

Veritas Storage Foundation™ Cluster File System High Availability Installation Guide

HP-UX

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 © 2015 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	21
Chapter 1 About Storage Foundation Cluster File System High Availability	22
About Veritas Storage Foundation Cluster File System High Availability	22
About I/O fencing	23
About Veritas Operations Manager	24
About Veritas Operations Manager	24
Chapter 2 Before you install	25
About planning for SFCFSHA installation	25
About installation and configuration methods	26
Assessing the system for installation readiness	27
Symantec Operations Readiness Tools	28
Prechecking your systems using the Veritas installer	28
Downloading the Veritas Storage Foundation Cluster File System High Availability software	29
Setting environment variables	30
Optimizing LLT media speed settings on private NICs	31
Guidelines for setting the media speed of the LLT interconnects	31
Creating the /opt directory	31
About using ssh or remsh with the Veritas installer	32
Setting up shared storage	32
Setting up shared storage: SCSI	33
Checking and changing SCSI Initiator IDs	33
Setting up shared storage: Fibre Channel	35
Cluster environment requirements	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	40
Chapter 3	System requirements	41
	Release notes	41
	Hardware compatibility list (HCL)	42
	Supported operating systems	42
	I/O fencing requirements	42
	Coordinator disk requirements for I/O fencing	42
	CP server requirements	43
	Non-SCSI-3 I/O fencing requirements	46
	Veritas File System requirements	47
	Storage Foundation Cluster File System High Availability memory requirements	47
	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	48
	Discovering product versions and various requirement information	48
	Number of nodes supported	49
Chapter 4	Licensing Veritas products	50
	About Veritas SFHA Solutions product licensing	50
	Setting or changing the Veritas SFHA Solutions product level for keyless licensing	51
	Installing Veritas SFHA Solutions product license keys	53
Section 2	Installation of Storage Foundation Cluster File System High Availability	54
Chapter 5	Installing Storage Foundation Cluster File System High Availability using the script-based installer	55
	Installation preparation overview	55
	About installing Veritas Storage Foundation Cluster File System High Availability on HP-UX	56

	Summary of Veritas Storage Foundation Cluster File System High Availability installation tasks	56
	Mounting the product disc	57
	About the Veritas installer	58
	Installing Storage Foundation Cluster File System High Availability using the product installer	59
Chapter 6	Installing Storage Foundation Cluster File System High Availability using the web-based installer	63
	About the Web-based installer	63
	Before using the Veritas Web-based installer	64
	Starting the Veritas Web-based installer	65
	Obtaining a security exception on Mozilla Firefox	65
	Performing a pre-installation check with the Veritas Web-based installer	66
	Installing SFCFSHA with the Web-based installer	66
Chapter 7	Installing Storage Foundation Cluster File System High Availability using other methods	69
	Installing SFCFSHA using Ignite-UX	69
	Creating the Software Distributor (SD) bundle for SFCFSHA or the operating system and SFCFSHA	69
	Using Ignite-UX to perform a standalone SFCFSHA installation	71
	Using Ignite-UX to install SFCFSHA and the HP-UX operating system	72
Section 3	Configuration of Veritas Storage Foundation Cluster File System High Availability	73
Chapter 8	Preparing to configure SFCFSHA	74
	About configuring SFCFSHA clusters for data integrity	74
	About I/O fencing for Storage Foundation Cluster File System High Availability in virtual machines that do not support SCSI-3 PR	75
	About I/O fencing components	76
	About data disks	76

About coordination points	76
About preferred fencing	77
About I/O fencing configuration files	78
About planning to configure I/O fencing	80
Typical SFCFSHA cluster configuration with server-based I/O fencing	84
Recommended CP server configurations	85
Setting up the CP server	88
Planning your CP server setup	88
Installing the CP server using the installer	89
Configuring the CP server cluster in secure mode	90
Setting up shared storage for the CP server database	91
Configuring the CP server using the configuration utility	92
Configuring the CP server manually	101
Verifying the CP server configuration	102

Chapter 9

Configuring Veritas Storage Foundation Cluster File System High Availability	104
Configuring SFCFSHA using the script-based installer	104
Overview of tasks to configure Storage Foundation Cluster File System High Availability using the script-based installer	104
Starting the software configuration	105
Specifying systems for configuration	106
Configuring the cluster name	107
Configuring private heartbeat links	107
Configuring the virtual IP of the cluster	111
Configuring the cluster in secure mode	113
Configuring a secure cluster node by node	114
Adding VCS users	117
Configuring SMTP email notification	118
Configuring SNMP trap notification	120
Configuring global clusters	122
Completing the SFCFSHA configuration	123
Verifying and updating licenses on the system	124
Configuring Storage Foundation Cluster File System High Availability using the Web-based installer	126
Configuring Veritas Storage Foundation Cluster File System High Availability manually	131
Configuring Veritas Volume Manager	131
Configuring Veritas File System	138
Configuring the SFDB repository database after installation	139

Chapter 10	Configuring SFCFSHA for data integrity	140
	Setting up disk-based I/O fencing using installsfcfsa	140
	Initializing disks as VxVM disks	140
	Configuring disk-based I/O fencing using installsfcfsa	141
	Checking shared disks for I/O fencing	144
	Setting up disk-based I/O fencing manually	148
	Identifying disks to use as coordinator disks	149
	Setting up coordinator disk groups	149
	Creating I/O fencing configuration files	150
	Modifying VCS configuration to use I/O fencing	151
	Verifying I/O fencing configuration	152
	Setting up server-based I/O fencing using installsfcfsa	153
	Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsfcfsa	153
	Configuring Storage Foundation Cluster File System High Availability for data integrity using the Web-based installer	155
	Setting up server-based I/O fencing manually	158
	Preparing the CP servers manually for use by the SFCFSHA cluster	159
	Configuring server-based fencing on the SFCFSHA cluster manually	162
	Configuring CoordPoint agent to monitor coordination points	168
	Verifying server-based I/O fencing configuration	169
	Setting up non-SCSI-3 fencing in virtual environments manually	170
	Sample /etc/vxfenmode file for non-SCSI-3 fencing	172
	Enabling or disabling the preferred fencing policy	175
Section 4	Upgrading Storage Foundation Cluster File System High Availability	178
Chapter 11	Preparing to upgrade Veritas Storage Foundation Cluster File System High Availability	179
	About upgrading	179
	About the different ways that you can upgrade	180
	Supported upgrade paths	180
	Preparing to upgrade	181
	Getting ready for the upgrade	182
	Creating backups	183
	Determining which 5.x release of Veritas File System and Veritas Volume Manager that you have installed	184

	Preparing to upgrade the Veritas software	185
	Pre-upgrade planning for Veritas Volume Replicator	188
	Preparing to upgrade VVR when VCS agents are configured	190
	Upgrading the array support	193
	Upgrading the disk layout versions	194
Chapter 12	Performing a typical Storage Foundation Cluster File System High Availability upgrade using the installer	196
	Performing a full upgrade from SFCFSHA versions on HP-UX 11i v2 to SFCFSHA 6.0	196
	Performing a full upgrade from SFCFSHA 5.x on HP-UX 11iv3 to 6.0 HP-UX 11iv3	199
	Performing a full upgrade from SFCFSHA 5.x on HP-UX 11iv3 to 6.0 on the latest HP-UX 11iv3	200
	Upgrading Veritas Storage Foundation Cluster File System High Availability software from a release prior to 6.0	202
	Upgrading from Storage Foundation or Storage Foundation and High Availability 5.0_11iv3, 5.0.1, 5.0.1 RP1, or 5.0.1 RP2, 5.1SP1, or 5.1SP1RP1	202
	Upgrading from previous versions of Storage Foundation on HP-UX 11i v2	204
	Upgrading from Storage Foundation 3.5 on 11i v1 to SFCFSHA 6.0 on HP-UX 11i v3	205
	Upgrading from VxVM 5.0 on HP-UX 11i v3 to VxVM 6.0 using integrated VxVM 6.0 package for HP-UX 11i v3	206
Chapter 13	Performing a rolling upgrade	207
	Performing a rolling upgrade using the installer	207
	About rolling upgrades	207
	Performing a rolling upgrade using the script-based installer	210
	Performing a rolling upgrade of SFCFSHA using the Web-based installer	211
Chapter 14	Performing a phased upgrade	214
	Performing a phased upgrade from version 5.x on HP-UX 11i v3 to Veritas Storage Foundation Cluster File System High Availability 6.0	214
	Performing phased upgrade of SFCFSHA from versions 4.x or 5.x on HP-UX 11i v2 to 5.1SP1	220

Chapter 15	Upgrading the operating system	225
	Upgrading the HP-UX operating system	225
Chapter 16	Upgrading Veritas Volume Replicator	227
	Upgrading Veritas Volume Replicator	227
	Upgrading VVR without disrupting replication	227
Chapter 17	Migrating from SFHA to SFCFSHA	229
	Migrating from SFHA to SFCFSHA 6.0	229
Section 5	Verification of the installation or the upgrade	232
Chapter 18	Verifying the Storage Foundation Cluster File System High Availability installation	233
	Performing a postcheck on a node	234
	Verifying that the products were installed	234
	Installation log files	234
	Using the installation log file	235
	Using the summary file	235
	About enabling LDAP authentication for clusters that run in secure mode	235
	Enabling LDAP authentication for clusters that run in secure mode	237
	Starting and stopping processes for the Veritas products	243
	Checking Veritas Volume Manager processes	243
	Checking Veritas File System installation	244
	Command installation verification	244
	Verifying agent configuration for Storage Foundation Cluster File System High Availability	244
	Synchronizing time on Cluster File Systems	245
	Configuring VCS for Storage Foundation Cluster File System High Availability	245
	main.cf file	246
	Storage Foundation Cluster File System HA Only	247
	Veritas Cluster Server application failover services	247
	Configuring the cluster UUID when creating a cluster manually	247
	About the cluster UUID	248

	About the LLT and GAB configuration files	248
	Verifying the LLT, GAB, and VCS configuration files	250
	Verifying LLT, GAB, and cluster operation	251
	Verifying LLT	251
	Verifying GAB	253
	Verifying the cluster	255
	Verifying the cluster nodes	255
Section 6	Adding and removing nodes	258
Chapter 19	Adding a node to a cluster	259
	About adding a node to a cluster	259
	Before adding a node to a cluster	260
	Adding a node to a cluster using the product installer	262
	Adding a node using the Web-based installer	265
	Adding the node to a cluster manually	266
	Starting Volume Manager on the new node	277
	Configuring cluster files on the new node	278
	Configuring cluster files on the new node	269
	Starting fencing on the new node	282
	Configuring CVM and CFS on the new node	284
	Starting VCS after adding the new node	283
	Configuring the ClusterService group for the new node	285
	Starting Volume Manager on the new node	277
	Configuring cluster files on the new node	278
	Setting up the node to run in secure mode	279
	Starting fencing on the new node	282
	Starting VCS after adding the new node	283
	Configuring CVM and CFS on the new node	284
	Configuring the ClusterService group for the new node	285
	Configuring server-based fencing on the new node	286
	Adding the new node to the vxfen service group	287
	Updating the Storage Foundation for Databases (SFDB) repository after adding a node	288
Chapter 20	Removing a node from Storage Foundation Cluster File System clusters	289
	About removing a node from a cluster	289
	Removing a node from a cluster	290
	Modifying the VCS configuration files on existing nodes	291
	Removing the node configuration from the CP server	293

	Removing security credentials from the leaving node	294
	Updating the Storage Foundation for Databases (SFDB) repository after removing a node	295
	Sample configuration file for removing a node from the cluster	295
Section 7	Setting up and configuring replicated global cluster	298
Chapter 21	Setting up a replicated global cluster	299
	Replication in the SFCFSHA environment	299
	About setting up a global cluster in an SFCFSHA environment	300
	Configuring an SFCFSHA global cluster at the primary site	302
	Configuring an SFCFSHA cluster at the secondary site	304
	Configuring replication for SFCFSHA clusters at both sites	306
	Modifying the ClusterService group for a global SFCFSHA cluster	306
	Modifying the global SFCFSHA cluster configuration using the wizard	307
	Defining the remote SFCFSHA cluster and heartbeat objects	308
	Configuring the VCS service groups for global SFCFSHA clusters	311
Chapter 22	Configuring a global cluster using VVR	312
	About configuring a global SFCFSHA cluster using VVR for replication	312
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the primary site	313
	Creating the data and SRL volumes on the primary site	313
	Setting up the Replicated Volume Group (RVG) on the primary site	315
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): setting up the secondary site	316
	Creating the data and SRL volumes on the secondary site	317
	Editing the /etc/vx/vras/.rdg files	317
	Setting up IP addresses for RLINKs on each cluster	318
	Setting up the disk group on secondary site for replication	319
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): starting replication	321
	Starting replication using automatic synchronization	321
	Starting replication using full synchronization with Storage Checkpoint	322

	Verifying replication status	322
	Configuring a global SFCFSHA cluster using Veritas Volume Replicator (VVR): configuring cluster resources	323
	About modifying the VCS configuration for replication	323
	Modifying the VCS Configuration on the Primary Site	325
	Modifying the VCS Configuration on the Secondary Site	329
	Managing a global SFCFSHA cluster using Veritas Volume Replicator (VVR)	333
	Migrating the role of primary site to the secondary site	335
	Taking over the primary role by the remote cluster	336
Section 8	Uninstallation of Storage Foundation Cluster File System High Availability	340
Chapter 23	Uninstalling Storage Foundation Cluster File System High Availability	341
	Shutting down cluster operations	341
	Disabling the agents on a system	342
	Removing the Replicated Data Set	343
	Uninstalling SFCFSHA depots using the script-based installer	344
	Uninstalling SFCFSHA with the Veritas Web-based installer	346
	Removing license files (Optional)	347
	Removing the CP server configuration using the removal script	347
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product	349
Section 9	Installation reference	351
Appendix A	Installation scripts	353
	About Veritas Storage Foundation and High Availability Solutions installation scripts	353
	Installation script options	354
	About using the postcheck option	359
Appendix B	Response files	361
	About response files	361
	Installing SFCFSHA using response files	362
	Configuring SFCFSHA using response files	362

	Upgrading SFCFSHA using response files	363
	Uninstalling SFCFSHA using response files	364
	Syntax in the response file	364
	Response file variables to install, upgrade, or uninstall Veritas Storage Foundation Cluster File System High Availability	365
	Response file variables to configure Veritas Storage Foundation Cluster File System High Availability	367
	Sample response file for Veritas Storage Foundation Cluster File System High Availability installation	376
	Sample response file for Veritas Storage Foundation Cluster File System High Availability configuration	377
Appendix C	Tunable files for installation	379
	About setting tunable parameters using the installer or a response file	379
	Setting tunables for an installation, configuration, or upgrade	380
	Setting tunables with no other installer-related operations	381
	Setting tunables with an un-integrated response file	382
	Preparing the tunables file	383
	Setting parameters for the tunables file	383
	Tunables value parameter definitions	384
Appendix D	Configuring I/O fencing using a response file	391
	Configuring I/O fencing using response files	391
	Response file variables to configure disk-based I/O fencing	392
	Sample response file for configuring disk-based I/O fencing	393
	Response file variables to configure server-based I/O fencing	394
	Sample response file for configuring server-based I/O fencing	396
	Response file variables to configure non-SCSI-3 server-based I/O fencing	396
	Sample response file for configuring non-SCSI-3 server-based I/O fencing	397
Appendix E	Configuring the secure shell or the remote shell for communications	399
	About configuring secure shell or remote shell communication modes before installing products	399
	Manually configuring and passwordless ssh	400
	Enabling remsh	404

Appendix F	Storage Foundation Cluster File System High Availability components	405
	Veritas Storage Foundation Cluster File System High Availability installation depots	405
	Veritas Cluster Server installation depots	408
	Veritas Cluster File System installation depots	409
	Veritas Storage Foundation obsolete and reorganized installation depots	409
Appendix G	High availability agent information	412
	About agents	412
	VCS agents included within SFCFSHA	413
	Enabling and disabling intelligent resource monitoring for agents manually	413
	Administering the AMF kernel driver	415
	CVMCluster agent	416
	Entry points for CVMCluster agent	416
	Attribute definition for CVMCluster agent	417
	CVMCluster agent type definition	418
	CVMCluster agent sample configuration	418
	CVMVxconfigd agent	418
	Entry points for CVMVxconfigd agent	419
	Attribute definition for CVMVxconfigd agent	419
	CVMVxconfigd agent type definition	421
	CVMVxconfigd agent sample configuration	421
	CVMVolDg agent	421
	Entry points for CVMVolDg agent	421
	Attribute definition for CVMVolDg agent	422
	CVMVolDg agent type definition	423
	CVMVolDg agent sample configuration	424
	CFSMount agent	424
	Entry points for CFSMount agent	425
	Attribute definition for CFSMount agent	425
	CFSMount agent type definition	427
	CFSMount agent sample configuration	428
	CFSfsckd agent	428
	Entry points for CFSfsckd agent	428
	Attribute definition for CFSfsckd agent	429
	CFSfsckd agent type definition	430
	CFSfsckd agent sample configuration	431

Appendix H	Troubleshooting information	432
	Restarting the installer after a failed connection	432
	What to do if you see a licensing reminder	432
	Storage Foundation Cluster File System High Availability installation	
	issues	433
	Incorrect permissions for root on remote system	433
	Resource temporarily unavailable	434
	Inaccessible system	435
	Storage Foundation Cluster File System High Availability	
	problems	436
	Unmount failures	436
	Mount failures	436
	Command failures	437
	Performance issues	438
	High availability issues	438
Appendix I	Troubleshooting cluster installation	440
	Installer cannot create UUID for the cluster	440
	The vxfentsthdw utility fails when SCSI TEST UNIT READY command	
	fails	441
	Troubleshooting CP server	441
	Troubleshooting issues related to the CP server service	
	group	442
	Checking the connectivity of CP server	442
	Troubleshooting server-based fencing on the SFCFSHA cluster	
	nodes	443
	Issues during fencing startup on SFCFSHA cluster nodes set up	
	for server-based fencing	443
	Issues during online migration of coordination points	444
Appendix J	Sample SFCFSHA cluster setup diagrams for CP	
	server-based I/O fencing	445
	Configuration diagrams for setting up server-based I/O fencing	445
	Two unique client clusters served by 3 CP servers	446
	Client cluster served by highly available CPS and 2 SCSI-3	
	disks	447
	Two node campus cluster served by remote CP server and 2	
	SCSI-3 disks	448
	Multiple client clusters served by highly available CP server and	
	2 SCSI-3 disks	450

Appendix K	Reconciling major/minor numbers for NFS shared disks	452
	Reconciling major/minor numbers for NFS shared disks	452
	Checking major and minor numbers for disk partitions	453
	Checking the major and minor number for VxVM volumes	454
Appendix L	Configuring LLT over UDP using IPv6	457
	Using the UDP layer of IPv6 for LLT	457
	When to use LLT over UDP	457
	Manually configuring LLT over UDP using IPv6	457
	The link command in the /etc/lfttab file	458
	The set-addr command in the /etc/lfttab file	459
	Selecting UDP ports	459
	Sample configuration: direct-attached links	460
	Sample configuration: links crossing IP routers	461
Appendix M	Configuring LLT over UDP using IPv4	464
	Using the UDP layer for LLT	464
	When to use LLT over UDP	464
	Manually configuring LLT over UDP using IPv4	464
	Broadcast address in the /etc/lfttab file	465
	The link command in the /etc/lfttab file	466
	The set-addr command in the /etc/lfttab file	467
	Selecting UDP ports	467
	Configuring the netmask for LLT	468
	Configuring the broadcast address for LLT	469
	Sample configuration: direct-attached links	469
	Sample configuration: links crossing IP routers	471
Appendix N	Compatibility issues when installing Storage Foundation Cluster File System High Availability with other products	473
	Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present	473
	Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present	474
	Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present	474
Index		475

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.

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

About I/O fencing

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

The fencing operation determines the following:

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

This decision prevents possible data corruption. The installer installs the I/O fencing driver, VRTSvxfen, when you install Storage Foundation Cluster File System High Availability. To protect data on shared disks, you must configure I/O fencing after you install and configure Storage Foundation Cluster File System High Availability.

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, Storage Foundation Cluster File System High Availability supports non-SCSI-3 server-based I/O fencing.

See [“About planning to configure I/O fencing”](#) on page 80.

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

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 remsh with the Veritas installer](#)
- [Setting up shared storage](#)
- [Cluster environment requirements](#)
- [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 361.</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 63.</p>
Manual installation and configuration	<p>Manual installation uses the HP-UX commands to install SFCFSHA. To retrieve a list of all depots and patches required for all products in the correct installation order, enter:</p> <pre># installer -allpkgs</pre> <p>Use the HP-UX 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 28.
- 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 28.

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
- Command checks

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 353.

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:

4 GB

See “[Disk space requirements](#)” on page 48.

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.

```
# df -b 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 remsh with the Veritas installer

The installer uses passwordless secure shell (`ssh`) or remote shell (`remsh`) communications among systems. The installer uses the `ssh` or `remsh` 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. Note that for security reasons, the installation program neither stores nor caches these passwords. The `ssh` or `remsh` 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 `remsh` configuration from the systems.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer can configure `ssh` or `remsh` on the target systems. In the following scenarios, you need to set up `ssh` or `remsh` manually:

- When the root broker is outside of the cluster that you plan to configure.
- 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 399.

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 80.


See also the *Storage Foundation Cluster File System High Availability Administrator's Guide* for a description of I/O fencing.

Setting up shared storage: SCSI

Perform the following steps to set up shared storage.

Figure 2-1 shows how to cable systems for shared storage.

Figure 2-1 Cabling the shared storage



To set up shared storage

- 1 Shut down the systems in the cluster.
- 2 Install the required SCSI host bus adapters and set up the external shared SCSI storage devices.
- 3 Cable the external shared storage devices. With cables connected to shared storage between two systems, you must terminate the two ends of the SCSI bus on the systems, as shown in the figure.

For more than two systems, disable SCSI termination on the systems that are not positioned at the ends of the SCSI chain.

Checking and changing SCSI Initiator IDs

The SCSI Initiator IDs for the host bus adapters (HBAs) on each of the systems that access the shared storage must be unique. So, you may have to change the HBA SCSI ID on one or more systems if these IDs are the same. Typically, the host bus adapters (HBAs) for the SCSI devices are shipped with a default SCSI ID of 7. Use the following procedure to check SCSI IDs and change them if necessary.

To check and change SCSI initiator IDs

- 1 For systems with PA-RISC architecture, turn on the power of the first system. During the boot process, the system delays for ten seconds, giving you the opportunity to stop the boot process and enter the boot menu:

To discontinue, press any key within 10 seconds.

Press any key. The boot process discontinues.

Boot terminated.

- 2 When you see the boot Main Menu, display the Information Menu by entering:

Main Menu: enter command or menu > **in**

- 3 From the Information Menu, enter "io" at the prompt for I/O interface information:

Information Menu: Enter command > **io**

The output shows information about the I/O interfaces and resembles:

Path	Bus	Slot	Vendor	Device	Id	Id
Description		(dec)	#	#		
-----		----	---	-----	-----	-----
.						
.						
SCSI bus cntlr		0/3/0/0	24	10	0x1000	0xf
.						

- 4 Return to the Main Menu:

Information Menu: Enter command > **main**

- 5 Go the Service Menu:

Main Menu: enter command or menu > **ser**

- 6 Display the host bus adapter's SCSI ID:

Service Menu: enter command or menu > **scsi**

The output displays information about the SCSI devices:

Path (dec)	Initiator	ID	SCSI Rate	Auto Term
-----	-----		-----	-----
0/3/0/0	7		Fast	Unknown

The output in this example shows the SCSI ID is 7, the preset default for the HBA as shipped.

- If you choose, you can leave the ID set at 7 and return to the Main Menu:

```
Service Menu: enter command or menu > main
```

- You can change the SCSI ID for the HBA. For example, to change the SCSI ID from 7 to 6, you would enter:

```
Service Menu: Enter command > SCSI init 0/3/0/0 6  

FAST
```

- To verify the change, enter "SCSI" at the prompt:

```
Service Menu: Enter command > SCSI
```

Path (dec)	Initiator ID	SCSI Rate	Auto
0/3/0/0	6	Fast	Unknown

7 Return to the Main Menu:

```
Service Menu: enter command or menu > main
```

8 At the Main Menu, enter the command to boot the system. Answer "n" when you are prompted to interact with IPL:

```
Menu: Enter command or menu > boot  

Interact with IPL (Y, N, or Cancel)?> n
```

Booting...

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up Fibre Channel shared storage

- 1 Shut down the cluster systems that must share the devices.
- 2 Install the required Fibre Channel host bus adapters on each system.
- 3 Cable the shared devices.

- 4 Reboot each system.
- 5 Verify that each system can see all shared devices. Use the command:

```
# ioscan -fnC disk
```

Where "disk" is the class of devices to be shared. For example, from a system galaxy type:

```
galaxy# ioscan -fnC disk
Class I H/W Path Driver S/W State H/W Type Description
=====
.
.
disk 4 0/4/0/0.1.16.255.13.4.0 sdisk CLAIMED DEVICE
SEAGATE ST318304 CLAR18
/dev/dsk/c4t4d0 /dev/rdisk/c4t4d0
disk 5 0/4/0/0.1.16.255.13.5.0 sdisk CLAIMED DEVICE
SEAGATE ST318304 CLAR18
/dev/dsk/c4t5d0 /dev/rdisk/c4t5d0
.
.
```

And on another system, nebula, enter:

```
nebula# ioscan -fnC disk
Class I H/W Path Driver S/W State H/W Type Description
=====
.
.
disk 4 0/4/0/0.1.16.255.13.4.0 sdisk CLAIMED DEVICE
SEAGATE ST318304 CLAR18
/dev/dsk/c4t4d0 /dev/rdisk/c4t4d0
disk 5 0/4/0/0.1.16.255.13.5.0 sdisk CLAIMED DEVICE
SEAGATE ST318304 CLAR18
/dev/dsk/c4t5d0 /dev/rdisk/c4t5d0
.
.
```

Cluster environment requirements

If you are configuring a cluster, which is a set of hosts that share a set of disks, set up the cluster environment.

To set up a cluster environment

- 1 If you plan to place the root disk group under VxVM control, decide into which disk group you want to configure it for each node in the cluster. The root disk group, usually aliased as `bootdg`, contains the volumes that are used to boot the system. VxVM sets `bootdg` to the appropriate disk group if it takes control of the root disk. Otherwise `bootdg` is set to `nodg`. To check the name of the disk group, enter the command:

```
# vxdg bootdg
```

- 2 Decide on the layout of shared disk groups. There may be one or more shared disk groups. Determine how many you wish to use.
- 3 If you plan to use Dirty Region Logging (DRL) with VxVM in a cluster, leave a small amount of space on the disk for these logs. The log size is proportional to the volume size and the number of nodes. Refer to the *Veritas Volume Manager Administrator's Guide* for more information on DRL.
- 4 Install the license that supports the clustering feature on every node in the cluster.

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 `remsh` (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 2-2](#) shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

Figure 2-2 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*.

Install the HP-UX 11i 64-bit operating system with the March 2011 HP-UX 11i Version 3.0 or later version of 11iv3 on each node and install a Fibre Channel host bus adapter to allow connection to the Fibre Channel switch.

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 42.
- CP servers
See “[CP server requirements](#)” on page 43.

If you have installed Storage Foundation Cluster File System High Availability 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 46.

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 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

Storage Foundation Cluster File System High Availability 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	<p>To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space:</p> <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
<p>CP server hosted on a VCS single-node cluster or on an SFHA cluster</p>	<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 *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:

- HP-UX Integrity Virtual Machines (IVM) Server 4.2 and 4.3

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:

- Storage Foundation Cluster File System High Availability must be configured with Cluster attribute UseFence set to SCSI3
- All coordination points must be CP servers

Veritas File System requirements

Complete the tasks in this section before installing Veritas File System.

Before installing Veritas File System, perform the following tasks:

- Review the *Veritas Storage Foundation Release Notes*.
- Ensure that the `/opt` directory exists and has write permissions for `root`.
- The Veritas File System does not support OmniStorage. Do not install VxFS without first retrieving any files archived using OmniStorage.
- Install all the latest required HP-UX patches.

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

All nodes in a Cluster File System must have the same operating system version and update level.

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 depots 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 51.

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 53.

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` depot 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](#)
- [About installing Veritas Storage Foundation Cluster File System High Availability on HP-UX](#)
- [Summary of Veritas Storage Foundation Cluster File System High Availability installation tasks](#)
- [Mounting the product disc](#)
- [About the Veritas installer](#)
- [Installing Storage Foundation Cluster File System High Availability using the product installer](#)

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 50.
Download the software, or insert the product DVD.	See “Downloading the Veritas Storage Foundation Cluster File System High Availability software” on page 29. See “Mounting the product disc” on page 57.
Set environment variables.	See “Setting environment variables” on page 30.
Create the <code>/opt</code> directory, if it does not exist.	See “Creating the /opt directory” on page 31.
Configure the secure shell (ssh) on all nodes.	See “About configuring secure shell or remote shell communication modes before installing products” on page 399.
Verify that hardware, software, and operating system requirements are met.	See “Release notes” on page 41.
Check that sufficient disk space is available.	See “Disk space requirements” on page 48.
Use the installer to install the products.	See “About the Veritas installer” on page 58.

About installing Veritas Storage Foundation Cluster File System High Availability on HP-UX

For an initial installation on a new system, you can use one of the installation procedures described in this section. If you have an existing installation of Storage Foundation Cluster File System that you are upgrading, you must perform an upgrade to move to the 6.0 versions of the Veritas products.

Summary of Veritas Storage Foundation Cluster File System High Availability installation tasks

Installation of Veritas Storage Foundation Cluster File System High Availability consists of the following tasks:

- Obtain a license key, if required.

- If the operating system is not at the required OS fusion level, upgrade the operating system to the latest release.
The operating system is bundled with Veritas Volume Manager and Veritas File System. If the Veritas Volume Manager or Veritas File System is in use, follow the steps in the upgrade chapter to upgrade the Storage Foundation and the operating system.
- If patches for the operating system are required, install the patches before upgrading the product.
- Mount the disk.
- Install Veritas Storage Foundation Cluster File System High Availability 6.0. Start the installer and select 'I' for install, or run the appropriate installation script.
- Reboot the system.

```
# /usr/sbin/shutdown -r now
```
- Configure the Veritas software.
Start the installer and select 'C' for configure, or run the appropriate installation script with the `-configure` option.

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 in the appropriate drive on your local system.
- 3 Determine the block device file for the DVD drive:

```
# iocan -fnC disk
```

Make a note of the device file as it applies to your system.

- 4 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom  
# mount /dev/rdisk/c0t0d0 /dvdrom
```

- 5 Verify that the disc is mounted:

```
# mount
```

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 353.

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 354.

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 399.

- 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 `1` to install and press Return.
- 6 From the Installation menu, choose the `1` 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 depots: installs only the basic functionality for the selected product.
 - Recommended depots: installs the full feature set without optional depots.
 - All depots: installs all available depots.

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 remsh. After the installation, the installer cleans up the ssh or remsh as needed.
- 11 After the system checks complete, the installer displays a list of the depots 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 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 `xprtlowid` 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/xprtlowid.conf`.

See [“Before using the Veritas Web-based installer”](#) on page 64.

See [“Starting the Veritas Web-based installer”](#) on page 65.

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 65.
- 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 66.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 65.
- 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 depots. 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 remsh. 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 SFCFSHA using Ignite-UX](#)

Installing SFCFSHA using Ignite-UX

You can install SFCFSHA or the HP-UX operating system and SFCFSHA using Ignite-UX.

The following procedures describe:

- See [“Creating the Software Distributor \(SD\) bundle for SFCFSHA or the operating system and SFCFSHA”](#) on page 69.
- See [“Using Ignite-UX to perform a standalone SFCFSHA installation”](#) on page 71.
- See [“Using Ignite-UX to install SFCFSHA and the HP-UX operating system”](#) on page 72.

Creating the Software Distributor (SD) bundle for SFCFSHA or the operating system and SFCFSHA

You can use the installer to create SD bundles.

You must run the following commands from an Ignite-UX Server. The `-ignite` option cannot run with other installation options.

Note: When you create the SD bundle for SFCFSHA, the Veritas product disc must be mounted on the Ignite-UX Server.

To create an SD bundle using the installer

- 1 Log in to a configured and running Ignite-UX Server and mount the Veritas installation disc.
- 2 From the prompt, run the **installer** command with the **-ignite** option.

```
# installer -ignite
```

- 3 Select the product to create its SD bundle.
- 4 The installer prompts you for the directory name to place the bundle.

```
Enter the file directory to create the VCS bundle:
(/var/opt/ignite/depots)
Checking the free space of file system ..... Done
Enter a name for the bundle which holds all the VCS depots:
(VCS60_bundle)
```

- 5 Accept the default bundle name or give the bundle a new name.
- 6 The installer copies the depots of the selected product from the disc to the Ignite-UX Server and creates the bundle. It then generates configuration files for the bundle.
- 7 The bundle is ready for a standalone installation of the specific product. To quit the installer choose the last option, **None of the above**.

Continue to the next step if you plan to create an SD bundle for both the operating system and SFCFSHA.

- 8 The installer checks the `/var/opt/ignite/data/INDEX` file to determine if the HP-UX operating system configuration files are available on the Ignite-UX Server. If the file is available, the installer prompts you to add the newly created bundle `cfg` into the HP-UX operating system `cfg` clause. You need to add it so that you can choose the bundle during the HP-UX operating system installation.

Answer **y** to add the bundle `cfg` into the HP-UX operating system `cfg` clause.

Using Ignite-UX to perform a standalone SFCFSHA installation

You can use Ignite-UX to install SFCFSHA on a standalone system.

To use Ignite-UX to install SFCFSHA

- 1 Make sure that the following bundles or packages are not on the target systems; remove them if they exist.
 - Operating system bundles:
 - Base-VxTools-50
 - Base-VxVM-50
 - B3929FB
 - Base-VxFS-50
 - Base-VxVM
 - Operating system bundle packages:
 - AVXTOOL
 - AVXVM
 - AONLINEJFS
 - OnlineJFS01
 - AVXFS
- 2 Create the SD bundle. You should be able to install this bundle to HP-UX systems on your network.

See [“Creating the Software Distributor \(SD\) bundle for SFCFSHA or the operating system and SFCFSHA”](#) on page 69.

- 3 On the system where you want to install the Veritas product, run the following command.

```
# swinstall -x autoreboot=true -s \  
ignite_server_ipadd:/var/opt/ignite/depots/  
product_bundle product_bundle
```

Where *ignite_server_ipadd* is the IP address of the Ignite-UX Server and where */var/opt/ignite/depots* is the directory path.

For example:

```
# swinstall -x autoreboot=true -s \  
10.198.92.81:/var/opt/ignite/depots/SFCFSHA60_bundle SFCFSHA60_bundle
```

- 4 After you install the bundle, reboot the system.
- 5 Configure the product. See the configuration chapter of this guide.

Using Ignite-UX to install SFCFSHA and the HP-UX operating system

You can use Ignite-UX to install SFCFSHA and the operating system.

To use Ignite-UX to install SFCFSHA and the operating system

- 1 Create the SD bundle. You should be able to install this bundle to HP-UX systems on your network.

See [“Creating the Software Distributor \(SD\) bundle for SFCFSHA or the operating system and SFCFSHA”](#) on page 69.
- 2 Install the operating system. See the appropriate HP-UX documentation for details.
- 3 If you use the Ignite-UX screen GUI, switch to the **Software** tab on the configuration page of the operating system installation. On the **Software** tab, select and enable the Veritas product bundle that you want to install.
- 4 On the **Software** tab, deselect the following operating system bundles if they are selected:
 - Base-VxTools-50
 - Base-VxVM-50
 - B3929FB
 - Base-VxFS-50
 - Base-VxVM
- 5 After you have installed the bundle, you need to configure the product. See the configuration chapter of this guide.

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 Storage Foundation Cluster File System High Availability 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 80.

About I/O fencing for Storage Foundation Cluster File System High Availability 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.

In virtualized environments that do not support SCSI-3 PR, Storage Foundation Cluster File System High Availability attempts to provide reasonable safety for the data disks. Storage Foundation Cluster File System High Availability 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 153.

See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 170.

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 76.
- Coordination points—Act as a global lock during membership changes
See [“About coordination points”](#) on page 76.

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.

Dynamic Multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. On cluster nodes with HP-UX 11i v3, you must use DMP devices or iSCSI devices

for I/O fencing. The following changes in HP-UX 11i v3 require you to not use raw devices for I/O fencing:

- Provides native multipathing support
- Does not provide access to individual paths through the device file entries
The metanode interface that HP-UX provides does not meet the SCSI-3 PR requirements for the I/O fencing feature. You can configure coordinator disks to use Veritas Volume Manager Dynamic Multi-pathing (DMP) feature. 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 [“Enabling or disabling the preferred fencing policy”](#) on page 175.

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/rc.config.d/vxfenconf	<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 Storage Foundation Cluster File System High Availability 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 <p>You must configure the vxfen module to use DMP devices or iSCSI devices, and set the SCSI-3 disk policy as dmp.</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>An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:</p> <pre data-bbox="319 586 561 670">/dev/vx/rdmp/c1t1d0 /dev/vx/rdmp/c2t1d0 /dev/vx/rdmp/c3t1d0</pre> <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 Storage Foundation Cluster File System High Availability 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 83.

[Figure 8-1](#) illustrates a high-level flowchart to configure I/O fencing for the Storage Foundation Cluster File System High Availability cluster.

Figure 8-1 Workflow to configure I/O fencing

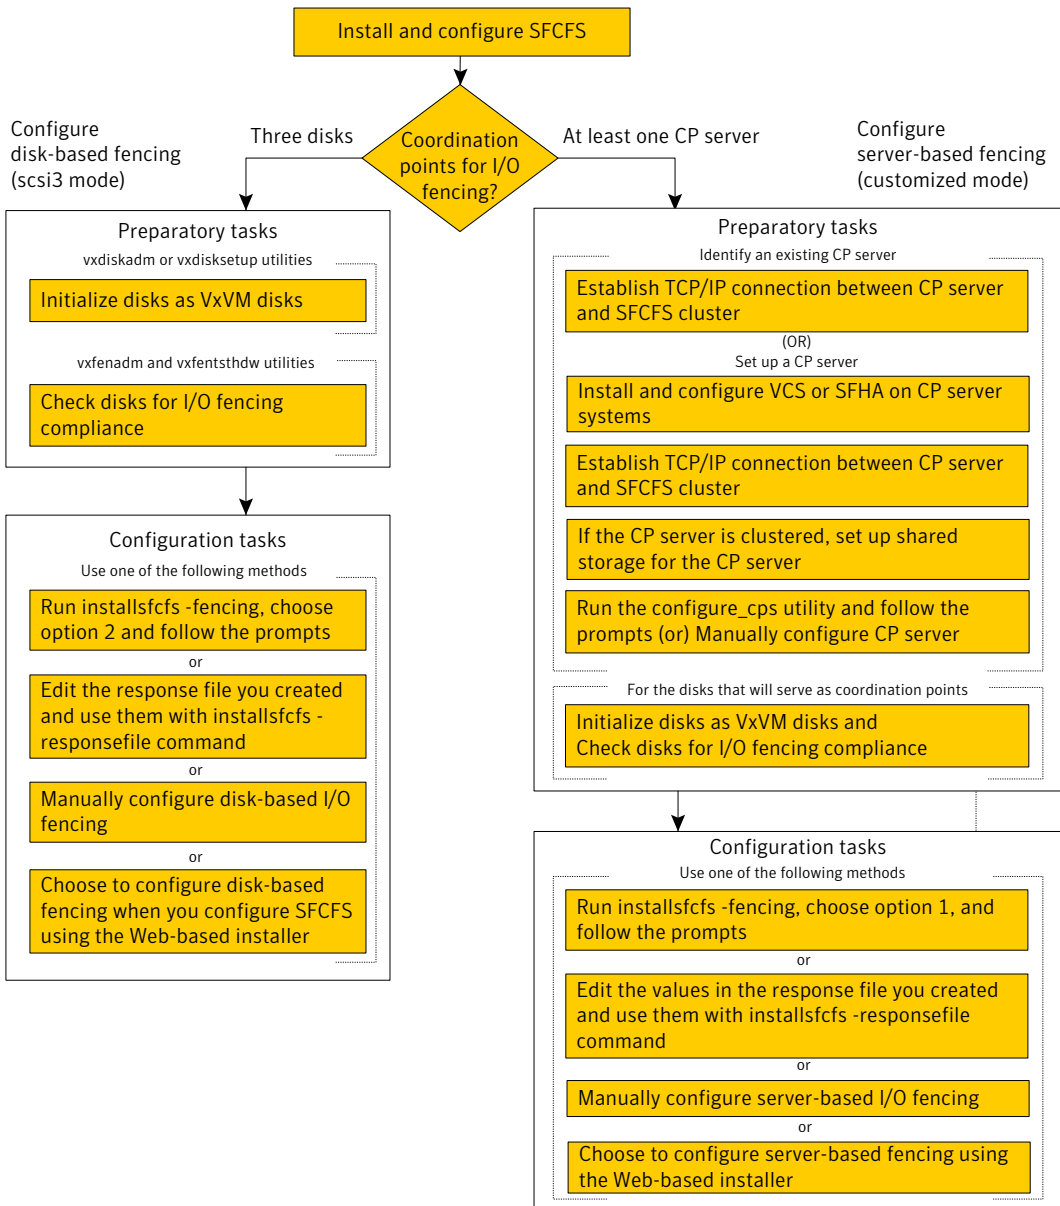
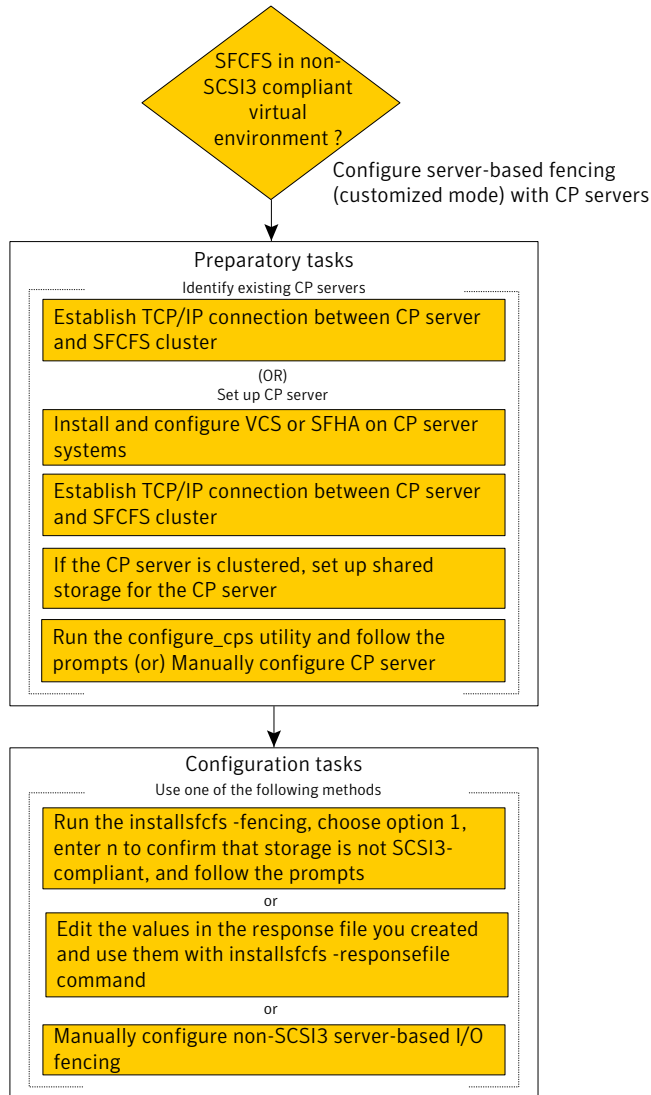


Figure 8-2 Illustrates a high-level flowchart to configure non-SCSI-3 server-based I/O fencing for the Storage Foundation Cluster File System High Availability cluster in virtual environments that do not support SCSI-3 PR.

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 140.
 See [“Setting up server-based I/O fencing using `installsfcfs`”](#) on page 153.
 See [“Setting up non-SCSI-3 server-based I/O fencing in virtual environments using `installsfcfs`”](#) on page 153.
- Using the Web-based installer See [“Configuring Storage Foundation Cluster File System High Availability for data integrity using the Web-based installer”](#) on page 155.
- Using response files See [“Response file variables to configure disk-based I/O fencing”](#) on page 392.
 See [“Response file variables to configure server-based I/O fencing”](#) on page 394.
 See [“Response file variables to configure non-SCSI-3 server-based I/O fencing”](#) on page 396.
 See [“Configuring I/O fencing using response files”](#) on page 391.
- Manually editing configuration files See [“Setting up disk-based I/O fencing manually”](#) on page 148.
 See [“Setting up server-based I/O fencing manually”](#) on page 158.
 See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 170.


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

See the *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 86.
- 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 87.
- 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 87.

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 445.

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 88.
Install the CP server	See “Installing the CP server using the installer” on page 89.
Configure the CP server cluster in secure mode	See “Configuring the CP server cluster in secure mode” on page 90.
Set up shared storage for the CP server database	See “Setting up shared storage for the CP server database” on page 91.
Configure the CP server	See “Configuring the CP server using the configuration utility” on page 92. See “Configuring the CP server manually” on page 101.
Verify the CP server configuration	See “Verifying the CP server configuration” on page 102.

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.

- Decide whether you want to configure server-based fencing for the SFCFSHA cluster (application cluster) with a single CP server as coordination point or with at least three coordination points. Symantec recommends using at least three coordination points.
- 3 Decide whether you want to configure the CP server cluster in secure mode. Symantec recommends configuring the CP server cluster in secure mode to secure the communication between the CP server and its clients (SFCFSHA clusters). It also secures the HAD communication on the CP server cluster.
 - 4 Set up the hardware and network for your CP server.
See “[CP server requirements](#)” on page 43.
 - 5 Have the following information handy for CP server configuration:
 - Name for the CP server
The CP server name should not contain any special characters. CP server name can include alphanumeric characters, underscore, and hyphen.
 - Port number for the CP server
Allocate a TCP/IP port for use by the CP server. Valid port range is between 49152 and 65535. The default port number is 14250.
 - Virtual IP address, network interface, netmask, and networkhosts for the CP server
You can configure multiple virtual IP addresses for the CP server.

Installing the CP server using the installer

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

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

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

CP server setup uses a single system	<p>Install and configure VCS to create a single-node VCS cluster.</p> <p>During installation, make sure to select all depots for installation. The VRTScps depot is installed only if you select to install all depots.</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 92.</p> <p>See “Configuring the CP server manually” on page 101.</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 depots for installation. The VRTScps depot is installed only if you select to install all depots. ■ 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 depot.

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 92.

For CP servers on an SFHA cluster: See [“To configure the CP server on an SFHA cluster”](#) on page 96.

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

- 1 Verify that the VRTScps depot 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: lan0
Enter a valid network interface for virtual IP 10.209.83.87
on mycps1.symantecexample.com: lan0
```

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
lan0 on system mycps1? [y/n] : y
Enter a valid IP address to configure NetworkHosts for
NIC lan0 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 *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` depot is installed on each node.
- 2 Make sure that you have configured passwordless ssh or remsh 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 remsh 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 : lan0
```

- 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 lan0 on system mycps1? [y/n] : y

Enter a valid IP address to configure NetworkHosts for NIC lan0 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 *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:

```
# hstop -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 Storage Foundation Cluster File System High Availability using the Web-based installer](#)
- [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 Storage Foundation Cluster File System High Availability using the script-based installer

[Table 9-1](#) lists the tasks that are involved in configuring Storage Foundation Cluster File System High Availability using the script-based installer.

Table 9-1 Tasks to configure Storage Foundation Cluster File System High Availability using the script-based installer

Task	Reference
Start the software configuration	See “Starting the software configuration” on page 105.
Specify the systems where you want to configure Storage Foundation Cluster File System High Availability	See “Specifying systems for configuration” on page 106.
Configure the basic cluster	See “Configuring the cluster name” on page 107. See “Configuring private heartbeat links” on page 107.
Configure virtual IP address of the cluster (optional)	See “Configuring the virtual IP of the cluster” on page 111.
Configure the cluster in secure mode (optional)	See “Configuring the cluster in secure mode” on page 113.
Add VCS users (required if you did not configure the cluster in secure mode)	See “Adding VCS users” on page 117.
Configure SMTP email notification (optional)	See “Configuring SMTP email notification” on page 118.
Configure SNMP email notification (optional)	See “Configuring SNMP trap notification” on page 120.
Configure global clusters (optional) Note: You must have enabled Global Cluster Option when you installed Storage Foundation Cluster File System High Availability.	See “Configuring global clusters” on page 122.
Complete the software configuration	See “Completing the SFCFSHA configuration” on page 123.

Starting the software configuration

You can configure Storage Foundation Cluster File System High Availability using the Veritas product installer or the `installsfcfsha` command.

Note: If you want to reconfigure Storage Foundation Cluster File System High Availability, before you start the installer you must stop all the resources that are under VCS control using the `hastop` command or the `hagrps -offline` command.

To configure Storage Foundation Cluster File System High Availability 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 Storage Foundation Cluster File System High Availability 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 Storage Foundation Cluster File System High Availability. 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 Storage Foundation Cluster File System High Availability.

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
- Checks whether Storage Foundation Cluster File System High Availability 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 80.

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 464.

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.

You must not enter the network interface card that is used for the public network (typically lan0.)

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

[b,q,?] **lan0**

lan0 has an IP address configured on it. It could be a public NIC on galaxy.

Are you sure you want to use lan0 for the first private heartbeat link? [y,n,q,b,?] (n) **y**

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,?] **lan1**

lan1 has an IP address configured on it. It could be a public NIC on galaxy.

Are you sure you want to use lan1 for the second private heartbeat link? [y,n,q,b,?] (n) **y**

Would you like to configure a third private heartbeat link?

[y,n,q,b,?] (n)

Do you want to configure an additional low priority heartbeat link? [y,n,q,b,?] (n)

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

```

Enter the NIC for the first private heartbeat
link on 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: lan0
Enter the NIC for Virtual IP of the Cluster to use on galaxy:
[b,q,?] (lan0)
```

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 lan0 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)
```

- Enter the NetworkHosts IP addresses that are separated with spaces for checking the connections.

```
Enter the NetworkHosts IP addresses, separated
by spaces: [b,q,?] 192.168.1.17
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

```
NIC: lan0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

```
NetworkHosts: 192.168.1.17
```

```
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
```

- Enter the NetworkHosts IP addresses that are separated with spaces for checking the connections.

```
Enter the NetworkHosts IP addresses, separated
by spaces: [b,q,?] 2001:db8::1 2001:db8::2
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

```
NIC: lan0
IP: 2001:454e:205a:110:203:baff:feee:10
Prefix: 64
```

```
NetworkHosts: 2001:db8::1 2001:db8::2
```

```
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 remsh, you cannot use the `-security` option to enable secure mode for your cluster. Instead, you can use the `-securityonemode` 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 114.
Configure security on the remaining nodes	See “Configuring the remaining nodes” on page 115.
Complete the manual configuration steps	See “Completing the secure cluster configuration” on page 116.

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 120.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: lan0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (lan0)
Is lan0 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: lan0
```

```
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 122.

3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: lan0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (lan0)
Is lan0 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: lan0

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 *Storage Foundation Cluster File System High Availability Installation Guide* for instructions to set up Storage Foundation Cluster File System High Availability 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, value for the netmask, and value for the network hosts.

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: lan0
IP: 192.168.1.16
Netmask: 255.255.240.0

NetworkHosts: 192.168.1.17
```

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

For IPv6 Global Cluster Option configuration verification:

```
NIC: lan0
IP: 2001:454e:205a:110:203:baff:feee:10
Prefix: 64

NetworkHosts: 2001:db8::1 2001:db8::2
```

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

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 Storage Foundation Cluster File System High Availability 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 362.

Verifying and updating licenses on the system

After you install Storage Foundation Cluster File System High Availability, 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 124.

See [“Updating product licenses using vxlicinst”](#) on page 125.

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 Storage Foundation Cluster File System High Availability license key on each node. If you have Storage Foundation Cluster File System High Availability already installed and configured and you use a demo license, you can replace the demo license.

See [“Replacing a Storage Foundation Cluster File System High Availability demo license with a permanent license”](#) on page 125.

To update product licenses

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

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

Replacing a Storage Foundation Cluster File System High Availability demo license with a permanent license

When a Storage Foundation Cluster File System High Availability 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 Storage Foundation Cluster File System High Availability on all nodes in the cluster:

```
# hastop -all -force
```

This command does not shut down any running applications.

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

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

- 4 Make sure demo licenses are replaced on all cluster nodes before starting Storage Foundation Cluster File System High Availability.

```
# vxlicrep
```

- 5 Start Storage Foundation Cluster File System High Availability on each node:

```
# hstart
```

Configuring Storage Foundation Cluster File System High Availability using the Web-based installer

Before you begin to configure Storage Foundation Cluster File System High Availability 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 Storage Foundation Cluster File System High Availability on a cluster

- 1 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 65.

- 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 Storage Foundation Cluster File System High Availability, 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 Storage Foundation Cluster File System High Availability for data integrity using the Web-based installer”](#) on page 155.

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.
Enter the value for the networkhosts.
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 *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. Enter the value for the networkhosts. 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 `installsfsha`.

Click **Next**.

- 8 On the NetworkHosts Configuration page, enter the details of the network hosts and click **Next**.
- 9 On the Stop Processes page, click **Next** after the installer stops all the processes successfully.

- 10 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 12. Go to step 11 to configure fencing.

- 11 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 Storage Foundation Cluster File System High Availability for data integrity using the Web-based installer”](#) on page 155.

- 12 Click **Next** to complete the process of configuring Storage Foundation Cluster File System High Availability.

On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.

- 13 Select the checkbox to specify whether you want to send your installation information to Symantec.

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

Configuring 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*.

Configuring Veritas Volume Manager with the `installvm` script

If you deferred configuring VxVM during installation, you can configure it by running the `installvm` script with the `-configure` option.

To configure VxVM using the `installvm` script

- 1 Enter the following commands.

```
# cd /dvdrom/volume_manager
# ./installvm -configure
```

- 2 The script runs an initial system check, and will tell you that you cannot configure already configured features such as enclosure-based naming and default disk groups.

- 3 Decide whether you want to set up the enclosure-based naming scheme:

```
Do you want to set up the enclosure-based naming scheme? [y, n, q]
```

- 4 You are then asked if you want to set up a default disk group for each system:

```
Do you want to set up a default disk group for each system? [y, n, q]
```

- 5 If you have a VVR license installed, the next phase concerns configuration of VVR:

```
Do you want to change any of the VVR ports ... [y, n, q]
```

- 6 You are now asked questions regarding the frequency of VVR statistics collection.

- 7 The next phase of the configuration procedure consists of setting up a centrally managed host:

```
Enable Centralized Management? [y,n,q]
```

- 8 If you selected centralized management, you will be asked a series of questions relating to hostnames.

After configuration is complete, the following message displays:

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

- 9 After the installation and configuration of VxVM is complete, you can use the `vxdiskadm` command and the VEA GUI to create disk groups, and to populate these with disks.

See the *Veritas Volume Manager Administrator's Guide* and the VEA online help for details.

Converting to a VxVM root disk

It is possible to select VxVM as a choice for your root disk when performing a new installation using Ignite-UX. Alternatively, you can use the following procedure to achieve VxVM rootability by cloning your LVM root disk using the `vxcp_lvmroot` command.

To convert to a VxVM root disk

- 1 Select the disk to be used as your new VxVM root disk. It is recommended that this disk is internal to the main computer cabinet. If this is currently an LVM disk, then it must be removed from LVM control as follows:
 - Use the `lvremove` command to remove any LVM volumes that are using the disk.
 - Use the `vgreduce` command to remove the disk from any LVM volume groups to which it belongs.
 - Use the `pvremove` command to erase the LVM disk headers

If the disk to be removed is the last disk in the volume group, use the `vgremove` command to remove the volume group, and then use `pvremove` to erase the LVM disk headers.

If the disk is not currently in use by any volume or volume group, but has been initialized by `pvcreate`, you must still use the `pvremove` command to remove LVM disk headers.

If you want to mirror the root disk across multiple disks, make sure that all the disks are free from LVM control.

- 2 While booted on the newly upgraded LVM root disk, invoke the `vxcp_lvmroot` command to clone the LVM root disk to the disk(s) you have designated to be the new VxVM root disks. In the following example, `c1t0d0` is used for the target VxVM root disk:

```
# /etc/vx/bin/vxcp_lvmroot -v c1t0d0
```

To additionally create a mirror of the root disk on `c2t0d0`:

```
# /etc/vx/bin/vxcp_lvmroot -v -m c2t0d0 c1t0d0
```

Use of the `-v` (verbose) option is highly recommended. The cloning of the root disk is a lengthy operation, and this option gives a time-stamped progress indication as each volume is copied, and other major events.

- 3 Use the `setboot` (1M) command to save the hardware path of the new VxVM root disk in the system NVRAM. The disk hardware paths can be found using this command:

```
# ioscan -kfnC disk
```

- 4 Reboot from the new VxVM root disk. If you created a mirrored root disk, then there is nothing more to do. The LVM root disk safely co-exists with your VxVM root disk, and provides a backup boot target.
- 5 If desired, you can convert the original LVM root disk into a mirror of your VxVM root disk by using the following commands:

```
# /etc/vx/bin/vxdestroy_lvmroot -v c2t0d0
```

```
# /etc/vx/bin/vxrootmir -v c2t0d0
```

Once this operation is complete, the system is running on a completely mirrored VxVM root disk.

- 6 If later required, you can use the `vxres_lvmroot` command to restore the LVM root disk.

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 `vxctl(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.

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

Use this procedure if you are upgrading from VxVM 3.x to VxVM 6.0 (or Storage Foundation 3.x to a Storage Foundation product at the 6.0 level) and you want to convert existing 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.

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/fstab
```

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

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

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: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 Storage Foundation Cluster File System High Availability 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:

```
# ioscan -nfc disk
# insf -e
```

Warning: The HP-UX man page for the `insf` command instructs you to run the command in single-user mode only. You can run `insf -e` in multiuser mode only when no other user accesses any of the device files. This command can change the mode, owner, or group of an existing special (device) file, or unlink and recreate a file. The special files that are currently open may be left in an indeterminate state.

- 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 Storage Foundation Cluster File System High Availability to complete I/O fencing configuration. Make sure to unfreeze any frozen VCS service groups in the cluster for the installer to successfully stop Storage Foundation Cluster File System High Availability.

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 Storage Foundation Cluster File System High Availability 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 144.
- 7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided. 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
- 8 Verify and confirm the I/O fencing configuration information that the installer summarizes.
- 9 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 Storage Foundation Cluster File System High Availability is cleanly configured to use the I/O fencing feature.
- 10 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.
- 11 Configure the Coordination Point agent to monitor the coordinator disks. See [“Configuring CoordPoint agent to monitor coordination points”](#) on page 168.

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 `vxfsentsthdw` utility. The two nodes must have `ssh` (default) or `remsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfsenadm` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

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

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

See the *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 144.
- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 146.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfsentsthdw utility”](#) on page 146.

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.sl	3PARdata	VV
libvxCLARiiON.sl	DGC	All
libvxFJTSYe6k.sl	FUJITSU	E6000
libvxFJTSYe8k.sl	FUJITSU	All
libvxautoraid.sl	HP	C3586A, C5447A, A5257A
libvxcompellent.sl	COMPELNT	Compellent Vol
libvxcopan.sl	COPANSYS	8814, 8818
libvxdds2a.sl	DDN	S2A 9550, S2A 9900, S2A 9700
libvxdothill.sl	DotHill	R/Evo 2730-2R, R/Evo 2530-2R, R/Evo 2330-2R, R/Evo 2130-2RX, R/Evo 2130-2J, R/Evo 5730-2R
libvxemc.sl	EMC	SYMMETRIX
libvxelogic.sl	EQLOGIC	100E-00
libvxfc60.sl	HP	A5277A
libvxfje3k4ka.sl	FUJITSU	E3000, E400A
libvxfjtsye2k.sl	FUJITSU	E2000, ETERNUS_DXL

- 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 `vxfcntlsthdw` utility, you must verify that the systems see the same disk.

To verify that the nodes have access to the same disk

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed 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/c1t1d0` path on node A and the `/dev/rdisk/c2t1d0` path on node B.

From node A, enter:

```
# vxfenadm -i /dev/rdisk/c1t1d0
```

```
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/c2t1d0` path.

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

```
# vxfenadm -i /dev/rdisk/c3t1d2
```

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

Testing the disks using vxfcntlsthdw utility

This procedure uses the `/dev/rdisk/c1t1d0` disk in the steps.

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

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

```
The disk /dev/rdisk/clt1d0 is ready to be configured for I/O Fencing on
node galaxy
```

To test the disks using `vxfcntlsthdw` 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 `remsh` 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:

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/c1t1d0 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 140.
Identifying disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 149.
Checking shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 144.
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 149.
Creating I/O fencing configuration files	See “Creating I/O fencing configuration files” on page 150.
Modifying Storage Foundation Cluster File System High Availability configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 151.
Configuring CoordPoint agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 168.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 152.

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 140.

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 144.

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 `c1t1d0`, `c2t1d0`, and `c3t1d0`.

To create the `vxencoorddg` disk group

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

```
# vxdg init vxencoorddg c1t1d0 c2t1d0 c3t1d0
```

- 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 Update the `/etc/vxfenmode` file to specify to use the SCSI-3 dmp disk policy. On all cluster nodes, type:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /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/rc.config.d/vxfenconf
```

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.

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

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

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

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

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

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`.

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

- Start VCS.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

- 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 153.

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 Storage Foundation Cluster File System High Availability 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 Storage Foundation Cluster File System High Availability on each node, starts I/O fencing on each node, updates the VCS configuration file `main.cf`, and restarts Storage Foundation Cluster File System High Availability 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 Storage Foundation Cluster File System High Availability for data integrity using the Web-based installer

After you configure Storage Foundation Cluster File System High Availability, you must configure the cluster for data integrity. Review the configuration requirements.

See [“Configuring Storage Foundation Cluster File System High Availability using the Web-based installer”](#) on page 126.

See [“About planning to configure I/O fencing”](#) on page 80.

To configure Storage Foundation Cluster File System High Availability for data integrity

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 65.
- 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, the default fencing disk policy for the disk group is dmp. 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** The default fencing disk policy for the disk group is dmp.

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.
 - The default fencing disk policy for the disk group is dmp.
- 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 Storage Foundation Cluster File System High Availability (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 Storage Foundation Cluster File System High Availability cluster	See “Preparing the CP servers manually for use by the SFCFSHA cluster” on page 159.
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 162.
Modifying Storage Foundation Cluster File System High Availability configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 151.
Configuring Coordination Point agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 168.

Table 10-2 Tasks to set up server-based I/O fencing manually (*continued*)

Task	Reference
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 169.

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 *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 (vxencoordg). You must specify this disk group in the `/etc/vxfenmode` file.

See [“Setting up coordinator disk groups”](#) on page 149.

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/rc.config.d/vxfenconf
```

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 163.

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

For example:

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

- 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=vxfencoorddg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
```

[Table 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. <code>cps<number>=[virtual_ip_address/virtual_host_name]:port</code> 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: <code>cps1=[192.168.0.23], [192.168.0.24]:58888, [mycps1.company.com]</code> 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 *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 158.

- 2 Make sure that the Storage Foundation Cluster File System High Availability 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/vxenvron` file as follows:

```
data_disk_fencing=off
```

- 5 Enter the following command to change the `vxfen_min_delay` parameter value:

```
# /usr/sbin/kctune vxfen_vxfnd_tmt=25
```

- 6 On each node, edit the `/etc/vxfenmode` file as follows:

```
loser_exit_delay=55
vxfen_script_timeout=25
```

Refer to the sample `/etc/vxfenmode` file.

- 7 On each node, set the value of the LLT `sendhbcap` timer parameter value as follows:

- Run the following command:

```
lltconfig -T sendhbcap:3000
```

- Add the following line to the `/etc/llttab` file so that the changes remain persistent after any reboot:

```
set-timer senhbcap:3000
```

- 8 On any one node, edit the VCS configuration file as follows:

- Make the VCS configuration file writable:

```
# haconf -makerw
```

- For each resource of the type `DiskGroup`, set the value of the `MonitorReservation` attribute to 0 and the value of the `Reservation` attribute to `NONE`.

```
# hares -modify <dg_resource> MonitorReservation 0
```

```
# hares -modify <dg_resource> Reservation "NONE"
```

- Run the following command to verify the value:

```
# hares -list Type=DiskGroup MonitorReservation!=0
```

```
# hares -list Type=DiskGroup Reservation!="NONE"
```

The command should not list any resources.

- Modify the default value of the `Reservation` attribute at type-level.

```
# haattr -default DiskGroup Reservation "NONE"
```

- Make the VCS configuration file read-only

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

- 9 Make sure that the UseFence attribute in the VCS configuration file main.cf is set to SCSI3.
- 10 To make these VxFEN changes take effect, stop and restart VxFEN and the dependent modules

- On each node, run the following command to stop VCS:

```
/sbin/init.d/vcs stop
```

- After VCS takes all services offline, run the following command to stop VxFEN:

```
/sbin/init.d/vxfen stop
```

- On each node, run the following commands to restart VxFEN and VCS:

```
# /sbin/init.d/vxfen start
# /sbin/init.d/vcs start
```

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 vxfsend process wait for a customized fencing
# script to complete. Only used with vxfsen_mode=customized
vxfsen_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=vxfencoorddg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
cps1=[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 77.

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:

```
# hagr -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. Migrating from SFHA to SFCFSHA](#)

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](#)
- [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 180.

Review the supported upgrade paths that are available for the different methods of upgrading.

See [“Supported upgrade paths”](#) on page 180.

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 HP-UX.

Table 11-2 HP-UX upgrades using the script- or Web-based installer

Veritas software versions	HP-UX 11.11	HP-UX 11.23	HP-UX 11.31
3.5	Upgrade the operating system to 11.23, then upgrade it to 11.31. Use the installer to upgrade the Veritas product to 6.0.SFCFS requires additional manual changes."	N/A	N/A
3.5_11iv2	N/A	Upgrade to 4.1. Upgrade the operating system to 11.31. Use the installer to upgrade to 6.0.	N/A
4.1 4.1 MP1 4.1 MP2 5.0 5.0 MP1 5.0 MP2 5.0 MP2 RPx	N/A	Upgrade the operating system to 11.31. Use the installer to upgrade to 6.0.	N/A
5.0_11iv3 5.0.1 5.0.1 RPx 5.1 SP1 5.1 SP1 RPx	N/A	N/A	Use the installer to upgrade to 6.0.

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/>
- 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 depots, 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 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.

- 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/fstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/fstab` 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.
- Upgrade arrays (if required).
- If replication using VVR is configured, Symantec recommends that the disk group version is at least 110 prior to upgrading.

```
# vxdg 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.
- 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 failback point.

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.
- 3 Copy the `fstab` file to `fstab.orig`:

```
# cp /etc/fstab /etc/fstab.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 which 5.x release of Veritas File System and Veritas Volume Manager that you have installed

If you are upgrading to this release and have a previously-installed release of Veritas File System (VxFS) and Veritas Volume Manager (VxVM), you must determine which release you have installed. Determining which 5.x release that you have installed can be difficult due to the binary path names being the same for both releases. Use the following procedures to determine which release you have installed.

See [“Discovering product versions and various requirement information”](#) on page 48.

To determine which release of VxFS that you have installed

- ◆ To determine which release of VxFS that you have installed, enter the following command:

```
# swlist -l product VRTSvxfs
```

If you have the 5.0 release installed, the command output includes the following information:

```
VRTSvxfs          5.0.31.0          VERITAS File System
```

If you have the 5.0.1 release installed, the command output includes the following information:

```
VRTSvxfs          5.0.31.5          VERITAS File System
```

If you have the 5.1 SP1 release installed, the command output includes the following information:

```
VRTSvxfs          5.1.100.000       VERITAS File System
```


To determine which release of VxVM that you have installed

- ◆ To determine which release of VxVM that you have installed, enter the following command:

```
# swlist -l product VRTSvxvm
```

If you have the 5.0 release installed, the command output includes the following information:

```
VRTSvxvm          5.0.31.1          Veritas Volume Manager by Symantec
```

If you have the 5.0.1 release installed, the command output includes the following information:

```
VRTSvxvm          5.0.31.5          Veritas Volume Manager by Symantec
```

If you have the 5.1 SP1 release installed, the command output includes the following information:

```
VRTSvxvm          5.1.100.000       Veritas Volume Manager by Symantec
```

Preparing to upgrade the Veritas software

Ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in files such as `/etc/fstab`. You should also run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands, and record the output from these. You may need this information to reconfigure your system after the upgrade.

If you are upgrading an HA cluster, follow the guidelines given in the *Veritas Cluster Server (VCS) Installation Guide* for information on preserving your VCS configuration across the upgrade procedure. In particular, you should take care to make backups of configuration files, such as `main.cf` and `types.cf`, in the `/etc/VRTSvcs/conf/config` directory. Additional configuration files, such as `OracleTypes.cf`, may also be present in this directory if you have installed any VCS agents. You should also back up these files.

To prepare for the Veritas software upgrade

- 1 Log in as superuser.
- 2 Perform any necessary preinstallation checks and configuration.
 See [“About planning for SFCFSHA installation”](#) on page 25.
- 3 Use the `vxlicrep` command to make a record of the currently installed Veritas licenses. Print the output or save it on a different system.

- 4 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes.
- 5 If you are upgrading a high availability (HA) product, take all service groups offline.

List all service groups:

```
# /opt/VRTSvcs/bin/hagrp -list
```

For each service group listed, take it offline:

```
# /opt/VRTSvcs/bin/hagrp -offline service_group -sys system_name
```

- 6 Use the following command to check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -F vxfs
```

- 7 Unmount all Storage Checkpoints and non-system VxFS file systems:

```
# umount /checkpoint_name
```

```
# umount /filesystem
```

- 8 Verify that all file systems have been cleanly unmounted:

```
# echo "8192B.p S" | fsdb -F vxfs filesystem | grep clean
flags 0 mod 0 clean clean_value
```

A `clean_value` value of `0x5a` indicates the file system is clean, `0x3c` indicates the file system is dirty, and `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.

- 9 (Optional) If a file system is not clean, enter the following commands for that file system:

```
# fsck -F vxfs filesystem
# mount -F vxfs filesystem mountpoint
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.

There may be a pending large fileset clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```

An extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system
file system still in progress.
```

- 10 (Optional) If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
- 11 (Optional) Repeat step 8 to verify that the unclean file system is now clean.
- 12 Stop all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

- 13 Comment out the non-system local VxFS mount points from the `/etc/fstab`. Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to uncomment these entries in the `/etc/fstab` file on the upgraded system.
- 14 If Veritas Volume Replicator (VVR) is configured, do the following steps in the order shown:
- Verify that all of the Primary RLINKs are up to date:


```
# vxrlink -g diskgroup status rlink_name
```
 - Detach the RLINKs.

- Disassociate the SRL.

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 11-3](#), 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 11-3 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

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 192.

- 11 For shared disk groups, run the following command on any node in the CVM cluster:

```
# vxctl -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.

Upgrading the array support

The Storage Foundation 6.0 release includes all array support in a single depot, VRTSaslapm. The array support depot includes the array support previously included in the VRTSvxvm depot. The array support depot 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 42.

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 depot 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` depot.

For more information about array support, see the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

Upgrading the disk layout versions

SFCFS 6.0 supports disk layouts Versions 5, 7, or later for locally mounted file systems and disk layouts Versions 7 or later for cluster mounted file systems. If you have cluster mounted file systems with disk layout versions lower than 7, then after upgrading to SFCFS 6.0, perform the following additional steps to prepare the file system for being mounted on all nodes of the cluster:

To upgrade the disk layout versions

- 1 Select one of the nodes of the cluster and `mount` the file system locally on this node. For example, mount it without the `-o cluster` option. Enter,

```
# mount -F vxfs block_device_path /mnt1
```

- 2 Current disk layout version on a file system can be found using

```
# fstyp -v char_device_path | grep version | \
  awk '{print $2}'
```

- 3 On the node selected in 1, incrementally upgrade the disk layout of this file system to disk layout Version 7 or later. For example, if you had a cluster mounted file system of disk layout Version 4 while running with SFCFS 3.5 on HP-UX 11i Version 1, after upgrading to SFCFS 6.0, you would need to upgrade the disk layout to Version 7 or later incrementally as follows:

```
# vxupgrade -n 5 /mnt1
# vxupgrade -n 6 /mnt1
# vxupgrade -n 7 /mnt1
```

- 4 On the node selected in 1, after the disk layout has been successfully upgraded, unmount the file system.

```
# umount /mnt1
```

- 5 This file system can be mounted on all nodes of the cluster using `cfsmount`.

Performing a typical Storage Foundation Cluster File System High Availability upgrade using the installer

This chapter includes the following topics:

- [Performing a full upgrade from SFCFSHA versions on HP-UX 11i v2 to SFCFSHA 6.0](#)
- [Performing a full upgrade from SFCFSHA 5.x on HP-UX 11iv3 to 6.0 HP-UX 11iv3](#)
- [Performing a full upgrade from SFCFSHA 5.x on HP-UX 11iv3 to 6.0 on the latest HP-UX 11iv3](#)
- [Upgrading Veritas Storage Foundation Cluster File System High Availability software from a release prior to 6.0](#)

Performing a full upgrade from SFCFSHA versions on HP-UX 11i v2 to SFCFSHA 6.0

Use these steps to perform a full upgrade from SFCFSHA 4.x or 5.x on HP-UX 11i v2 to SFCFSHA 6.0:

To upgrade from SFCFSHA versions on HP-UX 11i v2:

- 1 Log in as superuser to one of the nodes, *system01* for example, in the cluster.
- 2 Create a backup of the existing cluster configuration. Back up the *main.cf* and *types.cf* on all cluster nodes:

```
# cp /etc/VRTSvcs/conf/config/main.cf \  
/etc/VRTSvcs/conf/config/main.cf.save  
# cp /etc/VRTSvcs/conf/config/types.cf \  
/etc/VRTSvcs/conf/config/types.cf.save
```

- 3 If you created local VxFS mount points on VxVM volumes, added them to the */etc/fstab* file, and comment out the mount point entries in the */etc/fstab* file.
- 4 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.
If the applications are under VCS control:

```
# hagr -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 5 Take offline all VCS groups that contain CFSMount and CVMVolDg.

```
# hagr -offline group -sys system01  
# hagr -offline group -sys system02
```

- 6 Unmount all the VxFS file system which is not under VCS control.

```
# umount /mount_point
```

- 7 Make sure that no processes are running which make use of mounted shared file system or shared volumes.

```
# fuser -cu mount-point
```

- 8 Stop all VCS service groups.

To view the current state of the service groups:

```
# hagr -state
```

To stop each group:

```
# hagr -offline servicegroup -sys node_name
```

- 9 Freeze the VCS service groups. Run the following commands:

```
# haconf -makerw
# hagrps -freeze servicegroup -persistent
# haconf -dump -makero
```

- 10 Stop VCS on all nodes:

```
# hastop -all -force
```

- 11 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.

- 12 On each node, edit the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cat /etc/vxfenmode
vxfen_mode=disabled
```

- 13 On each node, change `LLT_START=0` in the file `/etc/rc.config.d/lltconf`.

- 14 On each node, remove the following device files:

```
# rm -f /dev/llt
# rm -f /dev/gab*
# rm -f /dev/vxfen
```

- 15 Upgrade the operating system from HP-UX 11i v2 to HP-UX 11i v3.

To upgrade from Veritas 5.0 releases on 11i v2, select the Base-VxFS-50, Base-VxVM-50 and Base-VxTools-50 bundles while using `update-ux (1M)`.

- 16 If any patches to the HP-UX 11i v3 are required, install all the prerequisite patches on all nodes before upgrading the Veritas products.

- 17 Install SFCFSHA 6.0.

```
# ./installer
```

From the installation menu, choose the G option for install and enter the number for VERITAS Storage Foundation Cluster File System.

- 18 Uncomment the VxFS mount point entries in the `/etc/fstab` file.

- 19 Set the `LLT_START` attribute to 1 in the `/etc/rc.config.d/lltconf` file:

```
LLT_START=1
```

- 20 Reboot all the nodes.

```
# /usr/sbin/shutdown -r now
```

- 21 Check the status of the cluster.

```
# hastatus -sum
```

- 22 Post Upgrade Tasks: Enable fencing.

```
# hastop -all  
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode  
# /sbin/init.d/vxfen stop  
# /sbin/init.d/vxfen start
```

- 23 Set the clusterwide attribute "UseFence" to use SCSI3. Add the following line to the `/etc/VRTSvcs/conf/config/main.cf` file:

```
UseFence=SCSI3
```

- 24 Start the VCS engine on each system:

```
# hstart
```

Performing a full upgrade from SFCFSHA 5.x on HP-UX 11iv3 to 6.0 HP-UX 11iv3

Use this full upgrade procedure if the operating system upgrade is not required.

To perform a full upgrade from SFCFS 5.x on HP-UX 11iv3 to 6.0 HP-UX 11iv3

- 1 Take offline all VCS groups that contain `CFSMount` and `CVMVolDg`

```
# hagr -offline group -sys system01  
# hagr -offline group -sys system02
```

- 2 Unmount all the non-system VxFS file systems which are not under VCS control.

```
# umount /mount_point
```

Note: Here the CVM is up.

- 3 Upgrade the stack following the installation of all the required patches.
- 4 Reboot the cluster.

```
# /usr/sbin/shutdown -r -y now
```

Performing a full upgrade from SFCFSHA 5.x on HP-UX 11iv3 to 6.0 on the latest HP-UX 11iv3

Use this full upgrade procedure if the operating system upgrade is required.

To perform a full upgrade from SFCF 5.x on HP-UX 11iv3 to 6.0 on the latest HP-UX 11iv3

- 1 Log in as superuser to one of the nodes in the cluster.
- 2 If you have created VxFS mount points on VxVM volumes, added them to the `/etc/fstab` file, and comment out the mount point entries in the `/etc/fstab` file.
- 3 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS. If the applications are under VCS control:

```
# hagr -offline grp_name -any
```

If the applications are not under VCS control:

Use native application commands to stop the application.

- 4 Take offline all VCS groups that contain `CFSMount` and `CVMVolDg`

```
# hagr -offline group -sys system01  
# hagr -offline group -sys system02
```
- 5 Unmount all the non-system VxFS file systems which are not under VCS control.

```
# umount /mount_point
```
- 6 Make sure that no processes which make use of mounted shared file system or shared volumes are running.

```
# fuser -cu mount-point
```


- 7 Stop all VCS service groups.

To view the current state of the service groups:

```
# hagrps -state
```

To stop each group:

```
# hagrps -offline servicegroup -sys node_name
```

- 8 Freeze all the VCS service groups by running the following commands:

```
# haconf -makerw
```

```
# hagrps -freeze servicegroup -persistent
```

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

- 9 Stop VCS on all the nodes:

```
# hastop -all
```

- 10 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

- 11 On each node, edit the `etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cat /etc/vxfenmode
```

```
vxfen_mode=disabled
```

- 12 If the HP-UX 11i v3 10 OEUR release is not already installed, you can now upgrade the HP-UX operating system to the latest available HP-UX 11i version 3 fusion release.

- 13 Install all the prerequisite patches on all the nodes.

- 14 Install SFCFS 6.0.

```
# ./installer
```

From the Installation menu, choose the G option for install and enter the number for Veritas Storage Foundation Cluster File System.

- 15 Uncomment the entries for the non-system VxFS mounts from the `/etc/fstab` file.

16 Reboot all the nodes

```
# /usr/sbin/shutdown -r now
```

17 Enable I/O fencing:

```
# /opt/VRTS/bin/hastop -all
```

18 Execute the following steps on all the nodes:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
# /sbin/init.d/vxfen stop
# /sbin/init.d/vxfen start
```

19 Set the clusterwide attribute "UseFence" to use SCSI3. Add the following line to the `/etc/VRTSvcs/conf/config/main.cf` file:

```
UseFence=SCSI3
```

20 Start the VCS engine on each system:

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

Upgrading Veritas Storage Foundation Cluster File System High Availability software from a release prior to 6.0

This section contains procedures for the Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) upgrade.

Upgrading from Storage Foundation or Storage Foundation and High Availability 5.0_11iv3, 5.0.1, 5.0.1 RP1, or 5.0.1 RP2, 5.1SP1, or 5.1SP1RP1

This procedure describes upgrading from Storage Foundation or Storage Foundation and High Availability 5.0_11iv3, 5.0.1, or 5.0.1 RP1 on HP-UX 11i v3 to 6.0 on HP-UX 11i v3.

After successful completion of the upgrade, disk groups that you created in previous versions are accessible.

To upgrade from Storage Foundation or Storage Foundation and High Availability on HP-UX 11i v3

- 1 Perform the necessary pre-upgrade tasks such as resynchronizing existing database snapshots.
See “[Preparing to upgrade the Veritas software](#)” on page 185.
- 2 Upgrade the HP-UX operating system to the latest available HP-UX 11i v3 fusion release.
See “[Upgrading the HP-UX operating system](#)” on page 225.
- 3 If patches to HP-UX 11i v3 are required, apply the patches before upgrading the product.

- 4 Install Storage Foundation 6.0 for HP-UX 11i v3 using the installer script.

```
# ./installer
```

- 5 Enter **G** to upgrade and press **Return**.
- 6 You are prompted to enter the system names on which the software is to be installed. Enter the system name or names and then press Return.
Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
- 7 You are prompted to agree with the End User License Agreement. Enter **y** and press Return.

```
Do you agree with the terms of the End User License Agreement as
specified in the storage_foundation/EULA/en/EULA_SF_Ux_6.0.pdf
file present on media? [y,n,q,?] y
```

- 8 The installer lists the packages to install or upgrade. You are prompted to confirm that you are ready to stop SF processes.

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

If you select **y**, the installer stops the product processes and makes some configuration updates before upgrading.

- 9 The installer uninstalls and reinstalls the listed packages.
- 10 Uncomment the entries in the `/etc/fstab` file which were commented as part of the pre-upgrade steps.

- 11 Reboot all the nodes.

```
# /usr/sbin/shutdown -r now
```

- 12 Check if the VEA service was restarted:

```
# /opt/VRTS/bin/vxsvcctl status
```

- 13 If the VEA service is not running, restart it:

```
# /opt/VRTS/bin/vxsvcctl start
```

Upgrading from previous versions of Storage Foundation on HP-UX 11i v2

Use this procedure to upgrade from Storage Foundation that uses a previous version on HP-UX 11i v2. You can upgrade to Storage Foundation 6.0 from Storage Foundation 4.1, 4.1 MP1, 4.1 MP2, 5.0, 5.0 MP1, or 5.0 MP2 on HP-UX 11i v2.

After successful completion of the upgrade, any disk groups that were created in Storage Foundation 4.1, 4.1 MP1, 4.1 MP2, 5.0, 5.0 MP1, or 5.0 MP2 are accessible by Storage Foundation 6.0.

Note: If you are upgrading from Storage Foundation for Oracle:

To upgrade from Storage Foundation on HP-UX 11i v2 or Storage Foundation for Oracle on HP-UX 11i v2

- 1 Perform the necessary preupgrade tasks such as resynchronizing existing database snapshots.

See [“Preparing to upgrade the Veritas software”](#) on page 185.

- 2 If you have any external Array Policy Modules (APMs) installed, uninstall the APMs. The following warning message displays during the OS upgrade and also when you issue an administrative command for HP-UX kernel modules after the upgrade, until SF 5.0 on HP-UX 11i v3 is installed:

```
WARNING: The file '/usr/conf/mod/dmpXXX.1' does not  
contain valid kernel code. It will be ignored.
```

This message can be ignored and does not affect the functionality of SF.

- 3 Upgrade the HP-UX operating system to the latest available HP-UX 11i v3 fusion release.
- 4 If patches to HP-UX 11i v3 are required, apply the patches before upgrading the product.
- 5 Install Storage Foundation 6.0 for HP-UX 11i v3 using the installer script.

```
# ./installer
```

- 6 Enter **G** to upgrade and press **Return**.
- 7 You are prompted to enter the system names on which the software is to be installed. Enter the system name or names and then press Return.
Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
- 8 You are prompted to agree with the End User License Agreement. Enter **y** and press Return.

```
Do you agree with the terms of the End User License Agreement as
specified in the storage_foundation/EULA/en/EULA_SF_Ux_6.0.pdf
file present on media? [y,n,q,?] y
```

- 9 The installer lists the packages that will be installed or updated. You are prompted to confirm that you are ready to stop SF processes.

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

If you select y, the installer stops the product processes and makes some configuration updates before upgrading.

- 10 The installer uninstalls and reinstalls the listed packages.
- 11 Reboot all the nodes.

```
# /usr/sbin/shutdown -r now
```

Upgrading from Storage Foundation 3.5 on 11i v1 to SFCFSHA 6.0 on HP-UX 11i v3

This procedure describes upgrading Storage Foundation 3.5 on HP-UX 11i v1 to Storage Foundation 6.0 on HP-UX 11i v3. Upgrading from HP-UX 11i v1 requires

an intermediate upgrade to HP-UX 11i v2 for Storage Foundation and Storage Foundation Clustered File System.

Veritas Volume Manager 3.5 and Veritas Volume Manager 6.0 both support disk group version 90. Therefore, any disk groups with version 90 are accessible by Storage Foundation 6.0 after the upgrade. However, certain features in Storage Foundation 6.0 may require the latest disk group version. Therefore, we recommend upgrading the disk group.

To upgrade from Storage Foundation 3.5 on HP-UX 11i v1 to SFCFSA 6.0 on HP-UX 11i v3

- 1 Stop activity to all Storage Foundation volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
- 2 Upgrade from HP-UX 11i v1 to HP-UX 11i v2 using practices recommended by HP. HP-UX 11i v2 includes VxVM 4.1 by default. All disk groups created using Storage Foundation 3.5 on HP-UX 11i v1 would be accessible.
- 3 Upgrade from HP-UX 11i v2 to the latest available HP-UX 11i v3 fusion release, using practices recommended by HP. The HP-UX 11i v3 fusion release includes VxVM 5.0 by default.
- 4 If patches to HP-UX 11i v3 are required, apply the patches before upgrading the product.
- 5 Install SFCFSA 6.0 for HP-UX 11i v3.

Upgrading from VxVM 5.0 on HP-UX 11i v3 to VxVM 6.0 using integrated VxVM 6.0 package for HP-UX 11i v3

You can upgrade from VxVM 5.0 on HP-UX 11i v3 to VxVM 6.0 on HP-UX 11i v3. Use the integrated VxVM 6.0 package for HP-UX 11i v3 from the Ignite depot.

To upgrade using the integrated VxVM 6.0 package from the Ignite depot

- 1 Use HP recommended steps to integrate VxVM 6.0 package with the latest 11i v3 fusion.
- 2 Stop activity to all Storage Foundation volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
- 3 Upgrade the HP-UX 11i v3 operating system to the latest fusion using the integrated VxVM 6.0 package. This process installs VxVM 6.0 package with the operating system. All disk groups that were created using VxVM 5.0 11i v3 are accessible.

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 depots in phase 1 and VCS and VCS agent depots 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.

- Perform rolling upgrades from 5.1 SP1 to 6.0.

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 Log in as superuser and mount the SFCFSHA Veritas Storage Foundation Cluster File System High Availability 6.0 installation media.
- 3 From root, start the installer.


```
# ./installer
```
- 4 From the menu select **G) Upgrade a product**; then select **2) Rolling Upgrade**;
- 5 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 6 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.
- 7 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.
- 8 Review the EULA, and type **y** if you agree to its terms.
- 9 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.

- 10 The installer stops relevant processes, uninstalls old kernel depots, and installs the new depots.
 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.
- 11 Complete the preparatory steps on the nodes that you have not yet upgraded.
- 12 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 6 through step 10.
 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.
- 13 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.
- 14 The installer determines the remaining depots to upgrade. Press **Enter** to continue.
- 15 The installer stops Veritas Cluster Server (VCS) processes. Type **y** to continue.
 The installer performs prechecks, uninstalls old depots, and installs the new depots. It performs post-installation tasks, and the configuration for the upgrade.
- 16 Type **y** or **n** to help Symantec improve the automated installation.
- 17 If you have network connection to the Internet, the installer checks for updates.
 If updates are discovered, you can apply them now.
- 18 Upgrade application.

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 207.

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 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 65.
- 4 In the Task pull-down menu, select G) Upgrade a product; then select 2)Rolling Upgrade;.
Click the **Next** button to proceed.
- 5 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.
- 6 Review the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.
- 7 The installer stops all processes. Click **Next** to proceed.
- 8 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.
- 9 Complete the preparatory steps on the nodes that you have not yet upgraded.

- 10 When prompted, perform step 4 through step 8 on the nodes that you have not yet upgraded.

After upgrading the kernelpkgs on each node, reboot the system and restart the Web-based installer.

- 11 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 65.

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 from version 5.x on HP-UX 11i v3 to Veritas Storage Foundation Cluster File System High Availability 6.0](#)
- [Performing phased upgrade of SFCFSHA from versions 4.x or 5.x on HP-UX 11i v2 to 5.1SP1](#)

Performing a phased upgrade from version 5.x on HP-UX 11i v3 to Veritas Storage Foundation Cluster File System High Availability 6.0

Perform the following procedures to upgrade SFCFSHA clusters from version 5.x on HP-UX 11i v3 to Veritas Storage Foundation Cluster File System High Availability 6.0.

The phased upgrade involves the following steps:

- Upgrading the first half of the cluster, system01 and system02.

Note: Your downtime starts after you complete the upgrade of the first half of the cluster.

- Stopping the second half of the cluster, system03 and system04.
- Bringing online the first half of the cluster, system01 and system02.

Note: Your downtime ends after you bring the first half of the cluster online.

- Upgrading the second half of the cluster, system03 and system04.

Perform the following steps on the first half of the cluster, system01 and system02.

To upgrade the first half of the cluster

- 1 Stop all the applications on the nodes that are not under VCS control. Use native application commands to stop the applications.
- 2 Switch the failover groups from the first half of the cluster to one of the nodes in the second half of the cluster.

```
# hagrpswitch failover_group -to system03
```

- 3 Stop all VCS service groups.

```
# hagrpsvc -offline group_name -sys system01
# hagrpsvc -offline group_name -sys system02
```

- 4 Freeze the nodes in the first half of the cluster

```
# haconf makerw
# hasys -freeze -persistent system01
# hasys -freeze -persistent system02
# haconf -dump -makero
```

- 5 Stop VCS on the first half of the cluster:

```
# hastop -local -force
```

- 6 If you created local VxFS mount points on VxVM volumes, added them to the `/etc/fstab` file, and comment out the mount point entries in the `/etc/fstab` file.

- 7 Set the `LLT_START` attribute to 0 in the `/etc/rc.config.d/lltconf` file:

```
LLT_START=0
```

- 8 On each node of the first half of the cluster, edit the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cat /etc/vxfenmode
vxfen_mode=disabled
```

9 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.

10 Stop all the modules on the first half of the cluster.

```
# /sbin/init.d/odm stop
# /sbin/init.d/vxfen stop
# /sbin/gabconfig -U
# kcmodule vxfen=unused
# kcmodule odm=unused
# kcmodule gab=unused
# lltconfig -U
# kcmodule llt=unused
```

11 On each node of the first half of the cluster, remove the following device files:

```
# rm -f /dev/llt
# rm -f /dev/gab*
# rm -f /dev/vxfen
```

12 Upgrade the HP-UX operating system to HP-UX 11.31 Mar 10 fusion.

See [“Upgrading the HP-UX operating system”](#) on page 225.

13 Upgrade SFCFSHA:

```
# ./installer
```

Note: DO NOT reboot the cluster.

After the installation completes, perform the following steps on the second half of the cluster.

Note: Your downtime starts now.

Perform the following steps on the second half of the cluster, `system03` and `system04`, to stop the second half of the cluster.

To stop the second half of the cluster

- 1 Stop all the applications on the node that are not under VCS control. Use native application commands to stop the applications.

- 2 Stop all VCS service groups.

```
# hagrps -offline group_name -sys system03
# hagrps -offline group_name -sys system04
```

- 3 Freeze the nodes in the second half of the cluster

```
# haconf makerw
# hasys -freeze group_name -persistent
# haconf -dump -makero
```

- 4 Stop VCS on the second half of the cluster:

```
# hastop -local -force
```

- 5 If you created local VxFS mount points on VxVM volumes, added them to the `/etc/fstab` file, and comment out the mount point entries in the `/etc/fstab` file.

- 6 Set the `LLT_START` attribute to 0 in the `/etc/rc.config.d/lltconf` file:

```
LLT_START=0
```

- 7 On each node of the second half of the cluster, edit the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cat /etc/vxfenmode
vxfen_mode=disabled
```

- 8 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.

- 9 Stop all the modules on the second half of the cluster:

```
# /sbin/init.d/odm stop
# /sbin/init.d/vxfen stop
# /sbin/gabconfig -U
# kcmodule vxfen=unused
# kcmodule odm=unused
# kcmodule gab=unused
# /sbin/lltconfig -U
# kcmodule llt=unused
```

- 10 On each node of the second half of the cluster, remove the following device files:

```
# rm -f /dev/llt
# rm -f /dev/gab*
# rm -f /dev/vxfen
```

Perform the following steps on the first half of the cluster, system01 and system02, to bring the first half of the cluster online.

To bring the first half of the cluster online

- 1 Uncomment the VxFS mount point entries in the `/etc/fstab` file.
- 2 Mount the VxFS file systems.
- 3 Enable fencing:

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

- 4 Remove the following line from the `/etc/VRTSvcs/conf/config/main.cf` file:

```
Frozen=1
```

- 5 Set the clusterwide attribute `UseFence` to use SCSI3. Add the following line to the `/etc/VRTSvcs/conf/config/main.cf` file:

```
UseFence=SCSI3
```

- 6 Set the `LLT_START` attribute to 1 in the `/etc/rc.config.d/lltconf` file:

```
LLT_START=1
```

- 7 Reboot the first half of the cluster:

```
# /usr/sbin/shutdown -r now
```

- 8 After the nodes come up, seed the cluster membership:

```
# gabconfig -x
```

The first half of the cluster is now up and running.

Note: The downtime ends here.

Perform the following steps on the second half of the cluster, system03 and system04, to upgrade the second half of the cluster.

To upgrade the second half of the cluster

- 1 Upgrade the HP-UX operating system to HP-UX 11.31 Mar 11.

See [“Upgrading the HP-UX operating system”](#) on page 225.

- 2 Upgrade SFCFSHA:

```
# ./installer
```

Note: DO NOT reboot the cluster.

- 3 Uncomment the VxFS mount point entries in the `/etc/fstab` file on the second half of the cluster.

- 4 Mount the VxFS file systems.

- 5 Enable fencing:

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

- 6 Set the `LLT_START` attribute to 1 in the `/etc/rc.config.d/lltconf` file:

```
LLT_START=1
```

- 7 Reboot the second half of the cluster:

```
# /usr/sbin/shutdown -r now
```

The nodes system03 and system04 now join the first half of the cluster.

Performing phased upgrade of SFCFSHA from versions 4.x or 5.x on HP-UX 11i v2 to 5.1SP1

The phased upgrade involves the following steps:

- Upgrading the first half of the cluster, system01 and system02.

Note: Your downtime starts after you complete the upgrade of the first half of the cluster.

- Stopping the second half of the cluster, system03 and system04.
- Bringing online the first half of the cluster, system01 and system02.

Note: Your downtime ends after you bring the first half of the cluster online.

- Upgrading the second half of the cluster, system03 and system04.

Perform the following steps on the first half of the cluster, system01 and system02, to upgrade the first half of the cluster.

To upgrade the first half of the cluster

- 1 Stop all the applications that are not configured under VCS.
- 2 Switch the failover groups from the first half of the cluster to one of the nodes in the second half of the cluster:

```
# hagrpswitch failover_group -to system03
```

- 3 Stop all VCS service groups.

```
# hagrpswitch -offline group_name -sys system01
# hagrpswitch -offline group_name -sys system02
```

- 4 Freeze the nodes in the first half of the cluster

```
# haconf makerw
# hasys -freeze -persistent system01
# hasys -freeze -persistent system02
# haconf -dump -makero
```

- 5 Stop VCS on the first half of the cluster:

```
# hstop -local -force
```

- 6 If you created local VxFS mount points on VxVM volumes, added them to `/etc/fstab` file, and comment out the mount point entries in the `/etc/fstab` file.

- 7 Set the `LLT_START` attribute to 0 in the `/etc/rc.config.d/lltconf` file:

```
LLT_START=0
```

- 8 On each node of the first half of the cluster, edit the `/etc/vxfenmode` file to configure I/O fencing in disabled mode.

```
# cat /etc/vxfenmode
vxfen_mode=disabled
```

- 9 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.

- 10 Stop all the modules on the first half of the cluster.

```
# /sbin/init.d/odm stop
# /sbin/init.d/vxfen stop
# /sbin/gabconfig -U
# kcmodule vxfen=unused
# kcmodule odm=unused
# kcmodule gab=unused
# lltconfig -U
# kcmodule llc=unused
```

- 11 On each node of the first half of the cluster, remove the following device files:

```
# rm -f /dev/llt
# rm -f /dev/gab*
# rm -f /dev/vxfen
```

- 12 Upgrade the operating system choosing three Base Bundles "Base-VxFS-50, Base-VxVM-50, Base-VxTools-50"..
- 13 Upgrade SFCFSHA:

```
# ./installer
```

Choose the upgrade option "G" when the installer prompts you.

Note: DO NOT reboot the cluster.

Perform the following steps on the second half of the cluster, system03 and system04, to stop the second half of the cluster.

Note: The downtime starts now.

To stop the second half of the cluster

- 1 Stop all the applications that are not configured under VCS.
- 2 Stop all VCS service groups.

```
# hagrps -offline group_name -sys system03
# hagrps -offline group_name -sys system04
```

- 3 Freeze the VCS service groups on the second half of the cluster:

```
# haconf -makerw
# hagrps -freeze group_name -persistent
# haconf -dump -makero
```

- 4 Stop VCS on the second half of the cluster:

```
# hastop -local -force
```

- 5 If you created local VxFS mount points on VxVM volumes, added them to `/etc/fstab` file, and comment out the mount point entries in the `/etc/fstab` file.
- 6 Set the `LLT_START` attribute to 0 in the `/etc/rc.config.d/lltconf` file:

```
LLT_START=0
```

- 7 On each node of the second half of the cluster, edit the `/etc/vxfenmode` file to configure I/O fencing in disabled mode:

```
# cat /etc/vxfenmode
vxfen_mode=disabled
```

- 8 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.
- 9 Stop all the modules on the second half of the cluster:

```
# /sbin/init.d/odm stop
# /sbin/init.d/vxfen stop
# /sbin/gabconfig -U
# kcmodule vxfen=unused
# kcmodule odm=unused
# kcmodule gab=unused
# /sbin/lltconfig -U
# kcmodule llt=unused
```

- 10 On each node of the second half of the cluster, remove the following device files:

```
# rm -f /dev/llt
# rm -f /dev/gab*
# rm -f /dev/vxfen
```

Perform the following steps on the first half of the cluster, `system01` and `system02`, to bring the first half of the cluster online.

To bring the first half of the cluster online

- 1 Uncomment the VxFS mount point entries in the `/etc/fstab` file.
- 2 Mount the VxFS file systems.
- 3 Enable fencing:

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

- 4 Remove the following line from the `/etc/VRTSvcs/conf/config/main.cf` file:

```
Frozen=1
```

- 5 Set the clusterwide attribute `UseFence` to use SCSI3. Add the following line to the `/etc/VRTSvcs/conf/config/main.cf` file:

```
UseFence=SCSI3
```

- 6 Reboot the first half of the cluster:

```
# /usr/sbin/shutdown -r now
```

- 7 After the nodes come up, seed the cluster membership:

```
# gabconfig -x
```

The first half of the cluster is now up and running.

Note: The downtime ends here.

Perform the following steps on the second half of the cluster, system03 and system04, to upgrade the second half of the cluster.

To upgrade the second half of the cluster

- 1 Upgrade the operating system.

See [“Upgrading the HP-UX operating system”](#) on page 225.

- 2 Upgrade SFCFSHA:

```
# ./installsfcfsha -upgrade system03 system04
```

Note: DO NOT reboot the cluster.

- 3 Uncomment the VxFS mount point entries in the `/etc/fstab` file.

- 4 Mount the VxFS file systems.

- 5 Enable fencing:

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

- 6 Reboot the second half of the cluster:

```
# /usr/sbin/shutdown -r now
```

The nodes system03 and system04 now join the first half of the cluster.

Upgrading the operating system

This chapter includes the following topics:

- [Upgrading the HP-UX operating system](#)

Upgrading the HP-UX operating system

If you are on an unsupported version of the operating system, you need to upgrade it to HP-UX B.11.31.1103, HP-UX 11i Version 3 March 2011 or later.

If you are upgrading the operating system from HP-UX 11i v2, make sure that you choose the following depots along with the HP-UX 11i v3 March 2011 OEUR release depots:

- `Base-VxFS-version`
where *version* is the base VxFS version bundled with the operating system.
- `Base-VxTools-version`
where *version* is the base VxTools version bundled with the operating system.
- `Base-VxVM-version`
where *version* is the base VxVM version bundled with the operating system.

To upgrade the operating system from HP-UX 11i v2, run the `update-ux` command specifying the Veritas depots along with the HP-UX operating system depots:

```
# swinstall -s os_path Update-UX
# update-ux -s os_path HPUX11i-DC-OE \
Base-VxFS-version Base-VxTools-version \
Base-VxVM-version
```

where *os_path* is the full path of the directory containing the operating system depots.

where *version* is the the base version of Veritas depots bundled with the operating system.

To upgrade the operating system from HP-UX 11i v3, run the `update-ux` command as follows:

```
# update-ux -s os_path HPUX11i-DC-OE
```

where *os_path* is the full path of the directory containing the operating system depots.

For detailed instructions on upgrading the operating system, see the operating system documentation.

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.

See [“Upgrading VVR without disrupting replication”](#) on page 227.

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 188.

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 4.1 or later to VVR 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 4.1 or later to VVR 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 188.

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 existing SFHA `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.

On the nodes that have any mounted VxFS local file systems that are not under VCS control:

```
# umount -F vxfs -a
```

- 5 Stop all 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 dg_name stopall  
# vxdg deport dg_name
```

- 6 Upgrade the existing SFHA to SFCFSHA 6.0:

For SFCFSHA:

```
# ./installsfcfsha
```

- 7 After installation is completed, reboot all the nodes.
- 8 After all nodes are rebooted, bring up CVM and the resources.
- 9 Verify that all SFHA processes have started. You can verify using the following commands:

```
# gabconfig -a  
# hastatus -sum
```

- 10 Configure CVM and the resources:

```
# /opt/VRTS/bin/cfscluster config
```

This automatically detects the cluster configuration such as node-names, cluster name, and cluster=id and brings up CVM resources on all the nodes in the cluster.

To verify:

```
# gabconfig -a  
# hastatus -sum
```

- 11 Find out which node in the cluster, is the master node:

```
# /opt/VRTS/bin/vxclustadm nidmap
```

- 12 On the master node, import disk groups:

```
# vxdg -s import dg_name
```

This release supports certain commands to be executed from the slave node such as `vxdg -s import dg_name`.

See the *Storage Foundation Cluster File System Administrator's Guide* for more information.

- 13 Start all the volumes on the imported disk group , run the following command:

```
# vxvol -g dg_name name startall
```

- 14 To have the VxFS file system to be under VCS control, run the following command

```
# cfsmntadm add shared_diskgroup_name volume_name \  
mount_point all=cluster_mount_options
```

This command creates Parallel service groups in VCS comprising of the supplied parameters of Diskgroup, Volume & Mountpoint.

- 15 Mount the CFS file system on all the nodes in the cluster:

```
# cfsmount mount_point
```

- 16 On the CVM Master node, re-import all the required disk groups which must be in shared mode:

Import all other local disk groups which have not been imported in shared mode in step 12.

```
# vxdg -s import dg_name
```

- 17 Start all the volumes whose disk groups have been imported as shared. Use the following command:

```
# vxdg -g dg_name startall
```

- 18 Repeat steps 14 and 15 for any of the VxFS file systems which VCS needs to monitor through Failover service groups.

Verification of the installation or the upgrade

- [Chapter 18. 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 359.

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 `swlist` command to check which depots have been installed:

```
# swlist -l product | grep VRTS
```

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 depots, and the status (success or failure) of each depot. 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 237.

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 18-1](#) depicts the SFCFSHA cluster communication with the LDAP servers when clusters run in secure mode.

Figure 18-1 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 `vxiod`, `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxpall`, `vxcached`, `vxconfigbackupd`, and `vxsvc` processes should appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

If you installed Storage Foundation and High Availability, the `vxodmd` and `vxdbd_11.31` processes are also displayed.

Checking Veritas File System installation

The Veritas File System package consists of a kernel component and administrative commands.

Command installation verification

The Veritas File System commands are installed in the `/opt/VRTS/bin` directory. To verify, determine that the subdirectory is present:

```
# ls /opt/VRTS/bin
```

Make sure you have adjusted your environment variables accordingly.

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 18-1](#) lists the LLT configuration files and the information that these files contain.

Table 18-1 LLT configuration files

File	Description
/etc/rc.config.d/lltconf	<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>

Table 18-1 LLT configuration files (*continued*)

File	Description
/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> 0 galaxy 1 nebula </pre>
/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 LLT network links that correspond to the specific system.</p> <p>For example, the file <code>/etc/llttab</code> contains the entries that resemble:</p> <pre> set-node galaxy set-cluster 2 link lan1 /dev/lan:1 - ether - - link lan2 /dev/lan: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 18-2](#) lists the GAB configuration files and the information that these files contain.

Table 18-2 GAB configuration files

File	Description
<p>/etc/rc.config.d/ gabconf</p>	<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 Storage Foundation Cluster File System High Availability configuration.</p>
<p>/etc/gabtab</p>	<p>After you install SFCFSHA, the file /etc/gabtab contains a <code>gabconfig(1)</code> command that configures the GAB driver for use.</p> <p>The file /etc/gabtab 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 248.

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 251.
- 4 Verify GAB operation.
See [“Verifying GAB”](#) on page 253.
- 5 Verify the cluster operation.
See [“Verifying the cluster”](#) on page 255.

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:

```
Node           State   Link   Status   Address
*0 galaxy     OPEN
                lan1 UP    08:00:20:93:0E:34
                lan2 UP    08:00:20:93:0E:38
  1 nebula     OPEN
                lan1 UP    08:00:20:8F:D1:F2
                lan2 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 `gabconfig` 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 Storage Foundation Cluster File System High Availability 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.

```
#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  141
galaxy       ConfigCheckSum    33975
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     5.0.31.0
galaxy       Frozen            0
galaxy       GUIIPAddr
galaxy       HostUtilization   CPU 0 Swap 0
```


#System	Attribute	Value
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
galaxy	SysInfo	HP-UX:galaxy,U,B.11.31,ia64
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 19. Adding a node to a cluster](#)
- [Chapter 20. 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](#)

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 19-1 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 260.

Table 19-1 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 262. See “Adding a node using the Web-based installer” on page 265. See “Adding the node to a cluster manually” on page 266.
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 288.

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 27.
- 2 Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
- 3

- 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:


```
# vxdctl protocolversion
Cluster running at protocol 110
```

- 5 If the cluster protocol is below 110, upgrade it on the master node using:

```
# vxdctl 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 19-1](#).

Figure 19-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the SFCFSHA private Ethernet controllers.

[Figure 19-1](#) 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
- 2 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 26.

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/clusuid
 - /etc/rc.config.d/lltconf
 - /etc/rc.config.d/gabconf
 - /etc/rc.config.d/vxfenconf
 - /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
# ./installsfsha -addnode
```

The installer displays the copyright message and the location where it stores the temporary installation logs.

- 3 Review and confirm the cluster information.
- 4 Enter the names of the systems that you want to add as new nodes to the cluster.

The installer checks the installed products and packages on the nodes and discovers the network interfaces.

- 5 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.

Note: At least two private heartbeat links must be configured for high availability of the cluster.

- 6 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.
- 7 Review and confirm the information.
- 8 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.
- 9 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 indicates the location of the log file with details of the actions performed.

- 10 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.
- 11 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 19-2 Tasks for adding a cluster manually

Task	Description
Start the Volume Manager.	
Add the cluster files.	Configure LLT, GAB, and VCSMM. Configure LLT and GAB.
Generate security credentials.	If the CP server of the existing cluster is secure, generate security credentials on the new node. See “Setting up the node to run in secure mode” on page 279.
Set up fencing.	Configure fencing for the new node to match the fencing configuration on the existing cluster. If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node. See “Starting fencing on the new node” on page 282.
Start VCS.	See “Starting VCS after adding the new node” on page 283.
Configure Cluster Volume Manager (CVM) and Cluster File System (CFS).	If CVM and CFS are configured on the existing cluster, configure them on the new node. See “Configuring CVM and CFS on the new node” on page 284.
Configure the ClusterService group.	If the ClusterService group is configured on the existing cluster, add the node to the group. See “Configuring the ClusterService group for the new node” on page 285.

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 VxVM uses license keys to control access. As you run the utility, answer "n" when prompted about licensing; you installed the appropriate license when you ran the `installsfcfsha` utility.
- 3 Enter **n** when prompted to set up a system wide disk group for the system. The installation completes.
- 4 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 lan1 /dev/lan:1 - ether --
link lan2 /dev/lan: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 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
```

- 7 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
```

- 8 Start the LLT, GAB, and ODM drivers on the new node:

```
# /sbin/init.d/llt start  
  
# sbin/init.d/gab start  
  
# /sbin/init.d/vxfen start  
  
# kcmodule vxgms=loaded  
  
# kcmodule odm=loaded  
  
# /sbin/init.d/odm stop  
  
# /sbin/init.d/odm start
```

- 9 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:

```
link lan1 /dev/lan:1 - ether --
link lan2 /dev/lan:2 - ether --
```

For Linux:

- 4 For Solaris:

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/rc.config.d/lltconf
/etc/rc.config.d/gabconf
/etc/rc.config.d/vcsconf
/etc/rc.config.d/vxfenconf
/etc/vcsmmtab
```

```
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:

```
# /sbin/init.d/llt start
```

For Linux:

For Solaris:

2 Start GAB.

For AIX:

For HP-UX:

```
# sbin/init.d/gab start
```

For Linux:

For Solaris:

3 Start disk-based fencing.

For AIX:

For HP-UX:

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

For Linux:

For Solaris:

4 Start VCSMM.

For AIX:

For HP-UX:

```
# kcmodule vxgms=loaded
```

For Linux:

For Solaris:

5 Start LMX.

For AIX:

For HP-UX:

For Solaris:

6 Start ODM.

For AIX:

For HP-UX:

```
# kcmodule odm=loaded
# /sbin/init.d/odm stop
# /sbin/init.d/odm start
```

For Linux:

For Solaris:

7 Start the LLT, GAB, and ODM drivers on the new node:

For AIX:

For HP-UX:

```
# /sbin/init.d/llt start
# sbin/init.d/gab start
# /sbin/init.d/vxfen start
# kcmodule vxgms=loaded
# kcmodule odm=loaded
# /sbin/init.d/odm stop
# /sbin/init.d/odm start
```

For Linux:

For Solaris:

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:

```
/etc/rc.config.d/vxfenconf
/etc/vxfendg
/etc/vxfenmode
```

For Linux:

For Solaris:

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:

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

- 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:

```
/etc/rc.config.d/vxfenconf
/etc/vxfendg
/etc/vxfenmode
```

For Linux:

For Solaris:

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:

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

- 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/CVMTTypes.cf \  
saturn:/etc/VRTSvcs/conf/config/CVMTTypes.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.

```
# haress -modify gcoip Device lan0 -sys saturn
```

```
# haress -modify gconic Device lan0 -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 `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 VxVM uses license keys to control access. As you run the utility, answer "n" when prompted about licensing; you installed the appropriate license when you ran the `installsfcfsha` utility.

- 3 Enter `n` when prompted to set up a system wide disk group for the system. The installation completes.

- 4 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 lan1 /dev/lan:1 - ether --
link lan2 /dev/lan: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 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
```

- 7 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
```

- 8 Start the LLT, GAB, and ODM drivers on the new node:

```
# /sbin/init.d/llt start  
  
# sbin/init.d/gab start  
  
# /sbin/init.d/vxfen start  
  
# kcmodule vxgms=loaded  
  
# kcmodule odm=loaded  
  
# /sbin/init.d/odm stop  
  
# /sbin/init.d/odm start
```

- 9 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 19-3](#) uses the following information for the following command examples.

Table 19-3 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:

```
/etc/rc.config.d/vxfenconf
/etc/vxfendg
/etc/vxfenmode
```

For Linux:

For Solaris:

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:

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

- 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 lan0 -sys saturn
# hares -modify gconic Device lan0 -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.

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*.

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 290.
- 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 290.
- Modify the VCS configuration files on the existing nodes.
See [“Modifying the VCS configuration files on existing nodes”](#) on page 291.
- 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 295.

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 291.

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 295.

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 *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 288.

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 349.

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 21. Setting up a replicated global cluster](#)
- [Chapter 22. 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 302.
- Configure a cluster at the secondary site.
See [“Configuring an SFCFSHA cluster at the secondary site”](#) on page 304.
- Configure replication on clusters at both sites.
See [“Configuring replication for SFCFSHA clusters at both sites”](#) on page 306.
- Configure VCS service groups for replication.
See [“Modifying the ClusterService group for a global SFCFSHA cluster”](#) on page 306.
See [“Defining the remote SFCFSHA cluster and heartbeat objects”](#) on page 308.
See [“Configuring the VCS service groups for global SFCFSHA clusters”](#) on page 311.
- 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 50.

- 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 21-1 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 21-1 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.

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 25.

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 304.

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 63.

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:

```
# hagrps -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 63.

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 313.](#)
- 2 Start replication between the sites.
[See “Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): starting replication” on page 321.](#)

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 325.
- 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 335.

- 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 50.

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 300.
- 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 313.
See [“Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): setting up the secondary site”](#) on page 316.
- Starting replication of the database.
See [“Configuring a global SFCFSHA cluster using Veritas Volume Replicator \(VVR\): starting replication”](#) on page 321.
- 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 323.

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 313.
- Creating the Replicated Volume Group (RVG) on the primary site.
See [“Setting up the Replicated Volume Group \(RVG\) on the primary site”](#) on page 315.
- 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 317.

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
```

- 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 317.
- Edit the `/etc/vx/vras/.rdg` file on the secondary site.
See [“Editing the /etc/vx/vras/.rdg files”](#) on page 317.
- 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 318.

- Create the replication objects on the secondary site.
See [“Setting up the disk group on secondary site for replication”](#) on page 319.

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 lan0:1
- The virtual IP address is 10.10.9.101
- The net mask is 255.255.255.0
- ```
ifconfig lan0:1 plumb
ifconfig lan0:1 inet 10.10.9.101 netmask 255.255.255.0
ifconfig lan0:1 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 lan0:1

**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 22-1** 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 421.</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).

- 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 = lan0
 Address = "10.10.9.101"
 NetMask = "255.255.255.0"
)

NIC nic (
 Device = lan0
 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 vxfsend (
 PathName = "/sbin/vxfsend"
 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 `jupitar`. 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, jupitar = 1 }
 AutoStartList = { mercury, jupitar }
)

IP logowner_ip (
 Device = lan0
 Address = "10.11.9.102"
 NetMask = "255.255.255.0"
)

NIC nic (
 Device = lan0
 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, jupitar = 1 }
 Parallel = 1
 AutoStartList = { mercury, jupitar }
)

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, jupitar = 1 }
 ClusterList = { clus2 = 0, clus1 = 1 }
 Parallel = 1
 OnlineRetryInterval = 300
 ClusterFailOverPolicy = Manual
 Authority = 1
 AutoStartList = { mercury, jupitar }
)

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 22-2**

| 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 22-3**

| 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. 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 22-4** Options for the remote cluster to take over the primary role

| Takeover options | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Disaster         | <p>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.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| 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       | <p>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.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Replica          | <p>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".</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

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 23. Uninstalling Storage Foundation Cluster File System High Availability](#)

# Uninstalling Storage Foundation Cluster File System High Availability

This chapter includes the following topics:

- [Shutting down cluster operations](#)
- [Disabling the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling SFCFSHA depots using the script-based installer](#)
- [Uninstalling SFCFSHA with the Veritas Web-based installer](#)
- [Removing license files \(Optional\)](#)
- [Removing the CP server configuration using the removal script](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product](#)

## 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:

```
hagrpl -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:

```
hagrpl -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:

```
vxrlink -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.

Go on to uninstalling Volume Manager to uninstall VVR.

- 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
```

- 6 Uninstall the VVR packages.

## Uninstalling SFCFSHA depots using the script-based installer

Use the following procedure to remove SFCFSHA products.

Not all depots 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 399.



**To shut down and remove the installed SFCFSHA depots**

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/fstab`. 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 depot (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

- 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 depots are uninstalled.

The uninstall script creates log files and displays the location of the log files.

- 9 Most depots have kernel components. In order to ensure complete removal, a system reboot is recommended after all depots have been removed.

# 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 65.
- 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**.

The Web-based installer prompts you to reboot the system.

You see a prompt recommending that you reboot the system, and then return to the Web page to complete additional tasks.

## Removing license files (Optional)

Optionally, you can remove the license files.

### To remove the VERITAS license files

- 1 To see what license key files you have installed on a system, enter:

```
/sbin/vxlicrep
```

The output lists the license keys and information about their respective products.

- 2 Go to the directory containing the license key files and list them:

```
cd /etc/vx/licenses/lic
ls -a
```

- 3 Using the output from step 1, identify and delete unwanted key files listed in step 2. Unwanted keys may be deleted by removing the license key file.

## 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
vxcpserv 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
```

- 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 29.

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>installvm</code>	Veritas Volume Manager
<code>installfs</code>	Veritas File System
<code>installdmp</code>	Veritas Dynamic Multi-Pathing

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 353.

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.
<code>-allpkgs</code>	Displays all depots and patches required for the specified product. The depots 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.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<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>-ignite</code>	The <code>-ignite</code> option allows you to create SD bundles for a product. When you create the SD bundle for the product, the Veritas product disc must be mounted on the Ignite-UX Server.
<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 depots are installed.
<code>-installminpkgs</code>	Specifies that the minimum depot set is installed.
<code>-installrecpkgs</code>	Specifies that the required depot set is installed.
<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.
<code>-makeresponsefile</code>	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.

Table A-2 Available command line options (*continued*)

Command Line Option	Function
<code>-minpkgs</code>	Displays the minimal depots and patches required for the specified product. The depots and patches are listed in correct installation order. Optional depots 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 depots 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 depots 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 depots.
<code>-pkgpath <i>package_path</i></code>	Designates the path of a directory that contains all depots 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 depot group (minimum, recommended, all) and depots that are installed on the specified systems.
<code>-pkgtable</code>	Displays product's depots 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 depots and patches required for the specified product. The depots and patches are listed in correct installation order. Optional depots 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 depots 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 depots 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 depots 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 depots. On HP-UX operating systems, <code>-rootpath</code> passes <code>-I path</code> to <code>swinstall</code> .
<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 399.
<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 depots 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 depots and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing depots 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 `VRTSllt` pkg version is not consistent on the nodes.
- The `llt-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 234.

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 354.

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 depots 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 367.

- 4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfcfsha -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 depots. 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 depots. 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 depots 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} {depot}	<p>Instructs the installation to not install the optional depots in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{donotremove} {depot}	<p>Instructs the uninstallation to not remove the optional depots 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 depots installed, without configuration. List or scalar: list Optional or required: optional
CFG{opt}{uninstall}	Uninstalls SFCFSHA depots. 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 depots 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>remsh</i> must be used instead of ssh as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
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 Storage Foundation Cluster File System High Availability cluster.

Table B-3 Response file variables specific to configuring a basic Storage Foundation Cluster File System High Availability 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 Storage Foundation Cluster File System High Availability 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.

Response file variables to configure Veritas Storage Foundation Cluster File System High Availability

Table B-4 Response file variables specific to configuring private LLT over Ethernet

Variable	List or Scalar	Description
CFG{vcs_lltlink#} {"system"}	Scalar	Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. You must enclose the system name within double quotes. (Required)
CFG{vcs_lltlinklowpri#} {"system"}	Scalar	Defines a low priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication. If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on. You must enclose the system name within double quotes. (Optional)

[Table 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	Indicates whether to configure heartbeat link using LLT over UDP. (Required)

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)

Response file variables to configure Veritas Storage Foundation Cluster File System High Availability

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 Storage Foundation Cluster File System High Availability cluster.

Table B-6 Response file variables specific to configuring virtual IP for Storage Foundation Cluster File System High Availability 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)

Response file variables to configure Veritas Storage Foundation Cluster File System High Availability

Table B-6 Response file variables specific to configuring virtual IP for Storage Foundation Cluster File System High Availability cluster (*continued*)

Variable	List or Scalar	Description
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 Storage Foundation Cluster File System High Availability cluster in secure mode.

Table B-7 Response file variables specific to configuring Storage Foundation Cluster File System High Availability cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonenode}	Scalar	Specifies that the securityonenode option is being used.
CFG{securityonenode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. <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)
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_smtprecip}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)

Table B-9 Response file variables specific to configuring VCS notifications using SMTP *(continued)*

Variable	List or Scalar	Description
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_snmpsev}	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 Storage Foundation Cluster File System High Availability global clusters.

Table B-11 Response file variables specific to configuring Storage Foundation Cluster File System High Availability 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/installsfcfs-xxxxxx/
```



```
installsfcfs-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;
$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{opt}{noextrapkgs}=1;
$CFG{vcs_clusterid}=127;
$CFG{vcs_clustername}="uxrt6_sol";
$CFG{vcs_username}=[ qw(admin operator) ];
$CFG{vcs_userenpw}=[ qw(JlmElgLimHmKumGlj
bQOsOUnVqOOUnTQsOSnUQuOUnPqTOS) ];
$CFG{vcs_userpriv}=[ qw(Administrators Operators) ];
$CFG{vcs_lltlink1}{system01}="lan1";
$CFG{vcs_lltlink2}{system01}="lan2";
$CFG{vcs_lltlink1}{system02}="lan1";
$CFG{vcs_lltlink2}{system02}="lan2";
$CFG{vcs_enabled}=1;
$CFG{opt}{logpath}="/opt/VRTS/install/logs/installsfcfs-xxxxxx/
installsfcfs-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 380.

- 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 381.

- 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 382.

See [“About response files”](#) on page 361.

You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 384.

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 384.

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 383.
- 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 384.

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 383.
- 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 384.

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 383.
- 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 384.

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_evm_handling	(Veritas Dynamic Multi-Pathing) Whether EVM should be handled or not.
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_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.
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.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
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.
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 tune(4) manual page for setting this tunable for a specified block device.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
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_checkpt_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.
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).

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
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.
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.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
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.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires system reboot to take effect.
voliomem_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_jobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_jobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires system reboot to take effect.
voliot_jobuf_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.
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.

Table C-1 Supported tunable parameters (*continued*)

Tunable	Description
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 Storage Foundation Cluster File System High Availability.

To configure I/O fencing using response files

- 1 Make sure that Storage Foundation Cluster File System High Availability 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 80.

- 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 393.
 See “[Sample response file for configuring server-based I/O fencing](#)” on page 396.
- 4 Edit the values of the response file variables as necessary.
 See “[Response file variables to configure disk-based I/O fencing](#)” on page 392.
 See “[Response file variables to configure server-based I/O fencing](#)” on page 394.
- 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 392.

```
#
# 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(c1t1d0 c2t1d0 c3t1d0) ];
$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 Storage Foundation Cluster File System High Availability in virtual machines that do not support SCSI-3 PR”](#) on page 75.

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](#)
- [Enabling remsh](#)

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 `remsh` (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 (`remsh`). Symantec recommends that you use `ssh` as it is more secure than `remsh`.

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.

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 `/opt/ssh/etc/sshd_config` file must contain the following two lines:

```
PermitRootLogin          yes
Subsystem sftp           /opt/ssh/libexec/sftp-server
```

- 2 If the lines are not there, add them and restart ssh:

```
system1 # /sbin/init.d/secsh start
```

- 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.

Enabling remsh

Remote shell functionality is enabled automatically after installing HP-UX .

Typically, the only requirement to enable remote installations is to modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. You must modify this file for each user who remotely accesses the system using `remsh`. Each line of the `.rhosts` file must contain a fully qualified domain name or IP address for each remote system that has access to the local system. For example, if the root user must remotely access `system1` from `system2`, add an entry for `system2.companyname.com` to the `.rhosts` file on `system1`.

```
# echo "system2.companyname.com" >> $HOME/.rhosts
```

After you complete an installation procedure, delete the `.rhosts` file from each user's `$HOME` directory to ensure security:

```
# rm -f $HOME/.rhosts
```

For more information on configuring the remote shell, see the operating system documentation and the `remsh(1M)` manual page.

Storage Foundation Cluster File System High Availability components

This appendix includes the following topics:

- [Veritas Storage Foundation Cluster File System High Availability installation depots](#)
- [Veritas Cluster Server installation depots](#)
- [Veritas Cluster File System installation depots](#)
- [Veritas Storage Foundation obsolete and reorganized installation depots](#)

Veritas Storage Foundation Cluster File System High Availability installation depots

[Table F-1](#) shows the depot name and contents for each English language depot for Veritas Storage Foundation Cluster File System High Availability. The table also gives you guidelines for which depots 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) depots, the combined functionality is called Veritas Storage Foundation Cluster File System High Availability and High Availability.

See [“Veritas Cluster Server installation depots”](#) on page 408.

Table F-1 Veritas Storage Foundation Cluster File System High Availability depots

depots	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 depots (*continued*)

depots	Contents	Configuration
VRTSsfpci60	<p>Veritas Storage Foundation Common Product Installer</p> <p>The Storage Foundation Common Product installer depot 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 depot 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 depots

[Table F-2](#) shows the depot name and contents for each English language depot for Veritas Cluster Server (VCS). The table also gives you guidelines for which depots to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and VCS depots, the combined functionality is called Storage Foundation and High Availability.

See [“Veritas Storage Foundation Cluster File System High Availability installation depots”](#) on page 405.

Table F-2 VCS installation depots

depot	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
VRTSvcsvcs	Veritas Cluster Server	Minimum
VRTSvcsvcsag	Veritas Cluster Server Bundled Agents	Minimum
VRTSvxfen	Veritas I/O Fencing	Minimum
VRTSvcsea	Consolidated database and enterprise agent depots	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 depots

[Table F-3](#) shows the depot name and contents for each English language depot for Veritas Cluster File System (CFS). The table also gives you guidelines for which depots to install based whether you want the minimum, recommended, or advanced configuration.

When you install all CFS depots and all the depots 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 depots”](#) on page 405.

See [“Veritas Cluster Server installation depots”](#) on page 408.

Table F-3 CFS installation depots

depot	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

Veritas Storage Foundation obsolete and reorganized installation depots

[Table F-4](#) lists the depots that are obsolete or reorganized for Veritas Storage Foundation Cluster File System High Availability.

Table F-4 Veritas Storage Foundation obsolete and reorganized depots

depot	Description
Obsolete and reorganized for 6.0	
VRTSat	Obsolete
Obsolete and reorganized for 5.1	
Infrastructure	

Table F-4 Veritas Storage Foundation obsolete and reorganized depots
(continued)

depot	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
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product depots	
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, VRTSacclib is not uninstalled. ■ For uninstalls, VRTSacclib is not uninstalled.
VRTSalloc	Obsolete
VRTScmccc	Obsolete
VRTScmcs	Obsolete
VRTScscm	Obsolete

Table F-4 Veritas Storage Foundation obsolete and reorganized depots
(continued)

depot	Description
VRTScscw	Obsolete
VRTScsocw	Obsolete
VRTScssim	Obsolete
VRTScutil	Obsolete
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSdsa	Obsolete
VRTSfsman	Included in the product's main depot.
VRTSfsmnd	Included in the product's main depot.
VRTSfspro	Included in VRTSsfmh
VRTSvcldb	Included in VRTSvcsea
VRTSvcsor	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvcsc
VRTSvdid	Obsolete
VRTSvmman	Included in the product's main depot.
VRTSvmpro	Included in VRTSsfmh
VRTSvrpro	Included in VRTSob
VRTSvrw	Obsolete
VRTSvxmsa	Obsolete
Documentation	All Documentation depots 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](#)
- [CFMount 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
- CFSSMount 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 CFSSMount 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 419.

See [“Attribute definition for CFSSMount agent”](#) on page 425.

See [“Attribute definition for CFSfsckd agent”](#) on page 429.

- 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/rc.config.d/amf
```

- 2 Start the AMF kernel driver. Run the following command:

```
/sbin/init.d/amf start
```

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/rc.config.d/amf
```

- 2 Stop the AMF kernel driver. Run the following command:

```
/sbin/init.d/amf stop
```

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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> 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. <ul style="list-style-type: none"> Default: 3. ■ Type and dimension: integer-association <p>See “Enabling and disabling intelligent resource monitoring for agents manually” on page 413.</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
Primary (Not set by user)	<p>Information only. Stores the primary node name for a VxFS file system. The value is automatically modified in the configuration file when an unmounted file system is mounted or another node becomes the primary node. The user does not set this attribute and user programs do not rely on this attribute.</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.

Table G-9 CFSfsckd agent entry points (*continued*)

Entry Points	Description
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.
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 413.</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 vxfscd (  
)
```

Troubleshooting information

This appendix includes the following topics:

- [Restarting the installer after a failed connection](#)
- [What to do if you see a licensing reminder](#)
- [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 53. 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>

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.

```
Checking ssh communication with system01 ..... permission denied
installer requires that ssh commands used between systems execute without
prompting for passwords or confirmations. Please run installer again with
the ssh configured for password free logins, or configure rsh and use the
-rsh option.
```

```
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 399.

Note: Remove remote shell permissions after completing the SFCFSHA installation and configuration.

Resource temporarily unavailable

If the installation fails with the following error message on the console:

```
fork() failed: Resource temporarily unavailable
```

The value of `nkthread` tunable parameter may not be large enough. The `nkthread` tunable requires a minimum value of 600 on all systems in the cluster. To determine the current value of `nkthread`, enter:

```
# kctune -q nkthread
```

If necessary, you can change the value of `nkthread` using the SAM (System Administration Manager) interface, or by running the `kctune` command. If you change the value of `nkthread`, the kernel must be rebuilt for the new value to take effect. It is easier to change the value using SAM because there is an option to process the new kernel immediately.

See the `kctune(1M)` and `sam(1M)` manual pages.

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.

```
Checking communication with system01 ..... FAILED
  System not accessible : system01

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.

If a system cannot access the software source depot, either `swagentd` is not running on the target system or the `swlist` command cannot see the source depot.

```
Correct /etc/{hosts, nsswitch.conf} and continue from here
Continue? [Y/N] :
```

Suggested solutions: check that `swagentd` is running. Check whether there is an entry for the target system in `/etc/hosts`. If there is no entry, then ensure the `hosts` file is not the primary lookup for the "hosts" entry.

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 30.
- 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/littab` 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 vxfsentsthdw 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 442.

See [“Checking the connectivity of CP server”](#) on page 442.

See [“Issues during fencing startup on SFCFSHA cluster nodes set up for server-based fencing”](#) on page 443.

See [“Issues during online migration of coordination points”](#) on page 444.

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 159.</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:

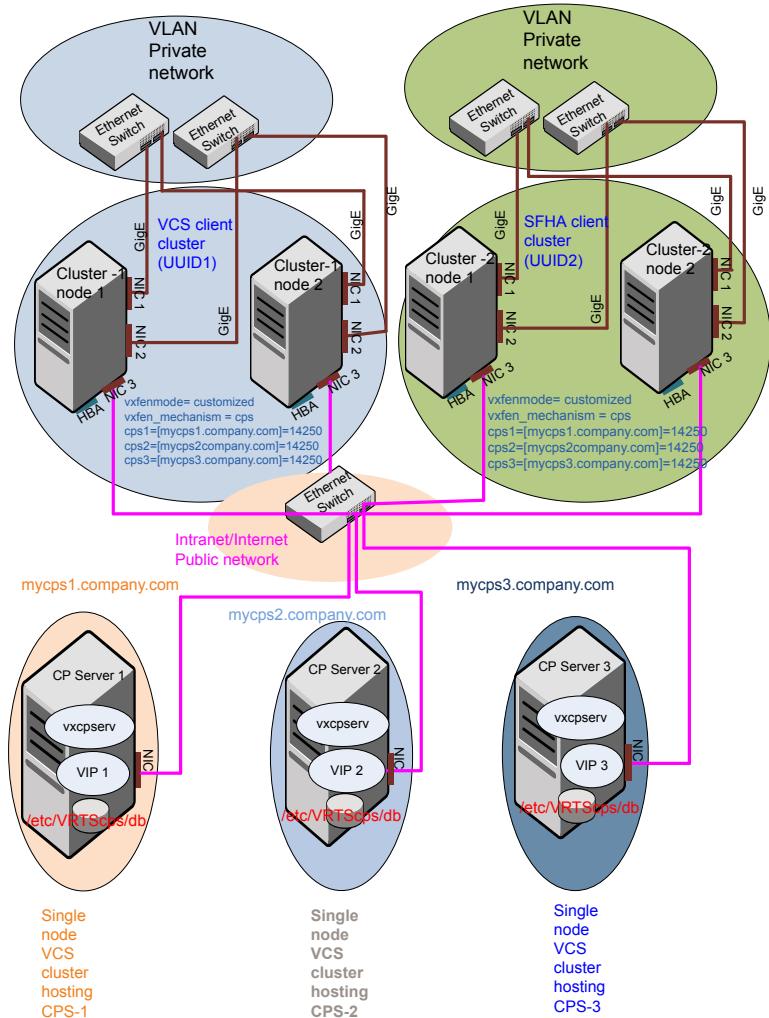
- Two unique client clusters that are served by 3 CP servers:
See [Figure J-1](#) on page 446.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
[Figure J-2](#)
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
[Figure J-3](#)
- Multiple 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



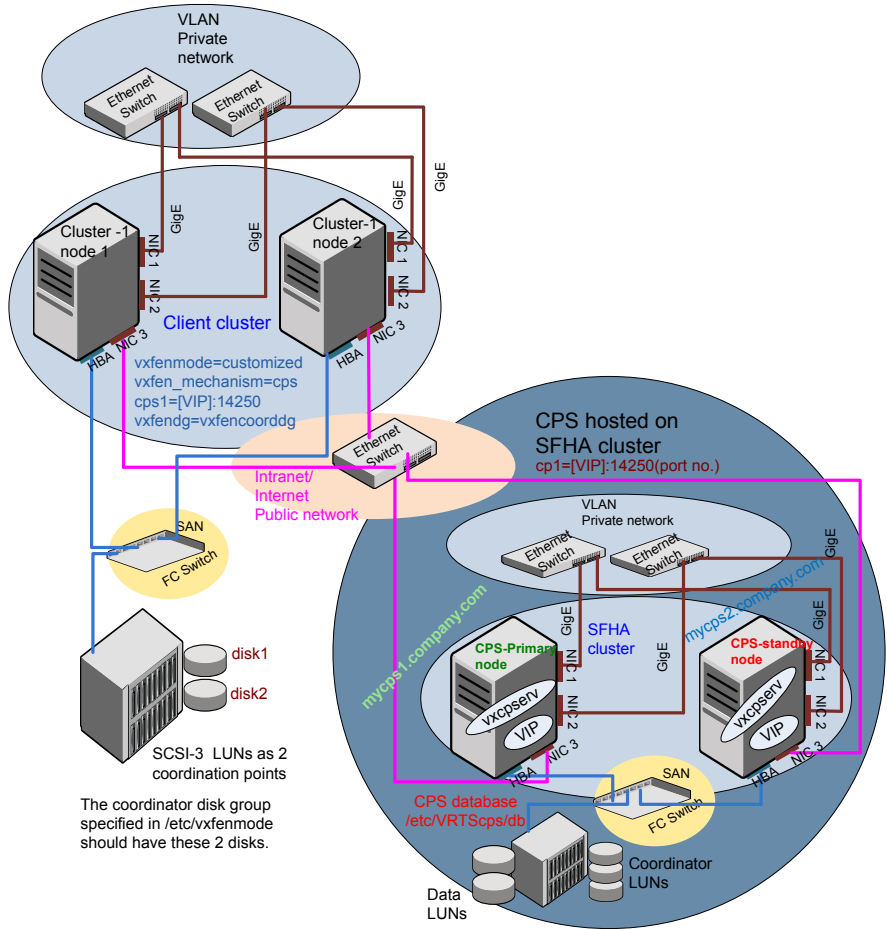
Client cluster served by highly available CPS and 2 SCSI-3 disks

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 `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 `vxfencoordg`. 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 vxencoordg.
 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



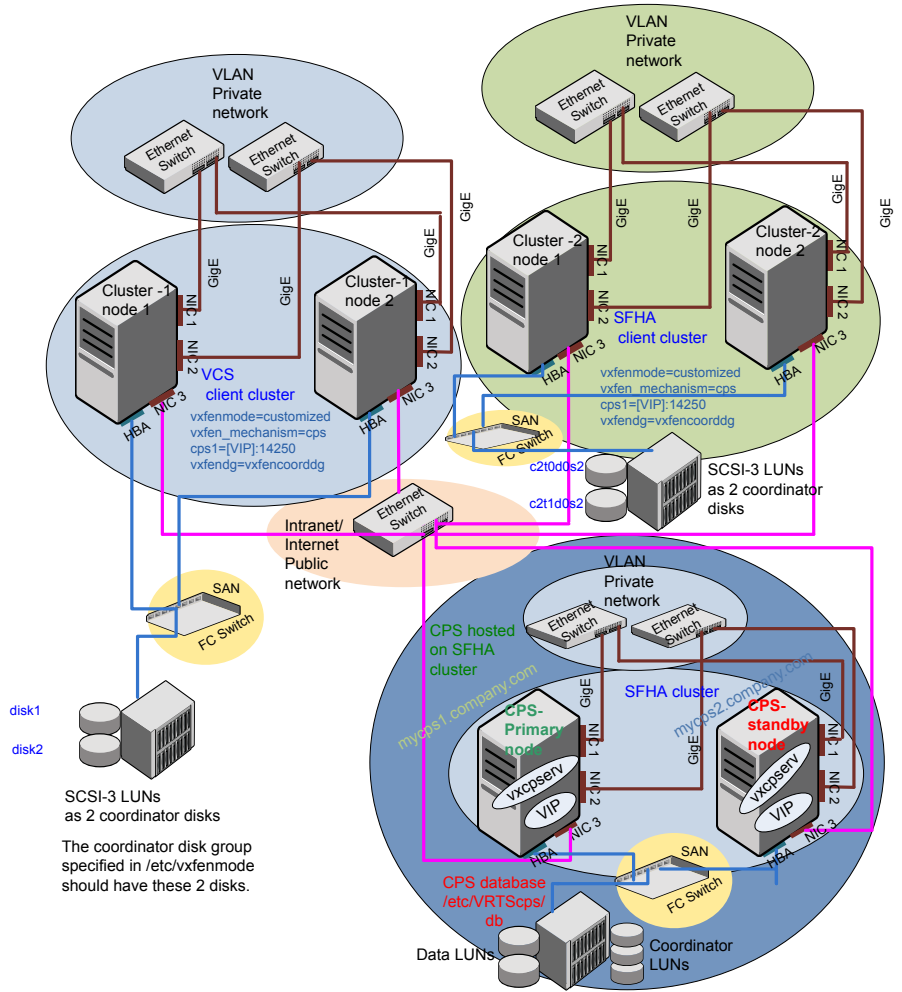
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

[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/c1t1d0`.

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 HP-UX 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/clt1d0
```

Output on Node A resembles:

```
crw-r----- 1 root sys 32,1 Dec 3 11:50 /dev/dsk/clt1d0
```

Output on Node B resembles:

```
crw-r----- 1 root sys 32,1 Dec 3 11:55 /dev/dsk/clt1d0
```

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/clt1d0
```

Output on Node B:

```
crw-r----- 1 root sys 36,1 Dec 3 11:55 /dev/dsk/clt1d0
```

- 2 Place the VCS command directory in your path.

- 3 Attempt to change the major number on System B (now 36) to match that of System A (32). Use the command:

```
# haremajord -sd major_number
```

For example, on Node B, enter:

```
# haremajord -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 `haremajord` command on Node B and change it to 128,

```
# haremajord -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.

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 Determine the major numbers for `vxio` and `vxspec` that Veritas Volume Manager uses on each node that is exporting an NFS file system.
- 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 459.
- 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 461.

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 460.
- See [“Sample configuration: links crossing IP routers”](#) on page 461.

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 459.
<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 461.

[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
<code>node-id</code>	The ID of the cluster node; for example, 0.
<code>link tag-name</code>	The string that LLT uses to identify the link; for example link1, link2,....
<code>address</code>	IPv6 address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file `/etc/services`. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -a | head -2 ; netstat -a | grep udp
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address Foreign Address (state)
udp 0 0 *.ntalk *.*
udp 0 0 *.* *.*
udp 0 0 *.49193 *.*
udp 0 0 *.49152 *.*
udp 0 0 *.portmap *.*
udp 0 0 *.* *.*
udp 0 0 *.135 *.*
udp 0 0 *.2121 *.*
udp 0 0 *.xdmcp *.*
```

```

udp      0      0 *.49196          *. *
udp      0      0 *.*              *. *
udp      0      0 *.snmp           *. *
udp      0      0 *.*              *. *
udp      0      0 *.49153          *. *
udp      0      0 *.echo           *. *
udp      0      0 *.discard        *. *
udp      0      0 *.daytime        *. *
udp      0      0 *.chargen        *. *
udp      0      0 *.syslog         *. *


```

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 interface_name` command to verify that the IPv6 address is set correctly.

You can also use the `lanscan` command to verify the IPv6 address.

```
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 465.

- 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 467.
- Set the broadcast address correctly for direct-attached (non-routed) links. See [“Sample configuration: direct-attached links”](#) on page 469.
- 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 471.

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.

```
galaxy # ifconfig lan1
lan1: flags=1843<UP,BROADCAST,RUNNING,MULTICAST,CKO>
       inet 192.168.9.1 netmask ffffffff broadcast 192.168.9.255
galaxy # ifconfig lan2
lan2: flags=1843<UP,BROADCAST,RUNNING,MULTICAST,CKO>
       inet 192.168.10.1 netmask ffffffff broadcast 192.168.10.255
```

- 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.

```
nebula # ifconfig lan1
lan1: flags=1843<UP,BROADCAST,RUNNING,MULTICAST,CKO>
       inet 192.168.9.2 netmask ffffffff broadcast 192.168.9.255
nebula # ifconfig lan2
lan2: flags=1843<UP,BROADCAST,RUNNING,MULTICAST,CKO>
       inet 192.168.10.2 netmask ffffffff broadcast 192.168.10.255
```

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 469.
- See “[Sample configuration: links crossing IP routers](#)” on page 471.

[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 <code>/dev/udp</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 "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 467.
<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.

Table M-1 Field description for link command in `/etc/llttab` (*continued*)

Field	Description
<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 471.

[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.
<i>link tag-name</i>	The string that LLT uses to identify the link; for example link1, link2,....
<i>address</i>	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file `/etc/services`. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -a | head -2 ; netstat -a | grep udp
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address Foreign Address (state)
udp      0      0 *.ntalk          *.*
```

```

udp      0      0  *.*                               *.*
udp      0      0  *.49193                           *.*
udp      0      0  *.49152                           *.*
udp      0      0  *.portmap                          *.*
udp      0      0  *.*                               *.*
udp      0      0  *.135                             *.*
udp      0      0  *.2121                             *.*
udp      0      0  *.xdmcp                            *.*
udp      0      0  *.49196                            *.*
udp      0      0  *.*                               *.*
udp      0      0  *.snmp                             *.*
udp      0      0  *.*                               *.*
udp      0      0  *.49153                            *.*
udp      0      0  *.echo                             *.*
udp      0      0  *.discard                          *.*
udp      0      0  *.daytime                          *.*
udp      0      0  *.chargen                          *.*
udp      0      0  *.syslog                           *.*

```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# set_parms ip_address
```

For example:

- For the first network interface on the node galaxy:

```
IP address=192.168.9.1, Broadcast address=192.168.9.255,
Netmask=255.255.255.0
```

For the first network interface on the node nebula:

```
IP address=192.168.9.2, Broadcast address=192.168.9.255,
Netmask=255.255.255.0
```

- For the second network interface on the node galaxy:

```
IP address=192.168.10.1, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

For the second network interface on the node nebula:

```
IP address=192.168.10.2, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.

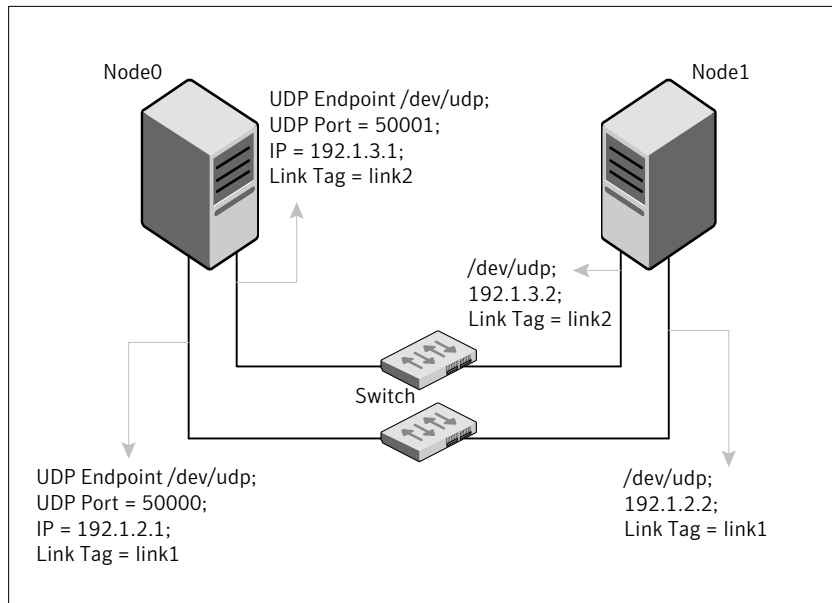
An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100  
  
link link1 /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 interface_name` 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:

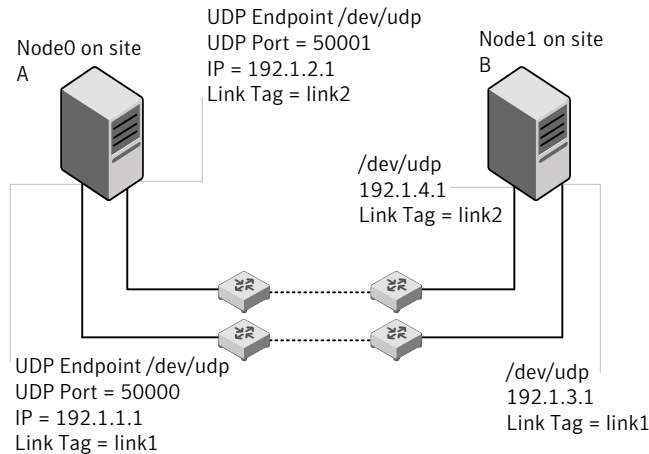
```
set-node Node1
set-cluster 1
```

```
#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 depots 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 VRTSspbx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspbx, VRTSicsco, and VRTSat.

Index

A

- adding
 - users 117
- agents
 - about 412
 - CFSfsckd 428
 - CFSMount 424, 428
 - CVMCluster 416
 - CVMVolDg 421
 - CVMVxconfigd 418
 - disabling 342
 - of VCS 413
- application
 - database replication 321
- applications, stopping 192
- attributes
 - about agent attributes 412
 - CFSMount agent 425, 429
 - CVMCluster agent 417
 - CVMVolDg agent 417, 422
 - CVMVxconfigd agent 419
 - UseFence 151

B

- block device
 - partitions
 - example file name 452
 - volumes
 - example file name 452

C

- cables
 - cross-over Ethernet 262
 - for SCSI devices 33
- CFS
 - mount and unmount failures 436
 - synchronization 245
 - troubleshooting 436
- CFSfsckd agent 428
 - attributes 429

- CFSMount agent 424, 428
 - attributes 425
 - entry points 425
 - sample configuration 427–428
 - type definition 427
- CFSTypes.cf 427
- cluster
 - removing a node from 290
 - verifying operation 255
- cluster functionality
 - enabling 136
 - environment requirements 36
 - shared disks 137
- command failures 437
- commands
 - gabconfig 253
 - gcoconfig 307
 - hastatus 255
 - hasys 255
 - lltconfig 248
 - lltstat 251
 - vradmin 322
 - vxassist 314, 317
 - vxdisksetup (initializing disks) 140
 - vxlicinst 125
 - vxlicrep 124
 - vxvol 314
- configuration daemon (vxconfigd)
 - starting 134
- configuration file
 - main.cf 246
- configuring
 - rsh 32
 - shared disks 137
 - ssh 32
- configuring Storage Foundation Cluster File System High Availability
 - script-based installer 104
- configuring VCS
 - adding users 117
 - event notification 118, 120
 - global clusters 122

configuring VCS *(continued)*

- secure mode 113
- starting 105

coordinator disks

- DMP devices 76
- for I/O fencing 76
- setting up 149

creating SD bundle 69

CVM

- CVMTypes.cf file 418

CVMCluster agent 416

- attributes 417
- entry points 416
- sample configuration 418
- type definition 418

CVMTypes.cf

- definition, CVMCluster agent 418
- definition, CVMVolDg agent 423
- definition, CVMVxconfigd agent 421

CVMVolDg agent 421

- attributes 422
- entry points 421
- sample configuration 424
- type definition 423

CVMVxconfigd agent 418

- attributes 419
- CVMTypes.cf 421
- entry points 419
- sample configuration 421
- type definition 421

D

data disks

- for I/O fencing 76

disabling the agents 342

disk groups

- rootdg 131

disks

- adding and initializing 140
- coordinator 149
- testing with vxfcntlshdw 144
- verifying node access 146

E

Ethernet controllers 262

F

Fibre Channel fabric 38

files

- main.cf 246
- freezing service groups 192

G

GAB

- port membership information 253
- verifying 253

gabconfig command 253

- a (verifying GAB) 253

gabtab file

- verifying after installation 248

gcoconfig command 307

global cluster

- return of primary role to original primary site 336
- takeover of primary site role by secondary 336
- troubleshooting VVR component 339

global clusters

- configuration 122
 - using VVR 306
- configuring VCS service groups 306
- illustration of dependencies 323
- migration of primary to secondary site 335
- replicating database volumes 323
- requirements 300
- return migration to original primary site 335

groups

- log owner 323
- RVG 323

H

hastatus -summary command 255

hasys -display command 255

high availability issues 438

- low memory 438
- network partition 438

hubs

- independent 262

I

I/O daemon (vxiod)

- starting 135

I/O fencing

- checking disks 144
- setting up 148
- shared storage 144

I/O fencing requirements

- non-SCSI-3 46

- Ignite
 - installing 69
- Ignite-UX 72
 - installing standalone 71
- Installing
 - SFCFSHA with the Web-based installer 66
- installing
 - Ignite 69
 - Ignite-UX 71–72
 - post 123
 - standalone 71
- intelligent resource monitoring
 - disabling manually 413
 - enabling manually 413

J

- jeopardy 438

K

- kctune command 435

L

- license keys
 - adding with vxlicinst 125
 - replacing demo key 125
- licenses
 - information about 124
- links
 - private network 248
- LLT
 - interconnects 31
 - verifying 251
- litconfig command 248
- llthosts file
 - verifying after installation 248
- lltstat command 251
- llttab file
 - verifying after installation 248
- log files 441

M

- main.cf file 246
- major and minor numbers
 - checking 453, 455
 - shared devices 452
- manual pages
 - potential problems 437
 - troubleshooting 437

- media speed 31
 - optimizing 31
- membership information 253
- mount command
 - potential problems 436
- mounting
 - software disc 57

N

- network partition 438
- NFS services
 - shared storage 452
- nodes
 - adding application nodes
 - configuring GAB 267, 278
 - configuring LLT 267, 278
 - configuring VXFEN 267, 278
 - starting Volume Manager 266, 277
 - adding new nodes
 - configuring GAB 269
 - configuring LLT 269
 - configuring VXFEN 269
 - preparing application nodes
 - configuring CVM 274, 284
 - removing nodes
 - GAB configuration 292
 - LLT configuration 291
 - removing nodes from VCS configuration files 291
 - removing Oracle nodes
 - modifying VCS configuration 292
 - tasks for removing nodes 289
- non-SCSI-3 fencing
 - manual configuration 170
 - setting up 170
- non-SCSI-3 I/O fencing
 - requirements 46
- non-SCSI3 fencing
 - setting up 153
 - using installsfcsfsha 153
- NTP
 - network time protocol daemon 245

O

- optimizing
 - media speed 31

P

- PATH variable
 - VCS commands 251
- persistent reservations
 - SCSI-3 32
- planning to upgrade VVR 188
- port a
 - membership 253
- port h
 - membership 253
- port membership information 253
- preinstallation 188
- preparing to upgrade VVR 192
- primary site
 - creating SRL volume 313
 - setting up 302
 - setting up replication objects 315
 - VCS configuration 325, 329
- problems
 - accessing manual pages 437
 - executing file system commands 437
 - mounting and unmounting file systems 436

Q

- Quick I/O
 - performance on CFS 438

R

- removing
 - the Replicated Data Set 343
- removing a node from a cluster 290
- remsh 106
- Replicated Data Set
 - removing the 343
- replication
 - automatic synchronization 321
 - configuring on both sites 306
 - full synchronization with Storage Checkpoint 322
 - modifying VCS configuration 323
 - options 302
 - setting up primary site 302
 - setting up secondary site 304
 - supported hardware 300
 - supported software 300
 - using VVR 313
 - verifying status 322
- resources
 - CVMVolDg 323

- resources *(continued)*
 - RVGSharedPri 323
- rolling upgrade 207, 210
 - versions 207
- root disk group 131
- rsh
 - configuration 32

S

- sam command 435
- SAN
 - see Storage Area Network 38
- script-based installer
 - Storage Foundation Cluster File System High Availability configuration overview 104
- SCSI
 - changing initiator IDs 33
- SCSI-3
 - persistent reservations 32
- SCSI-3 persistent reservations
 - verifying 148
- SD bundle 69
- secondary site
 - configuring replication 316
 - creating SRL volume 317
 - setting up 304
 - setting up disk groups 319
 - setting up RLINKs 318
- service groups
 - freezing 192
 - VCS, for global clusters 306
- SFCFSHA
 - coordinator disks 149
 - illustration of global cluster dependencies 323
- SFCFSHA installation
 - verifying
 - cluster operations 251
 - GAB operations 251
 - LLT operations 251
- SFCFSHA upgrade
 - preparation 196
 - preparing 197
- shared disks, configuring 137
- Shared storage
 - Fibre Channel 35
- shared storage 33
 - NFS services 452
 - SCSI 33
- SMTP email notification 118

- SNMP trap notification 120
- split brain 438
- ssh 106
 - configuration 32
- starting configuration
 - installvcs program 106
 - Veritas product installer 106
- starting vxconfigd configuration daemon 134
- starting vxiod daemon 135
- stopping
 - applications 192
- Storage Area Network 38
- Storage Foundation Cluster File System High Availability
 - configuring 104
- Symantec Product Authentication Service 113
- system state attribute value 255

T

- troubleshooting
 - accessing manual pages 437
 - executing file system commands 437
 - mounting and unmounting file systems 436
- tunables file
 - about setting parameters 379
 - parameter definitions 384
 - preparing 383
 - setting for configuration 380
 - setting for installation 380
 - setting for upgrade 380
 - setting parameters 383
 - setting with no other operations 381
 - setting with un-integrated response file 382

U

- upgrading
 - rolling 207
- upgrading VVR
 - from 4.1 188
 - planning 188
 - preparing 192

V

- VCS
 - command directory path variable 251
 - configuration, for database volume replication 323
 - configuring service groups 306

- VCS configuration
 - for replication 323
- VCS Global cluster option. *See* GCO
- verifying installation
 - kernel component 244
- Veritas Operations Manager 24
- Veritas Volume Replicator. *See* VVR
- vradmin
 - delpri 344
 - stoprep 343
- vradmin command 322
- VVR
 - application database replication 321
 - configuring global clusters 306
 - configuring on both sites 306
 - database volume replication
 - configuring VCS 323
 - defining heartbeat cluster objects 308
 - defining remote clusters 308
 - global cluster overview 312
 - primary site
 - creating SRL volume 313
 - setting up replication objects 315
 - replication
 - using automatic synchronization 321
 - using full synchronization with Storage Checkpoint 322
 - replication agents 302
 - secondary site
 - configuring replication 316
 - creating SRL volume 317
 - setting up disk groups 319
 - setting up RLINKs 318
 - setting up primary site 302
 - setting up replication 313
 - setting up secondary site 304
 - types of replication 302
 - VCS configuration
 - application database service group 323
 - CVMOLDG resource 323
 - log owner group 323
 - primary site 325
 - RVG group 323
 - RVGSharedPri resource 323
 - secondary site 329
 - verifying replication 322
- VVR 4.1
 - planning an upgrade from 188
- vxassist command 314, 317

- vxconfigd configuration daemon
 - starting 134
- vxctl mode command 135
- vxdisksetup command 140
- vxiod I/O daemon
 - starting 135
- vxlicinst command 125
- vxlicrep command 124
- vxvol command 314

W

- Web-based installer 66