you must meet the prerequisites to uninstall on all nodes in the cluster at one time.

The following prerequisites are required for remote uninstallation:

- Communication protocols must exist between systems. By default, the uninstall scripts use ssh.
- You must be able to execute ssh or rsh commands as superuser on all systems.
- The ssh or rsh must be configured to operate without requests for passwords or passphrases.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 446.

### Preparing to remove Veritas Volume Manager

This section describes the steps you need to take before removing Veritas Volume Manager (VxVM) to preserve the contents of the volumes.

---

**Warning:** Failure to follow the preparations in this section might result in unexpected behavior.

---

### Moving volumes from an encapsulated root disk

Use the following procedure to move volumes from an encapsulated root disk.

**To uninstall VxVM if root, swap, usr, or var is a volume under Volume Manager control**

- 1 Ensure that the `rootvol`, `swapvol`, `usr`, and `var` volumes have only one associated plex each.

The plex must be contiguous, non-striped, non-spanned, and non-sparse. To obtain this information, enter the following:

```
vxprint -ht rootvol swapvol usr var
```

If any of these volumes have more than one associated plex, remove the unnecessary plexes using the following command:

```
vxplex -o rm dis plex_name
```

- 2 Run the `vxunroot` command:

```
/etc/vx/bin/vxunroot
```

The `vxunroot` command changes the volume entries in `/etc/vfstab` to the underlying disk partitions for `rootvol`, `swapvol`, `usr`, and `var`. It also modifies `/etc/system` and prompts for a reboot so that disk partitions are mounted instead of volumes for `root`, `swap`, `usr`, and `var`.

- 3 Once you have changed the `root`, `swap`, `usr`, and `var` volumes, move all remaining volumes to disk partitions.

You can do this using one of the following procedures:

- Back up the entire system to tape and then recover from tape.
- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Move volumes incrementally to disk partitions.  
 See [“Moving volumes to disk partitions”](#) on page 381.  
 Otherwise, shut down VxVM.

**Moving volumes to disk partitions**

Use the following procedure to move volumes incrementally to disk partitions.

### To move volumes incrementally to disk partitions

- 1 Evacuate disks using `vxdiskadm`, the VOM GUI, or the `vxevac` utility.

Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.

- 2 Remove the evacuated disks from VxVM control by entering:

```
vxdg rmdisk diskname
vxdisk rm devname
```

- 3 Decide which volume to move first, and if the volume is mounted, unmount it.
- 4 If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that you have applied the `sync` command to the data on the volume.
- 5 Create a partition on free disk space of the same size as the volume using the `format` command.

If there is not enough free space for the partition, add a new disk to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this first volume.

- 6 Copy the data on the volume onto the newly created disk partition using a command such as `dd`.

```
dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/dsk/c2t2d2s7
```

where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly created partition.

- 7 Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.
- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Stop and remove the volume from VxVM using the commands.

```
vxvol -g diskgroup stop volume_name
vxedit -rf rm volume_name
```

- 10** Remove any free disks (those having no subdisks defined on them) by removing the volumes from VxVM control.

To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.

```
vxprint -g diskgroup -F '%snum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VxVM control.

```
vxdg rmdisk diskname
vxdisk rm devname
```

Use the free space created for adding the data from the next volume you want to remove.

- 11** After you successfully convert all volumes into disk partitions, reboot the system.  
**12** After the reboot, make sure none of the volumes are open by using the `vxprint` command.

```
vxprint -Aht -e v_open
```

- 13** If any volumes remain open, repeat the steps listed above.

### Example of moving volumes to disk partitions on Solaris

This example shows how to move the data on a volume to a disk partition. In the example, there are three disks: `disk1` and `disk2` are subdisks on volume `vol01` and `disk3` is a free disk. The data on `vol01` is copied to `disk3` using `vxevac`.

These are the contents of the disk group `vol01dg` before the data on `vol01` is copied to `disk3`.

```
vxprint -g vol01dg -ht
DG NAME NCONFIG NLOG MINORS GROUP-ID
DM NAME DEVICE TYPE PRIVLEN PUBLLEN STATE
RV NAME RLINK_CNT KSTATE STATE PRIMARY DATAVOLS SRL
RL NAME RVG KSTATE STATE REM_HOST REM_DG REM_RLNK
V NAME RVG KSTATE STATE LENGTH READPOL PREFPLEX UTYPE
PL NAME VOLUME KSTATE STATE LENGTH LAYOUT NCOL/WID MODE
SD NAME PLEX DISK DISKOFFS LENGTH [COL/]OFF DEVICE MODE
SV NAME PLEX VOLNAME NVOLLAYR LENGTH [COL/]OFF AM/NM MODE
DC NAME PARENTVOL LOGVOL
SP NAME SNAPVOL DCO

dg vol01dg default default 115000
```

```
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2 sliced 2591 17900352 -
dm disk2 c1t14d0s2 sliced 2591 17899056 -
dm disk3 c1t3d0s2 sliced 2591 17899056 -

v vol1 - ENABLED ACTIVE 4196448 ROUND - fsgen
pl pl1 vol1 ENABLED ACTIVE 4196448 CONCAT - RW
sd sd1 pl1 disk1 0 2098224 0 c1t12d0 ENA
sd sd2 pl1 disk2 0 2098224 2098224 c1t14d0 ENA
```

**Evacuate disk1 to disk3.**

```
/etc/vx/bin/vxevac -g voldg disk1 disk3
vxprint -g voldg -ht
```

| DG NAME | NCONFIG   | NLOG    | MINORS   | GROUP-ID |           |          |       |  |
|---------|-----------|---------|----------|----------|-----------|----------|-------|--|
| DM NAME | DEVICE    | TYPE    | PRIVLEN  | PUBLEN   | STATE     |          |       |  |
| RV NAME | RLINK_CNT | KSTATE  | STATE    | PRIMARY  | DATAVOLS  | SRL      |       |  |
| RL NAME | RVG       | KSTATE  | STATE    | REM_HOST | REM_DG    | REM_RLNK |       |  |
| V NAME  | RVG       | KSTATE  | STATE    | LENGTH   | READPOL   | PREFPLEX | UTYPE |  |
| PL NAME | VOLUME    | KSTATE  | STATE    | LENGTH   | LAYOUT    | NCOL/WID | MODE  |  |
| SD NAME | PLEX      | DISK    | DISKOFFS | LENGTH   | [COL/]OFF | DEVICE   | MODE  |  |
| SV NAME | PLEX      | VOLNAME | NVOLLAYR | LENGTH   | [COL/]OFF | AM/NM    | MODE  |  |
| DC NAME | PARENTVOL | LOGVOL  |          |          |           |          |       |  |
| SP NAME | SNAPVOL   | DCO     |          |          |           |          |       |  |

```
dg voldg default default 115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2 sliced 2591 17900352 -
dm disk2 c1t14d0s2 sliced 2591 17899056 -
dm disk3 c1t3d0s2 sliced 2591 17899056 -

v vol1 - ENABLED ACTIVE 4196448 ROUND - fsgen
pl pl1 vol1 ENABLED ACTIVE 4196448 CONCAT - RW
sd disk3-0111 disk3 0 2098224 0 c1t3d0 ENA
sd sd2 pl1 disk2 0 2098224 2098224 c1t14d0 ENA
```

**Evacuate disk2 to disk3.**

```
/etc/vx/bin/vxevac -g voldg disk2 disk3
vxprint -g voldg -ht
```

```

DG NAME NCONFIG NLOG MINORS GROUP-ID
DM NAME DEVICE TYPE PRIVLEN PUBLEN STATE
RV NAME RLINK_CNT KSTATE STATE PRIMARY DATAVOLS SRL
RL NAME RVG KSTATE STATE REM_HOST REM_DG REM_RLNK
V NAME RVG KSTATE STATE LENGTH READPOL PREFPLEX UTYPE
PL NAME VOLUME KSTATE STATE LENGTH LAYOUT NCOL/WID MODE
SD NAME PLEX DISK DISKOFFS LENGTH [COL/]OFF DEVICE MODE
SV NAME PLEX VOLNAME NVOLLAYR LENGTH [COL/]OFF AM/NM MODE
DC NAME PARENTVOL LOGVOL
SP NAME SNAPVOL DCO

```

```

dg voldg default default 115000
1017856044.1141.hostname.veritas.com

```

```

dm disk1 c1t12d0s2 sliced 2591 17900352 -
dm disk2 c1t14d0s2 sliced 2591 17899056 -
dm disk3 c1t3d0s2 sliced 2591 17899056 -

```

```

v voll - ENABLED ACTIVE 4196448 ROUND - fsgen
pl pl1 voll ENABLED ACTIVE 4196448 CONCAT - RW
sd disk3-01 pl1 disk3 0 2098224 0 c1t3d0 ENA
sd disk3-02 pl1 disk3 2098224 2098224 2098224 c1t3d0 ENA

```

Remove the evacuated disks from VxVM control.

```

vxdisk -g voldg list
DEVICE TYPE DISK GROUP STATUS
c1t3d0s2 sliced disk3 voldg online
c1t12d0s2 sliced disk1 voldg online
c1t14d0s2 sliced disk2 voldg online

vxdg rmdisk disk1
vxdg rmdisk disk2
vxdisk rm c1t12d0
vxdisk rm c1t14d0

```

Verify that the evacuated disks have been removed from VxVM control.

```

vxdisk -g voldg list
DEVICE TYPE DISK GROUP STATUS
c1t3d0s2 sliced disk3 voldg online

```

Check to see whether the volume you want to move first is mounted.

```
mount | grep voll
/voll on /dev/vx/dsk/voldg/voll
read/write/setuid/log/nolargefiles/dev=12dc138 on Wed Apr
3 10:13:11 2002
```

Create a partition on free disk space of the same size as the volume. In this example, a 2G partition is created on disk1 (c1t12d0s1).

```
format
Searching for disks...done
```

AVAILABLE DISK SELECTIONS:

0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>  
/sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
1. c1t3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
2. c1t9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
3. c1t10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
4. c1t11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
5. c1t12d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@c,0
6. c1t14d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@e,0
7. c1t15d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>  
/sbus@1f,0/SUNW,fas@2,8800000/sd@f,0

Specify disk (enter its number): 5

selecting c1t12d0

[disk formatted]

FORMAT MENU:

- disk - select a disk
- type - select (define) a disk type
- partition - select (define) a partition table
- current - describe the current disk
- format - format and analyze the disk
- repair - repair a defective sector
- label - write label to the disk
- analyze - surface analysis
- defect - defect list management
- backup - search for backup labels
- verify - read and display labels

```

save - save new disk/partition definitions
inquiry - show vendor, product and revision
volname - set 8-character volume name
!<cmd> - execute <cmd>, then return
quit
format> p

PARTITION MENU:
0 - change '0' partition
1 - change '1' partition
2 - change '2' partition
3 - change '3' partition
4 - change '4' partition
5 - change '5' partition
6 - change '6' partition
7 - change '7' partition
select - select a predefined table
modify - modify a predefined partition table
name - name the current table
print - display the current table
label - write partition map and label to the disk
!<cmd> - execute <cmd>, then return
quit

partition> 1
Part Tag Flag Cylinders Size Blocks
 1 unassigned wm 0 0 (0/0/0) 0
Enter partition id tag[unassigned]:
Enter partition permission flags[wm]:
Enter new starting cyl[0]:
Enter partition size[0b, 0c, 0.00mb, 0.00gb]: 2.00gb
partition> 1
Ready to label disk, continue? y

partition> p
Current partition table (unnamed):
Total disk cylinders available: 13814 + 2 (reserved cylinders)
Part Tag Flag Cylinders Size Blocks
 0 unassigned wm 0 0 (0/0/0) 0
 1 unassigned wm 0 - 3236 2.00GB (3237/0/0) 4195152
partition> q

```

Copy the data on `vol101` to the newly created disk partition.

```
dd if=/dev/vx/dsk/voldg/vol101 of=/dev/dsk/c1t12d0s1
```

In the `/etc/vfstab` file, remove the following entry.

```
/dev/vx/dsk/voldg/voll /dev/vx/rdisk/voldg/voll /voll vxfs 4 yes rw
```

Replace it with an entry for the newly created partition.

```
/dev/dsk/c1t12d0s1 /dev/rdisk/c1t12d0s1 /vol101 vxfs 4 yes rw
```

Mount the disk partition.

```
mount -F vxfs /dev/dsk/c1t12d0s1 /vol101
```

Remove `vol101` from VxVM.

```
vxedit -rf rm /dev/vx/dsk/voldg/vol101
```

To complete the procedure, follow the remaining steps.

## Preparing to remove Veritas File System

The `VRTSvxfs` package cannot be removed if there are any mounted VxFS file systems or Storage Checkpoints. Unmount the VxFS file systems and Storage Checkpoints before uninstalling Veritas Storage Foundation. After you remove the `VRTSvxfs` package, VxFS file systems are not mountable or accessible until another `VRTSvxfs` package is installed.

### To unmount a file system

- 1 Check if any VxFS file systems are mounted.

```
cat /etc/mnttab | grep vxfs
```

- 2 Unmount any file systems.

```
umount special | mount_point
```

Specify the file system to be unmounted as a *mount\_point* or *special* (the device on which the file system resides). See the `umount_vxfs(1M)` manual page for more information about this command and its available options.

You can use the `-a` option to unmount all file systems except `/`, `/usr`, `/usr/kvm`, `/var`, `/proc`, `/dev/fd`, and `/tmp`.

### To unmount a Storage Checkpoint

- 1 Check if any Storage Checkpoints are mounted.

```
cat /etc/mnttab | grep vxfs
```

- 2 Unmount any Storage Checkpoints.

```
umount /checkpoint_name
```

## Shutting down cluster operations

If the systems are running as an HA cluster, you have to take all service groups offline and shutdown VCS.

### To take all service groups offline and shutdown VCS

- ◆ Use the `hastop` command as follows:

```
/opt/VRTSvcs/bin/hastop -all
```

---

**Warning:** Do not use the `-force` option when executing `hastop`. This will leave all service groups online and shut down VCS, causing undesired results during uninstallation of the packages.

---

## Disabling the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

### To disable the agents

- 1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
hagrps -state service_group -sys system_name
```

If none of the service groups is online, skip to 3.

- 2 If the service group is online, take it offline.

To take the service group offline without bringing it online on any other system in the cluster, enter:

```
hagrps -offline service_group -sys system_name
```

- 3 Stop the agent on the system by entering:

```
haagent -stop agent_name -sys system_name
```

When you get the message `Please look for messages in the log file, check the file /var/VRTSvcs/log/engine_A.log for a message confirming that each agent has stopped.`

You can also use the `ps` command to confirm that the agent is stopped.

- 4 Remove the system from the `SystemList` of the service group. If you disable the agent on all the systems in the `SystemList`, you can also remove the service groups and resource types from the VCS configuration.

Read information on administering VCS from the command line.

Refer to the *Veritas Cluster Server User's Guide*.

## Removing the Replicated Data Set

If you use VVR, you need to perform the following steps. This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

## To remove the Replicated Data Set

- 1 Verify that all RLINKs are up-to-date:

```
vxlink -g diskgroup status rlink_name
```

If the Secondary is not required to be up-to-date, proceed to [2](#) and stop replication using the `-f` option with the `vradmin stoprep` command.

- 2 Stop replication to the Secondary by issuing the following command on any host in the RDS:

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

```
vradmin -g diskgroup stoprep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 3 Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

```
vradmin -g diskgroup delsec local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 4 Remove the Primary from the RDS by issuing the following command on the Primary:

```
vradmin -g diskgroup delpri local_rvgname
```

When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

The RDS is removed.

- 5 If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

```
vxedit -r -g diskgroup rm srl_name
```

# Uninstalling SFCFSHA packages using the script-based installer

Use the following procedure to remove SFCFSHA products.

Not all packages may be installed on your system depending on the choices that you made when you installed the software.

---

**Note:** After you uninstall the product, you cannot access any file systems you created using the default disk layout version in SFCFSHA 6.0 with a previous version of SFCFSHA.

---

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 446.

Language packages are uninstalled when you uninstall the English language packages.

## To shut down and remove the installed SFCFSHA packages

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems.

```
umount /mount_point
```

- 3 If the VxVM package (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

See [“Preparing to remove Veritas Volume Manager”](#) on page 380.

- 4 Make sure you have performed all of the prerequisite steps.

- 5 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
hastop -local
```

To stop VCS processes on all systems:

```
hastop -all
```

- 6 Move to the `/opt/VRTS/install` directory and run the uninstall script.

```
cd /opt/VRTS/install
```

For Veritas Storage Foundation Cluster File System High Availability

```
./uninstallsfcfsha
```

- 7 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall SFCFSHA, for example, `host1`:

```
Enter the system names separated by spaces: [q?] host1 host2
```

- 8 The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the packages are uninstalled.

The uninstall script creates log files and displays the location of the log files.

- 9 Most packages have kernel components. In order to ensure complete removal, a system reboot is recommended after all packages have been removed.

- 10 To verify the removal of the packages, use the `pkginfo` command.

```
pkginfo | grep VRTS
```

## Uninstalling SFCFSHA with the Veritas Web-based installer

This section describes how to uninstall using the Veritas Web-based installer.

---

**Note:** After you uninstall the product, you cannot access any file systems you created using the default disk layout Version in SFCFSHA 6.0 with with a previous version of SFCFSHA.

---

### To uninstall SFCFSHA

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 63.
- 3 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.

- 4 Select **Storage Foundation Cluster File System High Availability** 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 SFCFSHA 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**.

You see a prompt recommending that you reboot the system, and then return to the Web page to complete additional tasks.

## Uninstalling Storage Foundation using the `pkgrm` command

Use the following procedure to uninstall Storage Foundation using the `pkgrm` command.

If you are uninstalling Veritas Storage Foundation using the `pkgrm` command, the packages must be removed in a specific order, or else the uninstallation will fail. Removing the packages out of order will result in some errors, including possible core dumps, although the packages will still be removed.

### To uninstall Storage Foundation

- 1 Unmount all VxFS file systems and Storage Checkpoints, and close all VxVM volumes.

Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems and Storage Checkpoints.

```
umount /mount_point
```

- 3 Stop all applications from accessing VxVM volumes, and close all VxVM volumes.
- 4 Stop various daemons, if applicable.

```
/opt/VRTS/bin/vxsvcctl stop
```

- 5 Remove the packages in the following order:

```
pkgrm VRTSodm VRTSgms VRTScavf VRTSglm VRTSdbed \
VRTSvcsesa VRTSvcsag VRTSscps VRTSvcs VRTSsamf \
VRTSvxfen VRTSgab VRTSllt VRTSfssdk VRTSfsadv \
VRTSvxfs VRTSsfmh VRTSob VRTSaslapm VRTSvxvm \
VRTSspt VRTSsfcp160 VRTSperl VRTSvlic
```

## Uninstalling the language packages using the pkgrm command

If you would like to remove only the language packages, you can do so with the `pkgrm` command.

If you use the product installer menu or the uninstallation script, you can remove the language packages along with the English packages.

### To remove the language packages

- ◆ Use the `pkgrm` command to remove the appropriate packages.

See [“Chinese language packages”](#) on page 458.

See [“Japanese language packages”](#) on page 458.

```
pkgrm package_name package_name ...
```

Because the packages do not contain any dependencies, you can remove them in any order.

## Removing the CP server configuration using the removal script

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 SFCFSHA cluster (application cluster) uses the CP server that you want to unconfigure.

---

You can use the CP server configuration utility (`configure_cps.pl`) to remove the CP server configuration. This utility performs the following tasks when you choose to unconfigure the CP server:

- Removes all CP server configuration files
- Removes the VCS configuration for CP server

After you run this utility, you can uninstall VCS from the node or the cluster.

---

**Note:** You must run the configuration utility only once per CP server (which can be on a single-node VCS cluster or an SFHA cluster), when you want to remove the CP server configuration.

---

### To remove the CP server configuration

- 1 To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2 Select option 3 from the menu to unconfigure the CP server.

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

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

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

- 3 Review the warning message and confirm that you want to unconfigure the CP server.

```
WARNING: Unconfiguring Coordination Point Server stops the  
vxcperv process. VCS clusters using this server for  
coordination purpose will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)  
(Default:n) :y
```

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

- 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

- 5 Answer **y** to delete the CP server database.

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

- 6 Answer **y** at the prompt to confirm the deletion of the CP server database.

```
Warning: This database won't be available if CP server
is reconfigured on the cluster. Are you sure you want to
proceed with the deletion of database? (y/n) (Default:n) :
```

- 7 Answer **y** to delete the CP server configuration file and log files.

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 8 Run the `hagrp -state` command to ensure that the CPSSG service group has been removed from the node. For example:

```
root@mycps1.symantecexample.com # hagrp -state CPSSG
```

```
VCS WARNING V-16-1-40131 Group CPSSG does not exist
in the local cluster
```

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file disables the SFDB tools.

To remove the SFDB repository

- 1 Identify the SFDB repositories created on the host.

```
# cat /var/vx/vxdba/rep_loc

{
  "sfae_rept_version" : 1,
  "oracle" : {
    "SFAEDB" : {
      "location" : "/data/sfaedb/.sfae",
      "old_location" : "",
      "alias" : [
        "sfaedb"
      ]
    }
  }
}
```

- 2 Remove the directory identified by the `location` key.

```
# rm -rf /data/sfaedb/.sfae
```

- 3 Remove the repository location file.

```
# rm -rf /var/vx/vxdba/rep_loc
```

This completes the removal of the SFDB repository.

Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Response files](#)
- [Appendix C. Tunable files for installation](#)
- [Appendix D. Configuring I/O fencing using a response file](#)
- [Appendix E. Configuring the secure shell or the remote shell for communications](#)
- [Appendix F. Storage Foundation Cluster File System High Availability components](#)
- [Appendix G. High availability agent information](#)
- [Appendix H. Troubleshooting information](#)
- [Appendix I. Troubleshooting cluster installation](#)
- [Appendix J. Sample SFCFSHA cluster setup diagrams for CP server-based I/O fencing](#)
- [Appendix K. Reconciling major/minor numbers for NFS shared disks](#)
- [Appendix L. Configuring LLT over UDP using IPv6](#)
- [Appendix M. Configuring LLT over UDP using IPv4](#)

- [Appendix N. Compatibility issues when installing Storage Foundation Cluster File System High Availability with other products](#)

Installation scripts

This appendix includes the following topics:

- [About Veritas Storage Foundation and High Availability Solutions installation scripts](#)
- [Installation script options](#)
- [About using the postcheck option](#)

About Veritas Storage Foundation and High Availability Solutions installation scripts

Veritas Storage Foundation and High Availability Solutions products 6.0 provides several installation scripts. You can find these scripts at the root of the product media in the scripts directory.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from the Symantec download site, which does not include the installer, use the appropriate product installation script.

See [“Downloading the Veritas Storage Foundation Cluster File System High Availability software”](#) on page 30.

The following product installation scripts are available:

Table A-1 Product installation scripts

Product installation script	Veritas product name
<code>installvcs</code>	Veritas Cluster Server (VCS)
<code>installsf</code>	Veritas Storage Foundation (SF)

Table A-1 Product installation scripts (*continued*)

Product installation script	Veritas product name
<code>installsfha</code>	Veritas Storage Foundation and High Availability (SFHA)
<code>installsfcfsha</code>	Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)
<code>installsfrac</code>	Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
<code>installsfsybasece</code>	Veritas Storage Foundation for Sybase ASE CE (SF Sybase CE)
<code>installvm</code>	Veritas Volume Manager
<code>installfs</code>	Veritas File System
<code>installdmp</code>	Veritas Dynamic Multi-Pathing
<code>installsvs</code>	Symantec VirtualStore

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

Installation script options

[Table A-2](#) 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 Veritas Storage Foundation product scripts, except where otherwise noted.

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 401.

Table A-2 Available command line options

Command Line Option	Function
<code>system1 system2...</code>	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
<code>--addnode</code>	Adds a node to a high availability cluster.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-allpkgs</code>	Displays all packages and patches required for the specified product. The packages and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
<code>-comcleanup</code>	The <code>-comcleanup</code> 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.
<code>-configure</code>	Configures the product after installation.
<code>-fencing</code>	Configures I/O fencing in a running cluster.
<code>-hostfile <i>full_path_to_file</i></code>	Specifies the location of a file that contains a list of hostnames on which to install.
<code>-install</code>	The <code>-install</code> option is used to install products on systems.
<code>-installallpkgs</code>	Specifies that all packages are installed.
<code>-installminpkgs</code>	Specifies that the minimum package set is installed.
<code>-installrecpkgs</code>	Specifies that the required package set is installed.
<code>-jumpstart <i>dir_path</i></code>	Produces a sample finish file for Solaris JumpStart installation. The <i>dir_path</i> indicates the path to the directory in which to create the finish file.
<code>-keyfile <i>ssh_key_file</i></code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i <i>ssh_key_file</i></code> to every SSH invocation.
<code>-license</code>	Registers or updates product licenses on the specified systems.
<code>-logpath <i>log_path</i></code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-makeresponsefile</code>	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.
<code>-minpkgs</code>	Displays the minimal packages and patches required for the specified product. The packages and patches 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 <code>allpkgs</code> option.
<code>-nolic</code>	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
<code>-pkginfo</code>	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 <code>-pkginfo</code> option with the <code>installvcs</code> script to display VCS packages.
<code>-pkgpath <i>package_path</i></code>	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.
<code>-pkgset</code>	Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.
<code>-pkgtable</code>	Displays product's packages in correct installation order by group.
<code>-postcheck</code>	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
<code>-precheck</code>	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-recpkgs</code>	Displays the recommended packages and patches required for the specified product. The packages and patches 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 <code>allpkgs</code> option.
<code>-redirect</code>	Displays progress details without showing the progress bar.
<code>-requirements</code>	The <code>-requirements</code> option displays required OS version, required packages and patches, file system space, and other system requirements in order to install the product.
<code>-responsefile <i>response_file</i></code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<code>-rolling_upgrade</code>	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.
<code>-rollingupgrade_phase1</code>	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel packages get upgraded to the latest version
<code>-rollingupgrade_phase2</code>	The <code>-rollingupgrade_phase2</code> 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."

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-rootpath</code> <i>root_path</i>	Specifies an alternative root directory on which to install packages. On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code> command.
<code>-rsh</code>	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See “About configuring secure shell or remote shell communication modes before installing products” on page 446.
<code>-serial</code>	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.
<code>-settunables</code>	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 <code>-tunablesfile</code> option.
<code>-start</code>	Starts the daemons and processes for the specified product.
<code>-stop</code>	Stops the daemons and processes for the specified product.
<code>-tmppath</code> <i>tmp_path</i>	Specifies a directory other than <code>/var/tmp</code> 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.
<code>-uninstall</code>	The <code>-uninstall</code> option is used to uninstall products from systems.
<code>-tunablesfile</code>	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-upgrade</code>	Specifies that an existing version of the product exists and you plan to upgrade it.
<code>-version</code>	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, hotfixes, and available updates for the installed product if an Internet connection is available.

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.
- The VRTSIlt pkg version is not consistent on the nodes.
- The Ilt-linkinstall value is incorrect.
- The llthosts(4) or llttab(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.

See [“Performing a postcheck on a node”](#) on page 264.

Response files

This appendix includes the following topics:

- [About response files](#)
- [Installing SFCFSHA using response files](#)
- [Configuring SFCFSHA using response files](#)
- [Upgrading SFCFSHA using response files](#)
- [Uninstalling SFCFSHA using response files](#)
- [Syntax in the response file](#)
- [Response file variables to install, upgrade, or uninstall Veritas Storage Foundation Cluster File System High Availability](#)
- [Response file variables to configure Veritas Storage Foundation Cluster File System High Availability](#)
- [Sample response file for Veritas Storage Foundation Cluster File System High Availability installation](#)
- [Sample response file for Veritas Storage Foundation Cluster File System High Availability configuration](#)

About response files

The installer or product installation script generates a response file during any installation, configuration, upgrade, or uninstall procedure. The response file contains the configuration information that you entered during the procedure. When the procedure completes, the installation script displays the location of the response files.

You can use the response file for future installation procedures by invoking an installation script with the `-responsefile` option. The response file passes arguments to the script to automate the installation of that product. You can edit the file to automate installation and configuration of additional systems.

You can generate a response file using the `-makeresponsefile` option.

See [“Installation script options”](#) on page 402.

Installing SFCFSHA using response files

Typically, you can use the response file that the installer generates after you perform SFCFSHA installation on one cluster to install SFCFSHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To install SFCFSHA using response files

- 1 Make sure the systems where you want to install SFCFSHA meet the installation requirements.
- 2 Make sure the preinstallation tasks are completed.
- 3 Copy the response file to one of the cluster systems where you want to install SFCFSHA.
- 4 Edit the values of the response file variables as necessary.
- 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  
  
# ./installsfcfsha -responsefile /tmp/response_file
```

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

Configuring SFCFSHA using response files

Typically, you can use the response file that the installer generates after you perform SFCFSHA configuration on one cluster to configure SFCFSHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure SFCFSHA using response files

- 1 Make sure the SFCFSHA packages are installed on the systems where you want to configure SFCFSHA.
- 2 Copy the response file to one of the cluster systems where you want to configure SFCFSHA.

- 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 Veritas Storage Foundation Cluster File System High Availability](#)” on page 415.

- 4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfcfsa -responsefile /tmp/response_file
```

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

Upgrading SFCFSHA using response files

Typically, you can use the response file that the installer generates after you perform SFCFSHA upgrade on one system to upgrade SFCFSHA on other systems. You can also create a response file using the `make_responsefile` option of the installer.

To perform automated SFCFSHA upgrade

- 1 Make sure the systems where you want to upgrade SFCFSHA 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 SFCFSHA.
- 4 Edit the values of the response file variables as necessary.

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

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

Uninstalling SFCFSHA using response files

Typically, you can use the response file that the installer generates after you perform SFCFSHA uninstallation on one cluster to uninstall SFCFSHA on other clusters.

To perform an automated uninstallation

- 1 Make sure that you meet the prerequisites to uninstall SFCFSHA.
- 2 Copy the response file to the system where you want to uninstall SFCFSHA.
- 3 Edit the values of the response file variables as necessary.
- 4 Start the uninstallation from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/uninstallsfcfsha -responsefile /tmp/response_file
```

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

Syntax in the response file

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

For example, in the case of a string value:

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

or, in the case of an integer value:

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

or, in the case of a list:

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

Response file variables to install, upgrade, or uninstall Veritas Storage Foundation Cluster File System High Availability

[Table B-1](#) lists the response file variables that you can define to configure SFCFSHA.

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SFCFSHA

Variable	Description
CFG{opt}{install}	Installs SFCFSHA packages. Configuration can be performed at a later time using the <code>-configure</code> option. 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
\$CFG{opt}{vxkeyless}	Installs the product with keyless license. List or scalar: scalar Optional or required: optional
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SFCFSA (*continued*)

Variable	Description
CFG{opt}{pkgpath}	<p>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.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{tmppath}	<p>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.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{rsh}	<p>Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{donotinstall} {package}	<p>Instructs the installation to not install the optional packages in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{donotremove} {package}	<p>Instructs the uninstallation to not remove the optional packages in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{opt}{logpath}	<p>Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
\$CFG{opt}{prodmode}	<p>List of modes for product</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SFCFSHA (*continued*)

Variable	Description
CFG{opt}{upgrade}	Upgrades all packages installed, without configuration. List or scalar: list Optional or required: optional
CFG{opt}{uninstall}	Uninstalls SFCFSHA packages. List or scalar: scalar Optional or required: optional
CFG{mirrordgname}{system}	Splits the target disk group name for a system. List or scalar: scalar Optional or required: optional
CFG{splitmirror}{system}	Indicates the system where you want a split mirror backup disk group created. List or scalar: scalar Optional or required: optional

Response file variables to configure Veritas Storage Foundation Cluster File System High Availability

[Table B-2](#) lists the response file variables that you can define to configure SFCFSHA.

Table B-2 Response file variables specific to configuring Veritas Storage Foundation Cluster File System High Availability

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 SFCFSHA.
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)

Table B-2 Response file variables specific to configuring Veritas Storage Foundation Cluster File System High Availability (*continued*)

Variable	List or Scalar	Description
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 VCS60 for VCS. (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> 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 <code>/opt/VRTS/install/logs</code> . Note: The installer copies the response files and summary files also to the specified <i>logpath</i> 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 Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group

(csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtprecp, and smtpsev), the SNMP trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

[Table B-3](#) lists the response file variables that specify the required information to configure a basic SFCFSHA cluster.

Table B-3 Response file variables specific to configuring a basic SFCFSHA 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 SFCFSHA configuration, defines if fencing is enabled. Valid values are 0 or 1. (Required)

[Table B-4](#) lists the response file variables that specify the required information to configure LLT over Ethernet.

Table B-4 Response file variables specific to configuring private LLT over Ethernet

Variable	List or Scalar	Description
CFG{vcs_lltlink#} {"system"}	Scalar	<p>Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links.</p> <p>You must enclose the system name within double quotes.</p> <p>(Required)</p>
CFG{vcs_lltlinklowpri#} {"system"}	Scalar	<p>Defines a low priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication.</p> <p>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.</p> <p>You must enclose the system name within double quotes.</p> <p>(Optional)</p>

Table B-5 lists the response file variables that specify the required information to configure LLT over UDP.

Table B-5 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltverudp}=1	Scalar	<p>Indicates whether to configure heartbeat link using LLT over UDP.</p> <p>(Required)</p>

Table B-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_address} {<system1>}	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)
CFG {vcs_udplinklowpri<n>_address} {<system1>}	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} {<system1>}	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} {<system1>}	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)

Table B-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_netmask} {<system1>}	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)
CFG{vcs_udplinklowpri<n>_netmask} {<system1>}	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 B-6 lists the response file variables that specify the required information to configure virtual IP for SFCFSHA cluster.

Table B-6 Response file variables specific to configuring virtual IP for SFCFSHA 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 B-7 lists the response file variables that specify the required information to configure the SFCFSHA cluster in secure mode.

Table B-7 Response file variables specific to configuring SFCFSHA 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}{securityonnode}	Scalar	Specifies that the securityonnode option is being used.
CFG{securityonnode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. <ul style="list-style-type: none"> ■ 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.

Table B-8 lists the response file variables that specify the required information to configure VCS users.

Table B-8 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 B-8 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 B-9](#) lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table B-9 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)
CFG{vcs_smtpsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

[Table B-10](#) lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table B-10 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 B-11](#) lists the response file variables that specify the required information to configure SFCFSHA global clusters.

Table B-11 Response file variables specific to configuring SFCFSHA 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 Veritas Storage Foundation Cluster File System High Availability installation

The following example shows a response file for installing Veritas Storage Foundation Cluster File System High Availability.

```
#####  
#Auto generated sfcfsha responsefile #  
#####  
  
our %CFG;  
$CFG{accepteula}=1;  
$CFG{opt}{rsh}=1;  
$CFG{opt}{trace}=0;  
$CFG{opt}{gco}=1;  
$CFG{opt}{vr}=1;  
$CFG{opt}{install}=1;  
$CFG{opt}{installallpkgs}=1;  
$CFG{prod}="SFCFSHA60";  
$CFG{systems}=[ qw( system01 system02 ) ];  
$CFG{opt}{logpath}="/opt/VRTS/install/logs/installsfdfs-xxxxxx/  
installsfdfs-xxxxxx.response";  
  
1;
```

Sample response file for Veritas Storage Foundation Cluster File System High Availability configuration

The following example shows a response file for configuring Veritas Storage Foundation Cluster File System High Availability.

```
#####  
#Auto generated sfcfsha responsefile #  
#####  
  
our %CFG;  
$CFG{accepteula}=1;  
$CFG{opt}{rsh}=1;  
$CFG{opt}{trace}=0;  
$CFG{vcs_allowcomms}=1;
```

Sample response file for Veritas Storage Foundation Cluster File System High Availability configuration

```
$CFG{opt}{gco}=1;
$CFG{opt}{vr}=1;
$CFG{opt}{configure}=1;
$CFG{prod}="SFCFSHA60";
$CFG{systems}=[ qw( system01 system02 ) ];
$CFG{sfcfs_cvtimeout}=200;
$CFG{sfcfs_fencingenabled}=0;
$CFG{vm_newnames_file}{system01}=0;
$CFG{vm_restore_cfg}{system01}=0;
$CFG{vm_newnames_file}{system02}=0;
$CFG{vm_restore_cfg}{system02}=0;
$CFG{vcs_clusterid}=127;
$CFG{vcs_clustername}="uxrt6_sol";
$CFG{vcs_username}=[ qw(admin operator) ];
$CFG{vcs_userenpw}=[ qw(J1mElgLimHmKumGlj
bQOsOUnVQoOUnTQsOSnUQuOUnPQtOS) ];
$CFG{vcs_userpriv}=[ qw(Administrators Operators) ];
$CFG{vcs_11tlink1}{system01}="bge1";
$CFG{vcs_11tlink2}{system01}="bge2";
$CFG{vcs_11tlink1}{system02}="bge1";
$CFG{vcs_11tlink2}{system02}="bge2";
$CFG{vcs_enabled}=1;
$CFG{opt}{logpath}="/opt/VRTS/install/logs/installsfdfs-xxxxxx/
installsfdfs-xxxxxx.response";

1;
```

Tunable files for installation

This appendix includes the following topics:

- [About setting tunable parameters using the installer or a response file](#)
- [Setting tunables for an installation, configuration, or upgrade](#)
- [Setting tunables with no other installer-related operations](#)
- [Setting tunables with an un-integrated response file](#)
- [Preparing the tunables file](#)
- [Setting parameters for the tunables file](#)
- [Tunables value parameter definitions](#)

About setting tunable parameters using the installer or a response file

You can set non-default product and system tunable parameters using a tunables file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing. The tunables file passes arguments to the installer script to set tunables. With the file, you can set the tunables for the following operations:

- When you install, configure, or upgrade systems.

```
# ./installer -tunablesfile tunables_file_name
```

See [“Setting tunables for an installation, configuration, or upgrade”](#) on page 427.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -setttunables [  
system1 system2 ...]
```

See [“Setting tunables with no other installer-related operations”](#) on page 428.

- When you apply the tunables file with an un-integrated response file.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name
```

See [“Setting tunables with an un-integrated response file”](#) on page 429.

See [“About response files”](#) on page 409.

You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 431.

Setting tunables for an installation, configuration, or upgrade

You can use a tunables file for installation procedures to set non-default tunables. You invoke the installation script with the `tunablesfile` option. The tunables file passes arguments to the script to set the selected tunables. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 431.

Note: Certain tunables only take effect after a system reboot.

To set the non-default tunables for an installation, configuration, or upgrade

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 430.
- 2 Make sure the systems where you want to install SFCFSHA meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 431.

Note: Certain tunables only take effect after a system reboot.

To set tunables with no other installer-related operations

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 430.
- 2 Make sure the systems where you want to install SFCFSHA meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-setttunables` option.

```
# ./installer -tunablesfile tunables_file_name -setttunables [
sys123 sys234 ...]
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 431.

Note: Certain tunables only take effect after a system reboot.

To set tunables with an un-integrated response file

- 1 Make sure the systems where you want to install SFCFSHA meet the installation requirements.
- 2 Complete any preinstallation tasks.
- 3 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 430.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-settunables` option.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name -settunables
```

Where *response_file_name* is the full path name for the response file and *tunables_file_name* is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Preparing the tunables file

A tunables file is a Perl module and consists of an opening and closing statement, with the tunables defined between. Use the hash symbol at the beginning of the line to comment out the line. The tunables file opens with the line "our %TUN;" and ends with the return true "1;" line. The final return true line only needs to appear once at the end of the file. Define each tunable parameter on its own line.

Format the tunable parameter as follows:

```
$TUN{"tunable_name"}{"system_name"|"*"}=value_of_tunable;
```

For the *system_name*, use the name of the system, its IP address, or a wildcard symbol. The *value_of_tunable* depends on the type of tunable you are setting. End the line with a semicolon.

The following is an example of a tunables file.

```
#  
# Tunable Parameter Values:  
#  
our %TUN;  
  
$TUN{"tunable1"}{"*"}=1024;  
$TUN{"tunable3"}{"sys123"}="SHA256";  
  
1;
```

Setting parameters for the tunables file

Each tunables file defines different tunable parameters. The values that you can use are listed in the description of each parameter. Select the tunables that you want to add to the tunables file and then configure each parameter.

See [“Tunables value parameter definitions”](#) on page 431.

Each line for the parameter value starts with \$TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{"node123"}=16;
```

In this example, you are changing the *dmp_daemon_count* value from its default of 10 to 16. You can use the wildcard symbol "*" for all systems. For example:

```
$TUN("dmp_daemon_count"){ "*" }=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Veritas Storage Foundation and High Availability Solutions Tuning Guide* for detailed information on product tunable ranges and recommendations .

[Table C-1](#) describes the supported tunable parameters that can be specified in a tunables file.

Table C-1 Supported tunable parameters

Tunable	Description
dmp_cache_open	(Veritas Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_daemon_count	(Veritas Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_delayq_interval	(Veritas Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_fast_recovery	(Veritas Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_health_time	(Veritas Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_log_level	(Veritas Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_low_impact_probe	(Veritas Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_lun_retry_timeout	(Veritas Dynamic Multi-Pathing) The retry period for handling transient errors. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_fabric	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Veritas Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_multipathing	(Veritas Dynamic Multi-Pathing) Whether DMP will intercept the I/Os directly on the raw OS paths or not. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_support	(Veritas Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_path_age	(Veritas Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_pathswitch_blks_shift	(Veritas Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_probe_idle_lun	(Veritas Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_threshold	(Veritas Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_cycles	(Veritas Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_interval	(Veritas Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_policy	(Veritas Dynamic Multi-Pathing) The policy used by DMP path restoration thread. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_state	(Veritas Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_retry_count	(Veritas Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_scsi_timeout	(Veritas Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_sfg_threshold	(Veritas Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_stat_interval	(Veritas Dynamic Multi-Pathing) The time interval between gathering DMP statistics. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer sets only the system default value of max_diskq. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer sets only the system default value of read_ahead. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer sets only the system default value of read_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer sets only the system default value of read_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
vol_checkpoint_default	(Veritas File System) Size of VxVM checkpoints (sectors). This tunable requires system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for VERITAS Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for VERITAS Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires system reboot to take effect.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of VERITAS Volume Replicator (bytes).
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (sectors). This tunable requires system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (sectors). This tunable requires system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Veritas Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by VERITAS Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions that can exist on a non-sequential DRL. This tunable requires system reboot to take effect.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (sectors). This tunable requires system reboot to take effect.
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliorem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires system reboot to take effect.
voliorem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires system reboot to take effect.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires system reboot to take effect.
vx_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. This tunable requires system reboot to take effect.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. This tunable requires system reboot to take effect.
vxfs_mbuf	(Veritas File System) Maximum memory used for VxFS buffer cache. This tunable requires system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer sets only the system default value of write_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer sets only the system default value of write_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.

Configuring I/O fencing using a response file

This appendix 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 SFCFSHA.

To configure I/O fencing using response files

- 1 Make sure that SFCFSHA 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 84.

- 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 440.
 See “[Sample response file for configuring server-based I/O fencing](#)” on page 443.
- 4 Edit the values of the response file variables as necessary.
 See “[Response file variables to configure disk-based I/O fencing](#)” on page 439.
 See “[Response file variables to configure server-based I/O fencing](#)” on page 441.
- 5 Start the configuration from the system to which you copied the response file.
 For example:

```
# /opt/VRTS/install/installsfcsfsha -responsefile /tmp/response_file
```

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

Response file variables to configure disk-based I/O fencing

[Table D-1](#) lists the response file variables that specify the required information to configure disk-based I/O fencing for SFCFSHA.

Table D-1 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. <ul style="list-style-type: none"> ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode ■ 4—Fencing migration when the cluster is online (Required)

Table D-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
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 <code>fencing_scsi3_disk_policy</code> variable and either the <code>fencing_dgname</code> variable or the <code>fencing_newdg_disks</code> variable. (Optional)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.
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 <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.

Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions.

See [“Response file variables to configure disk-based I/O fencing”](#) on page 439.

```
#  
# Configuration Values:
```

```
#
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;

$CFG{prod}="SFCFSHA60";

$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_dgname}="fendg";
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{fencing_newdg_disks}=
  [ qw(c1t1d0s2 c2t1d0s2 c3t1d0s2) ];
$CFG{fencing_option}=2;
```

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 D-2](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table D-2 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_config_cpagent}	<p>Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not.</p> <p>Enter "0" if you do not want to configure the Coordination Point agent using the installer.</p> <p>Enter "1" if you want to use the installer to configure the Coordination Point agent.</p>

Table D-2 Coordination point server (CP server) based fencing response file definitions (*continued*)

Response file field	Definition
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 <code>fencing_config_cpagent</code> 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 " <code>\$CFG{fencing_reusedg}=0</code> " or " <code>\$CFG{fencing_reusedg}=1</code> " before proceeding with a silent installation.
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_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"

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_ports>{"10.200.117.145"}=14250;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFCFSHA60";
$CFG{systems}=[ qw(galaxy nebula) ];
$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 D-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 SFCFSHA in virtual machines that do not support SCSI-3 PR”](#) on page 79.

Table D-3 Non-SCSI-3 server-based I/O fencing response file definitions

Response file field	Definition
CFG{non_scsi3_fencing}	<p>Defines whether to configure non-SCSI-3 server-based I/O fencing.</p> <p>Valid values are 1 or 0. Enter 1 to configure non-SCSI-3 server-based I/O fencing.</p>

Table D-3 Non-SCSI-3 server-based I/O fencing response file definitions
(continued)

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 <code>fencing_config_cpagent</code> field is given a value of '0'.
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_ports}	The port of the CP server that is denoted by <i>cps</i> .

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_ports}{"10.198.89.251"}=14250;
```

```
$CFG{fencing_ports}{"10.198.89.252"}=14250;  
$CFG{fencing_ports}{"10.198.89.253"}=14250;  
$CFG{non_scsi3_fencing}=1;  
$CFG{opt}{configure}=1;  
$CFG{opt}{fencing}=1;  
$CFG{prod}="SFCSHA60";  
$CFG{systems}=[ qw(galaxy nebula) ];  
$CFG{vcs_clusterid}=1256;  
$CFG{vcs_clustername}="clus1";  
$CFG{fencing_option}=1;
```

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

- [About configuring secure shell or remote shell communication modes before installing products](#)
- [Manually configuring and passwordless ssh](#)
- [Restarting the ssh session](#)
- [Enabling and disabling rsh for Solaris](#)

About configuring secure shell or remote shell communication modes before installing products

Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installer is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install Veritas software.

You can install products to remote systems using either secure shell (`ssh`) or remote shell (`rsh`). Symantec recommends that you use `ssh` as it is more secure than `rsh`.

This section contains an example of how to set up `ssh` password free communication. The example sets up `ssh` between a source system (`system1`) that contains the installation directories, and a target system (`system2`). This procedure also applies to multiple target systems.

Note: The script- and Web-based installers support establishing password less communication for you.

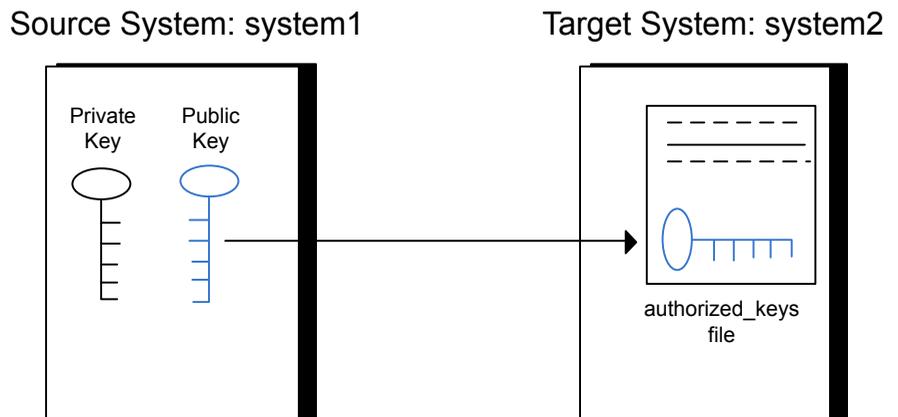
Manually configuring and passwordless ssh

The ssh program enables you to log into and execute commands on a remote system. ssh enables encrypted communications and an authentication process between two untrusted hosts over an insecure network.

In this procedure, you first create a DSA key pair. From the key pair, you append the public key from the source system to the `authorized_keys` file on the target systems.

Figure E-1 illustrates this procedure.

Figure E-1 Creating the DSA key pair and appending it to target systems



Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

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

To create the DSA key pair

- 1 On the source system (system1), log in as root, and navigate to the root directory.

```
system1 # cd /
```

- 2 To generate a DSA key pair on the source system, type the following command:

```
system1 # 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`.
- 4 When the program asks you to enter the passphrase, press the Enter key twice.

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

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 (system2 in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

```
system2 # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

```
system2 # chmod go-w /.ssh
```

To append the public key from the source system to the `authorized_keys` file on the target system, using secure file transfer

- 1 Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems (`system2` in this example).

To enable SFTP, the `/etc/ssh/sshd_config` file must contain the following two lines:

```
PermitRootLogin          yes
Subsystem                sftp          /usr/lib/ssh/sftp-server
```

- 2 If the lines are not there, add them and restart ssh.

To restart ssh on Solaris 10, type the following command:

```
system1 # svcadm restart ssh
```

- 3 From the source system (`system1`), move the public key to a temporary file on the target system (`system2`).

Use the secure file transfer program.

In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

Use the following command for secure file transfer:

```
system1 # sftp system2
```

If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

```
Connecting to system2 ...
The authenticity of host 'system2 (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
Are you sure you want to continue connecting (yes/no)?
```

- 4 Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'system2,10.182.00.00'
(DSA) to the list of known hosts.
root@system2 password:
```

- 5 Enter the root password of `system2`.

- 6 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
```

- 7 To quit the SFTP session, type the following command:

```
sftp> quit
```

- 8 To begin the `ssh` session on the target system (system2 in this example), type the following command on system1:

```
system1 # ssh system2
```

Enter the root password of system2 at the prompt:

```
password:
```

- 9 After you log in to system2, enter the following command to append the `id_dsa.pub` file to the `authorized_keys` file:

```
system2 # cat /id_dsa.pub >> /.ssh/authorized_keys
```

- 10 After the `id_dsa.pub` public key file is copied to the target system (system2), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, enter the following command on system2:

```
system2 # rm /id_dsa.pub
```

- 11 To log out of the `ssh` session, enter the following command:

```
system2 # exit
```

- 12 When you install from a source system that is also an installation target, also add the local system `id_dsa.pub` key to the local `authorized_keys` file. The installation can fail if the installation source system is not authenticated.

To add the local system `id_dsa.pub` key to the local `authorized_keys` file, enter the following command:

```
system1 # cat /.ssh/id_dsa.pub >> /.ssh/authorized_keys
```

- 13 Run the following commands on the source installation system. If your ssh session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and make the key globally available to the user `root`:

```
system1 # exec /usr/bin/ssh-agent $SHELL
system1 # ssh-add

Identity added: //./ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

To verify that you can connect to a target system

- 1 On the source system (`system1`), enter the following command:

```
system1 # ssh -l root system2 uname -a
```

where `system2` is the name of the target system.

- 2 The command should execute from the source system (`system1`) to the target system (`system2`) without the system requesting a passphrase or password.
- 3 Repeat this procedure for each target system.

Restarting the ssh session

After you complete this procedure, ssh can be restarted in any of the following scenarios:

- After a terminal session is closed
- After a new terminal session is opened
- After a system is restarted
- After too much time has elapsed, to refresh ssh

To restart ssh

- 1 On the source installation system (system1), bring the private key into the shell environment.

```
system1 # exec /usr/bin/ssh-agent $SHELL
```

- 2 Make the key globally available for the user `root`

```
system1 # ssh-add
```

Enabling and disabling rsh for Solaris

The following section describes how to enable remote shell on Solaris system.

Veritas recommends configuring a secure shell environment for Veritas product installations.

See [“Manually configuring and passwordless ssh”](#) on page 447.

See the operating system documentation for more information on configuring remote shell.

To enable rsh

- 1 To determine the current status of `rsh` and `rlogin`, type the following command:

```
# inetadm | grep -i login
```

If the service is enabled, the following line is displayed:

```
enabled online svc:/network/login:rlogin
```

If the service is not enabled, the following line is displayed:

```
disabled disabled svc:/network/login:rlogin
```

- 2 To enable a disabled `rsh/rlogin` service, type the following command:

```
# inetadm -e rlogin
```

- 3 To disable an enabled `rsh/rlogin` service, type the following command:

```
# inetadm -d rlogin
```

- 4 Modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. This file must be modified for each user who remotely accesses the system using rsh. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system having access to the local system. For example, if the root user must remotely access `system1` from `system2`, you must add an entry for `system2.companyname.com` in the `.rhosts` file on `system1`.

```
# echo "system2.companyname.com" >> $HOME/.rhosts
```

- 5 After you complete an installation procedure, delete the `.rhosts` file from each user's `$HOME` directory to ensure security:

```
# rm -f $HOME/.rhosts
```

Storage Foundation Cluster File System High Availability components

This appendix includes the following topics:

- [Veritas Storage Foundation Cluster File System High Availability installation packages](#)
- [Veritas Cluster Server installation packages](#)
- [Veritas Cluster File System installation packages](#)
- [Chinese language packages](#)
- [Japanese language packages](#)
- [Veritas Storage Foundation obsolete and reorganized installation packages](#)

Veritas Storage Foundation Cluster File System High Availability installation packages

[Table F-1](#) shows the package name and contents for each English language package for Veritas Storage Foundation Cluster File System High Availability. The table also gives you guidelines for which packages to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Veritas Storage Foundation Cluster File System High Availability and Veritas Cluster Server (VCS) packages, the combined functionality is called Veritas Storage Foundation Cluster File System High Availability and High Availability.

See [“Veritas Cluster Server installation packages”](#) on page 457.

Table F-1 Veritas Storage Foundation Cluster File System High Availability packages

packages	Contents	Configuration
VRTSaslapm	Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries Required for the support and compatibility of various storage arrays.	Minimum
VRTSperl	Perl 5.12.2 for Veritas	Minimum
VRTSvlic	Veritas License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Minimum
VRTSvxfv	Veritas File System binaries Required for VxFS file system support.	Minimum
VRTSvxvm	Veritas Volume Manager binaries, scripts, and utilities. Required for VxVM volume manager support.	Minimum
VRTSdbed	Veritas Storage Foundation for Databases	Recommended
VRTSob	Veritas Enterprise Administrator	Recommended
VRTSodm	Veritas ODM Driver for VxFS Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager allows Oracle 9i and 10g to improve performance and manage system bandwidth.	Recommended

Table F-1 Veritas Storage Foundation Cluster File System High Availability packages *(continued)*

packages	Contents	Configuration
VRTSsfpci60	<p>Veritas Storage Foundation Common Product Installer</p> <p>The Storage Foundation Common Product installer package contains the installer libraries and product scripts that perform the following:</p> <ul style="list-style-type: none"> ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ removing nodes ■ etc. <p>You can use these script to simplify the native operating system installations, configurations, and upgrades.</p>	Minimum
VRTSsfmh	<p>Veritas Storage Foundation Managed Host</p> <p>Discovers configuration information on a Storage Foundation managed host. This information is stored on a central database, which is not part of this release. You must download the database separately at:</p>	Recommended
VRTSspt	Veritas Software Support Tools	Recommended
VRTSfsadv	Minimum Veritas File System Advanced Solutions by Symantec (Solaris SPARC only).	Minimum
VRTSfssdk	<p>Veritas File System Software Developer Kit</p> <p>For VxFS APIs, the package contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.</p>	All

Veritas Cluster Server installation packages

[Table F-2](#) shows the package name and contents for each English language package for Veritas Cluster Server (VCS). The table also gives you guidelines for which packages to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and VCS packages, the combined functionality is called Storage Foundation and High Availability.

See [“Veritas Storage Foundation Cluster File System High Availability installation packages”](#) on page 454.

Table F-2 VCS installation packages

package	Contents	Configuration
VRTSgab	Veritas Cluster Server group membership and atomic broadcast services	Minimum
VRTSilt	Veritas Cluster Server low-latency transport	Minimum
VRTSamf	Veritas Cluster Server Asynchronous Monitoring Framework	Minimum
VRTSvc	Veritas Cluster Server	Minimum
VRTSvcsg	Veritas Cluster Server Bundled Agents	Minimum
VRTSvxfen	Veritas I/O Fencing	Minimum
VRTSvcsea	Consolidated database and enterprise agent packages	Recommended
VRTScps	Veritas Coordination Point Server The Coordination Point Server is an alternate mechanism for I/O fencing. It implements I/O fencing through a client/server architecture and can provide I/O fencing for multiple VCS clusters.	All

Veritas Cluster File System installation packages

[Table F-3](#) shows the package name and contents for each English language package for Veritas Cluster File System (CFS). The table also gives you guidelines for which packages to install based on whether you want the minimum, recommended, or advanced configuration.

When you install all CFS packages and all the packages that comprise Storage Foundation and Veritas Cluster Server, the resulting functionality is called Storage Foundation Cluster File System.

See [“Veritas Storage Foundation Cluster File System High Availability installation packages”](#) on page 454.

See [“Veritas Cluster Server installation packages”](#) on page 457.

Table F-3 CFS installation packages

package	Contents	Configuration
VRTScavf	Veritas Cluster Server Agents for Storage Foundation Cluster File System	Minimum
VRTSglm	Veritas Group Lock Manager for Storage Foundation Cluster File System	Minimum
VRTSgms	Veritas Group Messaging Services for Storage Foundation Cluster File System	Recommended

Chinese language packages

The following table shows the package name and contents for each Chinese language package.

Table F-4 Chinese language packages

package	Contents
VRTSzhvm	Chinese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages

Japanese language packages

The following table shows the package name and contents for each Japanese language package.

Table F-5 Japanese language packages

package	Contents
VRTSjacav	Japanese Veritas Cluster Server Agents for Storage Foundation Cluster File System – Manual Pages and Message Catalogs by Symantec
VRTSjacs	Veritas Cluster Server Japanese Message Catalogs by Symantec
VRTSjacse	Japanese Veritas High Availability Enterprise Agents by Symantec
VRTSjadba	Japanese Veritas Oracle Real Application Cluster Support package by Symantec
VRTSjadbe	Japanese Veritas Storage Foundation for Oracle from Symantec – Message Catalogs
VRTSjafs	Japanese Veritas File System – Message Catalog and Manual Pages
VRTSjaodm	Veritas Oracle Disk Manager Japanese Message Catalog and Manual Pages by Symantec
VRTSjavm	Japanese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages
VRTSmulic	Multi-language Symantec License Utilities

Veritas Storage Foundation obsolete and reorganized installation packages

[Table F-6](#) lists the packages that are obsolete or reorganized for Veritas Storage Foundation Cluster File System High Availability.

Table F-6 Veritas Storage Foundation obsolete and reorganized packages

package	Description
Obsolete and reorganized for 6.0	
VRTSat	Obsolete
VRTSatZH	Obsolete
VRTSatJA	Obsolete
Obsolete and reorganized for 5.1	
Infrastructure	

Table F-6 Veritas Storage Foundation obsolete and reorganized packages
(continued)

package	Description
SYMCIma	Obsolete
VRTSaa	Included in VRTSsfmh
VRTSccg	Included in VRTSsfmh
VRTSdbms3	Obsolete
VRTSicsco	Obsolete
VRTSjre	Obsolete
VRTSjre15	Obsolete
VRTSmh	Included in VRTSsfmh
VRTSobc33	Obsolete
VRTSobweb	Obsolete
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product packages	
VRTSacclib	<p>Obsolete</p> <p>The following information is for installations, upgrades, and uninstalls using the script- or Web-based installer.</p> <ul style="list-style-type: none"> ■ For fresh installations VRTSacclib is not installed. ■ For upgrades, the existing VRTSacclib is uninstalled and a new VRTSacclib is installed. ■ For uninstallation, VRTSacclib is not uninstalled.
VRTSalloc	Obsolete
VRTScmccc	Obsolete

Table F-6 Veritas Storage Foundation obsolete and reorganized packages
(continued)

package	Description
VRTScmcm	Obsolete
VRTScmcs	Obsolete
VRTScscm	Obsolete
VRTScscw	Obsolete
VRTScsocw	Obsolete
VRTScssim	Obsolete
VRTScutil	Obsolete
VRTSd2gui	Included in VRTSdbed
VRTSdb2ed	Included in VRTSdbed
VRTSdbcom	Included in VRTSdbed
VRTSdbed	Included in VRTSdbed
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSdsa	Obsolete
VRTSfas	Obsolete
VRTSfasag	Obsolete
VRTSfsman	Included in the product's main package.
VRTSfsmnd	Included in the product's main package.
VRTSfspro	Included in VRTSsfmh
VRTSgapms	Obsolete
VRTSmapro	Included in VRTSsfmh
VRTSorgui	Obsolete
VRTSsybed	Included in VRTSdbed
VRTSvail	Obsolete

Table F-6 Veritas Storage Foundation obsolete and reorganized packages
(continued)

package	Description
VRTSvcldb	Included in VRTSvcsea
VRTSvcsmn	Included in VRTSvc
VRTSvcsor	Included in VRTSvcsea
VRTSvcssy	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvc
VRTSvdid	Obsolete
VRTSvmman	Included in the product's main package.
VRTSvmpro	Included in VRTSsfmh
VRTSvrpro	Included in VRTSob
VRTSvrw	Obsolete
VRTSvxmsa	Obsolete
Documentation	All Documentation packages obsolete

High availability agent information

This appendix includes the following topics:

- [About agents](#)
- [Enabling and disabling intelligent resource monitoring for agents manually](#)
- [CVMCluster agent](#)
- [CVMVxconfigd agent](#)
- [CVMVolDg agent](#)
- [CFSMount agent](#)
- [CFSfsckd agent](#)

About agents

An agent is defined as a process that starts, stops, and monitors all configured resources of a type, and reports their status to Veritas Cluster Server (VCS). Agents have both entry points and attributes. Entry points are also known as agent functions and are referred to as "agent functions" throughout the document.

Attributes contain data about the agent. An attribute has a definition and a value. You change attribute values to configure resources, which are defined as the individual components that work together to provide application services to the public network. For example, a resource may be a physical component such as a disk or a network interface card, a software component such as Oracle or a Web server, or a configuration component such as an IP address or mounted file system.

Attributes are either optional or required, although sometimes the attributes that are optional in one configuration may be required in other configurations. Many optional attributes have predefined or default values, which you should change as required. A variety of internal use only attributes also exist. Do not modify these attributes—modifying them can lead to significant problems for your clusters. Attributes have type and dimension. Some attribute values can accept numbers, others can accept alphanumeric values or groups of alphanumeric values, while others are simple boolean on/off values.

The entry points and attributes for each SFCFSHA agent are described in this appendix.

VCS agents included within SFCFSHA

SFCFSHA includes the following VCS agents:

- CVMCluster agent
- CVMVxconfigd agent
- CVMVolDg agent
- CFSMount agent
- CFSfsckd
- Coordination Point agent

An SFCFSHA installation automatically configures the CVMCluster resource and the CVMVxconfigd resource.

You must configure the CVMVolDg agent for each shared disk group. If the database uses cluster file systems, configure the CFSMount agent for each volume in the disk group.

Use the information in this appendix about the entry points and attributes of the listed agents to make necessary configuration changes. For information on how to modify the VCS configuration:

See the *Veritas Cluster Server Administrator's Guide*

Enabling and disabling intelligent resource monitoring for agents manually

Review the following procedures to enable or disable intelligent resource monitoring manually. The intelligent resource monitoring feature is enabled by default. The IMF resource type attribute determines whether an IMF-aware agent must perform intelligent resource monitoring.

To enable intelligent resource monitoring

- 1 Make the VCS configuration writable.

```
# haconf -makerw
```

- 2 Run the following command to enable intelligent resource monitoring.

- To enable intelligent monitoring of offline resources:

```
# hatype -modify resource_type IMF -update Mode 1
```

- To enable intelligent monitoring of online resources:

```
# hatype -modify resource_type IMF -update Mode 2
```

- To enable intelligent monitoring of both online and offline resources:

```
# hatype -modify resource_type IMF -update Mode 3
```

- 3 If required, change the values of the MonitorFreq key and the RegisterRetryLimit key of the IMF attribute.

Review the agent-specific recommendations in the attribute definition tables to set these attribute key values.

See [“Attribute definition for CVMVxconfigd agent”](#) on page 470.

See [“Attribute definition for CFSSMount agent”](#) on page 476.

See [“Attribute definition for CFSfsckd agent”](#) on page 480.

- 4 Save the VCS configuration.

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

- 5 Restart the agent. Run the following commands on each node.

```
# haagent -stop agent_name -force -sys sys_name
```

```
# haagent -start agent_name -sys sys_name
```

To disable intelligent resource monitoring

- 1 Make the VCS configuration writable.

```
# haconf -makerw
```

- 2 To disable intelligent resource monitoring for all the resources of a certain type, run the following command:

```
# hatype -modify resource_type IMF -update Mode 0
```

- 3 To disable intelligent resource monitoring for a specific resource, run the following command:

```
# hares -override resource_name IMF  
# hares -modify resource_name IMF -update Mode 0
```

- 4 Save the VCS configuration.

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

Note: VCS provides `haimfconfig` script to enable or disable the IMF functionality for agents. You can use the script with VCS in running or stopped state. Use the script to enable or disable IMF for the IMF-aware bundled agents, enterprise agents, and custom agents.

Administering the AMF kernel driver

Review the following procedures to start or stop the AMF kernel driver.

To start the AMF kernel driver

- 1 Set the value of the `AMF_START` variable to 1 in the following file, if the value is not already 1:

```
/etc/default/amf
```

- 2 Start the AMF kernel driver. Run the following command:

```
svcadm enable amf
```

To stop the AMF kernel driver

- 1 Set the value of the AMF_START variable to 0 in the following file, if the value is not already 0:

```
/etc/default/amf
```

- 2 Stop the AMF kernel driver. Run the following command:

```
svcadm disable amf
```

To unload the AMF kernel driver

Unload the AMF kernel driver only after performing the following tasks:

- 1 Stop the agents that are registered with the AMF kernel driver.

The `amfstat` command output lists the agents that are registered with AMF under the Registered Reapers section. For more information about the `amfstat` command, see the manual page for the command.

- 2 Stop the AMF kernel driver by using the platform-specific command.
- 3 Start the agents.

To unload the AMF kernel driver with minimum downtime

- 1 Run the following command to disable the AMF driver even if agents are still registered with it.

```
amfconfig -Uof
```

- 2 Stop the AMF kernel driver by using the platform-specific command.

CVMCluster agent

The CVMCluster agent controls system membership on the cluster port that is associated with Veritas Volume Manager (VxVM).

The CVMCluster agent performs the following functions:

- Joins a node to the CVM cluster port.
- Removes a node from the CVM cluster port.
- Monitors the node's cluster membership state.

Entry points for CVMCluster agent

[Table G-1](#) describes the entry points used by the CVMCluster agent.

Table G-1 CVMCluster agent entry points

Entry Point	Description
Online	Joins a node to the CVM cluster port. Enables the Volume Manager cluster functionality by automatically importing the shared disk groups.
Offline	Removes a node from the CVM cluster port.
Monitor	Monitors the node's CVM cluster membership state.

Attribute definition for CVMCluster agent

[Table G-2](#) describes the user-modifiable attributes of the CVMCluster resource type.

Table G-2 CVMCluster agent attributes

Attribute	Description
CVMClustName	Name of the cluster. <ul style="list-style-type: none"> Type and dimension: string-scalar
CVMNodeAddr	List of host names and IP addresses. <ul style="list-style-type: none"> Type and dimension: string-association
CVMNodeId	Associative list. The first part names the system; the second part contains the LLT ID number for the system. <ul style="list-style-type: none"> Type and dimension: string-association
CVMTransport	Specifies the cluster messaging mechanism. <ul style="list-style-type: none"> Type and dimension: string-scalar Default = gab <p>Note: Do not change this value.</p>
PortConfigd	The port number that is used by CVM for vxconfigd-level communication. <ul style="list-style-type: none"> Type and dimension: integer-scalar
PortKmsgd	The port number that is used by CVM for kernel-level communication. <ul style="list-style-type: none"> Type and dimension: integer-scalar
CVMTimeout	Timeout in seconds used for CVM cluster reconfiguration. <ul style="list-style-type: none"> Type and dimension: integer-scalar Default = 200

CVMCluster agent type definition

The following type definition is included in the file, `CVMTypes.cf`:

```
type CVMCluster (
    static int InfoTimeout = 0
    static int NumThreads = 1
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMTransport, CVMClustName,
        CVMNodeAddr, CVMNodeId, PortConfigd, PortKmsgd,
        CVMTimeout }
    NameRule = ""
    str CVMClustName
    str CVMNodeAddr{}
    str CVMNodeId{}
    str CVMTransport
    int PortConfigd
    int PortKmsgd
    int CVMTimeout
)
```

Note: The attributes `CVMNodeAddr`, `PortConfigd`, and `PortKmsgd` are not used in an SFCFSHA environment. GAB, the required cluster communication messaging mechanism, does not use them.

CVMCluster agent sample configuration

The following is an example definition for the CVMCluster service group:

```
CVMCluster cvm_clus (
    Critical = 0
    CVMClustName = clus1
    CVMNodeId = { galaxy = 0, nebula = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)
```

CVMVxconfigd agent

The CVMVxconfigd agent starts and monitors the vxconfigd daemon. The vxconfigd daemon maintains disk and disk group configurations, communicates configuration

changes to the kernel, and modifies the configuration information that is stored on disks. CVMVxconfigd must be present in the CVM service group.

The CVMVxconfigd agent is an OnOnly agent; the agent starts the resource when the cluster starts up and VCS restarts the resource when necessary. The Operations attribute specifies these default aspects of startup.

Symantec recommends starting the vxconfigd daemon with the `syslog` option, which enables logging of debug messages. Note that the SFCFSHA installation configures the `syslog` option for the CVMVxconfigd agent.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

Entry points for CVMVxconfigd agent

[Table G-3](#) describes the entry points for the CVMVxconfigd agent.

Table G-3 CVMVxconfigd entry points

Entry Point	Description
Online	Starts the <code>vxconfigd</code> daemon
Offline	N/A
Monitor	Monitors whether <code>vxconfigd</code> daemon is running
<code>imf_init</code>	Initializes the agent to interface with the AMF kernel module. This function runs when the agent starts up.
<code>imf_getnotification</code>	Gets notification about the <code>vxconfigd</code> process state. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification. If the <code>vxconfigd</code> process fails, the function initiates a traditional CVMVxconfigd monitor entry point.
<code>imf_register</code>	Registers or unregisters the <code>vxconfigd</code> process id (pid) with the AMF kernel module. This function runs after the resource goes into steady online state.

Attribute definition for CVMVxconfigd agent

[Table G-4](#) describes the modifiable attributes of the CVMVxconfigd resource type.

Table G-4 CVMVxconfigd agent attribute

Attribute	Description
CVMVxconfigdArgs	<p>List of the arguments that are sent to the <code>online</code> entry point.</p> <p>Symantec recommends always specifying the <code>syslog</code> option.</p> <ul style="list-style-type: none"> ■ Type and dimension: keylist
IMF	<p>This resource-type level attribute determines whether the CVMVxconfigd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: <ul style="list-style-type: none"> ■ 0—Does not perform intelligent resource monitoring ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: <ul style="list-style-type: none"> ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the <code>imf_register</code> agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3. ■ Type and dimension: integer-association <p>For more details of IMF attribute for the agent type, refer to the <i>Veritas Cluster Server Administrator's Guide</i>.</p>

CVMVxconfigd agent type definition

The following type definition is included in the CVMTypes.cf file:

```
type CVMVxconfigd (
    static int FaultOnMonitorTimeouts = 2
    static int RestartLimit = 5
    static str ArgList[] { CVMVxconfigdArgs }
    static str Operations = OnOnly
    keylist CVMVxconfigdArgs
)
```

CVMVxconfigd agent sample configuration

The following is an example definition for the `CVMVxconfigd` resource in the CVM service group:

```
CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)
```

CVMVoIDg agent

The CVMVoIDg agent manages the CVM disk groups and CVM volumes and volume sets within the disk groups by performing the following functions:

- Imports the shared disk group from the CVM master node
- Starts the volumes and volume sets in the disk group
- Monitors the disk group, volumes, and volume sets
- Optionally, deports the disk group when the dependent applications are taken offline. The agent deports the disk group only if the appropriate attribute is set.

Configure the CVMVoIDg agent for each disk group used by a Oracle service group. A disk group must be configured to only one Oracle service group. If cluster file systems are used for the database, configure the CFMount agent for each volume or volume set in the disk group.

Entry points for CVMVoIDg agent

[Table G-5](#) describes the entry points used by the CVMVoIDg agent.

Table G-5 CVMVoIDg agent entry points

Entry Point	Description
Online	<p>Imports the shared disk group from the CVM master node, if the disk group is not already imported.</p> <p>Starts all volumes and volume sets in the shared disk group specified by the CVMVolume attribute.</p> <p>Sets the disk group activation mode to shared-write if the value of the CVMActivation attribute is sw. You can set the activation mode on both slave and master systems.</p>
Offline	<p>Removes the temporary files created by the online entry point.</p> <p>If the <code>CVMDeportOnOffline</code> attribute is set to 1 and if the shared disk group does not contain open volumes on any node in the cluster, the disk group is deported from the CVM master node.</p>
Monitor	<p>Determines whether the disk group, the volumes, and the volume sets are online.</p> <p>The agent takes a volume set offline if the file system metadata volume of a volume set is discovered to be offline in a monitor cycle.</p> <p>Note: If the CFSMount resource goes offline and the file system on the volume set is unmounted, the agent retains the online state of the volume set even if the file system metadata volume in the volume set is offline. This is because the CVMVoIDg agent is unable to determine whether or not the volumes that are offline are metadata volumes.</p>
Clean	Removes the temporary files created by the online entry point.

Attribute definition for CVMVoIDg agent

[Table G-6](#) describes the user-modifiable attributes of the CVMVoIDg resource type.

Table G-6 CVMVoIDg agent attributes

Attribute	Description
CVMDiskGroup (required)	<p>Shared disk group name.</p> <ul style="list-style-type: none"> ■ Type and dimension: string-scalar

Table G-6 CVMVoIDg agent attributes (*continued*)

Attribute	Description
CVMVolume (required)	<p>Name of shared volumes or volume sets. This list is used to check that the volumes or volume sets are in the correct state before allowing the resource to come online, and that the volumes remain in an enabled state.</p> <ul style="list-style-type: none"> Type and dimension: string-keylist
CVMActivation (required)	<p>Activation mode for the disk group.</p> <ul style="list-style-type: none"> Type and dimension: string-scalar Default = <code>sw</code> (<code>shared-write</code>) <p>This is a localized attribute.</p>
CVMVolumeIoTest(optional)	<p>List of volumes and volume sets that will be periodically polled to test availability. The polling is in the form of 4 KB reads every monitor cycle to a maximum of 10 of the volumes or volume sets in the list. For volume sets, reads are done on a maximum of 10 component volumes in each volume set.</p> <ul style="list-style-type: none"> Type and dimension: string-keylist
CVMDeportOnOffline (optional)	<p>Indicates whether or not the shared disk group must be deported when the last online CVMVoIDg resource for a disk group is taken offline.</p> <p>The value 1 indicates that the agent will deport the shared disk group from the CVM master node, if not already deported, when the last online CVMVoIDg resource for the disk group is taken offline.</p> <p>The value 0 indicates that the agent will not deport the shared disk group when the CVMVoIDg resource is taken offline.</p> <ul style="list-style-type: none"> Type and dimension: integer-scalar Default = 0 <p>Note: If multiple CVMVoIDg resources are configured for a shared disk group, set the value of the attribute to either 1 or 0 for all of the resources.</p> <p>The CVM disk group is deported based on the order in which the CVMVoIDg resources are taken offline. If the CVMVoIDg resources in the disk group contain a mixed setting of 1 and 0 for the <code>CVMDeportOnOffline</code> attribute, the disk group is deported only if the attribute value is 1 for the last CVMVoIDg resource taken offline. If the attribute value is 0 for the last CVMVoIDg resource taken offline, the disk group is not deported.</p> <p>The deport operation fails if the shared disk group contains open volumes.</p>

CVMVoIDg agent type definition

The `CVMTypes.cf` file includes the CVMVoIDg type definition:

```

type CVMVolDg (
    static keylist RegList = { CVMActivation, CVMVolume }
    static int OnlineRetryLimit = 2
    static int OnlineTimeout = 400
    static str ArgList[] = { CVMDiskGroup, CVMVolume, CVMActivation,
        CVMVolumeIoTest, CVMDGAction, CVMDeportOnOffline }
    str CVMDiskGroup
    str CVMDGAction
    keylist CVMVolume
    str CVMActivation
    keylist CVMVolumeIoTest
    int CVMDeportOnOffline
    temp int voldg_stat
)

```

CVMVolDg agent sample configuration

Each Oracle service group requires a CVMVolDg resource type to be defined. The following is a sample configuration:

```

CVMVolDg cvmvoldg1 (
    Critical = 0
    CVMDiskgroup = testdg
    CVMVolume = { vol1, vol2, mvoll, mvoll2, snapvol, vset1 }
    CVMVolumeIoTest = { snapvol, vset1 }
    CVMActivation @system1 = sw
    CVMActivation @system2 = sw
    CVMDeportOnOffline = 1
)

```

CFSMount agent

The CFSMount agent brings online, takes offline, and monitors a cluster file system mount point.

The agent executable is located in `/opt/VRTSvcs/bin/CFSMount/CFSMountAgent`.

The CFSMount type definition is described in the `/etc/VRTSvcs/conf/config/CFSTypes.cf` file.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

Entry points for CFSMount agent

Table G-7 provides the entry points for the CFSMount agent.

Table G-7 CFSMount agent entry points

Entry Point	Description
Online	Mounts a block device in cluster mode.
Offline	Unmounts the file system, forcing unmount if necessary, and sets primary to secondary if necessary.
Monitor	Determines if the file system is mounted. Checks mount status using the <code>fsclustadm</code> command.
Clean	Generates a null operation for a cluster file system mount.
imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for the agent. This function runs when the agent starts up.
imf_getnotification	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).

Attribute definition for CFSMount agent

Table G-8 lists user-modifiable attributes of the CFSMount Agent resource type.

Table G-8 CFSMount Agent attributes

Attribute	Description
MountPoint	Directory for the mount point. <ul style="list-style-type: none"> Type and dimension: string-scalar
BlockDevice	Block device for the mount point. <ul style="list-style-type: none"> Type and dimension: string-scalar

Table G-8 CFSMount Agent attributes (*continued*)

Attribute	Description
NodeList	<p>List of nodes on which to mount. If NodeList is NULL, the agent uses the service group system list.</p> <ul style="list-style-type: none"> ■ Type and dimension: string-keylist
IMF	<p>Resource-type level attribute that determines whether the CFSMount agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: <ul style="list-style-type: none"> ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: <ul style="list-style-type: none"> ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3. ■ Type and dimension: integer-association <p>See “Enabling and disabling intelligent resource monitoring for agents manually” on page 464.</p>

Table G-8 CFSMount Agent attributes (*continued*)

Attribute	Description
MountOpt (optional)	<p>Options for the mount command. To create a valid MountOpt attribute string:</p> <ul style="list-style-type: none"> Use the VxFS type-specific options only. Do not use the <code>-o</code> flag to specify the VxFS-specific options. Do not use the <code>-F vxfs</code> file system type option. Be aware the cluster option is not required. Specify options in comma-separated list: <pre>ro ro,cluster blkclear,mincache=closesync</pre> <ul style="list-style-type: none"> Type and dimension: string-scalar
Policy (optional)	<p>List of nodes to assume the primaryship of the cluster file system if the primary fails. If set to NULL or if none of the hosts specified in the list is active when the primary fails, a node is randomly selected from the set of active nodes to assume primaryship.</p> <ul style="list-style-type: none"> Type and dimension: string-scalar

CFSMount agent type definition

The `CFSTypes.cf` file includes the CFSMount agent type definition:

```
type CFSMount (
    static keylist RegList = { MountOpt, Policy, NodeList, ForceOff, SetPrimary }
    static keylist SupportedActions = { primary }
    static int FaultOnMonitorTimeouts = 1
    static int OnlineWaitLimit = 1
    static str ArgList[] = { MountPoint, BlockDevice, MountOpt, Primary, AMFMountType }
    str MountPoint
    str MountType
    str BlockDevice
    str MountOpt
    keylist NodeList
    keylist Policy
    temp str Primary
    str SetPrimary
    temp str RemountRes
    temp str AMFMountType
    str ForceOff
)
```

CFSMount agent sample configuration

Each Oracle service group requires a CFSMount resource type to be defined:

```
CFSMount ora_mount (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavoll"
    Primary = nebula;
)
```

To see CFSMount defined in a more extensive example:

CFSfsckd agent

The CFSfsckd agent starts, stops, and monitors the `vxfsckd` process. The CFSfsckd agent executable is `/opt/VRTSvcs/bin/CFSfsckd/CFSfsckdAgent`. The type definition is in the `/etc/VRTSvcs/conf/config/CFSTypes.cf` file. The configuration is added to the `main.cf` file after running the `cfsccluster config` command.

This agent is IMF-aware and uses asynchronous monitoring framework (AMF) kernel driver for IMF notification. For more information about the Intelligent Monitoring Framework (IMF) and intelligent resource monitoring, refer to the *Veritas Cluster Server Administrator's Guide*.

Entry points for CFSfsckd agent

[Table G-9](#) describes the CFSfsckd agent entry points.

Table G-9 CFSfsckd agent entry points

Entry Points	Description
Online	Starts the <code>vxfsckd</code> process.
Offline	Kills the <code>vxfsckd</code> process.
Monitor	Checks whether the <code>vxfsckd</code> process is running.
Clean	A null operation for a cluster file system mount.
imf_init	Initializes the agent to interface with the AMF kernel driver, which is the IMF notification module for the agent. This function runs when the agent starts up.

Table G-9 CFSfsckd agent entry points (*continued*)

Entry Points	Description
imf_getnotification	Gets notification about resource state changes. This function runs after the agent initializes with the AMF kernel module. This function continuously waits for notification and takes action on the resource upon notification.
imf_register	Registers or unregisters resource entities with the AMF kernel module. This function runs for each resource after the resource goes into steady state (online or offline).

Attribute definition for CFSfsckd agent

[Table G-10](#) lists user-modifiable attributes of the CFSfsckd Agent resource type.

Table G-10 CFSfsckd Agent attributes

Attribute	Description
IMF	<p>Resource-type level attribute that determines whether the CFSfsckd agent must perform intelligent resource monitoring. You can also override the value of this attribute at resource-level.</p> <p>This attribute includes the following keys:</p> <ul style="list-style-type: none"> ■ Mode: Define this attribute to enable or disable intelligent resource monitoring. Valid values are as follows: <ul style="list-style-type: none"> ■ 0—Does not perform intelligent resource monitoring ■ 1—Performs intelligent resource monitoring for offline resources and performs poll-based monitoring for online resources ■ 2—Performs intelligent resource monitoring for online resources and performs poll-based monitoring for offline resources ■ 3—Performs intelligent resource monitoring for both online and for offline resources Default: 0 ■ MonitorFreq: This key value specifies the frequency at which the agent invokes the monitor agent function. The value of this key is an integer. Default: 1 You can set this key to a non-zero value for cases where the agent requires to perform both poll-based and intelligent resource monitoring. If the value is 0, the agent does not perform poll-based process check monitoring. After the resource registers with the AMF kernel driver, the agent calls the monitor agent function as follows: <ul style="list-style-type: none"> ■ After every (MonitorFreq x MonitorInterval) number of seconds for online resources ■ After every (MonitorFreq x OfflineMonitorInterval) number of seconds for offline resources ■ RegisterRetryLimit: If you enable intelligent resource monitoring, the agent invokes the imf_register agent function to register the resource with the AMF kernel driver. The value of the RegisterRetyLimit key determines the number of times the agent must retry registration for a resource. If the agent cannot register the resource within the limit that is specified, then intelligent monitoring is disabled until the resource state changes or the value of the Mode key changes. Default: 3. ■ Type and dimension: integer-association <p>See “Enabling and disabling intelligent resource monitoring for agents manually” on page 464.</p>

CFSfsckd agent type definition

The CFSfsckd type definition:

```
type CFSfsckd (
    static int RestartLimit = 1
```

```
        str ActivationMode{}  
    )
```

CFSfsckd agent sample configuration

This is a sample of CFSfsckd configuration:

```
CFSfsckd vxfsckd (  
)
```

Troubleshooting information

This appendix includes the following topics:

- [Restarting the installer after a failed connection](#)
- [What to do if you see a licensing reminder](#)
- [Troubleshooting information](#)
- [Storage Foundation Cluster File System High Availability installation issues](#)
- [Storage Foundation Cluster File System High Availability problems](#)

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.

What to do if you see a licensing reminder

In this release, you can install without a license key. In order to comply with the End User License Agreement, you must either install a license key or make the host managed by a Management Server. If you do not comply with these terms within 60 days, the following warning messages result:

```
WARNING V-365-1-1 This host is not entitled to run Veritas Storage
Foundation/Veritas Cluster Server.As set forth in the End User
License Agreement (EULA) you must complete one of the two options
```

set forth below. To comply with this condition of the EULA and stop logging of this message, you have <nn> days to either:

- make this host managed by a Management Server (see <http://go.symantec.com/sfhakeyless> for details and free download), or
- add a valid license key matching the functionality in use on this host using the command 'vxlicinst'

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. See “[Installing Veritas SFHA Solutions product license keys](#)” on page 52. After you install the license key, you must validate the license key using the following command:

```
# /opt/VRTS/bin/vxkeyless
```

- 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>

Troubleshooting information

The VRTSpt package provides a group of tools for troubleshooting a system and collecting information on its configuration. The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. Although the tools are not required for the operation of any Veritas product, Symantec recommends installing them should a support case be needed to be opened with Symantec Support. If you are unfamiliar with their use and purpose, use caution when using them or use them in concert with Symantec Support.

Storage Foundation Cluster File System High Availability installation issues

If you encounter any issues installing SFCFSHA, refer to the following paragraphs for typical problems and their solutions:

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Failed to setup rsh communication on 10.198.89.241:
'rsh 10.198.89.241 <command>' failed
Trying to setup ssh communication on 10.198.89.241.
Failed to setup ssh communication on 10.198.89.241:
Login denied
```

```
Failed to login to remote system(s) 10.198.89.241.
Please make sure the password(s) are correct and superuser(root)
can login to the remote system(s) with the password(s).
If you want to setup rsh on remote system(s), please make sure
rsh with command argument ('rsh <host> <command>') is not
denied by remote system(s).
```

```
Either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

```
Would you like the installer to setup ssh/rsh communication
automatically between the nodes?
Superuser passwords for the systems will be asked. [y,n,q] (y) n
```

```
System verification did not complete successfully
```

```
The following errors were discovered on the systems:
```

```
The ssh permission denied on 10.198.89.241
rsh exited 1 on 10.198.89.241
either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

Suggested solution: You need to set up the systems to allow remote access using `ssh` or `rsh`.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 446.

Note: Remove remote shell permissions after completing the SFCFSHA installation and configuration.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
Verifying systems: 12% .....
Estimated time remaining: 0:10 1 of 8
Checking system communication ..... Done
System verification did not complete successfully
The following errors were discovered on the systems:
cannot resolve hostname host1
Enter the system names separated by spaces: q,? (host1)
```

Suggested solution: Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

Storage Foundation Cluster File System High Availability problems

If there is a device failure or controller failure to a device, the file system may become disabled cluster-wide. To address the problem, unmount file system on all the nodes, then run a full `fsck`. When the file system check completes, mount all nodes again.

Unmount failures

The `umount` command can fail if a reference is being held by an NFS server. Unshare the mount point and try the unmount again.

Mount failures

Mounting a file system can fail for the following reasons:

- The file system is not using disk layout Version 7 or later.
- The mount options do not match the options of already mounted nodes.
- A cluster file system is mounted by default with the `qio` option enabled if the node has a Quick I/O for Databases license installed, even if the `qio` mount option was not explicitly specified. If the Quick I/O license is not installed, a cluster file system is mounted without the `qio` option enabled. So if some nodes in the cluster have a Quick I/O license installed and others do not, a cluster mount can succeed on some nodes and fail on others due to different mount options. To avoid this situation, ensure that Quick I/O licensing is uniformly

applied, or be careful to mount the cluster file system with the `qio/noqio` option appropriately specified on each node of the cluster.
 See the `mount(1M)` manual page.

- A shared CVM volume was not specified.
- The device is still mounted as a local file system somewhere on the cluster. Unmount the device.
- The `fsck` or `mkfs` command is being run on the same volume from another node, or the volume is mounted in non-cluster mode from another node.
- The `vxfsckd` daemon is not running. This typically happens only if the `CFSfsckd` agent was not started correctly.
- If `mount` fails with an error message:

```
vxfs mount: cannot open mnttab
/etc/mnttab is missing or you do not have root privileges.
```

- If `mount` fails with an error message:

```
vxfs mount: device already mounted, ...
```

The device is in use by `mount`, `mkfs` or `fsck` on the same node. This error cannot be generated from another node in the cluster.

- If this error message displays:

```
mount: slow
```

The node may be in the process of joining the cluster.

- If you try to mount a file system that is already mounted without `-o cluster` option (that is, not in shared mode) on another cluster node,

```
# mount -F vxfs /dev/vx/dsk/share/vol01 /vol01
```

The following error message displays:

```
vxfs mount: /dev/vx/dsk/share/vol01 is already mounted,
/vol01 is busy, allowable number of mount points exceeded,
or cluster reservation failed for the volume
```

Command failures

This section describes command failures.

- Manual pages not accessible with the `man` command. Set the `MANPATH` environment variable appropriately. See “[Setting environment variables](#)” on page 31.
- The `mount`, `fsck`, and `mkfs` utilities reserve a shared volume. They fail on volumes that are in use. Be careful when accessing shared volumes with other utilities such as `dd`, it is possible for these commands to destroy data on the disk.
- Running some commands, such as `vxupgrade -n 7 /vol02`, can generate the following error message:

```
vxfs vxupgrade: ERROR: not primary in a cluster file system
```

This means that you can run this command only on the primary, that is, the system that mounted this file system first.

Performance issues

Quick I/O File system performance is adversely affected if a cluster file system is mounted with the `qio` option enabled, but the file system is not used for Quick I/O files. Because `qio` is enabled by default, if you do not intend to use a shared file system for Quick I/O, explicitly specify the `noqio` option when mounting.

High availability issues

This section describes high availability issues.

Network partition and jeopardy

Network partition (or split brain) is a condition where a network failure can be misinterpreted as a failure of one or more nodes in a cluster. If one system in the cluster incorrectly assumes that another system failed, it may restart applications already running on the other system, thereby corrupting data. CFS tries to prevent this by having redundant heartbeat links.

At least one link must be active to maintain the integrity of the cluster. If all the links go down, after the last network link is broken, the node can no longer communicate with other nodes in the cluster. Thus the cluster is in one of two possible states. Either the last network link is broken (called a network partition condition), or the last network link is okay, but the node crashed, in which case it is not a network partition problem. It is not possible to identify whether it is the first or second state, so a kernel message is issued to indicate that a network partition may exist and there is a possibility of data corruption.

Jeopardy is a condition where a node in the cluster has a problem connecting to other nodes. In this situation, the link or disk heartbeat may be down, so a jeopardy warning may be displayed. Specifically, this message appears when a node has only one remaining link to the cluster and that link is a network link. This is considered a critical event because the node may lose its only remaining connection to the network.

Warning: Do not remove the communication links while shared storage is still connected.

Low memory

Under heavy loads, software that manages heartbeat communication links may not be able to allocate kernel memory. If this occurs, a node halts to avoid any chance of network partitioning. Reduce the load on the node if this happens frequently.

A similar situation may occur if the values in the `/etc/llttab` files on all cluster nodes are not correct or identical.

Troubleshooting cluster installation

This appendix includes the following topics:

- [Installer cannot create UUID for the cluster](#)
- [The vxfsntsthdw utility fails when SCSI TEST UNIT READY command fails](#)
- [Troubleshooting CP server](#)
- [Troubleshooting server-based fencing on the SFCFSHA cluster nodes](#)

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 SFCFSHA configuration, upgrade, or when you add a node to the cluster using the installer.

Workaround: To start SFCFSHA, 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.

The vxfcntlshdw utility fails when SCSI TEST UNIT READY command fails

While running the vxfcntlshdw 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.

Troubleshooting CP server

All CP server operations and messages are logged in the /var/VRTSvcs/log directory in a detailed and easy to read format. The entries are sorted by date and time. The logs can be used for troubleshooting purposes or to review for any possible security issue on the system that hosts the CP server.

The following files contain logs and text files that may be useful in understanding and troubleshooting a CP server:

- /var/VRTSvcs/log/cpsrvr_[ABC].log
- /var/VRTSvcs/log/vcsauthserver.log (Security related)
- If the vxcpssrv process fails on the CP server, then review the following diagnostic files:
 - /var/VRTSvcs/diag/FFDC_CPS_pid_vxcpssrv.log
 - /var/VRTSvcs/diag/stack_pid_vxcpssrv.txt

Note: If the `vxcperv` process fails on the CP server, these files are present in addition to a core file. VCS restarts `vxcperv` process automatically in such situations.

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs that may be useful in understanding and troubleshooting fencing-related issues on a SFCFSHA cluster (client cluster) node.

See [“Troubleshooting issues related to the CP server service group”](#) on page 492.

See [“Checking the connectivity of CP server”](#) on page 492.

See [“Issues during fencing startup on SFCFSHA cluster nodes set up for server-based fencing”](#) on page 493.

See [“Issues during online migration of coordination points”](#) on page 494.

Troubleshooting issues related to the CP server service group

If you cannot bring up the CPSSG service group after the CP server configuration, perform the following steps:

- Verify that the CPSSG service group and its resources are valid and properly configured in the VCS configuration.
- Check the VCS engine log (`/var/VRTSvcs/log/engine_[ABC].log`) to see if any of the CPSSG service group resources are `FAULTED`.
- Review the sample dependency graphs to make sure the required resources are configured correctly.

Checking the connectivity of CP server

You can test the connectivity of CP server using the `cpsadm` command.

You must have set the environment variables `CPS_USERNAME` and `CPS_DOMAINTYPE` to run the `cpsadm` command on the SFCFSHA cluster (client cluster) nodes.

To check the connectivity of CP server

- ◆ Run the following command to check whether a CP server is up and running at a process level:

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

where `cp_server` is the virtual IP address or virtual hostname on which the CP server is listening.

Troubleshooting server-based fencing on the SFCFSHA cluster nodes

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs files that may be useful in understanding and troubleshooting fencing-related issues on a SFCFSHA cluster (application cluster) node.

Issues during fencing startup on SFCFSHA cluster nodes set up for server-based fencing

Table I-1 Fencing startup issues on SFCFSHA cluster (client cluster) nodes

Issue	Description and resolution
<p><code>cpsadm</code> command on the SFCFSHA cluster gives connection error</p>	<p>If you receive a connection error message after issuing the <code>cpsadm</code> command on the SFCFSHA cluster, perform the following actions:</p> <ul style="list-style-type: none"> ■ Ensure that the CP server is reachable from all the SFCFSHA cluster nodes. ■ Check that the SFCFSHA cluster nodes use the correct CP server virtual IP or virtual hostname and the correct port number. Check the <code>/etc/vxfenmode</code> file. ■ Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.
<p>Authorization failure</p>	<p>Authorization failure occurs when the CP server's nodes or users are not added in the CP server configuration. Therefore, fencing on the SFCFSHA 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.</p> <p>To resolve this issue, add the CP server node and user in the CP server configuration and restart fencing.</p> <p>See "Preparing the CP servers manually for use by the SFCFSHA cluster" on page 163.</p>
<p>Authentication failure</p>	<p>If you had configured secure communication between the CP server and the SFCFSHA cluster (client cluster) nodes, authentication failure can occur due to the following causes:</p> <ul style="list-style-type: none"> ■ Symantec Product Authentication Services (AT) is not properly configured on the CP server and/or the SFCFSHA cluster. ■ The CP server and the SFCFSHA cluster nodes use different root brokers, and trust is not established between the authentication brokers:

Issues during online migration of coordination points

During online migration of coordination points using the `vxfenswap` utility, the operation is automatically rolled back if a failure is encountered during validation of coordination points from any of the cluster nodes.

Validation failure of the new set of coordination points can occur in the following circumstances:

- The `/etc/vxfenmode` file is not updated on all the SFCFSHA cluster nodes, because new coordination points on the node were being picked up from an old `/etc/vxfenmode` file.
- The coordination points listed in the `/etc/vxfenmode` file on the different SFCFSHA cluster nodes are not the same. If different coordination points are listed in the `/etc/vxfenmode` file on the cluster nodes, then the operation fails due to failure during the coordination point snapshot check.
- There is no network connectivity from one or more SFCFSHA cluster nodes to the CP server(s).
- Cluster, nodes, or users for the SFCFSHA cluster nodes have not been added on the new CP servers, thereby causing authorization failure.

Vxfen service group activity after issuing the `vxfenswap` command

The Coordination Point agent reads the details of coordination points from the `vxfenconfig -l` output and starts monitoring the registrations on them.

Thus, during `vxfenswap`, when the `vxfenmode` file is being changed by the user, the Coordination Point agent does not move to FAULTED state but continues monitoring the old set of coordination points.

As long as the changes to `vxfenmode` file are not committed or the new set of coordination points are not reflected in `vxfenconfig -l` output, the Coordination Point agent continues monitoring the old set of coordination points it read from `vxfenconfig -l` output in every monitor cycle.

The status of the Coordination Point agent (either ONLINE or FAULTED) depends upon the accessibility of the coordination points, the registrations on these coordination points, and the fault tolerance value.

When the changes to `vxfenmode` file are committed and reflected in the `vxfenconfig -l` output, then the Coordination Point agent reads the new set of coordination points and proceeds to monitor them in its new monitor cycle.

Sample SFCFSHA 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:

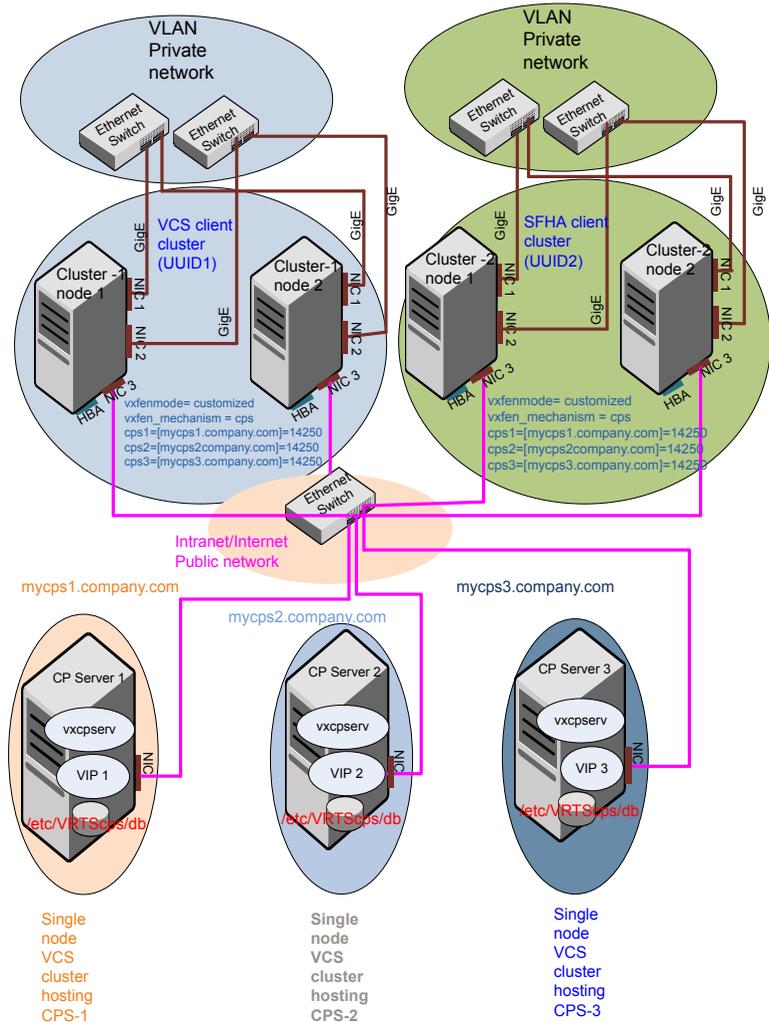
- Two unique client clusters that are served by 3 CP servers:
See [Figure J-1](#) on page 496.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
[Figure J-2](#)
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
[Figure J-3](#)
- Multiple 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



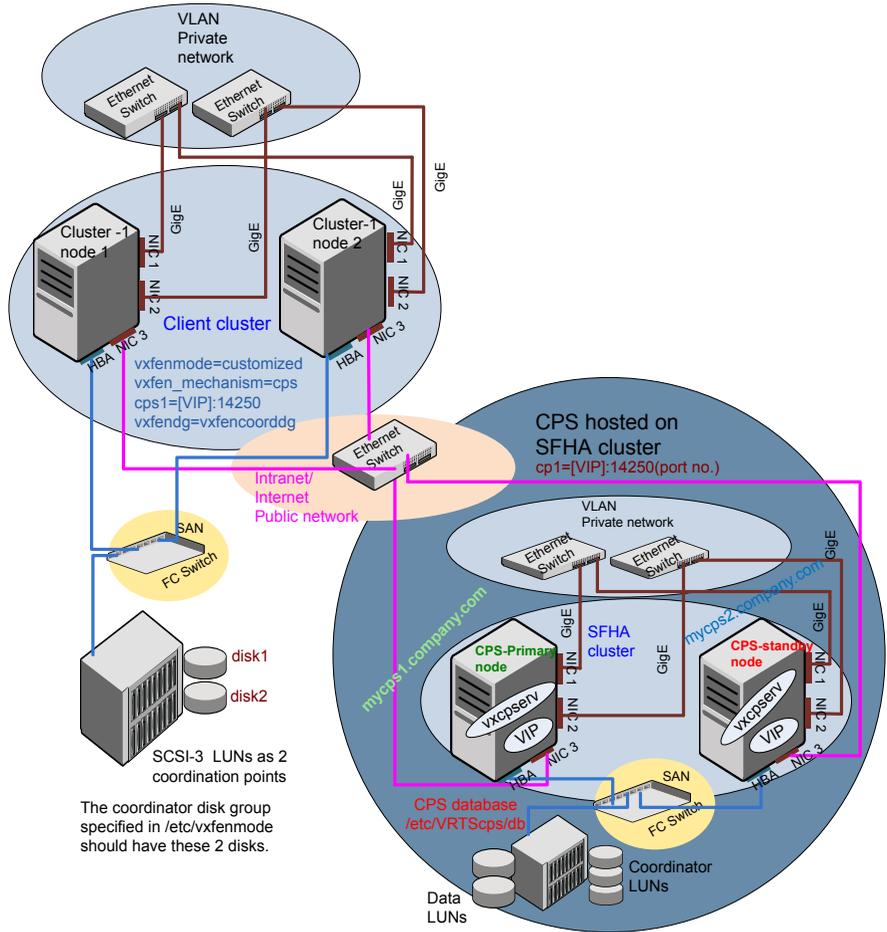
Client cluster served by highly available CPS and 2 SCSI-3 disks

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 `vxfsenmode` file on the client nodes, `vxfsenmode` is set to customized with `vxfsen` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfsencoordg`. 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



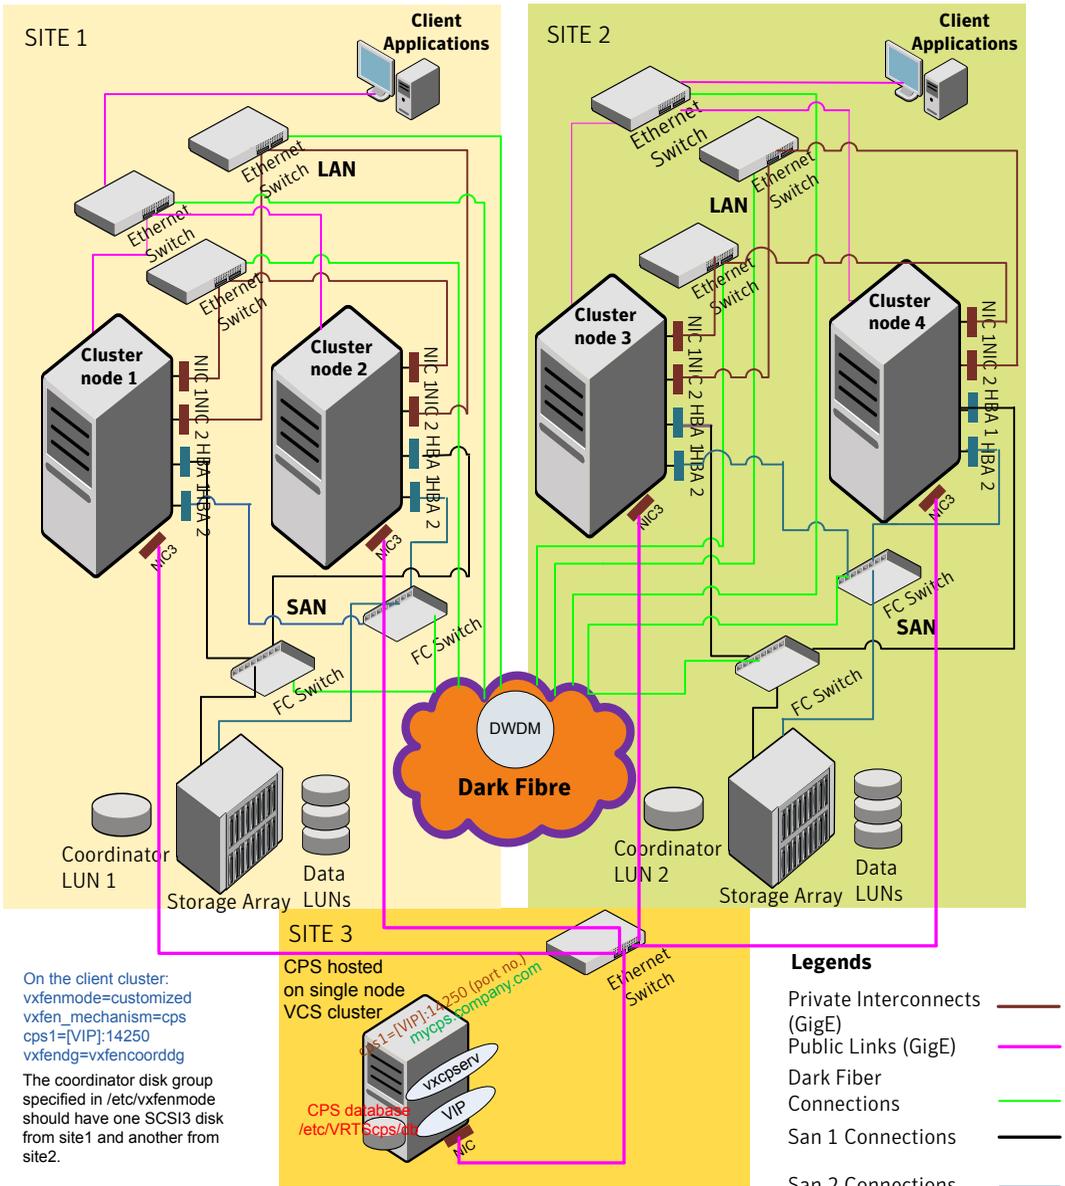
Two node campus cluster served by remote CP server and 2 SCSI-3 disks

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 vxfencoordg.
 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



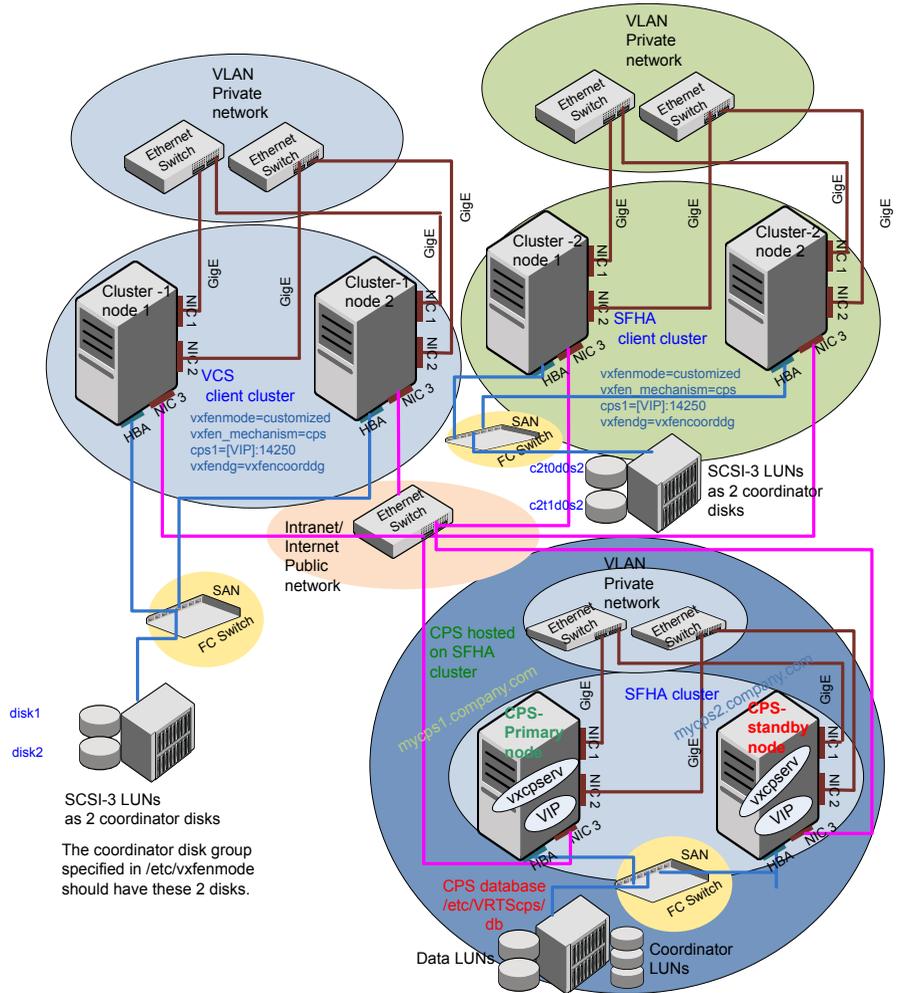
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

[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 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



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/shreddg/vol13`. 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 -l /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:

```
vxdump 30
vxio 32
vxspec 33
vxfen 87
vxglm 91
```

Output on Node B:

```
vxdump 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 `vxidg(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.

Configuring LLT over UDP using IPv6

This appendix includes the following topics:

- [Using the UDP layer of IPv6 for LLT](#)
- [Manually configuring LLT over UDP using IPv6](#)

Using the UDP layer of IPv6 for LLT

Veritas Storage Foundation Cluster File System High Availability 6.0 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

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

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/llttab` files are consistent with the IPv6 addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
See [“Selecting UDP ports”](#) on page 511.
- For the links that cross an IP router, disable multicast features and specify the IPv6 address of each link manually in the `/etc/llttab` file.
See [“Sample configuration: links crossing IP routers”](#) on page 514.

The link command in the `/etc/llttab` file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See [“Sample configuration: direct-attached links”](#) on page 512.
- See [“Sample configuration: links crossing IP routers”](#) on page 514.

Note that some of the fields in [Table L-1](#) differ from the command for standard LLT links.

[Table L-1](#) describes the fields of the link command that are shown in the `/etc/llttab` file examples.

Table L-1 Field description for link command in `/etc/llttab`

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example <code>/dev/udp6</code> .
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp6" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “Selecting UDP ports” on page 511.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IPv6 address</i>	IPv6 address of the link on the local node.
<i>mcast-address</i>	"-" is the default for clusters spanning routers.

The set-addr command in the /etc/llttab file

The `set-addr` command in the `/etc/llttab` file is required when the multicast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 514.

[Table L-2](#) describes the fields of the `set-addr` command.

Table L-2 Field description for `set-addr` command in `/etc/llttab`

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.
<i>link tag-name</i>	The string that LLT uses to identify the link; for example link1, link2,....
<i>address</i>	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

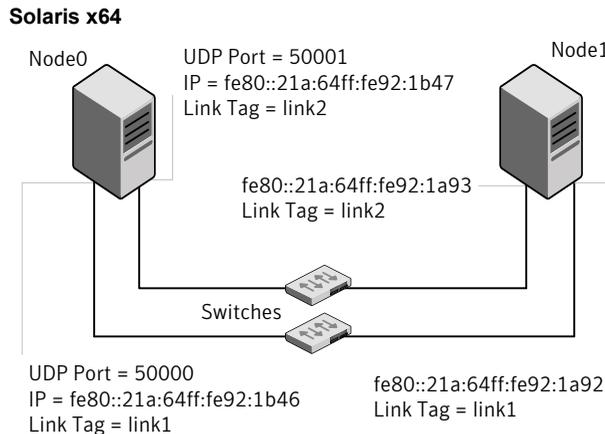
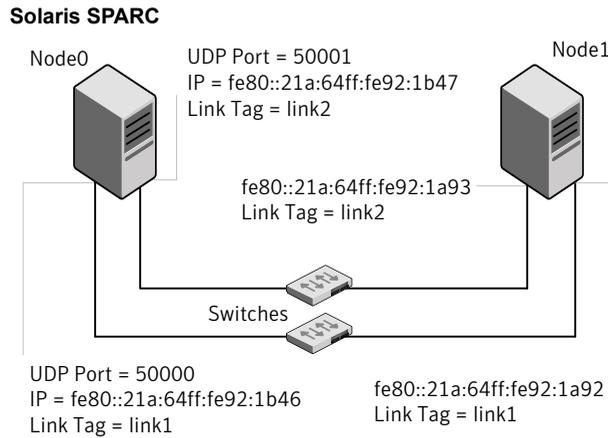
Local Address	Remote Address	State	If
*.servicetag		Idle	
*.177		Idle	

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 L-1](#) depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure L-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
```

```
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```

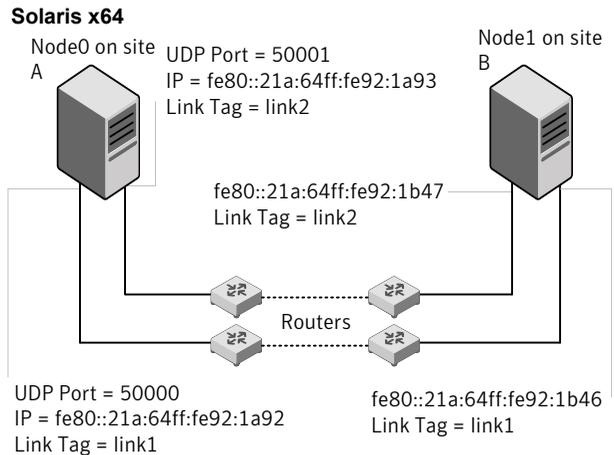
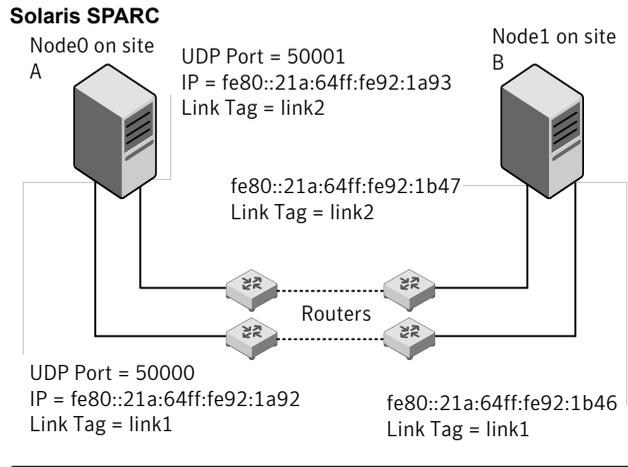
The file for Node 1 resembles:

```
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 L-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 L-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 /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
```

Configuring LLT over UDP using IPv4

This appendix includes the following topics:

- [Using the UDP layer for LLT](#)
- [Manually configuring LLT over UDP using IPv4](#)

Using the UDP layer for LLT

SFCFSHA 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/lfttab` explicitly depending on the subnet for each link. See [“Broadcast address in the /etc/lfttab file”](#) on page 518.

- 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 520.
- Set the broadcast address correctly for direct-attached (non-routed) links. See [“Sample configuration: direct-attached links”](#) on page 522.
- 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 523.

Broadcast address in the `/etc/llttab` file

The broadcast address is set explicitly for each link in the following example.

- Display the content of the `/etc/llttab` file on the first node galaxy:

```
galaxy # cat /etc/llttab

set-node galaxy
set-cluster 1
link link1 /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 nebula:

```
nebula # cat /etc/llttab

set-node nebula
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/llttab` file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See [“Sample configuration: direct-attached links”](#) on page 522.
- See [“Sample configuration: links crossing IP routers”](#) on page 523.

[Table M-1](#) describes the fields of the link command that are shown in the /etc/llttab file examples. Note that some of the fields differ from the command for standard LLT links.

Table M-1 Field description for link command in /etc/llttab

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example /dev/udp.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “Selecting UDP ports” on page 520.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IP address</i>	IP address of the link on the local node.
<i>bcast-address</i>	<ul style="list-style-type: none"> ■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the /etc/llttab file

The `set-addr` command in the /etc/llttab file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See [“Sample configuration: links crossing IP routers”](#) on page 523.

[Table M-2](#) describes the fields of the set-addr command.

Table M-2 Field description for set-addr command in /etc/llttab

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.

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 galaxy:

```
IP address=192.168.9.1, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

For the first network interface on the node nebula:

```
IP address=192.168.9.2, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

- For the second network interface on the node galaxy:

```
IP address=192.168.10.1, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

For the second network interface on the node nebula:

```
IP address=192.168.10.2, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.

An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

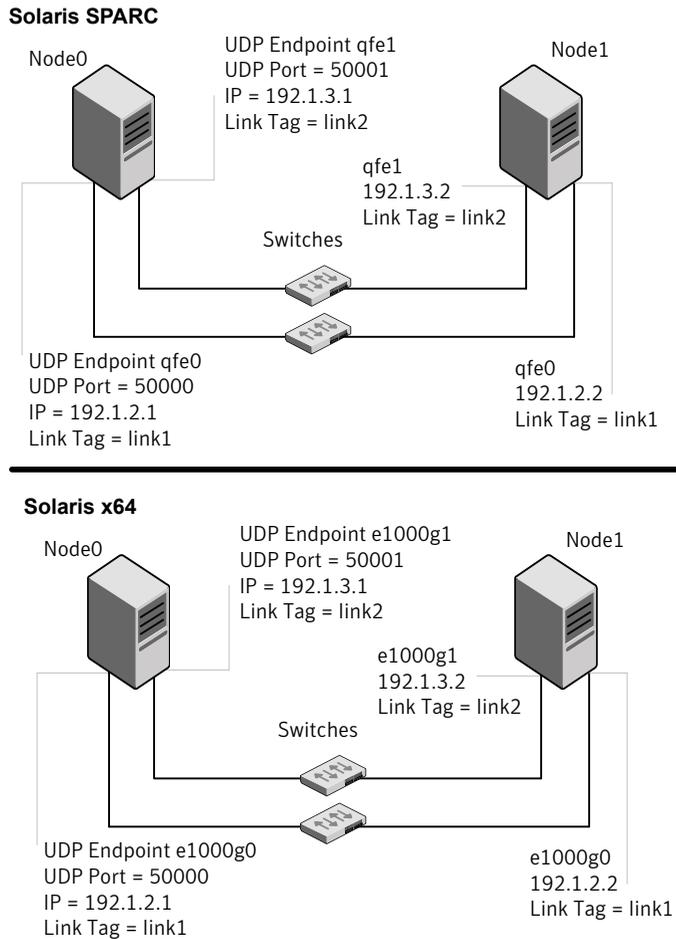
```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100
```

```
link link1 /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 M-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure M-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /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
```

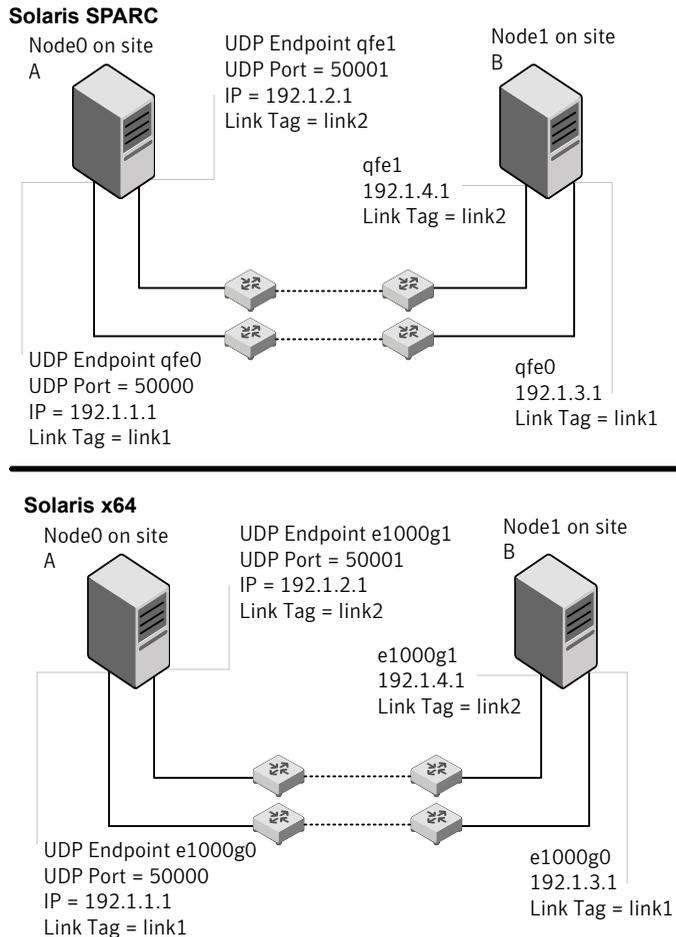
The file for Node 1 resembles:

```
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 M-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 M-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

```

Compatibility issues when installing Storage Foundation Cluster File System High Availability with other products

This appendix includes the following topics:

- [Installing, uninstalling, or upgrading Storage Foundation products when other Veritas 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 Veritas products are present

Installing Storage Foundation when other Veritas 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 SFM or VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the SFM Central Server and Managed Host packages as is.
- When uninstalling Storage Foundation products where SFM or VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where SFM or 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 VRTSspb and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspb, VRTSicsco, and VRTSat.

Index

A

- adding
 - users 120
- agents
 - about 463
 - CFSfsckd 479
 - CFSMount 475, 479
 - CVMCluster 467
 - CVMVolDg 472
 - CVMVxconfigd 469
 - disabling 389
 - of VCS 464
- application
 - database replication 359
- applications, stopping 195
- attributes
 - about agent attributes 463
 - CFSMount agent 476, 480
 - CVMCluster agent 468
 - CVMVolDg agent 468, 473
 - CVMVxconfigd agent 470
 - UseFence 155

B

- backup boot disk group 260–261
 - rejoining 260
- block device
 - partitions
 - example file name 502
 - volumes
 - example file name 502
- bootdg 137

C

- cables
 - cross-over Ethernet 293
- CFS
 - mount and unmount failures 486
 - synchronization 275
 - troubleshooting 486

- CFSfsckd agent 479
 - attributes 480
- CFSMount agent 475, 479
 - attributes 476
 - entry points 476
 - sample configuration 478–479
 - type definition 478
- CFSTypes.cf 478
- cluster
 - removing a node from 328
 - verifying operation 286
- cluster functionality
 - enabling 139
 - shared disks 140
- command failures 488
- commands
 - format 36
 - gabconfig 285
 - gcoconfig 345
 - hacf 202
 - hastatus 286
 - hasys 287
 - lltconfig 278
 - lltstat 282
 - vradmin 360
 - vxassist 352, 355
 - vxdisksetup (initializing disks) 144
 - vxlicinst 128
 - vxlicrep 127
 - vxvol 352
- configuration daemon (vxconfigd)
 - starting 135
- configuration file
 - main.cf 276
- configuring
 - rsh 33
 - shared disks 140
 - ssh 33
- configuring SFCFSHA
 - script-based installer 108
- configuring VCS
 - adding users 120

- configuring VCS (*continued*)
 - event notification 121, 123
 - global clusters 125
 - secure mode 116
 - starting 109

- controllers

- SCSI 34

- coordinator disks

- DMP devices 80
 - for I/O fencing 80
 - setting up 153

- creating

- Flash archive 72
 - post-deployment scripts 73

- CVM

- CVMTypes.cf file 469

- CVMCluster agent 467

- attributes 468
 - entry points 467
 - sample configuration 469
 - type definition 469

- CVMTypes.cf

- definition, CVMCluster agent 469
 - definition, CVMVolDg agent 474
 - definition, CVMVxconfigd agent 472

- CVMVolDg agent 472

- attributes 473
 - entry points 472
 - sample configuration 475
 - type definition 474

- CVMVxconfigd agent 469

- attributes 470
 - CVMTypes.cf 472
 - entry points 470
 - sample configuration 472
 - type definition 472

D

- data disks

- for I/O fencing 80

- default disk group 137

- defaultdg 137

- devices

- suppress devices 137

- disabling the agents 389

- disk groups

- bootdg 137
 - default 137
 - nodg 137

- disk groups (*continued*)

- root 137
 - rootdg 134, 137

- disks

- adding and initializing 144
 - coordinator 153
 - testing with vxfsntsthdw 147
 - verifying node access 149

E

- Ethernet controllers 293

F

- FC-AL controllers 36

- Fibre Channel fabric 38

- files

- main.cf 276

- flarcreate 72

- Flash archive 72

- post-deployment scripts 73

- freezing service groups 195

G

- GAB

- port membership information 285
 - verifying 285

- gabconfig command 285

- a (verifying GAB) 285

- gabtab file

- verifying after installation 278

- gcoconfig command 345

- global cluster

- return of primary role to original primary site 374
 - takeover of primary site role by secondary 374
 - troubleshooting VVR component 377

- global clusters

- configuration 125
 - using VVR 344
 - configuring VCS service groups 344
 - illustration of dependencies 361
 - migration of primary to secondary site 373
 - replicating database volumes 361
 - requirements 338
 - return migration to original primary site 373

- groups

- log owner 361
 - RVG 361

H

- hastatus -summary command 286
- hasys -display command 287
- high availability issues 489
 - low memory 489
 - network partition 488
- hubs
 - independent 293

I

- I/O daemon (vxiod)
 - starting 136
- I/O fencing
 - checking disks 147
 - setting up 152
 - shared storage 147
- I/O fencing requirements
 - non-SCSI-3 45
- Installing
 - SFCFSHA with the Web-based installer 64
- installing
 - JumpStart 67
 - post 126
- intelligent resource monitoring
 - disabling manually 464
 - enabling manually 464

J

- jeopardy 488–489
- JumpStart
 - installing 67

L

- language packages
 - removal 395
- license keys
 - adding with vxlicinst 128
 - replacing demo key 128
- licenses
 - information about 127
- links
 - private network 278
- Live Upgrade
 - preparing 242
 - upgrade paths 237
 - upgrading Solaris on alternate boot disk 246
- LLT
 - interconnects 32

LLT (*continued*)

- verifying 282
- lltconfig command 278
- llthosts file
 - verifying after installation 278
- lltstat command 282
- llttab file
 - verifying after installation 278
- localized environment settings for using VVR
 - settings for using VVR in a localized environment 192
- log files 491

M

- main.cf file 276
- major and minor numbers
 - checking 503, 506
 - shared devices 502
- manual pages
 - potential problems 488
 - troubleshooting 488
- media speed 32
 - optimizing 32
- membership information 285
- mount command
 - potential problems 487
- mounting
 - software disc 55

N

- network partition 488
- NFS services
 - shared storage 502
- nodes
 - adding application nodes
 - configuring GAB 300, 310
 - configuring LLT 300, 310
 - configuring VXFEN 300, 310
 - starting Volume Manager 299, 310
 - adding new nodes
 - configuring GAB 302
 - configuring LLT 302
 - configuring VXFEN 302
 - preparing application nodes
 - configuring CVM 307, 318
 - removing nodes
 - GAB configuration 330
 - LLT configuration 329

- nodes *(continued)*
 - removing nodes from VCS configuration files 329
 - removing Oracle 11g nodes
 - modifying VCS configuration 330
 - tasks for removing nodes 327
- nodg 137
- non-SCSI-3 fencing
 - manual configuration 174
 - setting up 174
- non-SCSI-3 I/O fencing
 - requirements 45
- non-SCSI3 fencing
 - setting up 157
 - using installsfcfsa 157
- NTP
 - network time protocol daemon 275

O

- optimizing
 - media speed 32

P

- PATH variable
 - VCS commands 282
- persistent reservations
 - SCSI-3 33
- planning to upgrade VVR 190
- port a
 - membership 285
- port h
 - membership 285
- port membership information 285
- post-deployment scripts 73
- preinstallation 190
- preparing
 - Live Upgrade 242
- preparing to upgrade VVR 195
- Prevent Multipathing/Suppress Devices from VxVM's
 - view 137
- primary site
 - creating SRL volume 351
 - setting up 340
 - setting up replication objects 353
 - VCS configuration 363, 367
- problems
 - accessing manual pages 488
 - executing file system commands 488
 - mounting and unmounting file systems 487

Q

- Quick I/O
 - performance on CFS 488

R

- rejoining
 - backup boot disk group 260
- removing
 - the Replicated Data Set 390
- removing a node from a cluster 328
- Replicated Data Set
 - removing the 390
- replication
 - automatic synchronization 359
 - configuring on both sites 344
 - full synchronization with Storage Checkpoint 360
 - modifying VCS configuration 361
 - options 340
 - setting up primary site 340
 - setting up secondary site 342
 - supported hardware 338
 - supported software 338
 - using VVR 351
 - verifying status 360
- resources
 - CVMVolDg 361
 - RVGSharedPri 361
- rolling upgrade 208, 211
 - versions 208
- root disk group 134, 137
- rootdg 137
- rsh 110
 - configuration 33

S

- SAN
 - see Storage Area Network 38
- script-based installer
 - SFCFSHA configuration overview 108
- SCSI driver
 - determining instance numbers 504
- SCSI-3
 - persistent reservations 33
- SCSI-3 persistent reservations
 - verifying 152
- secondary site
 - configuring replication 354
 - creating SRL volume 355

- secondary site *(continued)*
 - setting up 342
 - setting up disk groups 357
 - setting up RLINKs 356
- service groups
 - freezing 195
 - VCS, for global clusters 344
- settings for using VVR in a localized environment
 - localized environment settings for using VVR 192
- SFCFSHA
 - configuring 108
 - coordinator disks 153
 - illustration of global cluster dependencies 361
- SFCFSHA installation
 - verifying
 - cluster operations 282
 - GAB operations 282
 - LLT operations 282
- shared disks, configuring 140
- shared storage
 - Fibre Channel
 - setting up 36
 - NFS services 502
- SMTP email notification 121
- SNMP trap notification 123
- split brain 488
- ssh 110
 - configuration 33
- starting configuration
 - installvcs program 110
 - Veritas product installer 110
- starting vxconfigd configuration daemon 135
- starting vxiod daemon 136
- stopping
 - applications 195
- storage
 - setting up shared fibre 36
- Storage Area Network 38
- suppress devices 137
- Symantec Product Authentication Service 116
- system state attribute value 286

T

- troubleshooting
 - accessing manual pages 488
 - executing file system commands 488
 - mounting and unmounting file systems 487
- tunables file
 - about setting parameters 426

- tunables file *(continued)*
 - parameter definitions 431
 - preparing 430
 - setting for configuration 427
 - setting for installation 427
 - setting for upgrade 427
 - setting parameters 430
 - setting with no other operations 428
 - setting with un-integrated response file 429

U

- unsuccessful upgrade 261
- upgrade paths
 - Live Upgrade 237
- upgrading
 - rolling 208
- upgrading VVR
 - from 4.1 190
 - planning 190
 - preparing 195
- using Live Upgrade 237

V

- VCS
 - command directory path variable 282
 - configuration, for database volume replication 361
 - configuring service groups 344
- VCS configuration
 - for replication 361
- VCS Global cluster option. *See* GCO
- verifying installation
 - kernel component 274
- Veritas Operations Manager 25
- Veritas Volume Replicator. *See* VVR
- Volume Manager
 - Fibre Channel 36
- vradmin
 - delpri 391
 - stoprep 391
- vradmin command 360
- VVR
 - application database replication 359
 - configuring global clusters 344
 - configuring on both sites 344
 - database volume replication
 - configuring VCS 361
 - defining heartbeat cluster objects 346

VVR (*continued*)

- defining remote clusters 346
 - global cluster overview 350
 - primary site
 - creating SRL volume 351
 - setting up replication objects 353
 - replication
 - using automatic synchronization 359
 - using full synchronization with Storage Checkpoint 360
 - replication agents 340
 - secondary site
 - configuring replication 354
 - creating SRL volume 355
 - setting up disk groups 357
 - setting up RLINKs 356
 - setting up primary site 340
 - setting up replication 351
 - setting up secondary site 342
 - types of replication 340
 - VCS configuration
 - application database service group 361
 - CVMoldG resource 361
 - log owner group 361
 - primary site 363
 - RVG group 361
 - RVGSharedPri resource 361
 - secondary site 367
 - verifying replication 360
- VVR 4.1
- planning an upgrade from 190
- vxassist command 352, 355
- vxconfigd configuration daemon
 - starting 135
- vxctl mode command 135
- vxdisksetup command 144
- vxinstall program 136–138
- vxinstall program, running 136
- vxiod I/O daemon
 - starting 136
- vxlicinst command 128
- vxlicrep command 127
- vxvol command 352

W

- Web-based installer 64