

# Veritas Storage Foundation™ Cluster File System Installation Guide

AIX

5.1



# Veritas Storage Foundation™ Cluster File System Installation Guide

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

Product version: 5.1

Document version: 5.1.1

## Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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

This Symantec product may contain third party software for which Symantec is required to provide attribution to the third party ("Third Party Programs"). Some of the Third Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this Symantec product for more information on the Third Party Programs.

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  
20330 Stevens Creek Blvd.  
Cupertino, CA 95014

<http://www.symantec.com>

# Technical Support

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

Symantec's maintenance offerings include the following:

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

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

[www.symantec.com/techsupp/](http://www.symantec.com/techsupp/)

## Contacting Technical Support

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

[www.symantec.com/business/support/index.jsp](http://www.symantec.com/business/support/index.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/techsupp/](http://www.symantec.com/techsupp/)

## Customer service

Customer service information is available at the following URL:

[www.symantec.com/techsupp/](http://www.symantec.com/techsupp/)

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

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

## Maintenance agreement resources

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

Asia-Pacific and Japan	<a href="mailto:contractsadmin@symantec.com">contractsadmin@symantec.com</a>
Europe, Middle-East, and Africa	<a href="mailto:semea@symantec.com">semea@symantec.com</a>
North America and Latin America	<a href="mailto:supportsolutions@symantec.com">supportolutions@symantec.com</a>

## Additional enterprise services

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

Enterprise services that are available include the following:

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

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

[www.symantec.com](http://www.symantec.com)

Select your country or language from the site index.

# Contents

Technical Support .....	4
Chapter 1	About Storage Foundation Cluster File System ..... 19
	Veritas Storage Foundation Cluster File System suites ..... 19
	About I/O fencing ..... 20
	About Veritas product licensing ..... 21
Chapter 2	Before you install ..... 23
	About planning for a SFCFS installation ..... 23
	About installation and configuration methods ..... 24
	Assessing your system preparedness ..... 25
	Veritas Operations Services ..... 25
	About the installation simulator ..... 25
	Preinstallation or upgrade planning for Veritas Volume
	Replicator ..... 26
	Planning an upgrade from the previous VVR version ..... 27
	Downloading the Storage Foundation Cluster File System
	software ..... 28
	Setting environment variables ..... 28
	Configuring secure shell (ssh) or remote shell before installing
	products ..... 29
	Configuring and enabling ssh ..... 29
	Restarting ssh ..... 34
	Enabling rsh for AIX ..... 34
	Setting up shared storage ..... 35
	Setting the SCSI identifier value ..... 35
	Setting up Fiber Channel ..... 36
	Prerequisites for Storage Foundation Cluster File System ..... 37
	Hardware overview and requirements for Storage Foundation Cluster
	File System ..... 38
	Shared storage ..... 39
	Fibre Channel switch ..... 39
	Cluster platforms ..... 39

Chapter 3	System requirements .....	41
	Hardware and software requirements .....	41
	I/O fencing requirements .....	41
	Coordinator disk requirements for I/O fencing .....	42
	CP server requirements .....	42
	Release notes .....	44
	Supported AIX operating systems .....	44
	Storage Foundation Cluster File System node requirements .....	45
	Database requirements .....	45
	Disk space requirements .....	46
Chapter 4	Installing Storage Foundation Cluster File System using the common product installer .....	47
	Installation quick reference .....	47
	Mounting a software disc .....	48
	About the common product installer .....	49
	Installing Storage Foundation Cluster File System with the product installer .....	49
Chapter 5	Installing Storage Foundation Cluster File System using the web-based installer .....	55
	About the Web-based installer .....	55
	Features supported with Web-based installer .....	56
	Before using the Veritas Web-based installer .....	56
	Starting the Veritas Web-based installer .....	57
	Obtaining a security exception on Mozilla Firefox .....	57
	Performing a pre-installation check with the Veritas Web-based installer .....	58
	Installing SFCFS with the Veritas Web-based installer .....	58
Chapter 6	Installing Storage Foundation Cluster File System, other methods .....	61
	Installing SFCFS using NIM and the installer .....	61
	Preparing the bundle and script resources on NIM server .....	61
	Installing SFCFS on the NIM client using SMIT .....	63
	Installing SFCFS and the operating system on the NIM client using SMIT .....	63
	Installing SFCFS on an alternate disk .....	64
	Preparing to install SFCFS on an alternate disk .....	64
	Installing SFCFS on an alternate disk .....	65

	Verifying the installation .....	68
Chapter 7	Preparing to configure Storage Foundation Cluster File System .....	69
	Preparing to configure the clusters in secure mode .....	69
	Installing the root broker for the security infrastructure .....	73
	Creating authentication broker accounts on root broker system .....	74
	Creating encrypted files for the security infrastructure .....	75
	Preparing the installation system for the security infrastructure .....	77
	About configuring SFCFS clusters for data integrity .....	78
	About I/O fencing components .....	79
	About data disks .....	79
	About coordination points .....	79
	About I/O fencing configuration files .....	80
	About planning to configure I/O fencing .....	83
	Typical SFCFS cluster configuration with server-based I/O fencing .....	85
	Recommended CP server configurations .....	85
	Setting up the CP server .....	87
	Installing the CP server using the installer .....	88
	Configuring security on the CP server .....	89
	Setting up shared storage for the CP server database .....	90
	Configuring the CP server using the configuration utility .....	90
	Configuring the CP server manually .....	97
	Verifying the CP server configuration .....	98
Chapter 8	Configuring Storage Foundation Cluster File System .....	101
	Configuring the products using the common product installer .....	101
	Configuring Veritas Volume Manager .....	102
	Configuring Veritas File System .....	102
	Configuring Storage Foundation Cluster File System .....	102
	Configuring the SFDB repository database .....	106
	Setting or changing the product level for keyless licensing .....	107
	Installing Veritas product license keys .....	108

Chapter 9	Configuring Storage Foundation Cluster File System for data integrity .....	109
	Setting up disk-based I/O fencing using installsfcfs .....	109
	Initializing disks as VxVM disks .....	109
	Configuring disk-based I/O fencing using installsfcfs .....	110
	Checking shared disks for I/O fencing .....	113
	Setting up disk-based I/O fencing manually .....	117
	Identifying disks to use as coordinator disks .....	118
	Setting up coordinator disk groups .....	118
	Creating I/O fencing configuration files .....	119
	Modifying VCS configuration to use I/O fencing .....	120
	Verifying I/O fencing configuration .....	122
	Setting up server-based I/O fencing using installsfcfs .....	122
	Verifying security configuration on SFCFS cluster to use CP server coordination point .....	123
	Configuring server-based I/O fencing .....	125
	Setting up server-based I/O fencing manually .....	133
	Preparing the CP servers manually for use by the SFCFS cluster .....	133
	Configuring server-based fencing on the SFCFS cluster manually .....	137
	Configuring Coordination Point agent to monitor coordination points .....	141
	Verifying server-based I/O fencing configuration .....	143
Chapter 10	Upgrading Storage Foundation Cluster File System .....	145
	About upgrading Storage Foundation Cluster File System and High Availability .....	145
	Upgrading Storage Foundation Cluster File System and High Availability software from a release prior to 5.1 .....	146
	Upgrade paths for Veritas Storage Foundation Cluster File System 5.1 .....	146
	Planning the upgrade .....	146
	Preparing the system and backing up files before upgrading .....	147
	Preparing for upgrade of VVR in the presence of VCS agents .....	148
	Overview of procedures .....	151
	Performing a phased upgrade of SFCFSHA stack from version 5.0MP3 .....	152

	Performing a full upgrade .....	163
	Making the file systems clean .....	169
	Upgrading Veritas Volume Replicator .....	171
	Upgrading VVR without disrupting replication .....	172
	Upgrading from SFHA 5.1 to SFCFS 5.1 .....	173
Chapter 11	Upgrading SFCFS using an alternate disk .....	175
	About upgrading SFCFS using an alternate disk .....	175
	Supported upgrade scenarios .....	176
	Supported upgrade paths .....	176
	Preparing to upgrade SFCFS on an alternate disk .....	176
	Upgrading SFCFS on an alternate disk .....	178
	Upgrading a cluster that is in secure mode .....	185
	Verifying the upgrade .....	187
	Verifying that the cluster is in secure mode .....	187
Chapter 12	Verifying the Storage Foundation Cluster File System installation .....	189
	Verifying that the products were installed .....	189
	Installation log files .....	190
	Using the installation log file .....	190
	Using the summary file .....	190
	About enabling LDAP authentication for clusters that run in secure mode .....	190
	Enabling LDAP authentication for clusters that run in secure mode .....	192
	Starting and stopping processes for the Veritas products .....	198
	Checking Volume Manager processes .....	198
	Checking Veritas File System installation .....	199
	Verifying agent configuration for Storage Foundation Cluster File System .....	199
	Synchronizing time on Cluster File Systems .....	200
	Configuring VCS for Storage Foundation Cluster File System .....	200
	main.cf file .....	201
	Storage Foundation Cluster File System HA Only .....	202
	Veritas Cluster Server application failover services .....	202
	About the LLT and GAB configuration files .....	202
	Verifying the LLT, GAB, and VCS configuration files .....	205
	Verifying LLT, GAB, and cluster operation .....	205
	Verifying LLT .....	205
	Verifying GAB .....	208
	Verifying the cluster .....	209

	Verifying the cluster nodes .....	210
Chapter 13	Adding a node to Storage Foundation Cluster File System clusters .....	213
	About adding a node to an SFCFS cluster .....	213
	Preparing to add a node to an existing SFCFS cluster .....	214
	Enabling the DLPI driver .....	214
	Adding nodes to an existing SFCFS cluster .....	215
	Adding nodes to an existing SFCFS cluster using the SFCFS installer .....	216
	Adding the node to the SFCFS cluster manually .....	218
	Updating the Storage Foundation for Databases (SFDB) repository after adding a node .....	230
	Sample configuration file for adding a node to the cluster .....	231
Chapter 14	Removing a node from Storage Foundation Cluster File System clusters .....	237
	Removing nodes from an SFCFS cluster .....	237
	Modifying the VCS configuration files on existing nodes .....	238
	Editing the /etc/llhosts file .....	238
	Editing the /etc/gabtab file .....	239
	Modifying the VCS configuration to remove the node .....	239
	Removing the node configuration from the CP server .....	241
	Removing security credentials from the leaving node .....	242
	Updating the Storage Foundation for Databases (SFDB) repository after removing a node .....	242
	Sample configuration file for removing a node from the cluster .....	242
Chapter 15	Setting up a replicated global cluster .....	247
	Replication in the SFCFS environment .....	247
	Requirements for SFCFS global clusters .....	248
	Supported software and hardware for SFCFS .....	248
	Supported replication technologies for SFCFS .....	248
	About setting up a global cluster in an SFCFS environment .....	250
	Configuring an SFCFS cluster at the primary site .....	251
	Configuring an SFCFS cluster at the secondary site .....	252
	Setting up the cluster on the secondary site .....	253
	Setting up the database for the secondary site .....	254
	Configuring replication on clusters at both sites .....	254
	Modifying the ClusterService group for global clusters .....	255

	Modifying the global clustering configuration using the wizard .....	255
	Defining the remote cluster and heartbeat objects .....	256
	Configuring the VCS service groups for global clusters .....	260
Chapter 16	Configuring a global cluster using VVR .....	261
	About configuring global clustering using VVR .....	261
	Setting up replication using VVR on the primary site .....	262
	Creating the SRL volume on the primary site .....	262
	Setting up the Replicated Volume Group (RVG) on the primary site .....	263
	Setting up replication using VVR on the secondary site .....	265
	Creating the data and SRL volumes on the secondary site .....	265
	Editing the /etc/vx/vras/.rdg files .....	266
	Setting up IP addresses for RLINKs on each cluster .....	266
	Setting up the disk group on secondary site for replication .....	267
	Starting replication of application database volume .....	269
	Starting replication using automatic synchronization .....	270
	Starting replication using full synchronization with Checkpoint .....	270
	Verifying replication status .....	271
	Configuring VCS to replicate the database volume using VVR .....	272
	About modifying the VCS configuration for replication .....	272
	Modifying the VCS Configuration on the Primary Site .....	273
	Modifying the VCS Configuration on the Secondary Site .....	277
	Using VCS commands on SFCFS global clusters .....	281
	Using VVR commands on SFCFS global clusters .....	282
	About migration and takeover of the primary site role .....	282
	Migrating the role of primary site to the secondary site .....	282
	Migrating the role of new primary site back to the original primary site .....	283
	Taking over the primary role by the remote cluster .....	284
	VCS agents to manage wide-area failover .....	288
Chapter 17	Uninstalling Storage Foundation Cluster File System .....	289
	Preparing to uninstall a Storage Foundation product .....	290
	About removing Veritas Storage Foundation .....	291
	Shutting down cluster operations .....	291
	Moving volumes to physical disks .....	291
	Disabling the agents on a system .....	293
	Removing the Replicated Data Set .....	294

	Uninstalling SFCFS with the Veritas Web-based installer .....	295
	Uninstalling SFCFS filesets using the script-based installer .....	296
	Removing Storage Foundation products using SMIT .....	298
	Uninstalling Storage Foundation Cluster File System .....	298
	Removing the CP server configuration using the removal script .....	299
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product .....	302
Appendix A	Installation scripts .....	305
	About installation scripts .....	305
	Installation script options .....	306
Appendix B	Response files .....	311
	About response files .....	311
	Installing SFCFS using response files .....	312
	Configuring SFCFS using response files .....	312
	Upgrading SFCFS using response files .....	313
	Uninstalling SFCFS using response files .....	313
	Syntax in the response file .....	314
	Response file variable definitions .....	314
	Sample response file for SFCFS install .....	320
	Sample response file for SFCFS configure .....	321
Appendix C	Configuring I/O fencing using a response file .....	323
	Response file variables to configure disk-based I/O fencing .....	323
	Sample response file for configuring disk-based I/O fencing .....	324
	Configuring I/O fencing using response files .....	325
	Response file variables to configure server-based I/O fencing .....	326
	Sample response file for configuring server-based I/O fencing .....	328
Appendix D	Storage Foundation Cluster File System components .....	329
	Veritas Storage Foundation installation filesets .....	329
	Veritas Cluster Server installation filesets .....	331
	Veritas Cluster File System installation filesets .....	332
	Veritas Storage Foundation obsolete and reorganized installation filesets .....	333

Appendix E	High availability agent information .....	337
	About agents .....	337
	VCS agents included within SFCFS .....	338
	CVMCluster agent .....	338
	Entry points for CVMCluster agent .....	338
	Attribute definition for CVMCluster agent .....	339
	CVMCluster agent type definition .....	339
	CVMCluster agent sample configuration .....	340
	CVMVxconfigd agent .....	340
	Entry points for CVMVxconfigd agent .....	341
	Attribute definition for CVMVxconfigd agent .....	341
	CVMVxconfigd agent type definition .....	341
	CVMVxconfigd agent sample configuration .....	342
	CVMVolDg agent .....	342
	Entry points for CVMVolDg agent .....	342
	Attribute definition for CVMVolDg agent .....	343
	CVMVolDg agent type definition .....	343
	CVMVolDg agent sample configuration .....	343
	CFSMount agent .....	344
	Entry points for CFSMount agent .....	344
	Attribute definition for CFSMount agent .....	344
	CFSMount agent type definition .....	345
	CFSMount agent sample configuration .....	346
Appendix F	Troubleshooting information .....	347
	Restarting the installer after a failed connection .....	347
	What to do if you see a licensing reminder .....	347
	Troubleshooting an installation on AIX .....	348
	Storage Foundation Cluster File System installation issues .....	348
	Incorrect permissions for root on remote system .....	349
	Resource temporarily unavailable .....	349
	Inaccessible system .....	350
	Storage Foundation Cluster File System problems .....	350
	Unmount failures .....	350
	Mount failures .....	350
	Command failures .....	351
	Performance issues .....	352
	High availability issues .....	352

Appendix G	Troubleshooting cluster installation .....	355
	Installer cannot create UUID for the cluster .....	355
	The vxfcntl utility fails when SCSI TEST UNIT READY command fails .....	356
	Troubleshooting on the CP server .....	356
	CP server service group issues .....	357
	Testing the connectivity of the CP server .....	357
	Troubleshooting server-based I/O fencing on the SFCFS cluster .....	357
	Issues during server-based fencing start up on SFCFS cluster node .....	358
	Issues during online migration of coordination points .....	359
	Troubleshooting server-based I/O fencing in mixed mode .....	360
	Checking keys on coordination points when vxfen_mechanism value is set to cps .....	365
Appendix H	Sample SFCFS cluster setup diagrams for CP server-based I/O fencing .....	367
	Configuration diagrams for setting up server-based I/O fencing .....	367
	Two unique client clusters served by 3 CP servers .....	367
	Client cluster served by highly available CPS and 2 SCSI-3 disks .....	368
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks .....	370
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks .....	371
Appendix I	Changing NFS server major numbers for VxVM volumes .....	373
	Changing NFS server major numbers for VxVM volumes .....	373
Appendix J	Configuring LLT over UDP using IPv6 .....	375
	Using the UDP layer of IPv6 for LLT .....	375
	When to use LLT over UDP .....	375
	Configuring LLT over UDP using IPv6 .....	375
	The link command in the /etc/llttab file .....	376
	The set-addr command in the /etc/llttab file .....	377
	Selecting UDP ports .....	377
	Sample configuration: direct-attached links .....	378
	Sample configuration: links crossing IP routers .....	379

Appendix K	Configuring LLT over UDP using IPv4 .....	381
	Using the UDP layer for LLT .....	381
	When to use LLT over UDP .....	381
	Configuring LLT over UDP .....	381
	Broadcast address in the /etc/llttab file .....	382
	The link command in the /etc/llttab file .....	383
	The set-addr command in the /etc/llttab file .....	383
	Selecting UDP ports .....	384
	Configuring the netmask for LLT .....	385
	Configuring the broadcast address for LLT .....	385
	Sample configuration: direct-attached links .....	386
	Sample configuration: links crossing IP routers .....	387
Index .....		389



# About Storage Foundation Cluster File System

This chapter includes the following topics:

- [Veritas Storage Foundation Cluster File System suites](#)
- [About I/O fencing](#)
- [About Veritas product licensing](#)

## Veritas Storage Foundation Cluster File System suites

The following table lists the Symantec products and optionally licensed features available with each Veritas Storage Foundation Cluster File System (SFCFS) product suite.

**Table 1-1** Contents of Veritas Storage Foundation Cluster File System products

Storage Foundation Cluster File System version	Products and features
Storage Foundation Cluster File System	Veritas File System Veritas Volume Manager Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Optionally licensed features: Veritas Volume Replicator
Storage Foundation Cluster File System HA	Veritas File System Veritas Volume Manager Veritas Cluster Server Veritas Quick I/O option Veritas Extension for Oracle Disk Manager option Veritas Storage Checkpoint option Veritas Storage Mapping option Optionally licensed features: Veritas Volume Replicator

## About I/O fencing

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

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

The fencing operation determines the following:

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

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

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

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

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

## About Veritas product licensing

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

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

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

- Install a license key for the product and features that you want to install.  
When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
- Continue to install without a license key.  
The installer prompts for the product modes and options that you want to install, and then sets the required product level.

Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server. If you do not comply with the above terms, continuing to use the Veritas product is a violation of your end user license agreement, and results in warning messages. For more information about keyless licensing, see the following URL:  
<http://go.symantec.com/sfhakeyless>

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

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

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.  
See “[Setting or changing the product level for keyless licensing](#)” on page 107.  
See the `vxkeyless(1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the 5.1 products you have purchased.  
See “[Installing Veritas product license keys](#)” on page 108.  
See the `vxlicinst(1m)` manual page.

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

---

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

---

We recommend updating to keyless licensing for the following reasons:

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

# Before you install

This chapter includes the following topics:

- [About planning for a SFCFS installation](#)
- [About installation and configuration methods](#)
- [Assessing your system preparedness](#)
- [Preinstallation or upgrade planning for Veritas Volume Replicator](#)
- [Downloading the Storage Foundation Cluster File System software](#)
- [Setting environment variables](#)
- [Configuring secure shell \(ssh\) or remote shell before installing products](#)
- [Setting up shared storage](#)
- [Prerequisites for Storage Foundation Cluster File System](#)
- [Hardware overview and requirements for Storage Foundation Cluster File System](#)

## About planning for a SFCFS installation

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

[http://sfdoccentral.symantec.com/sf/5.1/aix/sfcfs\\_install.pdf](http://sfdoccentral.symantec.com/sf/5.1/aix/sfcfs_install.pdf)

This document is version 5.1.0.

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 SFCFS will be installed.

Follow the preinstallation instructions if you are installing one of the Veritas Storage Foundation Cluster File System products by Symantec.

The following Veritas Storage Foundation Cluster File System products by Symantec are installed with these instructions:

- Veritas Storage Foundation Cluster File System
- Veritas Storage Foundation Cluster File System High Availability (HA)

Several component products are bundled with each of these SFCFS products.

See [“Veritas Storage Foundation Cluster File System suites”](#) on page 19.

## About installation and configuration methods

You can install and configure SFCFS with Veritas installation programs or with native operating system methods.

Use one of the following methods to install and configure SFCFS:

- The Veritas product installer (Recommended)  
The common product installer displays a menu that simplifies the selection of installation options.  
See [“About the common product installer”](#) on page 49.
- The product-specific installation scripts  
The installation scripts provide a command-line interface to installing a specific product. The product-specific scripts enable you to specify some additional command-line options. Otherwise, installing with the installation script is identical to specifying SFCFS from the common product installer menu.
- The Web-based Veritas installer  
The installer provides an interface to manage the installation from a remote site using a standard Web browser.  
In this release, there are some limitations in the Web-based installer.  
See [“About the Web-based installer”](#) on page 55.
- Silent installation with response files  
You can use any of the above options to generate a response file. You can then customize the response file for another system. Run the product installation script with the response file to install silently on one or more other systems.  
See [“About response files”](#) on page 311.
- Network Installation Manager (NIM)

You can use the Veritas product installer or the product-specific installation script to generate a NIM configuration file. Use the generated script to install Veritas packages from your NIM server.

## Assessing your system preparedness

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 5.1.

### Veritas Operations Services

Veritas Operations Services (VOS) is a Web-based application that is designed specifically for Veritas Storage Foundation and High Availability products.

See “[Veritas Operations Services](#)” on page 25.

### Simulation option of the Veritas product installer

The Veritas product installer performs a complete simulation of the install process, including prechecks. The simulation provides you with a preview of the installation process, in addition to performing prechecks.

See “[About the installation simulator](#)” on page 25.

## Veritas Operations Services

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

Among its broad set of features, VOS evaluates the systems in your environment to determine if you are ready to install or upgrade Storage Foundation and High Availability products.

To access VOS, go to:

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

## About the installation simulator

The product installer includes an option to simulate installing, configuring, or uninstalling the selected Veritas product. The simulation option steps through

the installation script, including all of the preinstallation checks on the systems. However, the simulation does not actually install the packages, uninstall previously installed packages, or start or stop any processes.

The simulation process enables you to create a response file, that can be used as a template for installing or configuring a Veritas product. You can also use the simulator to view the installation questions or the configuration questions. The simulation lets you preview the steps for the installation or configuration, without disrupting your existing installation.

Use the installation simulator in the following situations:

- To understand the information that is required when you install, configure, or uninstall a Veritas product.

Because the simulator steps through the same code that is used by the installer, the simulation displays the exact prompts that the installer displays. The simulation includes running preinstallation checks on your system.

If the checks are not required, you can skip the preinstallation checks. For example, skip the preinstallation checks if you are running the simulator on a different system than the system on which you plan to install the Veritas product.

After viewing the prompts, you can gather any required information before performing the actual install, configure, or uninstall.

- To create a response file for your system.  
Response files store the values that are requested by the install program in the form of variables. The response file is a text file, which has comments defining what each variable represents. You can use the response file as a template for an installation or configuration. You can edit the response file with any text editor.

To simulate an installation or configuration, specify the `-makeresponsefile` option to the installer or product installation script at the command line.

To simulate an uninstallation, specify the `-makeresponsefile` option to the installer or the product uninstall script at the command line.

## Preinstallation or upgrade planning for Veritas Volume Replicator

Before installing or upgrading 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.

The following related documents are available:

*Veritas Volume Replicator Planning and Tuning Guide* Provides detailed explanation of VVR tunables

*Veritas Volume Replicator Administrator's Guide* Describes how to change tunable values

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 nodes. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS.

VVR supports replicating data between VVR 5.1 and VVR 4.0 or later.

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 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.

---

**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.

---

# Downloading the Storage Foundation Cluster File System software

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

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

See “[About installation scripts](#)” on page 305.

## To download the software

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

The estimated space that is needed for download is 5 GB.

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

See “[Disk space requirements](#)” on page 46.

- 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 -k 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. You must download the Veritas 5.0 software and the Veritas 5.1 software into separate 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, Veritas Storage Foundation commands are stored in `/opt/VRTS/bin` and HA commands are stored in `/opt/VRTSvcs/bin`. Storage Foundation HA manual pages are stored in `/opt/VRTS/man`.

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

If you are not installing an HA product, you can omit `/opt/VRTSvcs/bin`.

The `nroff` versions of the online manual pages are not readable using the `man` command if the `bos.txt.tfs` fileset is not installed; however, the `VRTSvxvm` and `VRTSvxfs` packages install ASCII versions in the `/opt/VRTS/man/cat*` and `/opt/VRTS/man/man*` directories that are readable without the `bos.txt.tfs` fileset.

## Configuring secure shell (ssh) or remote shell 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 installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. These utilities must run as `root` on all cluster nodes or remote systems.

You can install products to remote systems using either secure shell (`ssh`) or remote shell (`rsh`). `ssh` is the preferred method of remote communication because it provides a greater level of security than the `rsh` suite of protocols.

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.

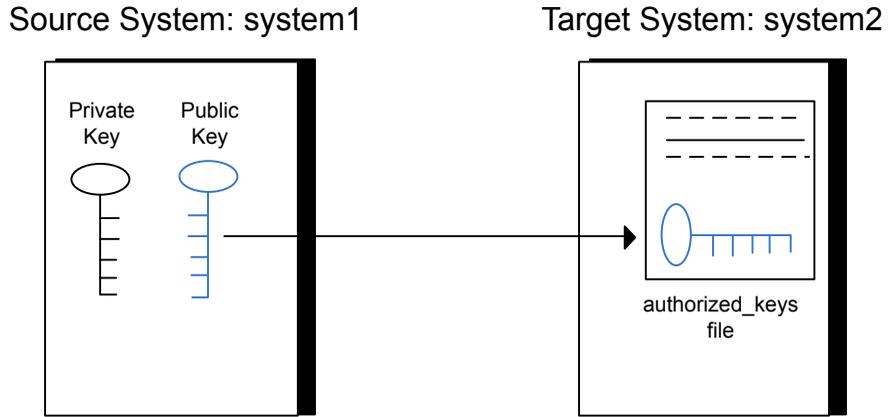
### Configuring and enabling ssh

The `ssh` program enables you to log into and execute commands on a remote system. `ssh` enables encrypted communications and an authentication process between two untrusted hosts over an insecure network.

In this procedure, you first create a DSA key pair. From the key pair, you append the public key from the source system to the `authorized_keys` file on the target systems.

Figure 2-1 illustrates this procedure.

**Figure 2-1** Creating the DSA key pair and appending it to target systems



Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

Visit the OpenSSH website that is located at: <http://openssh.org> to access online manuals and other resources.

#### To create the DSA key pair

- 1 On the source system (system1), log in as root, and navigate to the root directory.

```
system1 # cd /
```

- 2 To generate a DSA key pair on the source system, type the following command:

```
system1 # ssh-keygen -t dsa
```

System output similar to the following is displayed:

```
Generating public/private dsa key pair.  
Enter file in which to save the key (//.ssh/id_dsa):
```

- 3 Press Enter to accept the default location of `/.ssh/id_dsa`.

- 4 When the program asks you to enter the passphrase, press the Enter key twice.

```
Enter passphrase (empty for no passphrase):
```

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

- 5 Make sure the `/.ssh` directory is on all the target installation systems (system2 in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

```
system2 # cd /  
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 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)?
```

- 2 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:
```

- 3 Enter the root password of system2.
- 4 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
```

- 5 To quit the SFTP session, type the following command:

```
sftp> quit
```

- 6** 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:
```

- 7** After you log in to system2, enter the following command to append the `id_dsa.pub` file to the authorization key file:

```
system2 # cat /id_dsa.pub >> /.ssh/authorized_keys
```

- 8** 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, type the following command on system2:

```
system2 # rm /id_dsa.pub
```

- 9** To log out of the `ssh` session, type the following command:

```
system2 # exit
```

- 10** 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
```

- 11** 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 for 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), type the following command:

```
system1 # ssh -l root system2 uname -a
```

where system2 is the name of the target system.

- 2 The command should execute from the source system (system1) to the target system (system2) without the system requesting a passphrase or password.
- 3 Repeat this procedure for each target system.

## Restarting ssh

After you complete this procedure, ssh can be restarted in any of the following scenarios:

- After a terminal session is closed
- After a new terminal session is opened
- After a system is restarted
- After too much time has elapsed, to refresh ssh

### To restart ssh

- 1 On the source installation system (system1), bring the private key into the shell environment.

```
system1 # exec /usr/bin/ssh-agent $SHELL
```

- 2 Make the key globally available for the user root

```
system1 # ssh-add
```

## Enabling rsh for AIX

To enable rsh, create a `/.rhosts` file on each target system. Then add a line to the file specifying the full domain name of the source system. For example, add the line:

```
sysname.domainname.com root
```

Change permissions on the `/.rhosts` file to 600 by typing the following command:

```
# chmod 600 /.rhosts
```

After you complete an installation procedure, delete the `.rhosts` file from each target system to ensure security:

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

## Setting up shared storage

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

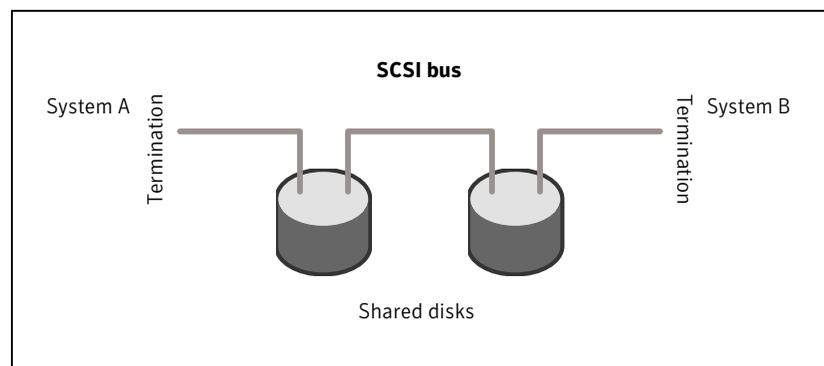
See also the *Storage Foundation Cluster File System Administrator's Guide* for a description of I/O fencing.

### Setting the SCSI identifier value

SCSI adapters are typically set with a default identifier value of 7. Each device on a SCSI bus must have a unique SCSI identifier value. When more than one system is connected to a SCSI bus, you must change the SCSI identifier to a unique number. You must make this change to one or more systems, usually the unique number is 5 or 6.

Perform the procedure if you want to connect to shared storage with shared SCSI devices.

**Figure 2-2** Cabling the shared storage



### To set the SCSI identifier value

- 1 Determine the SCSI adapters on each system:

```
north # lsdev -C -c adapter | grep scsi
scsi0   Available 11-08   Wide/Ultra-2 SCSI I/O Controller
scsi1   Available 11-09   Wide/Ultra-2 SCSI I/O Controller
south # lsdev -C -c adapter | grep scsi
scsi0   Available 11-08   Wide/Ultra-2 SCSI I/O Controller
scsi1   Available 11-09   Wide/Ultra-2 SCSI I/O Controller
```

- 2 Verify the SCSI ID of each adapter:

```
north # lsattr -E -l scsi0 -a id
id 7 Adapter card SCSI ID True
north # lsattr -E -l scsi1 -a id
id 7 Adapter card SCSI ID True
south # lsattr -E -l scsi0 -a id
id 7 Adapter card SCSI ID True
south # lsattr -E -l scsi1 -a id
id 7 Adapter card SCSI ID True
```

- 3 If necessary, change the SCSI identifier on each system so that it is unique:

```
south # chdev -P -l scsi0 -a id=5
scsi0 changed
south # chdev -P -l scsi1 -a id=5
scsi1 changed
```

- 4 Shut down all systems in the cluster.
- 5 Cable the shared storage as illustrated in [Figure 2-2](#).
- 6 Restart each system. After all systems have booted, use the `lspv` command to verify that each system can see all shared devices needed by the application.

## Setting up Fiber Channel

Perform the following steps to set up fiber channel.

### To set up fiber channel

- 1 Connect the Fiber Channel adapters and the shared storage devices to the same hub or switch.

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

- 2 Reboot each system:

```
shutdown -Fr
```

- 3 After all systems have booted, use the `lspv` command to verify that each system can see all shared devices needed by the application.

## Prerequisites for Storage Foundation Cluster File System

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 the SFCFS:

- The cluster name, beginning with a letter (a-z, A-Z).
- A unique ID from 0-65535 for the cluster. Within the public subnet, a new cluster using a duplicate cluster ID can cause existing clusters to fail.
- The host names of the cluster nodes.
- The device names of the network interface cards (NICs) used for the private networks among nodes.
- Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities as root on all cluster nodes or remote systems.
- Symantec recommends configuring the cluster with I/O fencing enabled. I/O fencing requires shared devices to support SCSI-3 Persistent Reservations (PR). Enabling I/O fencing prevents data corruption caused by a split brain scenario.

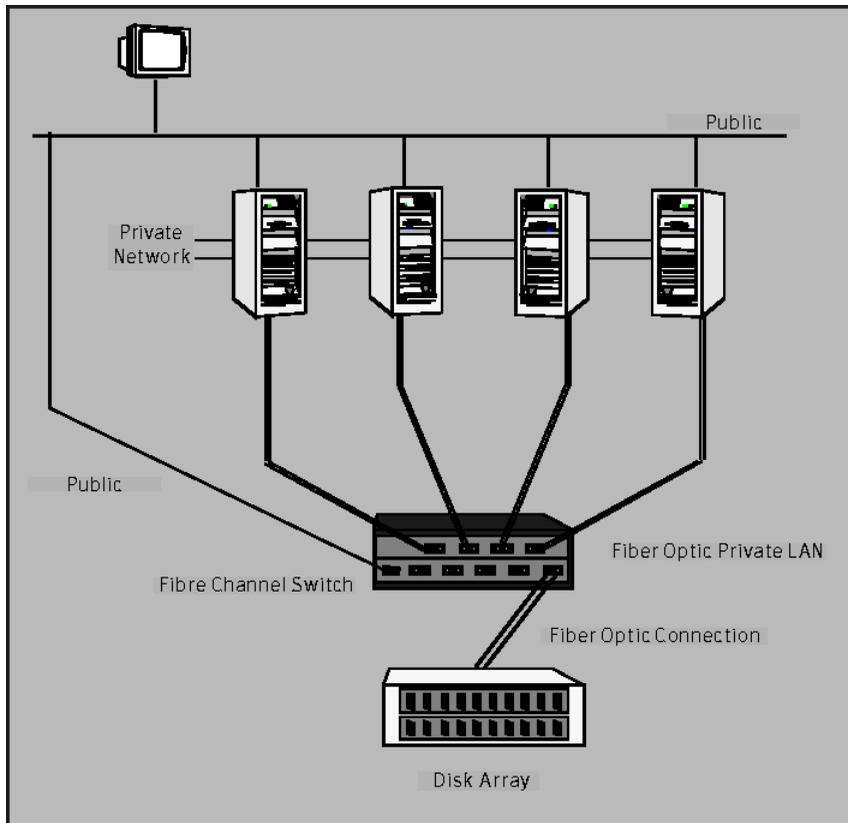
The Storage Foundation Cluster File System is supported without I/O fencing enabled. However, without I/O fencing enabled, split brain scenarios can result in data corruption.

# Hardware overview and requirements for Storage Foundation Cluster File System

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, SFCFS can be used in conjunction with the latest Veritas Storage Area Network (SAN) applications to provide a complete data storage and retrieval solution.

Figure 2-3 shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

Figure 2-3 Four Node SFCFS 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 Storage Foundation Cluster File System (SF CFS) cluster.

See the *Veritas Storage Foundation Release Notes*.

---

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

---



# System requirements

This chapter includes the following topics:

- [Hardware and software requirements](#)
- [I/O fencing requirements](#)
- [Release notes](#)
- [Supported AIX operating systems](#)
- [Storage Foundation Cluster File System node requirements](#)
- [Database requirements](#)
- [Disk space requirements](#)

## Hardware and software requirements

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

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

## 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 42.

## Coordinator disk requirements for I/O fencing

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

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

## CP server requirements

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

- CP server hardware-specific requirements
- OS requirements
- Networking requirements (and recommendations)
- Security requirements

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

[Table 3-1](#) lists additional requirements for hosting the CP server.

**Table 3-1** CP server hardware requirements

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

**Table 3-2** displays the CP server supported operating systems and versions.

**Table 3-2** CP server supported operating systems and versions

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

For networking requirements, Symantec recommends that network access from the SFCFS 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 SFCFS cluster(s) by these network paths. The CP server listens for messages from the SFCFS cluster(s) using TCP port 14250. This is the default port that can be changed during a CP server configuration.

---

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

---

When placing the CP server (s) within a specific network configuration, the number of hops from the different SFCFS cluster nodes to the CP server (s) should be taken

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

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

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

For additional information, see *Storage Foundation Cluster File System Administrator's Guide*.

## Release notes

Read the *Release Notes* for all products included with this product.

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

<http://www.symantec.com/business/support/index.jsp>

## Supported AIX operating systems

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

Product installation scripts verify the required update levels. The installation process terminates if the target systems do not meet the maintenance level requirements.

For any Veritas cluster product, all nodes in the cluster must have the same operating system version and update level.

The minimum system requirements for this release are as follows:

AIX 5.3 at one of the following levels:

- TL7 with SP6
- TL8 with SP4
- or any higher TLs.

AIX 6.1 at one of the following levels:

- TL0 with SP6
- TL1 with SP2
- or any higher TLs.

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

<http://entsupport.symantec.com/docs/334998>

## Storage Foundation Cluster File System 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://entsupport.symantec.com/docs/331625>

---

**Note:** SFCFS supports running Oracle, DB2, and Sybase on VxFS and VxVM.

SFCFS does not support running SFDB tools with DB2 and Sybase.

---

## Disk space requirements

Before installing any of the Veritas Storage Foundation products, confirm that your system has enough free disk space.

Use the "Perform a Preinstallation Check" (P) menu or the `-precheck` option of the product installer to determine whether there is sufficient space.

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

# Installing Storage Foundation Cluster File System using the common product installer

This chapter includes the following topics:

- [Installation quick reference](#)
- [Mounting a software disc](#)
- [About the common product installer](#)
- [Installing Storage Foundation Cluster File System with the product installer](#)

## Installation quick reference

The product installer displays a menu that simplifies the selection of installation and upgrade options. It is the recommended installation method. Select a product to install or upgrade from the menu to invoke that product's installation script.

[Table 4-1](#) provides a quick overview of a stand-alone installation using the product installer.

**Table 4-1** Installation overview

Installation task	For more information, refer to the following section:
Obtain product licenses.	

**Table 4-1** Installation overview (*continued*)

Installation task	For more information, refer to the following section:
Download the software, or insert the product DVD.	See <a href="#">“Downloading the Storage Foundation Cluster File System software”</a> on page 28. See <a href="#">“Mounting a software disc”</a> on page 48.
Set environment variables.	See <a href="#">“Setting environment variables”</a> on page 28.
Configure the secure shell (SSH) on all nodes.	See <a href="#">“Configuring secure shell (ssh) or remote shell before installing products”</a> on page 29.
Verify that hardware, software, and operating system requirements are met.	
Check that sufficient disk space is available.	See <a href="#">“Disk space requirements”</a> on page 46.
Use the installer to install the products.	

## Mounting a software disc

Veritas software is provided on a DVD format disc. If you have the media kit, then get the software disc from the media kit.

### To mount the software disc

- 1 Log in as superuser.
- 2 Place the Veritas software disc into a DVD drive connected to your system.
- 3 Mount the disc by determining the device access name of the DVD drive. The format for the device access name is `cdx` where `x` is the device number. Insert the disc and type the following commands:

```
# mkdir -p /mnt/cdrom
# mount -V cdrfs -o ro /dev/cdrom /mnt/cdrom
```

- 4 Change to the appropriate distribution directory and product subdirectory to view the product release notes and installation guides, or install the products.

## About the common product installer

The product installer is the recommended method to license and install the Veritas products. The installer also enables you to 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 installation scripts”](#) on page 305.

At most points during an installation, you can type `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. If an installation procedure hangs, use `Control-C` to stop and exit the program. After a short delay, the script exits. You can also enter `q` to quit the installer or `?` to display help information.

Default responses are in parentheses. Press Return to accept the defaults.

Additional options are available for the common product installer.

See [“Installation script options”](#) on page 306.

## Installing Storage Foundation Cluster File System with the product installer

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

---

**Note:** Save a copy of `/var/adm/ras/errtmpl` and `/etc/trcfmt` files before you install the product. If the packages fail to install due to the `template file is corrupted` error message, replace `/var/adm/ras/errtmpl` and `/etc/trcfmt` files with the ones that you had saved. Uninstall all the packages you installed and reinstall.

See [“Uninstalling Storage Foundation Cluster File System”](#) on page 298.

---

The following sample procedure is based on the installation of a Veritas Storage Foundation Cluster File System HA 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, use the `installsfdfs` script instead of the `installer` script.

---

### To install Storage Foundation Cluster File System products

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

See “[Configuring secure shell \(ssh\) or remote shell before installing products](#)” on page 29.

- 2 Load and mount the software disc.

See “[Mounting a software disc](#)” on page 48.

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

```
# cd /mnt/cdrom
```

- 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 (ssh) or remote shell (rsh) utilities are configured:

```
# ./installer
```

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

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

Do not select the "Storage Foundation for Oracle RAC" option unless you have the correct license and setup.

The example installation assumes you have selected SFCFS HA.

- 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 EULA.pdf file present on the media? [y,n,q,?] y
```

- 8 Select from one of the following install options:

- **Minimal filesets:** installs only the basic functionality for the selected product.

- Recommended filesets: installs the full feature set without optional filesets.
- All filesets: installs all available filesets.

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

For example, you should see output similar to the following:

```
SFCFS can be installed without optional filesets to conserve
disk space.
```

- 1) Install minimal Veritas Storage Foundation Cluster File System filesets - 721 MB required
- 2) Install recommended Veritas Storage Foundation Cluster File System filesets - 946 MB required
- 3) Install all Veritas Storage Foundation Cluster File System filesets - 994 MB required
- 4) Display filesets to be installed for each option.

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

- 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 on which to
install SFCFS: host1 host2
```

- 10 During the initial system check, the installer verifies that communication between systems has been set up.

If the installer hangs or asks for a login password, stop the installer and set up ssh or rsh. Then run the installer again.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 29.

- 11 After the system checks complete, the installer displays a list of the filesets that will be installed. Press Enter to continue with the installation.

## 12 You are prompted to choose your licensing method.

To ensure compliance 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 (see <http://go.symantec.com/vom> for details and free download)

- 1) Enter a valid license key
- 2) Enable keyless licensing

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

If you have a valid license key, select 1 and enter the license key at the prompt. Skip to step 17.

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 You are prompted to enter the Standard or Enterprise product mode.

## 14 Select **yes** to enable the Veritas Volume Replicator.

Would you like to enable Veritas Volume Replicator [y,n,q] (n) y

## 15 Select **yes** to enable the Global Cluster Option.

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

## 16 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) **n**

**17** The product installation completes.

Configure Storage Foundation Cluster File System when prompted.

Do you want to configure Storage Foundation Cluster File System on these systems at this time? [y,n,q] (y) **y**

If you select **y** to configure now, respond to the prompts to configure the cluster.

**18** If you select **n** to configure, the installation completes.

---

**Note:** You must configure Storage Foundation Cluster File System before you can use the product.

---

**19** View the log file, if needed, to confirm the installation.

Installation log files, summary file, and response file are saved at:

`/opt/VRTS/install/logs/installer-****`



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

This chapter includes the following topics:

- [About the Web-based installer](#)
- [Features supported with 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 SFCFS with the Veritas Web-based installer](#)

## About the Web-based installer

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

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

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

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

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

## Features supported with Web-based installer

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

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

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

## Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

**Table 5-1** Web-based installer requirements

System	Function	Requirements
Target system	The system(s) where the Veritas products will be installed.	Must be a supported platform for Veritas Storage Foundation Cluster File System 5.1
Installation server	The server from which to initiate the installation. The installation media is mounted and accessible from the installation server.	Must be the same OS as the system(s) on which to install.
Administrative system	The system on which you run the web browser to perform the installation.	Web browser

## Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

### To start the Web-based installer

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

```
# ./webinstaller start
```

The `webinstaller` script displays a URL.

- 2 On the administrative server, start the Web browser.
- 3 Navigate to the URL displayed from step 1.
- 4 The browser may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

- 5 When prompted, enter `root` and `root`'s password of the installation server.

## Obtaining a security exception on Mozilla Firefox

You may need to get a security exception on Mozilla Firefox.

### To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **Add Exception** button.

- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.
- 5 Click **Confirm Security Exception** button.
- 6 Enter root in User Name field and root password of the web server in the Password field.

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

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

### To perform a pre-installation check

- 1 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 57.
- 2 On the Select a task and a product page, select **Perform a Pre-installation check** from the **Task** drop-down list.
- 3 Select the product from the **Product** drop-down list, and click **Next**.
- 4 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Validate**.
- 5 The installer performs the precheck and displays the results.
- 6 If the validation completes successfully, click **Next**. The installer prompts you to begin the installation. Click **Ok** to install SFCFS on the selected system. Click **Cancel** to install later.
- 7 Click **Finish**. The installer prompts you for another task.

## Installing SFCFS with the Veritas Web-based installer

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

### To install SFCFS

- 1 Perform preliminary steps.
- 2 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 57.
- 3 On the Select a task and a product page, select **Install a Product** from the **Task** drop-down list.

- 4 Select SFCFS or SFCFS High Availability from the Product drop-down list, and click **Next**.
- 5 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 6 Choose minimal, recommended, or all packages. Click **Next**.
- 7 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.
- 8 After the validation completes successfully, click **Next** to install SFCFS on the selected system.
- 9 After the installation completes, you must choose your licensing method. On the license page, select one of the following tabs:

- Keyless licensing

---

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

For more information, go to the following website:

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

---

Complete the following information:

Choose whether you want to enable Veritas Volume 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**.

**10** The installer prompts you to configure the cluster.

If you select n, you can exit the installer. You must configure the product before you can use SFCFS.

After the installation completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

**11** Select the checkbox to specify whether you want to send your installation information to Symantec.

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

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

# Installing Storage Foundation Cluster File System, other methods

This chapter includes the following topics:

- [Installing SFCFS using NIM and the installer](#)
- [Installing SFCFS on an alternate disk](#)

## Installing SFCFS using NIM and the installer

You can use the product installer in concert with NIM to install the Veritas product, or to install the operating system and the Veritas product.

The instructions in this section assume a working knowledge of the Network Installation Management process. See the operating system documentation for detailed information on Network Installation Management.

In the following samples, the LPP resource uses LPP-6100-up2date and its relevant SPOT resource is spot-6100-up2date.

## Preparing the bundle and script resources on NIM server

You need to prepare the bundle and script resources on the NIM server before using NIM to install SFCFS packages. The following actions are executed on the NIM server.

### To prepare the bundle and script resources

- 1 Insert and mount the installation media.
- 2 At the command line enter the following command to list the LPP source and choose one for the next step.

```
# lsnim -t lpp_source
LPP-6100-up2date resources lpp_source
LPP-5300-up2date resources lpp_source
```

- 3 Navigate to the cluster\_server directory on the disc and run the `installsfcfs` command to prepare the bundle and script resources:

```
# ./installsfcfs -nim LPP-6100-up2date
```

- 4 When the installer asks you if you want to generate the NIM script resource, answer `y`.

Output resembles:

```
NIM script resource copy_cpi is created for copy installer
scripts to disk
```

The script configure file is created at `/opt/VRTS/nim/copy_cpi`.

- 5 Run the `lsnim -l` command to check that the `install_bundle` resource is created successfully.

Output resembles:

- 6 Run the `lsnim -l copy_cpi` command to check that the script resource is created successfully.

```
# lsnim -l copy_cpi
```

Output resembles:

```
copy_cpi:
class = resources
type = script
Rstate = ready for use
prev_state = unavailable for use
location = /opt/VRTS/nim/copy_cpi
alloc_count = 0
server = master
```

## Installing SFCFS on the NIM client using SMIT

You can install SFCFS on the NIM client using the SMIT tool.

Perform these steps on each node to have SFCFS installed in a cluster.

### To install SFCFS

- 1 On the NIM client, start smitty.

```
# smitty install
```

- 2 In the menu, select **Network Installation Management**.
- 3 In the menu, select **Install and Update Software**.
- 4 In the menu, select **Install Software Bundle**.
- 5 In the menu, select the LPP\_SOURCE. In this example, specify **LPP-6100-up2date**.
- 6 In the menu, select the bundle. In this example, specify the **SFCFS51\_bundle**.
- 7 For the customization script that you want to run after installation, specify **copy\_cpi**.
- 8 For the installp flags, specify that the ACCEPT new license agreements flag has a **yes** value.
- 9 Press the Enter key to start the installation. Note that it may take some time to finish.

## Installing SFCFS and the operating system on the NIM client using SMIT

You can install VCS and the operating system on the NIM client using the SMIT tool.

Perform these steps on each node to have SFCFS and AIX installed in a cluster.

### To install SFCFS and the operating system

- 1 On the NIM server, start smitty for a NIM and operating system installation.

```
# smitty nim_bosinst
```

- 2 In the menu, select the stand-alone target.
- 3 In the menu, select **rte- Install from installation images**.
- 4 In the menu, select the LPP\_SOURCE. In this example, select **LPP-6100-up2date**.

- 5 In the menu, select the SPOT. In this example, select **spot-6100-up2date**.
- 6 In the menu, select the following options:
  - For the Customization SCRIPT to run after installation option, specify **copy\_cpi**.
  - For the ACCEPT new license agreements option, specify **yes**.
- 7 For the installp flags, specify that the ACCEPT new license agreements flag has a **yes** value.

## Installing SFCFS on an alternate disk

Use the alternate disk installation process to install SFCFS on an alternate disk. Installing SFCFS on an alternate disk enables you to boot from the second disk instead of the default disk.

---

**Note:** The alternate disk installation is a manual procedure that does not require the Veritas installation program. As a result, the installation and uninstallation scripts are not available in the `/opt/VRTS/install` directory. If you need to access these scripts, find them on the installation media.

---

The installation process involves the following steps:

Preparing to install SFCFS on an alternate disk	See <a href="#">“Preparing to install SFCFS on an alternate disk”</a> on page 64.
Installing SFCFS on an alternate disk	See <a href="#">“Installing SFCFS on an alternate disk”</a> on page 65.
Verifying the installation	See <a href="#">“Verifying the installation”</a> on page 68.

The instructions in this section assume a working knowledge of the alternate disk installation process. See the operating system documentation for detailed information on alternate disk installations.

In the procedure examples, the primary or current boot environment resides on `hdisk0` and the alternate or inactive boot environment resides on `hdisk1`.

## Preparing to install SFCFS on an alternate disk

Complete the preparatory steps in the following procedure before you install SFCFS on an alternate disk.

### To prepare to install SFCFS on an alternate disk

- 1 Make sure that the SFCFS installation media is available.
- 2 On the nodes that you want to install SFCFS, create an alternate boot disk identical to the primary boot disk with the same size and partitions as the primary boot disk.
- 3 Check the status of the physical disks on your system.

---

**Note:** The alternate disk must have a physical identifier and must not contain any mounted volume groups.

---

```
# lspv
```

Output similar to the following displays:

```
hdisk0          0009710fa9c79877    rootvg    active
hdisk1          0009710f0b90db93    None
```

If the alternate disk does not have a physical identifier, set the physical identifier for the disk:

```
# chdev -l hdisk1 -a pv=yes
```

- 4 Make sure that the following filesets are installed on the primary disk:  
`bos.alt_disk_install.boot_images`, `bos.alt_disk.install.rte`
- 5 Mount the SFCFS installation media and copy the SFCFS filesets to a directory on the primary disk, for example `/usr`.
- 6 Create a bundle file `/usr/sys/inst.data/user_bundles/MyBundle.bnd` that contains the SFCFS filesets to be installed.

You can obtain the list of packages from the `pkginfo.txt` file that is available on the product disc in the corresponding Veritas product stack directory.

For instructions on creating a bundle, see the operating system documentation.

## Installing SFCFS on an alternate disk

This section provides instructions for cloning the primary boot environment to the alternate disk, installing SFCFS on the alternate disk, and rebooting the system to start from the alternate disk.

Use one of the following ways to install SFCFS on an alternate disk:

- SMIT interface See [“To install SFCFS on an alternate disk using the SMIT interface”](#) on page 66.
- Manual See [“To install SFCFS on an alternate disk manually”](#) on page 67.

### To install SFCFS on an alternate disk using the SMIT interface

- 1 Clone the primary boot disk `rootvg` to an alternate disk.

Start the SMIT menu :

```
# smit alt_clone
```

Provide information for the following fields in the SMIT menu.

Target disk to install Enter the name of the alternate boot disk.

```
hdisk1
```

Bundle to install Enter the name of the bundle that contains the SFCFS filesets to be installed.

```
MyBundle
```

Directory or Device with images Enter the full path of the bundle file that contains the SFCFS filesets.

```
/usr/
```

ACCEPT new license agreements? Enter **yes** to accept the license agreements.

Set bootlist to boot from this disk on next reboot? Enter **yes** to enable the system to boot from the alternate disk on rebooting.

- 2 Press **Enter** to start the alternate disk installation.

The installation process takes some time.

- 3 Verify that the alternate disk is created and the volume group `altinst_rootvg` is mounted:

```
# lspv
```

Output similar to the following displays:

```
hdisk0          0009710fa9c79877      rootvg
hdisk1          0009710f0b90db93      altinst_rootvg
```

- 4 Verify that the normal boot list includes the name of the alternate boot disk.

```
# bootlist -m normal -o
hdisk1
```

- 5 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

```
# shutdown -r
```

- 6 Verify the installation.

See [“Verifying the installation”](#) on page 68.

- 7 Configure SFCFS.

#### To install SFCFS on an alternate disk manually

- 1 Clone the primary boot disk `rootvg` to an alternate disk.

```
# /usr/sbin/alt_disk_copy -d "hdisk1" \
-b MyBundle -l /usr/
```

Where:

- `-d` indicates the name of the target disk on which you clone the primary disk.
  - `-b` indicates the name of the user bundle that contains the list of SFCFS filesets that you want to install on the alternate boot disk.
  - `-l` indicates the full path of the directory that contains the filesets
- 2 Verify that the alternate disk is created and the volume group `altinst_rootvg` is mounted:

```
# lspv
hdisk0          0009710fa9c79877    rootvg
hdisk1          0009710f0b90db93    altinst_rootvg
```

- 3 Verify that the normal boot list includes the name of the alternate boot disk. By default, the alternate disk installation process changes the boot list to enable the system to boot from the alternate disk.

```
# bootlist -m normal -o
hdisk1
```

- 4 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

```
# shutdown -r
```

- 5 Configure SFCFS.

## Verifying the installation

To ensure that alternate disk installation has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

### To verify the installation

- 1 Verify that the alternate boot environment is active:

```
# lspv
```

Output similar to the following displays:

```
hdisk0          0009710fa9c79877    old_rootvg
hdisk1          0009710f0b90db93    rootvg
```

- 2 Verify that the filesets in your user bundle are installed:

```
# ls1pp -Lb MyBundle
```

# Preparing to configure Storage Foundation Cluster File System

This chapter includes the following topics:

- [Preparing to configure the clusters in secure mode](#)
- [About configuring SFCFS clusters for data integrity](#)
- [About I/O fencing components](#)
- [About I/O fencing configuration files](#)
- [About planning to configure I/O fencing](#)
- [Setting up the CP server](#)

## Preparing to configure the clusters in secure mode

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

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

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

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

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

- A system in your enterprise that serves as root broker (RB).

You can either use an external system as root broker, or use one of the cluster nodes as root broker.

- To use an external root broker, identify an existing root broker system in your enterprise or install and configure root broker on a stable system. See [“Installing the root broker for the security infrastructure”](#) on page 73.
- To use one of the cluster nodes as root broker, the installer does not require you to do any preparatory tasks. When you configure the cluster in secure mode using the `installsfcfs`, choose the automatic mode and choose one of the nodes for the installer to configure as root broker.

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

- For external root broker, an authentication broker (AB) account for each node in the cluster is set up on the root broker system. See [“Creating authentication broker accounts on root broker system”](#) on page 74.
- The system clocks of the external root broker and authentication brokers must be in sync.

The `installsfcfs` provides the following configuration modes:

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

[Figure 7-1](#) depicts the flow of configuring SFCFS cluster in secure mode.



**Table 7-1** Preparatory tasks to configure a cluster in secure mode (with an external root broker)

Tasks	Who performs this task
<p>Decide one of the following configuration modes to set up a cluster in secure mode:</p> <ul style="list-style-type: none"> <li>■ Automatic mode</li> <li>■ Semi-automatic mode</li> <li>■ Manual mode</li> </ul>	VCS administrator
<p>Install the root broker on a stable system in the enterprise.</p> <p>See <a href="#">“Installing the root broker for the security infrastructure”</a> on page 73.</p>	AT administrator
<p>To use the semi-automatic mode or the manual mode, on the root broker system, create authentication broker accounts for each node in the cluster.</p> <p>See <a href="#">“Creating authentication broker accounts on root broker system”</a> on page 74.</p> <p>AT administrator requires the following information from the VCS administrator:</p> <ul style="list-style-type: none"> <li>■ Node names that are designated to serve as authentication brokers</li> <li>■ Password for each authentication broker</li> </ul>	AT administrator
<p>To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator.</p> <p>See <a href="#">“Creating encrypted files for the security infrastructure”</a> on page 75.</p> <p>AT administrator requires the following additional information from the VCS administrator:</p> <ul style="list-style-type: none"> <li>■ Administrator password for each authentication broker Typically, the password is the same for all nodes.</li> </ul>	AT administrator
<p>To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.</p>	AT administrator
<p>Copy the files that are required to configure a cluster in secure mode to the system from where you plan to install and configure SFCFS.</p> <p>See <a href="#">“Preparing the installation system for the security infrastructure”</a> on page 77.</p>	VCS administrator

## Installing the root broker for the security infrastructure

Install the root broker only if you plan to use AT to configure the cluster in secure mode. You can use a system outside the cluster or one of the systems within the cluster as root broker. If you plan to use an external broker, the root broker administrator must install and configure the root broker before you configure the Authentication Service for SFCFS. Symantec recommends that you install the root broker on a stable system that is outside the cluster.

You can also identify an existing root broker system in the data center to configure the cluster in secure mode. The root broker system can run AIX, HP-UX, Linux, or Solaris operating system.

See Symantec Product Authentication Service documentation for more information.

### To install the root broker

- 1 Mount the product disc and start the installer.

```
# ./installer
```

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

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

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

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

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

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

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

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

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

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

## Creating authentication broker accounts on root broker system

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

### To create authentication broker accounts on root broker system

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

```
venus> # vssat showalltrustedcreds
```

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

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

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

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

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

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

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

"Failed To Get Attributes For Principal"

Proceed to step 3.

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

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

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

## Creating encrypted files for the security infrastructure

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

### To create encrypted files

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

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

```
venus> # vssat showbrokerhash
```

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

```
venus> # vssat showalltrustedcreds
```

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

identity	<p>The value for the authentication broker identity, which you provided to create authentication broker principal on the root broker system.</p> <p>This is the value for the <code>--prplname</code> option of the <code>addprpl</code> command.</p> <p>See <a href="#">“Creating authentication broker accounts on root broker system”</a> on page 74.</p>
password	<p>The value for the authentication broker password, which you provided to create authentication broker principal on the root broker system.</p> <p>This is the value for the <code>--password</code> option of the <code>addprpl</code> command.</p> <p>See <a href="#">“Creating authentication broker accounts on root broker system”</a> on page 74.</p>
broker_admin_password	<p>The value for the authentication broker password for Administrator account on the node. This password must be at least five characters.</p>

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

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

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

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

start_broker=false
enable_pbx=false
```

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

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

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

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

For example:

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

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

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

## Preparing the installation system for the security infrastructure

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

### To prepare the installation system for the security infrastructure

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

Automatic mode Do the following:

- Gather the root broker system name from the AT administrator.
- During SFCFS configuration, choose the configuration option 1 when the installsfcfs prompts.

Semi-automatic mode Do the following:

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

Manual mode

Do the following:

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

## About configuring SFCFS clusters for data integrity

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

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

## About I/O fencing components

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

- Data disks—Store shared data  
See [“About data disks”](#) on page 79.
- Coordination points—Act as a global lock during membership changes  
See [“About coordination points”](#) on page 79.

### 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. Racing for control of the coordination points to fence data disks is the key to understand how fencing prevents split-brain.

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

- Coordinator disks  
Disks that act as coordination points are called coordinator disks. Coordinator disks are three standard disks or LUNs set aside for I/O fencing during cluster reconfiguration. Coordinator disks do not serve any other storage purpose in the SFCFS configuration.

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

removal capabilities of DMP. So, you can configure I/O fencing to use either DMP devices or the underlying raw character devices. I/O fencing uses SCSI-3 disk policy that is either raw or dmp based on the disk device that you use. The disk policy is dmp by default.

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

■ Coordination point servers

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

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

---

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

## About I/O fencing configuration files

[Table 7-2](#) lists the I/O fencing configuration files.

**Table 7-2** I/O fencing configuration files

File	Description
/etc/default/vxfen	<p>This file stores the start and stop environment variables for I/O fencing:</p> <ul style="list-style-type: none"> <li>■ <b>VXFEN_START</b>—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that I/O fencing is enabled to start up.</li> <li>0—Indicates that I/O fencing is disabled to start up.</li> </ul> </li> <li>■ <b>VXFEN_STOP</b>—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that I/O fencing is enabled to shut down.</li> <li>0—Indicates that I/O fencing is disabled to shut down.</li> </ul> </li> </ul> <p>The installer sets the value of these variables to 1 at the end of Storage Foundation Cluster File System 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 7-2** I/O fencing configuration files (*continued*)

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

**Table 7-2** I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	<p>When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfermode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.</p> <p><b>Note:</b> The /etc/vxfentab file is a generated file; do not modify this file.</p> <p>For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:</p> <ul style="list-style-type: none"> <li>■ Raw disk:           <ul style="list-style-type: none"> <li>/dev/rhdisk75</li> <li>/dev/rhdisk76</li> <li>/dev/rhdisk77</li> </ul> </li> <li>■ DMP disk:           <ul style="list-style-type: none"> <li>/dev/vx/rdmp/rhdisk75</li> <li>/dev/vx/rdmp/rhdisk76</li> <li>/dev/vx/rdmp/rhdisk77</li> </ul> </li> </ul> <p>For server-based fencing, the /etc/vxfentab file also includes the security settings information.</p>

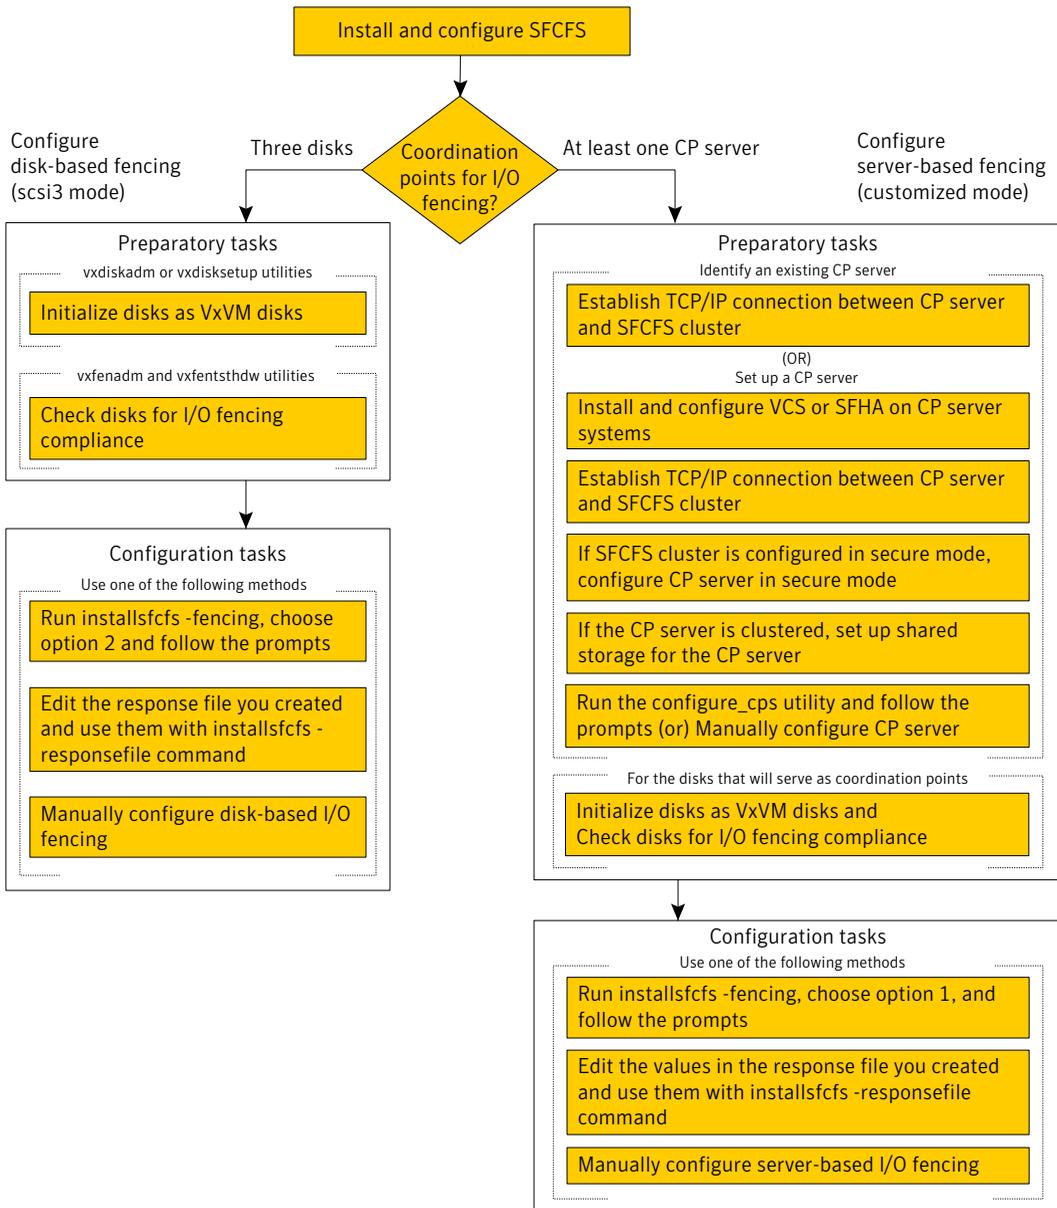
## About planning to configure I/O fencing

After you configure Storage Foundation Cluster File System with the installer, the installer starts Storage Foundation Cluster File System with I/O fencing in disabled mode. To use I/O fencing in the cluster for data integrity, you must configure I/O fencing.

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

**Figure 7-2** illustrates a high-level flowchart to configure I/O fencing for the Storage Foundation Cluster File System cluster.

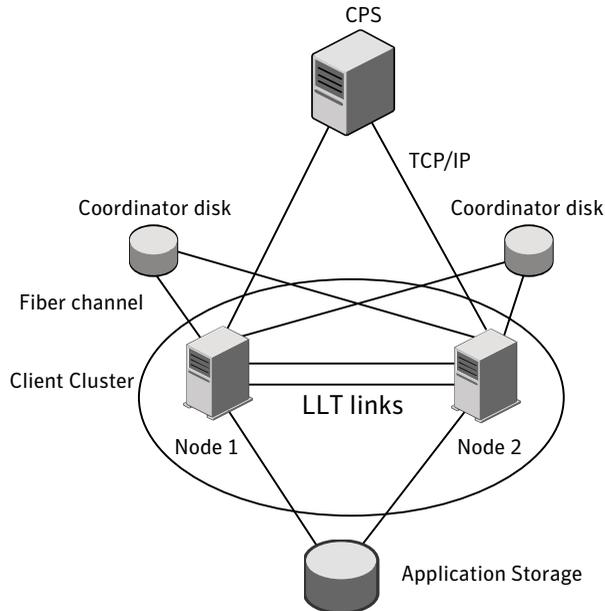
Figure 7-2 Workflow to configure I/O fencing



## Typical SFCFS cluster configuration with server-based I/O fencing

Figure 7-3 displays a configuration using a SFCFS cluster (with two nodes), a single CP server, and two coordinator disks. The nodes within the SFCFS cluster are connected to and communicate with each other using LLT links.

**Figure 7-3** CP server, SFCFS cluster, and coordinator disks



## Recommended CP server configurations

This section discusses the following recommended CP server configurations:

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

---

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

---

Figure 7-4 displays a configuration using a single CP server that is connected to multiple SFCFS clusters with each SFCFS cluster also using two coordinator disks.

Figure 7-4 Single CP server connecting to multiple SFCFS clusters

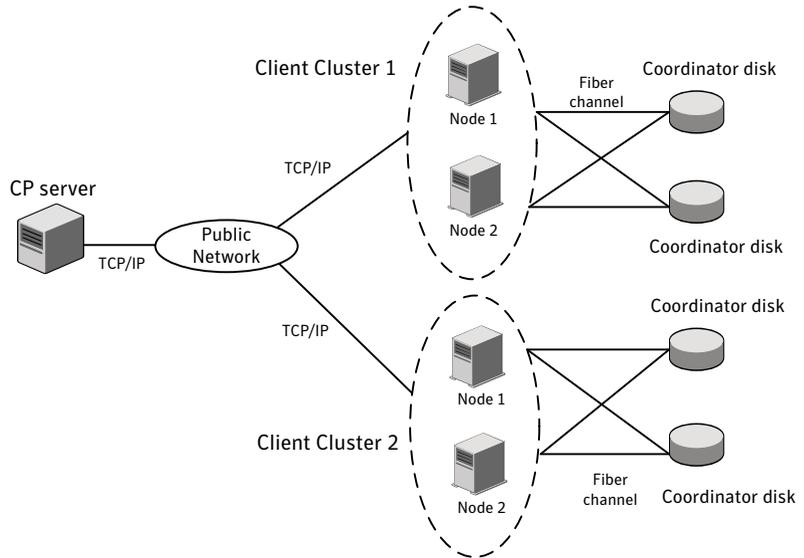
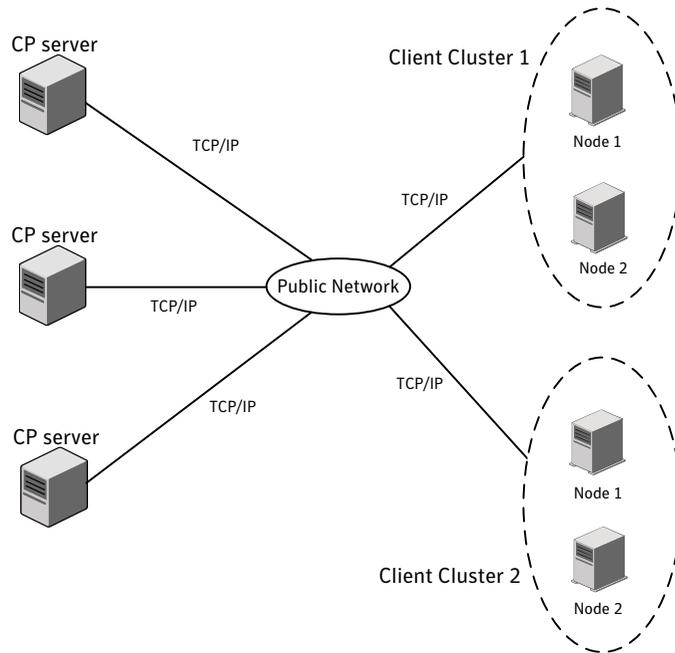


Figure 7-5 displays a configuration using 3 CP servers that are connected to multiple SFCFS clusters.

**Figure 7-5** Three CP servers connecting to multiple SFCFS clusters



For detailed deployment diagrams for server-based fencing:

See [“Configuration diagrams for setting up server-based I/O fencing”](#) on page 367.

## Setting up the CP server

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

### To prepare to configure the CP server

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

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

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

- 3 Decide if you want to secure the communication between the CP server and SFCFS clusters using the Symantec Product Authentication Service (AT).  
Symantec recommends setting up security for the CP server and SFCFS cluster communications.  
For information about configuring security on the CP server:  
See [“Configuring security on the CP server”](#) on page 89.
- 4 Choose a name for the CP server.  
The CP server name should not contain any special characters.
- 5 Choose a port number for the CP server.  
Allocate a TCP/IP port for use by the CP server.  
The default port number is 14250. Alternatively, the user can specify any other valid port from the following valid port range: 49152-65535.
- 6 If CP server is hosted on an SFHA cluster, then set up shared storage for the CP server database.  
For information about setting up shared storage for the CP server database:  
See [“Setting up shared storage for the CP server database”](#) on page 90.
- 7 Choose a valid virtual IP address, network interface, and netmask for the CP server.

## Installing the CP server using the installer

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

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

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

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

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

- 3 For installing CP server on a single node:
  - Install VCS 5.1 onto the system where you are installing the CP server. Installing VCS 5.1 also installs CP server on the system. Refer to the *Veritas™ Cluster Server Installation Guide, Version 5.1* for instructions on installing VCS 5.1.

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

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

## Configuring security on the CP server

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

---

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

---

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

For additional information:

See [“Preparing to configure the clusters in secure mode”](#) on page 69.

## Setting up shared storage for the CP server database

### To set up shared storage for the CP server database

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

For a command example:

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

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

For a command example:

```
# vxdg import cps_dg
```

- 3 Create a mirrored volume over the disk group.

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

For a command example:

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

- 4 Create a file system over the volume.

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

Symantec recommends the vxfs file system type.

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

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

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

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

## Configuring the CP server using the configuration utility

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

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

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

---

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

---

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

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

See “[Setting up the CP server](#)” on page 87.

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

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

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

Use the -n option for using rsh communication.

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

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

```
=====
```

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

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

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

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

```
Enter the option:
```

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

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

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

**5** Enter the name of the CP server.

For example:

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

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

For example:

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

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

For example:

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

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

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

For example:

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

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

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

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

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

For information about configuring security on the CP server:

See [“Configuring security on the CP server”](#) on page 89.

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

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

For a single node VCS configuration for CP server example:

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

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

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

For configuring CP server on an SFHA cluster example:

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

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

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

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

- 10 Review the displayed CP server configuration information.

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

For example:

```
Following is the CP Server configuration information:
```

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

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

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

For a single node VCS configuration for CP server example:

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

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

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

For configuring CP server on an SFHA cluster example:

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

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

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

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

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

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

For a single node VCS configuration for CP server example:

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

For configuring CP server on an SFHA cluster example:

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

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

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

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

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

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

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

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

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

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

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

For example:

```
# hagrp -state CPSSG

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

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

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

For information about the CPSSG, refer to the *Storage Foundation Cluster File System 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 Ensure that the CP server preparation procedures have been performed:
- 2 Stop VCS on each node by using the following command:

```
# hastop -local
```

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

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

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

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

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

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

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

For more information about configuring security on the CP server:

See [“Configuring security on the CP server”](#) on page 89.

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

- 6 Start VCS on all the cluster nodes.

Enter the following command:

```
# hastart
```

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

Enter the following command:

```
# hagrps -state CPSSG
```

Output similar to the following should appear:

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

## Verifying the CP server configuration

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

- /etc/vxcps.conf (CP server configuration file)
- /etc/VRTSvcs/conf/config/main.cf

- `/etc/VRTScps/db` (default location for CP server database)

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

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

where *cp\_server* is the virtual IP/ virtual hostname of the CP server.



# Configuring Storage Foundation Cluster File System

This chapter includes the following topics:

- [Configuring the products using the common product installer](#)
- [Configuring Veritas Volume Manager](#)
- [Configuring Veritas File System](#)
- [Configuring Storage Foundation Cluster File System](#)
- [Configuring the SFDB repository database](#)
- [Setting or changing the product level for keyless licensing](#)
- [Installing Veritas product license keys](#)

## Configuring the products using the common product installer

After installation, you must configure the product. To configure, run the Veritas product installer or the appropriate installation script using the `-configure` option.

To configure Storage Foundations and High Availability Solutions or cluster configurations, refer to that section.

## 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 Volume Manager Administrator's Guide*.

## 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/filesystems
```

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

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

## Configuring Storage Foundation Cluster File System

This section describes configuring Storage Foundation Cluster File System using the Veritas product installer. If you configured Storage Foundation Cluster File System during the installation process, you do not need to perform the procedure in this section.

To configure the product, run the Veritas product installer or the appropriate installation script using the `-configure` option.

### To configure Storage Foundation Cluster File System

- 1 To invoke the common installer, run the `installer` command with the `configure` option, as shown in this example:

```
# ./installer -configure
```

- 2 When the list of available products is displayed, select Veritas Storage Foundation Cluster File System (SFCFS), or Veritas Storage Foundation Cluster File System High Availability (SFCFS HA), enter the corresponding number, and press Return.

```
Select a product to configure:
```

- 3 You are prompted to enter the system names (in the following example, "system01" and "system02") 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 on which to
configure SFCFS: system01 system02
```

- 4 During the initial system check, the installer checks that communication between systems has been set up.

The installer requires that ssh commands used between systems execute without prompting for passwords or confirmations. If the installer hangs or asks for a login password, stop the installer and run it again with the ssh configured for password free logins, or configure rsh and use the -rsh option.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 29.

- 5 The procedure checks system licensing, and you can enter additional licenses, if needed.

```
Checking system licensing
```

```
SFCFS license registered on system01
```

```
Do you want to enter another license key for system01?
```

```
[y,n,q] (n) n
```

- 6 Any running SFCFS processes are stopped. Press **Return** to continue.

- 7** Starting I/O Fencing in enabled mode requires manual intervention after SFCFS Configuration. I/O Fencing can be configured in disabled mode now and it does not require any manual intervention after SFCFS Configuration.

Determine at this time if you plan to configure I/O Fencing in enabled mode or disabled mode, as well as the number of network interconnects (NICs) required on your systems. If you configure I/O Fencing in enabled mode only a single NIC is required, though at least two is recommended.

Enter `y` or `n` for configuring I/O Fencing in enabled mode.

```
Will you be configuring I/O Fencing in enabled mode?  
[y,n,q,?] (y) n
```

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

- 8** No configuration changes are made to the systems until all configuration questions are completed and confirmed. Press **Return** to continue.

All systems are configured to create one cluster.

Enter the unique cluster name and Cluster ID number.

```
Enter the unique cluster name: [?] cluster2  
Enter the unique Cluster ID number between 0-65535: [b,?] 76
```

- 9** The installer discovers the NICs available on the first system and reports them.

```
Discovering NICs on host1 ... discovered en0 en1 en2 en3 en4 en5
```

- 10** To use aggregated interfaces for private heartbeat, enter the name of an aggregated interface. To use a NIC for private heartbeat, enter a NIC which is not part of an aggregated interface.

Enter private heartbeat NIC information for each host.

```
Enter the NIC for the first private heartbeat link
on host1: [b,?] en1
en1 has an IP address configured on it. It could be a public NIC on
system01.
Are you sure you want to use en1 for the first private heartbeat link?
[y,n,q,b,?] (n) y
Is en1 a bonded NIC? [y,n,q] (n) n
Would you like to configure a second private heartbeat link? [y,n,q,b,?]
Enter the NIC for the second private heartbeat link on system01: [b,q,?]
en2 has an IP address configured on it. It could be a public NIC on
system01.
Are you sure you want to use en2 for the second private heartbeat link?
[y,n,q,b,?] (n) y
Is en2 a bonded NIC? [y,n,q] (n)

Would you like to configure a third private heartbeat
link? [y,n,q,b,?] (n) n
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) n
Are you using the same NICs for private heartbeat links
on all systems? [y,n,q,b,?] (y) y
```

---

**Warning:** When answering *y*, be sure that the same NICs are available on each system; the installer may not verify this.

---

Notice that in this example, *en0* is not selected for use as a private heartbeat NIC because it already in use as the public network interface.

- 11** A summary of the information you entered is given. When prompted, confirm that the information is correct.

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

If the information is correct, enter *y*. If the information is not correct, enter *n*. The installer prompts you to enter the information again.

- 12** 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) n
```

- 13** The Veritas Storage Foundation Cluster File System software is verified and configured.

Check the log file, if needed, to confirm the configuration.

Configuration log files, summary file, and response file are saved at:

```
/opt/VRTS/install/logs/installer-****
```

- 14** If you do not plan to use SFCFS file systems to store the Oracle database or binaries, you have completed the SFCFS installation and configuration.

Before installing Oracle binaries (`ORACLE_HOME`), consider these points:

- Local installations provide a comfort level using traditional installation methods and the possibility of improved protection against a single point of failure.
- CFS installations provide a single Oracle installation to manage, regardless of number of nodes. This scenario offers a necessary reduction in storage requirements and easy addition of nodes.

Select the location based on your high availability requirements. Symantec generally recommends using local installations.

See the *Veritas Storage Foundation Cluster File System Administrator's Guide* for more information on Oracle Disk Manager.

## Configuring the SFDB repository database

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

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

# Setting or changing the product level for keyless licensing

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

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

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

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

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

## To set or change the product level

- 1 View the current setting for the product level.

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

- 2 View the possible settings for the product level.

```
# vxkeyless displayall
```

- 3 Set the desired product level.

```
# vxkeyless [-q] set prod_levels
```

where *prod\_levels* is a comma-separated list of keywords, as shown in step 2

If you want to remove keyless licensing and enter a key, you must clear the keyless licenses. Use the NONE keyword to clear all keys from the system.

---

**Warning:** Clearing the keys disables the Veritas products until you install a new key or set a new product level.

---

### To clear the product license level

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

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

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

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

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

## Installing Veritas product license keys

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

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

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

### To install a new license

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

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

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

# Configuring Storage Foundation Cluster File System for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing using `installsfcs`](#)
- [Setting up disk-based I/O fencing manually](#)
- [Setting up server-based I/O fencing using `installsfcs`](#)
- [Setting up server-based I/O fencing manually](#)

## Setting up disk-based I/O fencing using `installsfcs`

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

### Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

**To initialize disks as VxVM disks**

- 1 Scan for the new hdisk devices.

```
# /usr/sbin/cfgmgr
```

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

```
# lsdev -Cc disk
```

**3** Determine the VxVM name by which a disk drive (or LUN) is known.

In the following example, VxVM identifies a disk with the AIX device name `/dev/rhdisk75` as `EMC0_17`:

```
# vxdmpadm getdmpnode nodename=hdisk75
NAME      STATE      ENCLR-TYPE  PATHS    ENBL     DSBL     ENCLR-NAME
=====
EMC0_17   ENABLED    EMC         1        1        0        EMC0
Notice that in the example command, the AIX device name for
the block device was used.
```

As an option, you can run the command `vxdisk list vxvm_device_name` to see additional information about the disk like the AIX device name. For example:

```
# vxdisk list EMC0_17
```

**4** To initialize the disks as VxVM disks, use one of the following methods:

- Use the interactive `vxdiskadm` utility to initialize the disks as VxVM disks. For more information see the *Veritas Volume Managers Administrator's Guide*.
- Use the `vxdisksetup` command to initialize a disk as a VxVM disk.

```
vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i EMC0_17
```

Repeat this command for each disk you intend to use as a coordinator disk.

## Configuring disk-based I/O fencing using installsfcfs

---

**Note:** The installer stops and starts Storage Foundation Cluster File System 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.

---

### To set up disk-based I/O fencing using the installsfcfs

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

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

The installsfcfs 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 5.1 is configured properly.

- 3 Review the I/O fencing configuration options that the program presents. Type **2** to configure disk-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster
[1-3,b,q] 2
```

- 4 Review the output as the configuration program checks whether VxVM is already started and is running.

- If the check fails, configure and enable VxVM before you repeat this procedure.
- If the check passes, then the program prompts you for the coordinator disk group information.

- 5 Choose whether to use an existing disk group or create a new disk group to configure as the coordinator disk group.

The program lists the available disk group names and provides an option to create a new disk group. Perform one of the following:

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.  
The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
- To create a new disk group, perform the following steps:
  - Enter the number corresponding to the **Create a new disk group** option.  
The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks.

Symantec recommends to use three disks as coordination points for disk-based I/O fencing.

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
  - Enter the disk group name.
- 6 Verify that the coordinator disks you chose meet the I/O fencing requirements. You must verify that the disks are SCSI-3 PR compatible using the `vxfcntlsthdw` utility and then return to this configuration program.

See [“Checking shared disks for I/O fencing”](#) on page 113.

- 7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.
- 8 Enter the I/O fencing disk policy that you chose to use. For example:

```
Enter fencing mechanism name (raw/dmp): [b,q,?] raw
```

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
  - Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information
- 9 Verify and confirm the I/O fencing configuration information that the installer summarizes.
- 10 Review the output as the configuration program does the following:
- Stops Storage Foundation Cluster File System 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 Storage Foundation Cluster File System on each node to make sure that the Storage Foundation Cluster File System is cleanly configured to use the I/O fencing feature.

- 11 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.
- 12 Configure the Coordination Point agent to monitor the coordinator disks.  
See [“Configuring Coordination Point agent to monitor coordination points”](#) on page 141.

## Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure SFCFS meets the I/O fencing requirements. You can test the shared disks using the `vxfcntlsthdw` utility. The two nodes must have `ssh` (default) or `rsh` communication. To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxflenadm` 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.

You can use the `vxfcntlsthdw` utility to test disks either in DMP format or in raw format.

- If you test disks in DMP format, use the VxVM command `vxdisk list` to get the DMP path name.
- If you test disks in raw format for Active/Passive disk arrays, you must use an active enabled path with the `vxfcntlsthdw` command. Run the `vxdkmpadm getsubpaths dmpnodename=enclosure-based_name` command to list the active enabled paths.

DMP opens the secondary (passive) paths with an exclusive flag in Active/Passive arrays. So, if you test the secondary (passive) raw paths of the disk, the `vxfcntlsthdw` command may fail due to DMP’s exclusive flag.

The `vxfcntlsthdw` 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 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 114.
- Verifying that nodes have access to the same disk  
See [“Verifying that the nodes have access to the same disk”](#) on page 115.

- Testing the shared disks for SCSI-3  
See “Testing the disks using `vxfcntlhdw` utility” on page 115.

## Verifying Array Support Library (ASL)

Make sure that the Array Support Library (ASL) for the array that you add is installed.

### To verify Array Support Library (ASL)

- 1 If the Array Support Library (ASL) for the array that you add is not installed, obtain and install it on each node before proceeding.

The ASL for the supported storage device that you add is available from the disk array vendor or Symantec technical support.

- 2 Verify that the ASL for the disk array is installed on each of the nodes. Run the following command on each node and examine the output to verify the installation of ASL.

The following output is a sample:

```
# vxddladm listsupport all
```

```
LIBNAME          VID          PID
=====
libvx3par.so     3PARdata    VV
libvxCLARiion.so DGC          All
libvxFJTSYe6k.so FUJITSU     E6000
libvxFJTSYe8k.so FUJITSU     All
libvxcompellent.so COMPELNT    Compellent Vol
libvxcopan.so   COPANSYS    8814, 8818
libvxdds2a.so   DDN         S2A 9550, S2A 9900,
                S2A 9700
```

- 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 SFCFS.
- 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
```

For A/P arrays, run the `vxfcntlsthdw` command only on secondary paths.

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

For example, an EMC disk is accessible by the `/dev/rhdisk75` path on node A and the `/dev/rhdisk76` path on node B.

From node A, enter:

```
vxfenadm -i /dev/rhdisk75
```

```
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/rhdisk76` path.

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

```
# vxfenadm -i /dev/rhdisk77
```

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

## Testing the disks using `vxfcntlsthdw` utility

This procedure uses the `/dev/rhdisk75` 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 `vxfcntlshdw` utility indicates a disk can be used for I/O fencing with a message resembling:

```
The disk /dev/rhdisk75 is ready to be configured for I/O Fencing on
node galaxy
```

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

### To test the disks using `vxfcntlshdw` utility

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

2 From one node, start the utility.

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

```
# vxfcntlshdw [-n]
```

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

---

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

---

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

Do you still want to continue : [y/n] (default: n) y
Enter the first node of the cluster: galaxy
Enter the second node of the cluster: nebula
```

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

```
Enter the disk name to be checked for SCSI-3 PGR on node
IP_adrs_of_galaxy in the format:
for dmp: /dev/vx/rdmp/DiskXX
for raw: /dev/rhdiskXX
Make sure it's the same disk as seen by nodes
IP_adrs_ofgalaxy and IP_adrs_of_nebula
/dev/rhdisk75
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
IP_adrs_of_nebula in the format:
for dmp: /dev/vx/rdmp/DiskXX
for raw: /dev/rhdiskXX
Make sure it's the same disk as seen by nodes
IP_adrs_ofgalaxy and IP_adrs_of_nebula
/dev/rhdisk75
```

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

- 5 Review the output as the utility performs the checks and report its activities.
- 6 If a disk is ready for I/O fencing on each node, the utility reports success:

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

ALL tests on the disk /dev/rhdisk75 have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

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

## Setting up disk-based I/O fencing manually

Tasks that are involved in setting up I/O fencing include:

**Table 9-1** Tasks to set up I/O fencing manually

Action	Description
Initializing disks as VxVM disks	See <a href="#">“Initializing disks as VxVM disks”</a> on page 109.

**Table 9-1** Tasks to set up I/O fencing manually (*continued*)

Action	Description
Identifying disks to use as coordinator disks	See <a href="#">“Identifying disks to use as coordinator disks”</a> on page 118.
Checking shared disks for I/O fencing	See <a href="#">“Checking shared disks for I/O fencing”</a> on page 113.
Setting up coordinator disk groups	See <a href="#">“Setting up coordinator disk groups”</a> on page 118.
Creating I/O fencing configuration files	See <a href="#">“Creating I/O fencing configuration files”</a> on page 119.
Modifying SFCFS configuration to use I/O fencing	See <a href="#">“Modifying VCS configuration to use I/O fencing”</a> on page 120.
Configuring Coordination Point agent to monitor coordination points	See <a href="#">“Configuring Coordination Point agent to monitor coordination points”</a> on page 141.
Verifying I/O fencing configuration	See <a href="#">“Verifying I/O fencing configuration”</a> on page 122.

## Identifying disks to use as coordinator disks

After you add and initialize disks, identify disks to use as coordinator disks.

See [“Initializing disks as VxVM disks”](#) on page 109.

### 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 113.

## Setting up coordinator disk groups

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

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

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

The following example procedure assumes that the disks have the device names EMC0\_12, EMC0\_16, and EMC0\_17.

#### To create the vxfencoorddg disk group

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

```
# vxdg init vxfencoorddg EMC0_12 EMC0_16 EMC0_17
```

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

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

## Creating I/O fencing configuration files

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

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

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

- 1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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

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

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

- 4 Edit the following file on each node in the cluster to change the values of the VXFEN\_START and the VXFEN\_STOP environment variables to 1:

```
/etc/default/vxfen
```

## Modifying VCS configuration to use I/O fencing

After you add coordinator disks and configure I/O fencing, add the UseFence = SCSI3 cluster attribute to the VCS configuration file /etc/VRTSvcs/conf/config/main.cf. If you reset this attribute to UseFence = None, VCS does not make use of I/O fencing abilities while failing over service groups. However, I/O fencing needs to be disabled separately.

### To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

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

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

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

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

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

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

- 6 Save and close the file.

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

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

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

For example, on each remaining node, enter:

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

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

- Start the I/O fencing driver.

The `vxfen` startup script also invokes the `vxfenconfig` command, which configures the `vxfen` driver to start and use the coordinator disks that are listed in `/etc/vxfentab`.

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

- Start VCS.

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

## Verifying I/O fencing configuration

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

### To verify I/O fencing configuration

- ◆ On one of the nodes, type:

```
# vxfenadm -d

I/O Fencing Cluster Information:
=====

Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:

* 0 (galaxy)
  1 (nebula)

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

## Setting up server-based I/O fencing using installsfcfs

If Storage Foundation Cluster File System cluster is configured to run in secure mode, then verify that the configuration is correct before you configure CP server-based I/O fencing.

See [“Verifying security configuration on SFCFS cluster to use CP server coordination point ”](#) on page 123.

See [“Configuring server-based I/O fencing”](#) on page 125.

## Verifying security configuration on SFCFS cluster to use CP server coordination point

After configuring security using the `installsfcfs -security` command, follow the procedure below on each SFCFS cluster node to confirm that security is correctly configured.

**To verify the security configuration on SFCFS cluster to use CP server coordination point**

- 1 Run the following command:

```
# /opt/VRTScps/bin/cpsat listpd -t local
```

```
Domain(s) Found 1
```

```
*****
```

```
Domain Name HA_SERVICES@galaxy.symantec.com
```

```
Expiry Interval 0
```

```
*****
```

- 2 There should be a domain name entry with the following format:

```
HA_SERVICES@hostname.domainname
```

or

```
HA_SERVICES@hostname
```

**3** There should not be duplicate entries for HA\_SERVICES domain.

An example of incorrect configuration is given below.

```
showdomains

Domain(s) Found :          3

*****

Domain Name:      HA_SERVICES@galaxy.symantec.com

Domain Type:     vx

*****

Domain Name:      broker@galaxy.symantec.com

Domain Type:     vx

*****

Domain Name:      HA_SERVICES@galaxy

Domain Type:     vx

*****
```

Proceed to reconfigure security in case duplicate entries appear as shown in the above example.

## Configuring server-based I/O fencing

This section describes how to configure server-based I/O fencing for the SFCFS cluster. With server-based I/O fencing, a combination of CP servers and SCSI-3 compliant coordinator disks can act as coordination points for I/O fencing.

### To configure the SFCFS cluster with server-based I/O fencing

- 1 Ensure that the CP server(s) are configured and reachable from the cluster. If coordinator disks are to be used as coordination points, ensure that they are SCSI-3 compliant.
- 2 Run the `installsfcfs -fencing` command to configure fencing.

For example:

```
/opt/VRTS/install/installsfcfs -fencing
```

The installer creates a `vxfenmode` file on each node. The file is located at `/etc/vxfenmode`.

The following procedure can be used as an example to configure server-based I/O fencing. In this procedure example, there is one CP server and two disks acting as the coordination points.

### To configure fencing configuration using the installer - CP client-based fencing

- 1 After installing and configuring VCS on the SFCFS cluster, the user issues the following command for configuring fencing:

```
/opt/VRTS/install/installsfcfs -fencing
```

- 2 After issuing the command, the installer displays Symantec copyright information and the location of log files for the configuration process.

Access and review these log files if there is any problem with the installation process. The following is an example of the command output:

```
Logs for installsfcfs are being created in /var/tmp/installsfcfs-LqwKwB.
```

- 3 Next, the installer displays the current cluster information for verification purposes. The following is an example of the command output:

```
Cluster information verification:
```

```
Cluster Name: clus1  
Cluster ID Number: 4445  
Systems: galaxy nebula
```

The cluster name, systems, and ID number are all displayed.

You are then asked whether you want to configure I/O fencing for the cluster. Enter "y" for yes. The rsh (or ssh) communication with the cluster nodes is then checked by the installer.

- 4** Next, you are prompted to select one of the following options for your fencing configuration:

```
Fencing configuration
```

- 1) Configure CP client based fencing
- 2) Configure disk based fencing
- 3) Configure fencing in disabled mode

```
Select the fencing mechanism to be configured in this  
Application Cluster [1-3,q]
```

Select the first option for CP client-based fencing.

- 5** Enter the total number of coordination points including both servers and disks. This number should be at least 3.

For example:

```
Enter the total number of co-ordination points including both  
CP servers and disks: [b] (3)
```

- 6** Enter the total number of coordinator disks among the coordination points. In this example, there are two coordinator disks.

For example:

```
Enter the total number of disks among these:  
[b] (0) 2
```

- 7** Enter the Virtual IP addresses or host names of the virtual IP address for each of the Coordination Point servers.

---

**Note:** The installer assumes these values to be the identical as viewed from all the client cluster nodes.

---

For example:

```
Enter the Virtual IP address/fully qualified host name  
for the Co-ordination Point Server #1::  
[b] 10.209.80.197
```

**8** Enter the port that the CP server would be listening on.

For example:

```
Enter the port in the range [49152, 65535] which the
Co-ordination Point Server 10.209.80.197
would be listening on or simply accept the default port suggested:
[b] (14250)
```

**9** Enter the fencing mechanism for the disk or disks.

For example:

```
Enter fencing mechanism for the disk(s) (raw/dmp):
[b,q,?] raw
```

**10** The installer then displays a list of available disks to choose from to set up as coordinator points.

```
Select disk number 1 for co-ordination point
```

- 1) c3t0d0s2
- 2) c3t1d0s3
- 3) c3t2d0s4

```
Please enter a valid disk which is available from all the
cluster nodes for co-ordination point [1-3,q] 1
```

Select a disk from the displayed list.

Ensure that the selected disk is available from all the SFCFS cluster nodes.

- 11** Read the displayed recommendation from the installer to verify the disks prior to proceeding:

```
It is strongly recommended to run the 'VxFen Test Hardware' utility
located at '/opt/VRTSvcs/vxfen/bin/vxfentsthdw' in another window
before continuing. The utility verifies if the shared storage
you intend to use is configured to support I/O
fencing. Use the disk you just selected for this
verification. Come back here after you have completed
the above step to continue with the configuration.
```

Symantec recommends that you verify that the disks you are using as coordination points have been configured to support I/O fencing. Press Enter to continue.

You are then prompted to confirm your disk selection after performing a 'vxfentsthdw' test.

Press Enter to accept the default (y) and continue.

- 12** The installer then displays a list of available disks to choose from to set up as coordinator points.

Select a disk from the displayed list for the second coordinator point.

Ensure that the selected disk is available from all the SFCFS cluster nodes.

- 13** Proceed to read the displayed recommendation from the installer to verify the disks prior to proceeding.

Press Enter to continue.

- 14** You are then prompted to confirm your disk selection after performing a 'vxfentsthdw' test.

Press Enter to accept the default (y) and continue.

- 15** Proceed to enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):
[b] (vxfencoorddg)
```

- 16** The installer now begins verification of the coordination points. At the end of the verification process, the following information is displayed:

- Total number of coordination points being used
- CP Server Virtual IP/hostname and port number
- SCSI-3 disks

- Disk Group name for the disks in customized fencing
- Disk mechanism used for customized fencing

For example:

```
Total number of coordination points being used: 3
CP Server (Port):
  1. 10.209.80.197 (14250)
SCSI-3 disks:
  1. c3t0d0s2
  2. c3t1d0s3
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk mechanism used for customized fencing: raw
```

You are then prompted to accept the above information. Press Enter to accept the default (y) and continue.

The disks and disk group are initialized and the disk group deported on the SFCFS cluster node.

- 17** The installer now automatically determines the security configuration of the CP server's side and takes the appropriate action:
- If the CP server's side is configured for security, then the SFCFS cluster's side will be configured for security.
  - If the CP server's side is not configured for security, then the SFCFS cluster's side will not be configured for security.

For example:

```
While it is recommended to have secure communication
configured between CP Servers and CP client cluster, the client
cluster must be in the same mode (secure or non-secure) as the
CP servers are.
```

```
Since the CP servers are configured in secure mode, the installer
will configure the client cluster also as a secure cluster.
```

```
Press [Enter] to continue:
```

```
Trying to configure Security on the cluster:
```

```
All systems already have established trust within the
```

```
Symantec Product Authentication Service domain  
root@galaxy.symantec.com
```

- 18** Enter whether you are using different root brokers for the CP servers and SFCFS clusters.

If you are using different root brokers, then the installer tries to establish trust between the authentication brokers of the CP servers and the SFCFS cluster nodes for their communication.

After entering "y" for yes or "n" for no, press Enter to continue.

- 19** If you entered "y" for yes in step 18, then you are also prompted for the following information:
- Hostname for the authentication broker for any one of the CP servers
  - Port number where the authentication broker for the CP server is listening for establishing trust
  - Hostname for the authentication broker for any one of the SFCFS cluster nodes
  - Port number where the authentication broker for the SFCFS cluster is listening for establishing trust

Press Enter to continue.

- 20** The installer then displays your I/O fencing configuration and prompts you to indicate whether the displayed I/O fencing configuration information is correct.

If the information is correct, enter "y" for yes.

For example:

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm  
Cluster ID: 2122  
Cluster Name: clus1  
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

- 21** The installer then updates the SFCFS cluster information on each of the CP Servers to ensure connectivity between them.

The installer then populates the file `/etc/vxfenmode` with the above details in each of the CP SFCFS cluster nodes.

For example:

```
Updating client cluster information on CP Server 10.210.80.199

Adding the client cluster to the CP Server 10.210.80.199 ..... Done

Registering client node galaxy with CP Server 10.210.80.199..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Registering client node nebula with CP Server 10.210.80.199 ..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Updating /etc/vxfenmode file on galaxy ..... Done
Updating /etc/vxfenmode file on nebula ..... Done
```

For additional information about the `vxfenmode` file in mixed disk and CP server mode, or pure server-based mode:

See [“About I/O fencing configuration files”](#) on page 80.

- 22** You are then prompted to configure the CP agent on the client cluster.

```
Do you want to configure CP Agent on the client cluster? [y,n,q]
(y)

Enter a non-existing name for the service group for CP Agent:
[b] (vxfen)

Adding CP Agent via galaxy ..... Done
```

- 23 The VCS and the fencing process are then stopped and restarted on each SFCFS cluster node, and the I/O configuration process then finished.

```
Stopping VCS on galaxy ..... Done
Stopping Fencing on galaxy ..... Done
Stopping VCS on nebula ..... Done
Stopping Fencing on nebula ..... Done
```

- 24 At the end of this process, the installer then displays the location of the configuration log files, summary files, and response files.

## Setting up server-based I/O fencing manually

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

**Table 9-2** Tasks to set up server-based I/O fencing manually

Action	Description
Preparing the CP servers for use by the Storage Foundation Cluster File System cluster	See <a href="#">“Preparing the CP servers manually for use by the SFCFS cluster”</a> on page 133.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See <a href="#">“Configuring server-based fencing on the SFCFS cluster manually”</a> on page 137.
Configuring Coordination Point agent to monitor coordination points	See <a href="#">“Configuring Coordination Point agent to monitor coordination points”</a> on page 141.
Verifying the server-based I/O fencing configuration	See <a href="#">“Verifying server-based I/O fencing configuration”</a> on page 143.

### Preparing the CP servers manually for use by the SFCFS cluster

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

[Table 9-3](#) displays the sample values used in this procedure.

**Table 9-3** Sample values in procedure

CP server configuration component	Sample name
CP server	mycps1.symantecexample.com

**Table 9-3** Sample values in procedure (*continued*)

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

**To manually configure CP servers for use by the SFCFS cluster**

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

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

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster clus1

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2}
```

- 2 Check whether the SFCFS cluster and nodes are present in the CP server.

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

ClusName  UUID                               Hostname(Node ID) Registered
clus1     {f0735332-1dd1-11b2} galaxy(0)          0
clus1     {f0735332-1dd1-11b2} nebula(1)         0
```

If the output does not show the cluster and nodes, then add them as described in the next step.

**3** Add the SFCFS cluster and nodes to each CP server.

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

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

Cluster clus1 added successfully

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

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

Node 0 (galaxy) successfully added

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

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

Node 1 (nebula) successfully added

**4** If security is to be enabled, check whether the `_HA_VCS_` users are created in the CP server.

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

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

Username/Domain Type	Cluster Name / UUID	Role
<code>_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com/vx</code>	<code>clus1/{f0735332-1dd1-11b2}</code>	Operator
<code>_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com/vx</code>	<code>clus1/{f0735332-1dd1-11b2}</code>	Operator

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

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

5 Add the users to the CP server.

First, determine the user@domain to be added.

The user for fencing should be of the form `_HA_VCS_`*short-hostname* and domain name is that of HA\_SERVICES user in the output of command:

```
# /opt/VRTScps/bin/cpsat listpd -t local
```

Next, issue the following commands on the CP server (mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com  
successfully added
```

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com  
successfully added
```

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

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

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

```
Cluster successfully added to user  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com privileges.
```

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

```
Cluster successfully added to user  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com privileges.
```

## Configuring server-based fencing on the SFCFS cluster manually

The configuration process for the client or SFCFS cluster to use CP server as a coordination point requires editing the `/etc/vxfenmode` file. You need to edit this file to specify the following information for your configuration:

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

Whenever coordinator disks are used as coordination points in your I/O fencing configuration, a disk group (vxfendg) has to be created. This disk group has to be

specified in the `/etc/vxfenmode` file. For information about creating the disk group, see the *Veritas™ Cluster Server Installation Guide*.

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

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

`/etc/default/vxfen`

Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

The following file output provides an example of what the `/etc/vxfenmode` file contains:

```
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps        - use a coordination point server with optional script
#              controlled scsi3 disks
#
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
```

```
scsi3_disk_policy=dmp

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

#
# Specify 3 or more odd number of coordination points in this file,
# one in each row. They can be all-CP servers, all-SCSI-3 compliant
# coordinator disks, or a combination of CP servers and SCSI-3
# compliant coordinator disks. Please ensure that the CP server
# coordination points are numbered sequentially and in the same
# order on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as:
#
# cps<number>=<Virtual IP/ Virtual hostname of cp server> in
# square brackets ([]), followed by ":" and CPS port number.
#
# Examples:
# cps1=[192.168.0.23]:14250
# cps2=[mycps.company.com]:14250
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoordg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
```

```
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
```

[Table 9-4](#) defines the vxfenmode parameters that must be edited.

**Table 9-4** 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”. <b>Note:</b> The configured disk policy is applied on all the nodes.
security	Security parameter 1 indicates that Symantec Product Authentication Service is used for CP server communications. Security parameter 0 indicates that communication with the CP server is made in non-secure mode. The default security value is 1. <b>Note:</b> Symantec only supports a configuration where both the CP server and client sides have the same security setting. The security setting on both sides must be either enabled or disabled.

**Table 9-4** vxfenmode file parameters (*continued*)

vxfenmode File Parameter	Description
cps1, cps2, cps3, or vxfendg	<p>Coordination point parameters.</p> <p>Enter either the Virtual IP address or FQHN (whichever is accessible) of the CP server.</p> <p><b>Note:</b> Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfendg) and specified in the /etc/vxfenmode file. Additionally, the customized fencing framework also generates the /etc/vxfentab file which specifies the security setting and the coordination points (all the CP servers and the disks from disk group specified in /etc/vxfenmode file).</p>

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

For example:

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

## Configuring Coordination Point agent to monitor coordination points

The following procedure describes how to manually configure the Coordination Point agent to monitor coordination points (CP server or SCSI-3 disks).

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

### To configure Configuration Point agent to monitor coordination points

- 1 Ensure that your SFCFS cluster has been properly installed and configured with fencing enabled.
- 2 Create a parallel service group vxfen and add a coordpoint resource to the vxfen service group as follows:

```
# haconf -makerw
# hagr -add vxfen
# hagr -modify vxfen SystemList system1 0 system2 1
# hagr -modify vxfen AutoFailOver 0
# hagr -modify vxfen Parallel 1
# hagr -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 1
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

- 3 Verify the status of the agent on the SFCFS cluster using the `hares` commands.

For example:

```
# hares -state coordpoint
```

The following is an example of the command and output::

```
# hares -state

# Resource      Attribute      System      Value
coordpoint     State          galaxy      ONLINE
```

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

The agent log contains detailed Coordination Point agent monitoring information; including information about whether the Coordination Point agent is able to access all the coordination points, information to check on which coordination points the Coordination Point agent is reporting missing keys, etc.

To view all such information in the engine log, change the `dbg` level for that node using the following command:

```
# hatype -modify coordpoint LogDbg 10
```

The agent log can now be viewed at the following location:

```
/var/VRTSvcs/log/engine_A.log
```

## Verifying server-based I/O fencing configuration

During the SFCFS cluster installation, the installer populates the following files based on inputs that are received during the configuration phase:

- `/etc/vxfenmode` (edited for CP server)
- `/etc/vxfentab` (edited for CP server)

Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

For troubleshooting server-based I/O fencing configuration issues, refer to the *Storage Foundation Cluster File System Administrator's Guide*.

Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```



# Upgrading Storage Foundation Cluster File System

This chapter includes the following topics:

- [About upgrading Storage Foundation Cluster File System and High Availability](#)
- [Upgrading Storage Foundation Cluster File System and High Availability software from a release prior to 5.1](#)
- [Upgrading Veritas Volume Replicator](#)
- [Upgrading from SFHA 5.1 to SFCFS 5.1](#)

## About upgrading Storage Foundation Cluster File System and High Availability

Perform the procedures in the following sections to upgrade Storage Foundation Cluster File System and High Availability. You can perform an upgrade to Storage Foundation using the Veritas product installer or product installation script if you already have Storage Foundation Cluster File System installed.

See [“Upgrading Storage Foundation Cluster File System and High Availability software from a release prior to 5.1”](#) on page 146.

# Upgrading Storage Foundation Cluster File System and High Availability software from a release prior to 5.1

This section contains procedures for the Veritas Storage Foundation Cluster File System upgrade.

## Upgrade paths for Veritas Storage Foundation Cluster File System 5.1

[Table 10-1](#) shows the upgrade paths for Veritas Storage Foundation Cluster File System.

**Table 10-1** Upgrade paths

From	Upgrade to	Tasks
Storage Foundation Cluster File System 4.1 MP4	Storage Foundation Cluster File System 5.1	See <a href="#">“Overview of procedures”</a> on page 151.
Storage Foundation Cluster File System 5.0, 5.0MP1, and 5.0MP3	Storage Foundation Cluster File System 5.1	

## Planning the upgrade

Complete the following tasks in advance of upgrading:

- Review the *Veritas Storage Foundation Release Notes* for any late-breaking information on upgrading your system.
- Be sure that the administrator doing the upgrade has root access and a working knowledge of system administration.
- Schedule sufficient outage time for the upgrade.
- Make sure you have upgraded all the file systems to disk layout Version 6, before you upgrade SFCFS to 5.1. Disk layout Version 7 is the recommended version for SFCFS 5.1.

See `vxupgrade(1M)`, `vxconvert(1M)`, and `fsadm(1M)` manual pages.

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

- Verify all the file systems are working fine and data is intact.  
See the `cfsmount(1M)` manual page.

## Preparing the system and backing up files before upgrading

Before upgrading an installed Veritas Storage Foundation Cluster File System, preserve the existing configuration information.

To preserve the existing configuration information, perform the following actions:

- Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/filesystems` file. You will need to recreate these entries in the `/etc/filesystems` file on the freshly installed system.
- Before upgrading, ensure that you have made backups of all data that you want to preserve. In particular, you will need the information in such as `/etc/filesystems`. 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.
- Use the `vxlicrep` command to make a record of the currently installed Veritas licenses.
- Back up the configuration files.

```
# cp -r /etc/VRTSvcS/conf/ backupdirectory/
```

- If Veritas Volume Replicator (VVR) is configured, do the following steps in the order shown:
  - Make sure that the disk groups that contain RVGs are at least at disk group version 110.

```
# vxdg list diskgroup
```

- Make sure the size of the SRL volume is greater than 110 MB.  
Refer to the *Veritas Volume Replicator Administrator's Guide*.
- Stop all the applications involved in replication. For example, if a data volume contains a file system, unmount it.
- Verify that all of 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.

---

## Preparing for upgrade of VVR in the presence of VCS agents

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:** Write down 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 150.

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

```
# vxdctl -c mode
```

Note the master and record it for future use.

### Determining the nodes on which disk groups are online

For private disk groups, determine and note down the hosts on which the disk groups containing RVG resources are imported. This information is required for restoring the configuration after the upgrade.

#### To determine the online disk groups

- 1** On any node in the cluster, list the disk groups in your configuration, and note down the disk group names listed in the output for future use:

```
# hares -display -type RVG -attribute DiskGroup
```

---

**Note:** Write down the list of the disk groups that are under VCS control.

---

- 2** For each disk group listed in the output in step 1, list its corresponding disk group resource name:

```
# hares -list DiskGroup=diskgroup Type=DiskGroup
```

- 3** For each disk group resource name listed in the output in step 2, get and note down the node on which the disk group is imported by typing the following command:

```
# hares -display dg_resname -attribute State
```

The output displays the disk groups that are under VCS control and nodes on which the disk groups are imported.

### Preparing for the upgrade when VCS agents are configured

If you have configured the VCS agents, it is recommended that you take backups of the configuration files, such as `main.cf` and `types.cf`, which are present in the `/etc/VRTSvcs/conf/config` directory.

### To prepare a configuration with VCS agents for an upgrade

- 1 List the disk groups on each of the nodes by typing the following command on each node:

```
# vxdisk -o alldgs list
```

The output displays a list of the disk groups that are under VCS control and the disk groups that are not under VCS control.

---

**Note:** The disk groups that are not locally imported are displayed in parentheses.

---

- 2 If any of the disk groups have not been imported on any node, import them. For disk groups in your VCS configuration, you can import them on any node. For disk groups that are not under VCS control, choose an appropriate node on which to import the disk group. Enter the following command on the appropriate node:

```
# vxdg -t import diskgroup
```

- 3 If a disk group is already imported, then recover the disk group by typing the following command on the node on which it is imported:

```
# vxrecover -bs
```

- 4 Verify that all the Primary RLINKs are up to date.

```
# vxrlink -g diskgroup status rlink_name
```

---

**Note:** Do not continue until the Primary RLINKs are up-to-date.

---

## Overview of procedures

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation Cluster File System: phased and full.

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation High Availability or Storage Foundation for Oracle High Availability: phased and full.

---

**Note:** If VVR is configured, phased upgrade is not supported. We recommend that the secondary cluster be upgraded before the primary cluster in the RDS.

---

The upgrade procedures apply to both the phased and full upgrade procedures unless otherwise noted. Occasionally, steps differ between the two procedures. Screen output is also common between both procedures unless otherwise noted.

---

**Note:** Both procedures automatically uninstall the previous version of the software.

---

## Phased upgrade

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time. Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

The stages of the phased upgrade procedure are:

- Select two or more nodes to upgrade.
- Install the new version.
- Shut down VCS on remaining non-upgraded nodes and ensure the file systems are clean.
- Reboot the upgraded nodes.
- Install the new version on each remaining node and reboot them.

## Full upgrade

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

The stages of the full upgrade procedure are:

- Ensure the file systems are clean.
- Install the new version.
- Reboot the upgraded nodes.

## Performing a phased upgrade of SFCFSHA stack from version 5.0MP3

Performing a phased upgrade involves the following tasks:

- Moving the service groups to the second subcluster

- Upgrading the SFCFSHA stack on the first subcluster
- Preparing the second subcluster
- Activating the first subcluster
- Upgrading the operating system on the second subcluster
- Upgrading the second subcluster
- Finishing the phased upgrade

Before you start the upgrade on the first half of the cluster, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.

## Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade.

## Planning for a phased upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group. Some basic guidelines follow:

- Split the cluster in half. If the cluster has an odd number of nodes, calculate  $(n+1)/2$ , and start the upgrade with the even number of nodes.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.

## Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules. Also, do not add or remove service groups to any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.

## Moving the service groups to the second subcluster

### To move the service groups to the second subcluster

- 1 Switch failover groups from the first half of the cluster to one of the nodes in the second half of the cluster. In this procedure, `galaxy` is a node in the first half of the cluster and `jupiter` is a node in the second half of the cluster. Enter the following:

```
# hagr -switch failover_group -to jupiter
```

- 2 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the applications.
- 3 On the first half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS.

```
# mount | grep vxfs
```

Verify that no processes use the VxFS or CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS or CFS mount point with the mechanism provided by the application.

Unmount the VxFS or CFS file system. Enter the following:

```
# umount /mount_point
```

- 4 On the first half of the cluster, bring all the VCS service groups offline including CVM group. Enter the following:

```
# hagr -offline group_name -sys galaxy
```

When the CVM group becomes OFFLINE, all the parallel service groups such as the CFS file system will also become OFFLINE on the first half of the cluster nodes.

- 5 Verify that the VCS service groups are offline on all the nodes in first half of the cluster. Enter the following:

```
# hagr -state group_name
```

- 6 Freeze the nodes in the first half of the cluster. Enter the following:

```
# haconf -makerw

# hasys -freeze -persistent galaxy

# haconf -dump -makero
```

- 7 If IO fencing is enabled, then on each node of the first half of the cluster, change the contents of the `/etc/vxfenmode` file to configure I/O fencing in disabled mode. Enter the following:

```
# cp /etc/vxfen.d/vxfenmode_disabled /etc/vxfenmode
# cat /etc/vxfenmode
[root@swlx08 ~]# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=disabled
```

- 8 If the cluster-wide attribute `UseFence` is set to `SCSI3`, then reset the value to `NONE` in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

- 9 Verify that only GAB ports a, b, d and h are open. Enter the following:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port h gen ada40f membership 01
```

Do not stop VCS. Port h should be up and running.

- 10 In the first half of the cluster, stop all VxVM and CVM volumes. Enter the following command for each disk group:

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

Verify that no volumes remain open. Enter the following:

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

- 11 On first half of the cluster, upgrade the operating system on all the nodes, if applicable. For instructions, see the upgrade paths for the operating system. See [“Supported AIX operating systems”](#) on page 44.

## Upgrading the SFCFSHA stack on the first subcluster

### To upgrade the SFCFSHA stack on the first subcluster

---

- ◆ **Note:** This procedure is based on an "in-place" upgrade path; that is, if the operating system is upgraded, the release will be the same, but only the path level will change. If you are moving from major operating system release to another, you must uninstall the SFCFSHA stack before you upgrade the operating system. After the operating system is upgraded, you must reinstall SFCFSHA.
- 

On the first half of the cluster, upgrade SFCFSHA by using the `installsfcfs` script. For example use the `installsfcfs` script as shown below:

```
# ./installsfcfs galaxy
```

After the upgrade for first half of the cluster is complete, no GAB ports will be shown in `gabconfig -a` output.

To upgrade your operating system, follow the normal procedures for your platform.

## Preparing the second subcluster

### To prepare the second subcluster

- 1 On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application.  
[Downtime starts now.]
- 2 On the second half of the cluster, unmount the VxFS and CFS file systems that are not managed by VCS. Enter the following:

```
# mount | grep vxfs
```

Verify that no processes use the VxFS and CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS and CFS mount point with the mechanism provided by the application.

Unmount the VxFS and CFS file system. Enter the following:

```
# umount /mount_point
```

- 3 On the second half of the cluster, unfreeze all the VCS service groups on all the nodes using the following commands:

```
# haconf -makerw  
# hagr -unfreeze group_name -persistent  
# haconf -dump -makero
```

- 4 On the second half of the cluster, bring all the VCS service groups offline, including CVM group. Enter the following:

```
# hagr -offline group_name -sys jupiter
```

- 5 On the second half of the cluster, verify that the VCS service groups are offline. Enter the following:

```
# hagr -state group_name
```

- 6 Stop VCS on the second half of the cluster. Enter the following:

```
# hstop -local
```

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

```
# cp /etc/vxfen.d/vxfenmode_disabled /etc/vxfenmode
# cat /etc/vxfenmode#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=disabled
```

- 8 If the cluster-wide attribute `UseFence` is set to `SCSI3`, reset the value to `NONE` in the `/etc/VRTSvcs/conf/config/main.cf` file, in second half of the cluster.
- 9 On the second half on cluster, stop the following SFCFSHA modules: `VCS`, `VxFEN`, `ODM`, `GAB`, and `LLT`. Enter the following:

```
# /etc/methods/glmkextadm unload
# /etc/rc.d/rc2.d/s99odm stop
# /etc/methods/gmskextadm status
# /etc/init.d/vxfen.rc stop
# /etc/init.d/gab.rc stop
# /etc/init.d/llt.rc stop
```

- 10** On each node in the first half of the cluster, enable fencing. Enter the following:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# sybase     - use scsi3 disks in kernel but coordinate membership with
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=scsi3
#
# scsi3_disk_policy determines the way in which I/O Fencing communicates
# the coordination disks.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
scsi3_disk_policy=dmp
```

- 11** If the cluster-wide attribute UseFence is set to NONE, reset the value to SCSI3 in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

## Activating the first subcluster

### To activate the first subcluster

- 1 Restart the upgraded nodes; that is, the nodes in first half of the cluster. Enter the following:

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

When the first half of the cluster nodes come up, no GAB ports are OPEN. The following command does not show any GAB ports:

```
# /sbin/gabconfig -a
GAB Port Memberships
```

```
=====
```

- 2 Force gab to form a cluster after the upgraded nodes are rebooted in first half of the cluster.

```
# /sbin/gabconfig -xc
```

GAB ports a, b, d and h appear in `gabconfig -a` command output.

- 3 On first half of the cluster, unfreeze all the upgraded nodes. Enter the following:

```
# haconf -makerw
# hasys -unfreeze -persistent node_name
# haconf -dump -makero
```

- 4 On the first half of the cluster, bring the VCS service groups online. Enter the following:

```
# hagrps -online group_name -sys node_name
```

After you bring the CVM service group ONLINE, all the GAB ports v, w and f come ONLINE and all the CFS mounts service groups also come ONLINE automatically. Only failover service groups need to be brought ONLINE manually.

- 5 Manually mount the VxFS and CFS file systems that are not managed by VCS. [Downtime ends now.]

## Upgrading the operating system on the second subcluster

### To upgrade the operating system on the second subcluster

- ◆ Enter the following.

On the second half of the cluster, upgrade the operating system, if applicable. For instructions, see the upgrade paths for the operating system.

See [“Supported AIX operating systems”](#) on page 44.

## Upgrading the second subcluster

### To upgrade the second subcluster

- ◆ Enter the following:

```
# ./installsfcfs node_name
```

## Completing the phased upgrade

### To complete the phased upgrade

- 1 Verify that the cluster UUID on the nodes in the second subcluster is the same as the cluster UUID on the nodes in the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -display nodename
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first subcluster to the nodes in the second subcluster. For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -copy -from_sys node01 -to_sys
```

- 2 On each node in the second half of the cluster, enable fencing. Enter the following:

```
# cp /etc/vxfen.d/vxfenmode_scsi3_dmp /etc/vxfenmode
# cat /etc/vxfenmode
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# sybase     - use scsi3 disks in kernel but coordinate membership with Syb
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=scsi3
#
# scsi3_disk_policy determines the way in which I/O Fencing communicates wi
# the coordination disks.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
scsi3_disk_policy=dmp
```

- 3 Restart the upgraded nodes; that is, the nodes in second half of the cluster. Enter the following:

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

When second half of the nodes come up, all the GAB ports a, b, d, h, v, w and f are ONLINE. Also all the CFS mounts service groups come online automatically.

- 4 Manually mount the VxFS and CFS file systems that are not managed by VCS in the second half of the cluster.

## Performing a full upgrade

Performing a full upgrade involves the following tasks:

- Ensuring that the file systems are clean
- Performing the upgrade
- Updating the configuration and confirming startup
- Upgrading the remaining nodes

### Ensuring the file systems are clean

Before upgrading to SFCFS 5.1, ensure that the file systems are clean. To ensure that the logs have been replayed and the file systems are marked clean:

**To ensure the file systems are clean**

- 1 Log in as superuser onto any node in the cluster.
- 2 Offline the group on each node of the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

Repeat this step for each SFCFS service group.

- 3 Check and repair each file system:

```
# fsck -V vxfs /dev/vx/dsk/diskgroup/volume
```

The `fsck` command in `/opt/VRTS/bin` accepts either the block or character device (`/dev/vx/dsk/dg/vol`) or (`/dev/vx/rdisk/dg/vol`). The operating system version of `fsck` may limit the device types it accepts.

For more information, see the `fsck` and `fsck_vxfs` man pages.

Repeat this step for each file system.

**Performing the upgrade (phased or full)**

This section describes how to upgrade for phased or full.

If you are performing a phased upgrade, select one or more nodes to upgrade.

**To perform the upgrade**

- 1 Log in as superuser.
- 2 Insert the appropriate media disc into your system's DVD-ROM drive.
- 3 If volume management software is running on your system, the software disc automatically mounts as `/mnt/cdrom`.

If volume management software is not available to mount the CD-ROM, you must mount it manually, enter:

```
# mkdir -p /mnt/cdrom
# mount -V cdrfs -o ro /dev/cd0 /mnt/cdrom
```

- 4 Change to the top-level directory on the disc:

```
# cd /mnt/cdrom
```

- 5 Verify there are no VxFS file systems mounted on the nodes being upgraded:

```
# mount | grep vxfs
```

If any VxFS file systems are mounted, offline the group on each node of the cluster:

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

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

```
# hagr -offline group -sys system03
```

```
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

Repeat this step for each SFCFS service group.

- 6 If performing a phased upgrade, start the installation procedure from the node selected in the cluster. In the following example the phased upgrade is performed on one node in a four-node cluster.

Ensure that the HAD daemon of VCS is running on all nodes selected for the upgrade. Enter the following command, and then press **Return**.

```
# ./installsfcfs system01 system02
```

If performing a full upgrade, start the installation from any node in the cluster. Enter the following command, and then press **y** to upgrade the cluster configuration.

```
# ./installsfcfs
```

- 7 Press **Return** to begin installing infrastructure packages.
- 8 Press **Return** to begin license verification.
- 9 Press **Return** to begin the uninstall. The uninstall utility checks the system's uninstallation requirements and identifies packages, patches and dependencies.

The output can be extensive and exceed 100 lines, depending on your configuration.

- 10 If VCS is running you are prompted to upgrade and reset the VCS password. To ensure security, passwords are not written to the install log.

```
installsfcfs must now make configuration updates and stop  
the cluster before upgrading VCS packages.
```

```
Are you ready to begin the Cluster Server upgrade at this  
time? [y,n,q]
```

- 11 Enter **y**.
- 12 At the prompt, enter your new password.
- 13 Reenter your new password.
- 14 Output shows information that Cluster Server must be stopped on a running system. Enter **y** to continue.
- 15 Press **Return** to begin removing the previous packages and installing the new.
- 16 Press **Return** again for summary information about logs and reboots.  
Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.
- 17 If performing a phased upgrade, proceed to shut down VCS.  
If performing a full upgrade, proceed to updating the configuration.  
See [“Updating the configuration and confirm startup \(phased or full\)”](#) on page 166.

### Updating the configuration and confirm startup (phased or full)

Perform the following steps on each upgraded node.

#### To update the configuration and confirm startup

- 1 Remove the `/etc/VRTSvcs/conf/config/.stale` file, if it exists.

```
# rm -f /etc/VRTSvcs/conf/config/.stale
```

- 2 Reboot the upgraded nodes.

```
# reboot
```

- 3 After the nodes reboot, verify that LLT is running:

```
# lltconfig
LLT is running
```

- 4 Verify GAB is configured:

```
# gabconfig -l | grep 'Driver.state' | \
grep Configured
Driver state : Configured
```

- 5 Verify VxVM daemon is started and enabled:

```
# /opt/VRTS/bin/vxdctl mode
mode: enabled
```

- 6 Confirm all upgraded nodes are in a running state.

```
# /opt/VRTSvc/bin/hasys -state | grep RUNNING
1
```

- 7 Log in as superuser.

- 8 Insert the appropriate media disc into your system's CD-ROM drive.

- 9 If volume management software is running on your system, the software disc automatically mounts as `/mnt/cdrom`.

If volume management software is not available to mount the CD-ROM, you must mount it manually, enter:

```
# mkdir -p /mnt/cdrom
# mount -V cdrfs -o ro /dev/cd0 /mnt/cdrom
```

- 10 Change to the top-level directory on the disc:

```
# cd /mnt/cdrom
```

- 11 Run `installsfefs` from any node in the cluster:

```
# ./installsfefs -configure system01 system02
```

- 12 After the configuration is complete, the CVM and SFCFS groups may come up frozen. To find out the frozen CVM and SFCFS groups, enter the following command:

```
# /opt/VRTS/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

- Make the configuration read/write:

```
# /opt/VRTS/bin/haconf -makerw
```

- Unfreeze the group.

```
# /opt/VRTS/bin/hagrp -unfreeze group -persistent
```

- Save the configuration.

```
# /opt/VRTS/bin/haconf -dump -makero
```

- 13** If VVR is configured, and the CVM and SFCFS groups are offline, bring the groups online in the following order:

Bring online the CVM groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group -sys system01
```

```
# /opt/VRTS/bin/hagrp -online group -sys system02
```

where *group* is the VCS service group that has the CVMVolDg resource.

Bring online the RVGShared groups and the virtual IP on the master node using the following commands:

```
# hagrp -online RVGShared -sys masterhost
```

```
# hares -online ip_name masterhost
```

Bring online the SFCFS groups on all systems.

```
# /opt/VRTS/bin/hagrp -online group -sys system01
```

```
# /opt/VRTS/bin/hagrp -online group -sys system02
```

where *group* is the VCS service group that has the CFSSMount resource.

If the SFCFS service groups do not come online then your file system could be dirty.

See [“Making the file systems clean”](#) on page 169.

- 14** If performing a phased upgrade, upgrade the remaining nodes.

If performing a full upgrade, proceed to upgrade the remaining nodes.

See [“Upgrading remaining nodes \(full only\)”](#) on page 169.

## Upgrading remaining nodes (full only)

This section describes how to upgrade the remaining nodes.

### To upgrade remaining nodes

- 1 If you are configuring SFCFS for a fenced environment.

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

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

- 2 To verify the cluster protocol version, enter the following command that displays the same on a running node:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

If the cluster protocol version is less than 70, then it needs to be upgraded to 70 for SFCFS 5.1.

The cluster protocol version can only be updated on the master node.

Upgrade the entire cluster using the following command on the master node.

```
# /opt/VRTS/bin/vxdctl upgrade
```

- 3 Type the following command on one upgraded node to enable membership:

```
# gabconfig -xc
```

## Making the file systems clean

If you upgrade to SFCFS 5.1 and the file systems are dirty, you have to deport the shared disk group and import it as non-shared. After the import, run `fsck`. `fsck` should succeed. Then deport the disk group and import it back as shared.

### To make the file systems clean

- 1 Log in as superuser onto the CVM master node.
- 2 If performing a full upgrade, offline the group on all the nodes in the cluster:

```
# hagrps -offline group -sys system01
# hagrps -offline group -sys system02
# hagrps -offline group -sys system03
# hagrps -offline group -sys system04
```

If performing a phased upgrade, offline the group:

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

where *group* is the VCS service group that has the CVMVolDg and CFSSMount resource.

- 3 Deport the disk group:

```
# vxvg deport diskgroup
```

where *diskgroup* is the SFCFS disk group.

- 4 Import the disk group:

```
# vxvg -C import diskgroup
```

- 5 Start all the volumes in the disk group:

```
# vxvol -g diskgroup startall
```

- 6 Check and repair each file system:

```
# fsck -V vxfs /dev/vx/dsk/diskgroup/volume
```

Repeat this step for each file system.

- 7 Deport the disk group:

```
# vxvg deport diskgroup
```

- 8 Import the disk group:

```
# vxvg -s import diskgroup
```

- 9 Start all the volumes in the disk group:

```
# vxvol -g diskgroup startall
```

- 10 If VVR is configured, upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

- 11 If performing a full upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01  
# hares -clear resource -sys system02  
# hares -clear resource -sys system03  
# hares -clear resource -sys system04
```

If performing a phased upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01  
# hares -clear resource -sys system02
```

- 12 If performing a full upgrade, online the group on all the nodes in the cluster:

```
# hagr -online group -sys system01  
# hagr -online group -sys system02  
# hagr -online group -sys system03  
# hagr -online group -sys system04
```

If performing a phased upgrade, online the group:

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

- 13 If VVR is configured, bring online the RVGLogowner group on the master:

```
# hagr -online RVGLogownerGrp -sys masterhost
```

Restart the applications that were stopped.

## Upgrading Veritas Volume Replicator

If a previous version of Veritas Volume Replicator (VVR) is configured, the product installer upgrades VVR automatically when you upgrade the Storage Foundation products.

When upgrading from 4.0 or later, you have the option to upgrade without disrupting replication.

See [“Upgrading VVR without disrupting replication”](#) on page 172.

## 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 27.

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 4.0 to VVR 5.1 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.

---

**Note:** Reduce application downtime while upgrading by planning your upgrade.

---

See “[Planning an upgrade from the previous VVR version](#)” on page 27.

## Upgrading from SFHA 5.1 to SFCFS 5.1

The product installer does not support direct upgrades from a previous version of Storage Foundation High Availability to Storage Foundation Cluster File System 5.1. First upgrade Storage Foundation High Availability to 5.1. Then use this procedure to upgrade the product from SFHA 5.1 to SFCFS 5.1. The same steps are used to upgrade to SFCFS 5.1 or SFCFS HA 5.1

### To upgrade from SFHA 5.1 to SFCFS 5.1 or SFCFS HA 5.1

- 1 Back up the `main.cf` file before beginning the upgrade.
- 2 Confirm that the storage disks are visible on all the nodes in the 5.1 SFHA cluster.
- 3 Bring all the failover service groups offline, using the following command:

```
# hagrps -offline group_name -any
```

The above command brings the service group offline on the node where the service group is currently online.

- 4 Unmount all the VxFS file systems which are not under VCS control. If the local file systems are under VCS control, then VCS unmounts the file systems when the failover service group is brought offline in step 3.

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

```
# umount -t 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 diskgroup_name stopall  
# vxdg deport diskgroup_name
```

- 6 Upgrade the existing 5.1 SFHA to 5.1 SFCFSHA using the SFCFSHA installation script as below:

```
# ./installsfdfs
```

- 7 After installation is completed, the install script asks you to install licenses. Enter the correct license key to register the key.

- 8 The installer prompts to reconfigure the VCS. Provide the same cluster name, cluster ID, and LLT link interfaces details that were used during configuration of the SFHA cluster.

See “[Configuring Storage Foundation Cluster File System](#)” on page 102.

- 9 Find out which node is the CVM master, using the following command:

```
# vxdctl -c mode
```

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

```
# vxdg -s import diskgroup_name
```

- 11 Start all the volumes whose disk groups have been imported as shared in step 10. Use the following command:

```
# vxdg -g diskgroup_name startall
```

- 12 Run the following command for each of the file systems you want to mount as CFS:

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

- 13 Run the following command to mount CFS file systems on all the nodes:

```
# cfsmount mount_point
```

- 14 Import all other local disk groups which have not been imported in shared mode in step 10.

```
# vxdg import diskgroup_name
```

Start all the volumes of these disk groups using:

```
# vxvol -g diskgroup_name startall
```

Mount these volumes.

- 15 For any of the file systems which need to be monitored by VCS through failover service groups, create these failover service groups by adding the Mount, Diskgroup & Volume resources for VxFS file systems under VCS control.

# Upgrading SFCFS using an alternate disk

This chapter includes the following topics:

- [About upgrading SFCFS using an alternate disk](#)
- [Supported upgrade scenarios](#)
- [Supported upgrade paths](#)
- [Preparing to upgrade SFCFS on an alternate disk](#)
- [Upgrading SFCFS on an alternate disk](#)
- [Verifying the upgrade](#)

## About upgrading SFCFS using an alternate disk

Use the alternate disk installation process to upgrade the operating system and SFCFS on a production server while the server runs. Perform the upgrade on an alternate or inactive boot environment. After the upgrade, you restart the system to use the updated environment. The instructions in this section assume a working knowledge of the alternate disk installation process. See the operating system documentation for detailed information on alternate disk installations.

---

**Note:** Only Technology Level (TL) and Service Pack (SP) releases of the operating system can be upgraded using this procedure.

---

Upgrading SFCFS on an alternate disk has the following advantages:

- The server remains active during the time the new boot environment is created and upgraded on the alternate boot device.

- The actual downtime for the upgrade is reduced to the period of time required for a single reboot.
- The original boot environment is still available for use if the updated environment fails to become active.

Upgrading SFCFS on an alternate disk involves the following steps:

Preparing to upgrade SFCFS on an alternate disk See [“Preparing to upgrade SFCFS on an alternate disk”](#) on page 176.

Upgrading SFCFS on an alternate disk See [“Upgrading SFCFS on an alternate disk”](#) on page 178.

Verifying the upgrade See [“Verifying the upgrade”](#) on page 187.

## Supported upgrade scenarios

The following upgrade scenarios are supported on an alternate disk:

- Upgrading only SFCFS
- Upgrading only the operating system (Technology Level (TL) and Service Pack (SP) releases)
- Upgrading the operating system (Technology Level (TL) and Service Pack (SP) releases) and SFCFS

## Supported upgrade paths

You can upgrade the operating system and SFCFS using an alternate disk from the following versions:

AIX version Technology Level and Service Pack releases of AIX 5.3 and later

SFCFS version 4.0 MP4 and later

## Preparing to upgrade SFCFS on an alternate disk

Complete the preparatory steps in the following procedure before you upgrade SFCFS on an alternate disk.

### To prepare to upgrade SFCFS on an alternate disk

- 1 Make sure that the SFCFS installation media is available.
- 2 On the nodes that you want to upgrade SFCFS, create an alternate boot disk identical to the primary boot disk with the same size and partitions as the primary boot disk.
- 3 Check the status of the physical disks on your system.

---

**Note:** The alternate disk must have a physical identifier and must not contain any mounted volume groups.

---

```
# lspv
```

Output similar to the following displays:

```
hdisk0          0009710fa9c79877    rootvg    active
hdisk1          0009710f0b90db93    None
```

If the alternate disk does not have a physical identifier, set the physical identifier for the disk:

```
# chdev -l hdisk1 -a pv=yes
```

- 4 Make sure that the following filesets are installed on the primary disk:  
`bos.alt_disk_install.boot_images`, `bos.alt_disk.install.rte`
- 5 Mount the SFCFS installation media.

Determine the filesets you want to install on the alternate disk by running the following command:

```
# ./installsfcfs -install_option
```

where `install_option` is one of the following:

- minpkgs: For installing the minimum set of packages
- recpkgs: For installing the recommended packages
- allpkgs: For installing all packages

Copy the required filesets to a directory on the primary boot disk, for example `/usr`

If you are upgrading the operating system along with SFCFS, copy the necessary operating system filesets and the SFCFS filesets to a directory on the primary disk, for example `/usr`.

## Upgrading SFCFS on an alternate disk

Use one of the following options to upgrade SFCFS on an alternate disk:

- SMIT interface    See [“To upgrade Storage Foundation Cluster File System on an alternate disk using the SMIT interface”](#) on page 179.
- Manual            See [“To upgrade Storage Foundation Cluster File System on an alternate disk manually”](#) on page 182.

The procedure provides instructions to clone the primary boot environment to the alternate disk, upgrade SFCFS on the alternate disk, and reboot the system to start from the alternate disk.

---

**Note:** The alternate disk upgrade is a manual procedure that does not require the Veritas installation program. As a result, the installation and uninstallation scripts are not available in the `/opt/VRTS/install` directory. You can access these scripts from the installation media.

---

---

**Note:** The instructions in the procedure are valid for all upgrade scenarios.

---

In the procedure examples, the primary or current boot environment resides on `hdisk0` and the alternate or inactive boot environment resides on `hdisk1`.

## To upgrade Storage Foundation Cluster File System on an alternate disk using the SMIT interface

- 1 Clone the primary boot disk `rootvg` to an alternate disk.

Start the SMIT menu :

```
# smit alt_clone
```

Provide information for the following fields in the SMIT menu.

Target disk to install      Enter the name of the alternate boot disk.

```
hdisk1
```

Fileset(s) to install      Enter the following:

```
all
```

Directory or Device with images      Enter the full path of the directory that contains the filesets to be upgraded.

```
/usr/
```

ACCEPT new license agreements?      Enter **yes** to accept the license agreements.

Set bootlist to boot from this disk on next reboot?      Enter **yes** to enable the system to boot from the alternate disk on rebooting.

- 2 Press **Enter** to start the upgrade on the alternate disk.

The upgrade process takes some time.

- 3 Use the following command to wake up the volume group on the alternate boot disk (`hdisk1`) that you cloned.

```
# /usr/sbin/alt_rootvg_op -W -d hdisk1
```

- 4 Verify that the alternate disk is created:

```
# lspv
```

Output similar to the following displays:

```
hdisk0            0009710fa9c79877      rootvg
```

```
hdisk1            0009710f0b90db93      altinst_rootvg
```

- 5 Run the following command to export the root path installation environment variable.

```
# export INSTALL_ROOT_PATH=/alt_inst
```

- 6 Run the following command on the alternate root path of any one node in the cluster to configure a Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure \  
-use_llthost
```

The `-use_llthost` option indicates that the `/etc/llthost` file is used to determine the names of the nodes in the cluster. Alternatively, you can specify the node names instead of the file name.

- 7 Confirm that the you have created the Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -display \  
-use_llthost
```

The output should resemble:

```
Finding existing UUID information ...
```

```
nodeA .... exist.
```

```
nodeB .... exist.
```

```
Done
```

```
Valid uuid exist on nodeA nodeB
```

```
{ef228450-1dd1-11b2-a7bb-5938100f2199} : nodeA nodeB
```

- 8 Change directory to `/alt_inst/etc/VRTSvcs/conf/config`.

```
# cd /alt_inst/etc/VRTSvcs/conf/config
```

- 9 Open the `main.cf` file and delete the `include vvrtypes.cf` line, which is deprecated. The VVR agents are now included in the updated `5.1 types.cf` file. The `main.cf` file is in the `/alt_inst/etc/VRTSvcs/conf/config` directory.

- 10** Back up a copy of the old `types.cf` file and copy the new one for SFCFS to use.

```
# mv types.cf types.cf.ORIG
```

```
# cp ../types.cf .
```

- 11** If you have a secure cluster, perform the instructions in the following section:  
See [“Upgrading a cluster that is in secure mode”](#) on page 185.

- 12** Move to root and run the `alt_rootvg_op -S` command to put the alternate root to sleep.

```
# cd /
```

```
# alt_rootvg_op -S
```

- 13** Verify that the normal boot list includes the name of the alternate boot disk.

```
# bootlist -m normal -o  
hdisk1
```

- 14** Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

```
# shutdown -r
```

**15** Verify the upgrade.

See “[Verifying the upgrade](#)” on page 187.

**16** Remove the obsolete filesets from the alternate disk to complete the upgrade.

You can determine the obsolete packages by comparing the list of installed filesets with the filesets that are displayed on running the following command:

```
# ./installsfcfs -install_option
```

where `install_option` is one of the following:

-minpkgs: Displays the minimum set of packages

-recpkgs: Displays the recommended packages

-allpkgs: Displays all packages

If you find obsolete packages, remove them by running the following command:

```
# installp -u pkg_name
```

---

**Note:** Make sure that you do not remove the filesets of other Veritas products, such as Veritas Cluster Server Management Console, that may be installed on your system. If you need assistance, contact Symantec Technical Support.

---

**To upgrade Storage Foundation Cluster File System on an alternate disk manually**

**1** Clone the primary boot disk `rootvg` to an alternate disk.

```
# /usr/sbin/alt_disk_copy -I "acNgXY" -P "all" -l "/usr" -w \
"all" -d "hdisk1"
```

Where:

- `-d` indicates the name of the target disk on which you clone the primary disk.
- `-l` indicates the full path of the directory that contains the filesets to be upgraded
- `-w` indicates the list of SFCFS filesets that you want to upgrade on the alternate boot disk. The option `all` indicates that all the filesets contained

in the directory you specified (using option `-l`) must be installed to the alternate boot disk.

- 2 Use the following command to wake up the volume group on the alternate boot disk (`hdisk1`) that you cloned.

```
# alt_rootvg_op -W -d hdisk1
```

- 3 Verify that the alternate disk is created:

```
# lspv
hdisk0          0009710fa9c79877    rootvg
hdisk1          0009710f0b90db93    altinst_rootvg
```

- 4 Run the following command to export the root path installation environment variable.

```
# export INSTALL_ROOT_PATH=/alt_inst
```

- 5 Run the following command on the alternate root path of any one node in the cluster to configure a Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure \
-use_llthost
```

The `-use_llthost` option indicates that the `/etc/llthost` file is used to determine the names of the nodes in the cluster. Alternatively, you can specify the node names instead of the file name.

- 6 Confirm that the you have created the Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -display \  
-use_llthost
```

The output should resemble:

```
Finding existing UUID information ...
```

```
nodeA .... exist.
```

```
nodeB .... exist.
```

```
Done
```

```
Valid uuid exist on nodeA nodeB
```

```
{ef228450-1dd1-11b2-a7bb-5938100f2199} : nodeA nodeB
```

- 7 Change directory to `cd /alt_inst/etc/VRTSvcs/conf/config`.

```
# cd /alt_inst/etc/VRTSvcs/conf/config
```

- 8 Open the `main.cf` file and delete the `include vvrtypes.cf` line, which is deprecated. The VVR agents are now included in the updated 5.1 `types.cf` file. The `main.cf` file is in the `/alt_inst/etc/VRTSvcs/conf/config` directory.
- 9 Back up a copy of the old `types.cf` file and copy the new one for SFCFS to use.

```
# mv types.cf types.cf.ORIG
```

```
# cp ../types.cf .
```

- 10 If you have a secure cluster, perform the instructions in the following section:  
See [“Upgrading a cluster that is in secure mode”](#) on page 185.

- 11 Move to root and run the `alt_rootvg_op -S` command to put the alternate root to sleep.

```
# cd /
```

```
# alt_rootvg_op -S
```

- 12 Verify that the normal boot list includes the name of the alternate boot disk. By default, the alternate disk installation process changes the boot list to enable the system to boot from the alternate disk.

```
# bootlist -m normal -o  
hdisk1
```

- 13 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

```
# shutdown -r
```

- 14 Verify the upgrade.

See [“Verifying the upgrade”](#) on page 187.

- 15 Remove the obsolete filesets from the alternate disk to complete the upgrade.

You can determine the obsolete packages by comparing the list of installed filesets with the filesets that are displayed on running the following command:

```
# ./installsfcfs -install_option
```

where `install_option` is one of the following:

-minpkgs: Displays the minimum set of packages

-recpkgs: Displays the recommended packages

-allpkgs: Displays all packages

If you find obsolete packages, remove them by running the following command:

```
# installp -u pkg_name
```

---

**Note:** Make sure that you do not remove the filesets of other Veritas products, such as Veritas Cluster Server Management Console, that may be installed on your system. If you need assistance, contact Symantec Technical Support.

---

## Upgrading a cluster that is in secure mode

If you had a secure cluster in your 5.0 MP3 cluster and want to keep it for 5.1, perform the following procedure.

## To enable security for the upgraded secure cluster

### 1 Change directory to VRTSat.

```
# cd /alt_inst/var/VRTSat
```

### 2 Edit the /alt\_inst/var/VRTSat/ABAuthSource file. Delete all HA\_SERVICES-related entries in it. Remove text similar to the following:

```
[HA_SERVICES@symantecexample]
"PD_state"=dword:00000001
"PD_expiryinterval"=dword:00000000
[HA_SERVICES@symantecexample\admin]
"PD_password"=hex:8d,ab,d2,a3,fe, . . . c4,17,5d,6f,35,3c,12,40
"IsDomainAdmin"=dword:00000001
[HA_SERVICES@symantecexample\HA_VCS_symantecexample]
"PD_principaltype"=dword:00000002
"PD_password"=hex:7f,31,af,c0,b2, . . . 6c,48,33,fe,13,2d,4e,56
"IsBrokerAdmin"=dword:00000000
"IsDomainAdmin"=dword:00000000
"CanAcceptProxyFlag"=dword:00000000
"CanProxyFlag"=dword:00000000
[HA_SERVICES@symantecexample\CMDSERVER_VCS_symantecexample]
"PD_principaltype"=dword:00000002
"PD_password"=hex:da,79,b1,9d,fe, . . . 24,54,e1,90,fb,fb,fb,82
"IsDomainAdmin"=dword:00000000
"IsBrokerAdmin"=dword:00000000
"CanProxyFlag"=dword:00000000
"CanAcceptProxyFlag"=dword:00000000
[HA_SERVICES@symantecexample\webserver_VCS_symantecexample.com]
"PD_principaltype"=dword:00000002
"PD_password"=hex:38,29,ba,6d,57, . . . d1,c1,1d,ca,34,0c,82,9f
"IsDomainAdmin"=dword:00000000
"IsBrokerAdmin"=dword:00000000
"CanProxyFlag"=dword:00000001
"CanAcceptProxyFlag"=dword:00000000
"PD_expiryinterval"=dword:00000000
```

- 3 Touch `/alt_inst/var/VRTSat/LocalAuthSource`.

```
# touch /alt_inst/var/VRTSat/LocalAuthSource
```

- 4 Ensure that DNS has the IPv6 address for localhost. Add the following lines into `/etc/hosts`.

```
127.0.0.1 localhost
::1 localhost
```

## Verifying the upgrade

To ensure that alternate disk installation has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

### To verify the upgrade

- ◆ Verify that the alternate boot environment is active:

```
# lspv
Output similar to the following displays:
hdisk0          0009710fa9c79877    old_rootvg
hdisk1          0009710f0b90db93    rootvg
```

## Verifying that the cluster is in secure mode

Perform the following procedure to verify that the cluster's security (VxAT) works after the reboot.

### To verify that VxAT works

- 1 Make sure that the vxatd process is running. Grep the vxatd process.

```
# ps -ef | grep vxatd
```

For a running process, output resembles:

```
root 139410      1    0 13:01:19      -   0:02 /opt/VRTSsat/bin/vxatd
root 176206 229686    0 17:22:39 pts/0    0:00 grep vxatd
```

If the process is not running, enter the following command to start it.

```
# /opt/VRTSvcs/bin/vxatd
```

- 2 Make sure that the CmdServer process is running. Grep the CmdServer process.

```
# ps -ef | grep CmdServer
```

For a running process, output resembles:

```
root 176142 229686    0 17:24:04 pts/0    0:00 grep CmdServer
root 262272      1    0 13:05:42      -   0:00 /opt/VRTSvcs
/bin/CmdServer
```

If the process is not running, enter the following command to start it.

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

- 3 For secure communication with Veritas Cluster Management console, you may have to restart HAD.

- Stop HAD, at the prompt type:

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

```
# ps -ef | grep had
```

```
root 426094 319550    0 15:06:44 pts/0    0:00 grep had
```

- Start HAD, at the prompt type:

```
# hastart
```

# Verifying the Storage Foundation Cluster File System installation

This chapter includes the following topics:

- [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 Volume Manager processes](#)
- [Checking Veritas File System installation](#)
- [Verifying agent configuration for Storage Foundation Cluster File System](#)
- [Synchronizing time on Cluster File Systems](#)
- [Configuring VCS for Storage Foundation Cluster File System](#)
- [About the LLT and GAB configuration files](#)
- [Verifying the LLT, GAB, and VCS configuration files](#)
- [Verifying LLT, GAB, and cluster operation](#)

## Verifying that the products were installed

Verify that the Veritas Storage Foundation products are installed.

Use the `ls1pp` command to check which packages have been installed.

```
# ls1pp -L | grep VRTS
```

The packages should be in the COMMITTED state.

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 common product installer or product installation scripts. The summary includes the list of the packages, and the status (success or failure) of each package. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

## About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as Sun Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

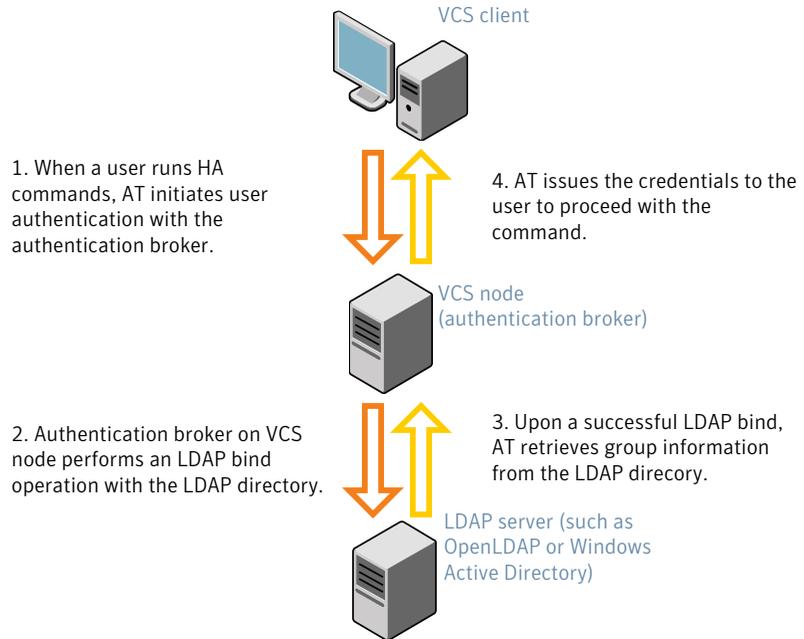
See [“Enabling LDAP authentication for clusters that run in secure mode”](#) on page 192.

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 12-1](#) depicts the SFCFS cluster communication with the LDAP servers when clusters run in secure mode.

**Figure 12-1** Client communication with LDAP servers



See the *Symantec Product Authentication Service Administrator's Guide*.

The LDAP schema and syntax for LDAP commands (such as `ldapadd`, `ldapmodify`, and `ldapsearch`) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)

- UserObjectClass (the default is posixAccount)
- UserObject Attribute (the default is uid)
- User Group Attribute (the default is gidNumber)
- Group Object Class (the default is posixGroup)
- GroupObject Attribute (the default is cn)
- Group GID Attribute (the default is gidNumber)
- Group Membership Attribute (the default is memberUid)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)
- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

## Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 5.0.32.0 or later.

```
# /opt/VRTSat/bin/vssat showversion  
vssat version: 5.0.32.0
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

**To enable OpenLDAP authentication for clusters that run in secure mode**

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
# /opt/VRTSsat/bin/vssat addldapdomain \  
--domainname "MYENTERPRISE.symantecdomain.com"\  
--server_url "ldap://my_openldap_host.symantecexample.com"\  
--user_base_dn "ou=people,dc=symantecdomain,dc=myenterprise,dc=com"\  
--user_attribute "cn" --user_object_class "account"\  
--user_gid_attribute "gidNumber"\  
--group_base_dn "ou=group,dc=symantecdomain,dc=myenterprise,dc=com"\  
--group_attribute "cn" --group_object_class "posixGroup"\  
--group_gid_attribute "member"\  
--admin_user "cn=manager,dc=symantecdomain,dc=myenterprise,dc=com"\  
--admin_user_password "password" --auth_type "FLAT"
```

- 2 Verify that you can successfully authenticate an LDAP user on the SFCFS nodes.

You must have a valid LDAP user ID and password to run the command. In the following example, authentication is verified for the MYENTERPRISE domain for the LDAP user, `vcsadmin1`.

```
galaxy# /opt/VRTSsat/bin/vssat authenticate  
--domain ldap:MYENTERPRISE.symantecdomain.com  
--prplname vcsadmin1 --broker galaxy:2821
```

Enter password for vcsadmin1: #####

```
authenticate  
-----  
-----
```

```
Authenticated User vcsadmin1  
-----
```

**3** Add the LDAP user to the main.cf file.

```
# haconf makerw
# hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
# haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
# hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

**4** Verify that the main.cf file has the following lines:

```
# cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster clus1 (
  SecureClus = 1
  Administrators = {
    "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/
    DC=com@myenterprise.symantecdomain.com" }
  AdministratorGroups = {
    "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/
    DC=com@myenterprise.symantecdomain.com " }
  )
...
...
```

**5** Set the VCS\_DOMAIN and VCS\_DOMAINTYPE environment variables as follows:

- VCS\_DOMAIN=myenterprise.symantecdomain.com
- VCS\_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=myenterprise.symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

**6** Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute  Value
galaxy       Attribute  RUNNING
nebula       Attribute  RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFCFS node using the VCS Cluster Manager (Java Console).

**7** To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

### To enable Windows Active Directory authentication for clusters that run in secure mode

- 1 Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSat/bin/atldapconf -d  
-s domain_controller_name_or_ipaddress  
-u domain_user -g domain_group
```

For example:

```
# /opt/VRTSat/bin/atldapconf -d -s 192.168.20.32 \  
-u Administrator -g "Domain Admins"
```

Search User provided is invalid or Authentication is required to proceed further.

Please provide authentication information for LDAP server.

Username/Common Name: **symantecdomain\administrator**

Password:

Attribute file created.

- 2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSat/bin/atldapconf -c -d windows_domain_name
```

For example:

```
# /opt/VRTSat/bin/atldapconf -c -d symantecdomain.com
```

Attribute list file not provided, using default AttributeList.txt.

CLI file name not provided, using default CLI.txt.

CLI for addldapdomain generated.

- 3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSat/bin/atldapconf -x
```

- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSat/bin/vssat listldapdomains
```

```
Domain Name :          symantecdomain.com
Server URL :          ldap://192.168.20.32:389
SSL Enabled :         No
User Base DN :       CN=people,DC=symantecdomain,DC=com
User Object Class :  account
User Attribute :     cn
User GID Attribute : gidNumber
Group Base DN :     CN=group,DC=symantecdomain,DC=com
Group Object Class : group
Group Attribute :   cn
Group GID Attribute : cn
Auth Type :         FLAT
Admin User :
Admin User Password :
Search Scope :      SUB
```

- 5 Set the VCS\_DOMAIN and VCS\_DOMAINTYPE environment variables as follows:

- VCS\_DOMAIN=symantecdomain.com
- VCS\_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
# export VCS_DOMAIN=symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute  Value
galaxy       Attribute  RUNNING
nebula       Attribute  RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFCFS 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 the product installation script.

For example, to stop VCS processes, enter the following command:

```
# ./installvcs -stop
```

### To start the processes

- ◆ Use the `-start` option to the product installation script.

For example: To start VCS processes, enter the following command:

```
# ./installvcs -start
```

## Checking Volume Manager processes

Use the following procedure to verify that Volume Manager processes are running.

**To confirm that key Volume Manager processes are running**

- ◆ Type the following command:

```
# ps -ef | grep vx
```

Entries for the `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxcached`, and `vxconfigbackupd` processes should appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

## Checking Veritas File System installation

After the Storage Foundation software has been successfully installed, you can confirm successful Veritas File System installation.

**To confirm the File System installation**

- ◆ Use the `lsvfs` command as follows:

```
# lsvfs vxfs
```

Entries for these processes appear in output similar to the following:

```
vxfs 32 /sbin/helpers/vxfs /sbin/helpers/vxfs
```

## Verifying agent configuration for Storage Foundation Cluster File System

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

SFCFS 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

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 SFCFS file resembles:

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleASMTypes.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 }
)

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

// 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.

## About the LLT and GAB configuration files

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

LLT and GAB also require the initialization configuration files:

- `/etc/default/llt`
- `/etc/default/gab`

The information that these LLT and GAB configuration files contain is as follows:

- The `/etc/default/llt` file

This file stores the start and stop environment variables for LLT:

- `LLT_START`—Defines the startup behavior for the LLT module after a system reboot. Valid values include:
  - 1—Indicates that LLT is enabled to start up.
  - 0—Indicates that LLT is disabled to start up.
- `LLT_STOP`—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include:
  - 1—Indicates that LLT is enabled to shut down.
  - 0—Indicates that LLT is disabled to shut down.

The installer sets the value of these variables to 1 at the end of Storage Foundation Cluster File System configuration.

- The `/etc/llthosts` file

The file `llthosts` is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.

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

```
0      galaxy
1      nebula
```

- The `/etc/llttab` file

The file `llttab` contains the information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the LLT network links that correspond to the specific system.

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

```
set-node galaxy
set-cluster 2
link en1 /dev/dlpi/en:1 - ether - -
link en2 /dev/dlpi/en:2 - ether - -

set-node galaxy
set-cluster 2
link en1 /dev/en:1 - ether - -
link en2 /dev/en:2 - ether - -
```

The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the `link` command. These lines identify the two network cards that the LLT protocol uses.

If you configured a low priority link under LLT, the file also includes a "link-lopri" line.

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

- The `/etc/default/gab` file

This file stores the start and stop environment variables for GAB:

- `GAB_START`—Defines the startup behavior for the GAB module after a system reboot. Valid values include:

- 1—Indicates that GAB is enabled to start up.

- 0—Indicates that GAB is disabled to start up.

- `GAB_STOP`—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include:

- 1—Indicates that GAB is enabled to shut down.

- 0—Indicates that GAB is disabled to shut down.

The installer sets the value of these variables to 1 at the end of Storage Foundation Cluster File System configuration.

- The `/etc/gabtab` file

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

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

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

The `-c` option configures the driver for use. The `-nN` specifies that the cluster is not formed until at least `N` nodes are ready to form the cluster. Symantec recommends that you set `N` to be the total number of nodes in the cluster.

---

**Note:** Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. Using `-c -x` can lead to a split-brain condition.

---

## 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 202.

## 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 205.
- 4 Verify GAB operation.  
See [“Verifying GAB”](#) on page 208.
- 5 Verify the cluster operation.  
See [“Verifying the cluster”](#) on page 209.

## Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The

command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

**To verify LLT**

- 1 Log in as superuser on the node galaxy.
- 2 Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node           State      Links
*0 galaxy      OPEN      2
 1 nebula      OPEN      2
```

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

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
Node           State      Links
* 0 galaxy      OPEN      2
 1 nebula      OPEN      2
 2 saturn       OPEN      1
```

- 3 Log in as superuser on the node nebula.
- 4 Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
Node           State      Links
 0 galaxy      OPEN      2
*1 nebula      OPEN      2
```

- 5 To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv | more
```

The output on galaxy resembles:

Node	State	Link	Status	Address
*0 galaxy	OPEN	en1	UP	08:00:20:93:0E:34
		en2	UP	08:00:20:93:0E:34
1 nebula	OPEN	en1	UP	08:00:20:8F:D1:F2
		en2	DOWN	
2	CONNWAIT	en1	DOWN	
		en2	DOWN	
3	CONNWAIT	en1	DOWN	
		en2	DOWN	
.				
.				
.				
31	CONNWAIT	en1	DOWN	
		/dev/en:2	DOWN	

Note that the output lists 32 nodes. The command reports the status on the two nodes in the cluster, galaxy and nebula, along with the details for the non-existent nodes.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- A MAC address for each link

However, the output in the example shows different details for the node nebula. The private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

- 6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node galaxy in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
  Port  Usage      Cookie
  ----  -
  0     gab        0x0
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects: 0 1
  7     gab        0x7
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects: 0 1
  31    gab        0x1F
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects: 0 1
```

## Verifying GAB

Verify the GAB operation using the `gabconfig -a` command. This command returns the GAB port membership information.

The ports indicate the following:

- |   |  |
|---|--|
| a | GAB  |
| b | I/O fencing  |
| d | Oracle Disk Manager (ODM)                              |
| f | Cluster File System (CFS)                              |
| h | Veritas Cluster Server (VCS: High Availability Daemon) |
| v | Cluster Volume Manager (CVM)                           |
| w | vxconfigd (module for CVM)                             |

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 v gen ada416 membership 01
Port w gen ada418 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 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
```

- 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
```

#System	Attribute	Value
galaxy	AvailableCapacity	100
galaxy	CPUUsage	0
galaxy	CPUUsageMonitoring	Enabled 0 ActionThreshold 0 ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
galaxy	Capacity	100
galaxy	ConfigBlockCount	117
galaxy	ConfigChecksum	10844
galaxy	ConfigDiskState	CURRENT
galaxy	ConfigFile	/etc/VRTSvcs/conf/config
galaxy	ConfigInfoCnt	0
galaxy	ConfigModDate	Wed 14 Oct 2009 17:22:48
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	
galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	5.1.00.0
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	HostUtilization	CPU 0 Swap 0
galaxy	LLTNodeId	0
galaxy	LicenseType	DEMO
galaxy	Limits	
galaxy	LinkHbStatus	link2 UP link3 UP
galaxy	LoadTimeCounter	0

#System	Attribute	Value
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	ShutdownTimeout	600
galaxy	SourceFile	./main.cf
galaxy	SysInfo	Aix:galaxy,5,2,00023BDA4C00
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 a node to Storage Foundation Cluster File System clusters

This chapter includes the following topics:

- [About adding a node to an SFCFS cluster](#)
- [Preparing to add a node to an existing SFCFS cluster](#)
- [Adding nodes to an existing SFCFS cluster](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node](#)
- [Sample configuration file for adding a node to the cluster](#)

## About adding a node to an SFCFS cluster

You can add multiple nodes to an existing cluster using the SFCFS installer or manually.

---

**Note:** SFCFS enables you to add nodes to an existing cluster while the cluster is running. The applications running on the cluster continue to be available as new nodes join the cluster.

---

The sample procedures in this chapter describe how to add a node (saturn) to an existing cluster (clus1) with two nodes (galaxy and nebula).

## Preparing to add a node to an existing SFCFS cluster

Complete the following preparatory steps on the new node before you add the node to an existing SFCFS cluster.

### To prepare the new node

- 1 Make sure that the existing cluster is an SFCFS cluster and that SFCFS is running on the cluster.
- 2 Enable the DLPI driver.

See “[Enabling the DLPI driver](#)” on page 214.

- 3 Add the node physically to the existing cluster.

Make sure that you meet the following requirements:

- The 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.  
For more information, see the *Veritas Cluster Server Installation Guide*.
- 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.

- 4 Install SFCFS on the new system.

---

**Note:** Use the `-install` option to install SFCFS. Do not configure SFCFS after the installation.

```
Would you like to configure SFCFS on saturn [y, n, q] (n)
```

You can configure the new node later using the configuration from the existing cluster nodes.

---

## Enabling the DLPI driver

The file `/etc/pse.conf` must be configured to enable the Streams DLPI driver. The following procedure describes checking the file and modifying it, if necessary, and loading the `dlpi` driver.

### Checking and modifying /etc/pse.conf

- 1 Check to see if the Ethernet driver is configured in the /etc/pse.conf file:

```
# egrep 'ethernet driver' /etc/pse.conf
```

- 2 In the output, examine the line containing the "ethernet driver" expression:

```
# d+ dlpi en /dev/dlpi/en # streams dlpi ethernet driver
```

- 3 If the comment symbol("#") precedes the line, the DLPI driver is not configured. Using vi or another text editor, edit the file:

```
# vi /etc/pse.conf
```

- 4 Find the section in the file labeled "#PSE drivers" and look for the line shown in step 2.

- 5 Uncomment the line by removing the initial "#" symbol.

- 6 Save and close the file.

- 7 To load the dlpi driver, enter:

```
# strload -f /etc/dlpi.conf
```

## Adding nodes to an existing SFCFS cluster

---

**Note:** Before you add the node, make sure that SFCFS is not configured on the node.

---

Use one of the following ways to add nodes to an existing SFCFS cluster:

Using the SFCFS installer	See <a href="#">"Adding nodes to an existing SFCFS cluster using the SFCFS installer"</a> on page 216.
---------------------------	--

Manual	See <a href="#">"Adding the node to the SFCFS cluster manually"</a> on page 218.
--------	--

## Adding nodes to an existing SFCFS cluster using the SFCFS installer

---

**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.

---

The SFCFS installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:
  - `/etc/llttab`
  - `/etc/VRTSvcs/conf/sysname`
- Updates the following configuration files and copies them on the new node:
  - `/etc/llthosts`
  - `/etc/gabtab`
  - `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files from the existing cluster to the new node:
  - `/etc/vxfenmode`
  - `/etc/vxfendg`
  - `/etc/vx/.uuids/clusuid`
- Configures security on the new node if the existing cluster is a secure cluster.
- 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.

---

- Starts SFCFS processes and configures CVM and CFS on the new node.

At the end of the process, the new node joins the SFCFS cluster.

**To add the node to an existing SFCFS cluster using the SFCFS installer**

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the SFCFS installer with the `-addnode` option.

```
# cd /opt/VRTS/install
# ./installsfdfs -addnode
```

The installer displays the copyright and the location where the temporary installation logs are stored.

- 3 Enter the name of a node in the existing SFCFS cluster. The installer uses the node information to identify the existing cluster.

```
Enter a node name in the SFCFS cluster to which
you want to add a node: galaxy
```

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

```
Enter the system names separated by spaces
to add to the cluster: saturn
```

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

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

---

**Note:** The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

---

```
Enter the NIC for the first private heartbeat
link on saturn: [b,q,?] en1
```

- 7 Enter **y** to configure a second private heartbeat link.

---

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

---

```
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
```

- 8 Enter the name of the network interface that you want to configure as the second private heartbeat link.

```
Enter the NIC for the second private heartbeat link
on saturn: [b,q,?] en2
```

- 9 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 10 Review and confirm the information.

- 11 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: en3
```

- 12 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.

- 13 The installer then starts all the required Veritas processes and joins the new node to cluster. Confirm using `lltstat -n` and `gabconfig -a`.

## Adding the node to the SFCFS cluster manually

Perform this step after you install SFCFS only if you plan to add the node manually to the cluster.

### To add the node manually to the SFCFS cluster

- 1 Start the Volume Manager.  
See “[Starting Volume Manager on the new node](#)” on page 219.
- 2 Configure LLT and GAB  
See “[Configuring LLT and GAB on the new node](#)” on page 220.
- 3 If the existing cluster is a secure cluster, set up the new node to run in secure mode.
- 4 If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node.  
See “[Configuring server-based fencing on the new node](#)” on page 224.
- 5 Start fencing.  
See “[Starting fencing on the new node](#)” on page 227.
- 6 Configure CVM and CFS.  
See “[Configuring CVM and CFS on the new node](#)” on page 228.
- 7 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 229.

### 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 `installsfcfs` program.

#### To start Volume Manager on the new node

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:  

```
# vxinstall
```
- 2 Enter **n** when prompted to set up a system wide disk group for the system.  
The installation completes.
- 3 Verify that the daemons are up and running. Enter the command:

```
# vxdisk list
```

Make sure the output displays the shared disks without errors.

## Configuring LLT and GAB 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 en1 /dev/dlpi/en:1 - ether - -
link en2 /dev/dlpi/en: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. 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:

```
# /etc/init.d/llt.rc start  
  
# /etc/init.d/gab.rc start  
  
# /etc/methods/gmskextadm load  
  
# /etc/rc.d/rc2.d/S99odm start
```

- 9** On the new node, verify that the GAB port memberships are a, b, d, h, v, w and f:

```
# gabconfig -a  
GAB Port Memberships  
=====
```

```
Port a gen df204 membership 012  
Port b gen df20a membership 012  
Port d gen df207 membership 012  
Port h 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 13-1](#) uses the following information for the following command examples.

**Table 13-1** The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
saturn	saturn.nodes.example.com	The new node that you are adding to the cluster.
RB1	RB1.brokers.example.com	The root broker for the cluster

**Table 13-1** The command examples definitions (*continued*)

Name	Fully-qualified host name (FQHN)	Function
RB2	RB2.brokers.example.com	Another root broker, not the cluster's RB

**To verify the existing security setup on the node**

- 1 If node saturn is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Else, proceed to configuring the authentication broker on node saturn.

See “[Configuring the authentication broker on node saturn](#)” on page 222.

- 2 Find out the root broker to which the node saturn belongs using the following command.

```
# vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
-k "BrokerName"
```

- 3 If the node saturn already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.

- 4 If the node saturn belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node saturn.

- Kill /opt/VRTSat/bin/vxatd process.
- Remove the credential that RB2 has given to AB on node saturn.

```
# vssat deletecred --domain type:domainname \
--prplname prplname
```

For example:

```
# vssat deletecred --domain vx:root@RB2.brokers.example.com \
--prplname saturn.nodes.example.com
```

**Configuring the authentication broker on node saturn**

Configure a new authentication broker (AB) on node saturn. This AB belongs to root broker RB1.

### To configure the authentication broker on node saturn

- 1 Create a principal for node saturn on root broker RB1. Execute the following command on root broker RB1.

```
# vssat addprpl --pdrtype root --domain domainname \  
  --prplname prplname --password password \  
  --prpltype service
```

For example:

```
# vssat addprpl --pdrtype root \  
  --domain root@RB1.brokers.example.com \  
  --prplname saturn.nodes.example.com \  
  --password flurbdicate --prpltype service
```

- 2 Ensure that there is no clock skew between the times on node saturn and RB1.
- 3 Copy the /opt/VRTSat/bin/root\_hash file from RB1 to node saturn.
- 4 Configure AB on node saturn to talk to RB1.

```
# vxatd -o -a -n prplname -p password -x vx -y domainname -q \  
  rootbroker -z 2821 -h roothash_file_path
```

For example:

```
# vxatd -o -a -n saturn.nodes.example.com -p flurbdicate \  
  -x vx -y root@RB1.brokers.example.com -q RB1 \  
  -z 2821 -h roothash_file_path
```

- 5 Verify that AB is configured properly.

```
# vssat showbrokermode
```

The command should return 1, indicating the mode to be AB.

### Adding a node in a VxSS group

Perform the following procedure when adding a node in a VxSS group.

### To add a node in the VxSS group using the CLI

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

```
# hasys -add saturn
```

- 4 Add the node saturn to the existing VxSS group.

```
# hagrps -modify VxSS SystemList -add saturn 2
# hagrps -modify VxSS AutoStartList -add saturn
```

- 5 Save the configuration by running the following command from any node in the cluster:

```
# 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
```

Perform the following procedure for a secure configuration.

**To configure server-based fencing with security on the new node**

- 1 As the root user, create the VCS user and the domain on the new node:
  - Create a dummy configuration file /etc/VRTSvcs/conf/config/main.cf that resembles the following example:

```
# cat main.cf  
  
include "types.cf"  
cluster clus1 {  
    SecureClus = 1  
}  
  
system saturn {  
}
```

- Verify the dummy configuration file:

```
# cd /etc/VRTSvcs/conf/config  
# /opt/VRTSvcs/bin/hacf -verify .
```

- Start VCS in one node mode on the new node:

```
# /opt/VRTSvcs/bin/hastart -onenode
```

- 2 Verify that the VCS user and the domain are created on the new node:

```
# /opt/VRTScps/bin/cpsat showcred | grep _HA_VCS_  
# /opt/VRTScps/bin/cpsat listpd -t local | grep HA_SERVICES
```

- 3 Stop VCS if the VCS user and domain are created successfully on the new node:

```
# /opt/VRTSvcs/bin/hastop -local
```

- 4 If the root broker for the CP server and the new node are different, run the following command to establish trust between the authentication broker of the CP Server and the new node:

```
# /usr/bin/echo y | /opt/VRTScps/bin/cpsat setuptrust \  
-b mycps1.symantecexample.com -s high
```

- 5 Log in to each CP server as the root user.

- 6 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 7 Verify that the new node is added to the CP server configuration:

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

The new node must be listed in the output.

- 8 Add the VCS user `_HA_VCS_saturn@HA_SERVICES@saturn.symantec.com` to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e _HA_VCS_saturn@HA_SERVICES@saturn.symantec.com \  
-f cps_operator -g vx
```

User `_HA_VCS_saturn@HA_SERVICES@saturn.symantec.com` successfully added

## 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 Copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/default/vxfen
```

```
/etc/vxfendg
```

```
/etc/vxfenmode
```

- 2 Start fencing on the new node:

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

- 3 On the new node, verify that the GAB port memberships are a, b, d, and o:

```
# gabconfig -a
```

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====  
Port a gen df204 membership 012
```

```
Port b gen df20d membership 012
```

```
Port d gen df20a membership 012
```

## 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 SFCFS 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 SFCFS cluster:

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

## 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, 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
```

## 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 en0 -sys saturn
```

```
# hares -modify gconic Device en0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

## Updating the Storage Foundation for Databases (SFDB) repository after adding a node

If you are using Database Checkpoints, Database Flashsnap, or Database Dynamic Storage Tiering 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 Run the following to change permission, owner, group of various SFDB directories:

```
# run sfua_db_config
```

- 2 Run the `dbed_update` command on any one node in the cluster. For example:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -G $ORACLE_SERVICE_GROUP
```

This completes the addition of the node to the SFDB repository.

For information on using SFDB tools features:

See the Storage Foundation guide: *Storage Foundation: Storage and Availability Management for Oracle Databases*.

## Sample configuration file for adding a node to the cluster

You may use this sample file as reference information to understand the configuration changes that take place when you add a node to a cluster.

The existing sample configuration before adding the node `saturn` is as follows:

- The existing cluster `rac_cluster101` comprises two nodes `galaxy` and `nebula` and hosts a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle. The agent starts, stops, and monitors the database.
- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

The following sample configuration file shows the changes (in **bold**) effected in the configuration after adding a node "saturn" to the cluster.

```
include "types.cf"  
include "CFSTypes.cf"  
include "CVMTypes.cf"
```

```
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)
system nebula (
)
system saturn (
)
```

---

**Note:** In the following group `oradb1_grp`, the saturn node has been added.

---

```
group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)
```

---

**Note:** In the following Oracle resource, the saturn node information has been added.

---

```
Oracle oral (
    Critical = 0
    Sid @galaxy = vrts1
    Sid @nebula = vrts2
    Sid @saturn = vrts3
    Owner = oracle
    Home = "/app/oracle/orahome"
    StartUpOpt = "SRVCTLSTART"
    ShutDownOpt = "SRVCTLSTOP"
)

CFSMount oradata_mnt (
    Critical = 0
```

```

MountPoint = "/oradata"
BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"
)

CVMVolDg oradata_voldg (
    Critical = 0
    CVMDiskGroup = oradatadg
    CVMVolume = { oradatavol }
    CVMActivation = sw
)

requires group cvm online local firm
ora1 requires oradata_mnt
oradata_mnt requires oradata_voldg

```

---

**Note:** In the following CVM and CVMCluster resources, the saturn node information has been added.

---

```

group cvm (
    SystemList = { galaxy = 0, nebula = 1, saturn =2}
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { galaxy, nebula, saturn }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
    OnlineRetryLimit = 20
)

CFSMount ocrvote_mnt (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
    MountOpt= "mincache=direct"
)

```

```
CVMVolDg ocrvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1, saturn =2 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0
    CVMVxconfigdArgs = { syslog }
)
```

---

**Note:** In the following PrivNIC resource, the saturn node information has been added.

---

```
PrivNIC ora_priv (
    Critical = 0
    Device@galaxy = { en1 = 0, en2 = 1 }
    Device@nebula = { en1 = 0, en2 = 1 }
    Device@saturn = { en1 = 0, en2 = 1 }
    Address@galaxy = "192.168.12.1"
    Address@nebula = "192.168.12.2"
    Address@saturn = "192.168.12.3"
    NetMask = "255.255.255.0"
)
```

```
cssd requires ocrvote_mnt
cssd requires ora_priv
```

```
ocrvote_mnt requires ocrvote_voldg  
ocrvote_mnt requires vxfsckd  
ocrvote_voldg requires cvm_clus  
vxfsckd requires cvm_clus  
cvm_clus requires cvm_vxconfigd
```



# Removing a node from Storage Foundation Cluster File System clusters

This chapter includes the following topics:

- [Removing nodes from an SFCFS 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](#)

## Removing nodes from an SFCFS cluster

Perform the following steps to remove a node from an SFCFS cluster.

### To remove a node from an SFCFS 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.

- 3 Remove your application database instance from the node.

For instructions, see your application documentation.

- 4 Remove your application database software from the node.

For instructions, see your application documentation.

- 5 Stop VCS on the node:

```
# hastop -local
```

- 6 Unmount the VxFS/CFS file systems that are not under configured under VCS.

```
# umount mount_point
```

- 7 Uninstall SFCFS from the node using the SFCFS installer.

```
# cd /opt/VRTS/install
```

```
# ./uninstallsfcfs saturn
```

The installer stops all SFCFS processes and uninstalls the SFCFS packages.

- 8 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 238.

## 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](#)

### Editing 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
```

## Editing 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
```

## Modifying the VCS configuration to remove the node

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 3 and step 4 for each of them.

- 7 Remove the node from the `NodeList` attribute of all CFS mount resources:

```
# hares -modify CFSMount NodeList -delete saturn
```

- 8 After deleting the removed node from all service groups in the configuration, delete the node from the cluster system list:

```
# hasys -delete saturn
```

- 9 Save the new configuration to disk:

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

- 10 Verify that the node is removed from the VCS configuration.

```
# grep -i saturn 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 SFCFS cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

### To remove the node configuration from the CP server

- 1 Log into the CP server as the root user.
- 2 View the list of VCS users on the CP server, using the following command:

```
# cpsadm -s cp_server -a list_users
```

Where *cp\_server* is the virtual IP/ virtual hostname of the CP server.

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e _HA_VCS_saturn@HA_SERVICES@saturn.nodes.example.com \  
-f cps_operator -g vx
```

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
cpsadm -s cp_server -a rm_node -h saturn -c clus1 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
cpsadm -s cp_server -a list_nodes
```

## Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node saturn. Perform the following steps.

### To remove the security credentials

- 1 Kill /opt/VRTSat/bin/vxatd process.
- 2 Remove the root credentials on node saturn.

```
# vssat deletecred --domain type:domainname --prplname prplname
```

## Updating the Storage Foundation for Databases (SFDB) repository after removing a node

If you are using Database Checkpoints, Database Flashsnap, or Database Dynamic Storage Tiering in your configuration, update the SFDB repository to remove the reference for the node after removing the node from the cluster.

---

**Note:** If you have not created an SFDB repository, you do not need to perform the following steps.

---

### To update the SFDB repository after removing a node

- 1 As Oracle user, list the nodes in the cluster:

```
$ /opt/VRTSdbed/bin/dbed_rept_node -S $ORACLE_SID -o list
```

- 2 Run the following command after physically removing the node from the cluster.

For example:

```
$ /opt/VRTSdbed/bin/dbed_rept_node -S $ORACLE_SID -n NODE -o remove
```

This completes the removal of the node from the SFDB repository.

## 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 saturn is as follows:

- The existing cluster `rac_cluster101` comprises three nodes `galaxy`, `nebula`, and `saturn` and hosts a single database.
- The Oracle database is stored on CFS.
- The database is managed by the VCS agent for Oracle. The agent starts, stops, and monitors the database.
- Only one private IP address is configured for Oracle Clusterware. The private IP address is managed by the PrivNIC agent for high availability.
- The Oracle Cluster Registry (OCR) and voting disk are stored on CFS.

---

**Note:** The following sample file shows in **bold** the configuration information that is removed when the node "saturn" is removed from the cluster.

---

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"

cluster rac_cluster101 (
    UserNames = { admin = bopHo }
    Administrators = { admin }
    UseFence = SCSI3
)

system galaxy (
)
system nebula (
)
system saturn (
)
```

---

**Note:** In the following group `oradb1_grp`, the saturn node must be removed.

---

```
group oradb1_grp (
    SystemList = { galaxy = 0, nebula = 1, saturn = 2 }
    AutoFailOver = 0
    Parallel = 1
```

```
AutoStartList = { galaxy, nebula, saturn }  
)
```

---

**Note:** In the following Oracle resource, the saturn node information must be removed.

---

```
Oracle oral (  
    Critical = 0  
    Sid @galaxy = vrts1  
    Sid @nebula = vrts2  
    Sid @saturn = vrts3  
    Owner = oracle  
    Home = "/app/oracle/orahome"  
    StartUpOpt = "SRVCTLSTART"  
    ShutDownOpt = "SRVCTLSTOP"  
)  
  
CFSMount oradata_mnt (  
    Critical = 0  
    MountPoint = "/oradata"  
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol"  
)  
  
CVMVolDg oradata_voldg (  
    Critical = 0  
    CVMDiskGroup = oradatadg  
    CVMVolume = { oradatavol }  
    CVMActivation = sw  
)  
  
requires group cvm online local firm  
oral requires oradata_mnt  
oradata_mnt requires oradata_voldg
```

---

**Note:** In the following CVM and CVMCluster resources, the saturn node information must be removed.

---

```
group cvm (  
    SystemList = { galaxy = 0, nebula = 1, saturn =2}  
    AutoFailOver = 0  
    Parallel = 1
```

```

AutoStartList = { galaxy, nebula, saturn }
)

Application cssd (
    Critical = 0
    StartProgram = "/opt/VRTSvcs/rac/bin/cssd-online"
    StopProgram = "/opt/VRTSvcs/rac/bin/cssd-offline"
    CleanProgram = "/opt/VRTSvcs/rac/bin/cssd-clean"
    MonitorProgram = "/opt/VRTSvcs/rac/bin/cssd-monitor"
    OnlineRetryLimit = 20
)

CFSMount ocrvote_mnt (
    Critical = 0
    MountPoint = "/ocrvote"
    BlockDevice = "/dev/vx/dsk/ocrvotedg/ocrvotevol"
    MountOpt= "mincache=direct"
)

CVMVolDg ocrvote_voldg (
    Critical = 0
    CVMDiskGroup = ocrvotedg
    CVMVolume = { ocrvotevol }
    CVMActivation = sw
)

CFSfscsd vxfsckd (
)

CVMcluster cvm_clus (
    CVMClustName = rac_cluster101
    CVMNodeId = { galaxy = 0, nebula = 1, saturn =2 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
    Critical = 0

```

```
CVMVxconfigdArgs = { syslog }  
)
```

---

**Note:** In the following PrivNIC resource, the saturn node information must be removed.

---

```
PrivNIC ora_priv (  
    Critical = 0  
    Device@galaxy = { en1 = 0, en2 = 1 }  
    Device@nebula = { en1 = 0, en2 = 1 }  
    Device@saturn = { en1 = 0, en2 = 1 }  
    Address@galaxy = "192.168.12.1"  
    Address@nebula = "192.168.12.2"  
    Address@saturn = "192.168.12.3"  
    NetMask = "255.255.255.0"  
)
```

```
cssd requires ocrvote_mnt  
cssd requires ora_priv  
ocrvote_mnt requires ocrvote_voldg  
ocrvote_mnt requires vxfsckd  
ocrvote_voldg requires cvm_clus  
vxfsckd requires cvm_clus  
cvm_clus requires cvm_vxconfigd
```

# Setting up a replicated global cluster

This chapter includes the following topics:

- [Replication in the SFCFS environment](#)
- [Requirements for SFCFS global clusters](#)
- [About setting up a global cluster in an SFCFS environment](#)
- [Configuring an SFCFS cluster at the primary site](#)
- [Configuring an SFCFS cluster at the secondary site](#)
- [Configuring replication on clusters at both sites](#)
- [Modifying the ClusterService group for global clusters](#)
- [Defining the remote cluster and heartbeat objects](#)
- [Configuring the VCS service groups for global clusters](#)

## Replication in the SFCFS environment

You can set up a primary SFCFS cluster for replication to a secondary SFCFS 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 SFCFS.
- Supported hardware-based replication technologies. Using hardware-based replication you can replicate data from a primary array to a secondary array.

## Requirements for SFCFS global clusters

Review the requirements information to make sure your configuration is supported for SFCFS.

For product licensing information:

See [“About Veritas product licensing”](#) on page 21.

## Supported software and hardware for SFCFS

For supported hardware and software:

- See [“Hardware overview and requirements for Storage Foundation Cluster File System”](#) on page 38.
- See the current compatibility list in the Veritas Technical Support website to confirm the compatibility of your hardware:  
<http://entsupport.symantec.com/docs/283161>

## Supported replication technologies for SFCFS

SFCFS supports the following replication technologies through the use of Veritas replication agents:

**Table 15-1** Supported replication options for SFCFS global clusters

Replication technology	Supported modes	Supported software
Veritas Volume Replicator (VVR) Supporting agents <ul style="list-style-type: none"> <li>■ RVGShared</li> <li>■ RVGSharedPri</li> <li>■ RVGLogOwner</li> </ul>	<ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication</li> </ul>	Host-based replication

**Table 15-1** Supported replication options for SFCFS global clusters (*continued*)

Replication technology	Supported modes	Supported software
EMC SRDF Supporting agent: SRDF	<ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication</li> </ul>	All versions of Solutions Enabler
Hitachi True Copy Supporting agent: HTC	<ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication</li> </ul>	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"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication</li> </ul>	SSH access to the SVC
EMC Mirror View Supporting agent: MirrorView	<ul style="list-style-type: none"> <li>■ Asynchronous replication</li> <li>■ Synchronous replication: only individual LUNs may be replicated</li> </ul>	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 SFCFS.

---

You can use the Veritas replication agents listed in the table above for global clusters that run SFCFS. 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 SFCFS nodes. The agents make sure that the system

which has the resource online also has safe and exclusive access to the configured devices.

This information is current at the time this document is released. For more current information on the replicated agents, see:

- *Veritas Cluster Server Agent for EMC SRDF Installation and Configuration Guide*
- *Veritas Cluster Server Agent for Hitachi TrueCopy Installation and Configuration Guide*
- *Veritas Cluster Server Agent for IBM Metro Mirror Installation and Configuration Guide*
- *Veritas Cluster Server Agent for IBM SVC Installation and Configuration Guide*
- *Veritas Cluster Server Agent for EMC MirrowView Installation and Configuration Guide*
- Technical Support TechNote for the latest updates or software issues for replication agents:  
<http://entsupport.symantec.com/docs/282004htm>

## About setting up a global cluster in an SFCFS environment

Configuring a global cluster for application requires the coordination of many component setup tasks. The procedures provided in this document are guidelines.

The tasks required to set up a global cluster:

- Configure an SFCFS cluster at the primary site
- Configure an SFCFS cluster at the secondary site
- Configure replication on clusters at both sites
- Configure VCS service groups for replication
- Test the HA/DR configuration
- Upon successful testing, bring the environment into production

Some SFCFS HA/DR configuration tasks may require adjustments depending upon your particular starting point, environment, and configuration. Review the installation requirements and sample cluster configuration files for primary and secondary clusters.

For requirements:

For instructions for configuring AT in a global cluster:

See the *Veritas Cluster Server User's Guide*

## Configuring an SFCFS cluster at the primary site

You can use an existing SFCFS cluster or you can install a new SFCFS cluster for your primary site.

For planning information:

See [“About planning for a SFCFS installation”](#) on page 23.

If you have an existing SFCFS cluster, and you want to set up a global cluster, skip the steps below and proceed to configure your secondary cluster.

If you do not have an existing cluster and you are setting up two new sites for an SFCFS 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 for a global cluster. If you are using VVR for replication, you must have it enabled.
- 4 Install and configure SFCFS. Prepare for your installation according to your configuration needs.

For preparation:

See [“Prerequisites for Storage Foundation Cluster File System”](#) on page 37.

For installation:

See [“Installation quick reference”](#) on page 47.

After verifying a successful SFCFS installation:

Verify the CVM group is online on all nodes in the primary cluster:

```
# hagr -state cvm
```

- 5 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
- 6 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.

- 7 Identify the disks that will be replicated, create the required CVM disk group, volume, and file system.
- 8 Create the database on the file system you created in the previous step.
- 9 Configure the VCS service groups for the database.
- 10 Verify that all VCS service groups are online.

## Configuring an SFCFS cluster at the secondary site

To set up a multi-node or single-node cluster on the secondary site:

- Set up the cluster
- Set up the database

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.
- Make sure that you use the same OS user and group IDs for for installation and configuration on both the primary and secondary clusters.  
Make sure that you use the same OS user and group IDs for for installation and configuration on both the primary and secondary clusters.

## Setting up the cluster on the secondary site

### 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 SFCFS. Prepare for your installation according to your configuration needs.

For preparation:

See [“Prerequisites for Storage Foundation Cluster File System”](#) on page 37.

For installation:

See [“Installation quick reference”](#) on page 47.

After verifying a successful SFCFS installation:

Verify the CVM group is online on all nodes in the primary cluster:

```
# hagrps -state cvm
```

- 5 For a multi-node cluster, configure I/O fencing.
  - Verify the shared storage on the secondary site supports SCSI-3 reservations.
  - Set up coordinator disks
  - Configure I/O fencing
- 6 For a single-node cluster, do not enable I/O fencing. Fencing will run in disabled mode.
- 7 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
- 8 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.

## Setting up the database for the secondary site

### 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.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 on 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 [“Setting up replication using VVR on the primary site”](#) on page 262.
- 2 Start replication between the sites.  
See [“Starting replication of application database volume”](#) on page 269.

## Modifying the ClusterService group for global clusters

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 User’s Guide* for complete details on global clustering.

## Modifying the global clustering 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 group online.

## Defining the remote 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 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` enables 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 through step 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

# hastatus -sum
.....
-- WAN HEARTBEAT STATE
-- Heartbeat      To                State

L  Icmp           clus2                ALIVE

-- REMOTE CLUSTER STATE
-- Cluster        State

M  clus2          RUNNING

-- REMOTE SYSTEM STATE
-- cluster:system State                Frozen

N  clus2:  RUNNING           0
N  clus2:  RUNNING           0
```

## 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 (
)

.
.
```

See the *Veritas Cluster Server User's Guide* for details for configuring the required and optional attributes of the heartbeat object.

# Configuring the VCS service groups for global 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 273.
- 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 282.  
  
See [“Migrating the role of new primary site back to the original primary site”](#) on page 283.
- 3 Upon successful testing, bring the environment into production.

For complete details on VVR in a shared disk environment:

See the *Veritas Volume Replicator Administrator’s Guide*.

# Configuring a global cluster using VVR

This chapter includes the following topics:

- [About configuring global clustering using VVR](#)
- [Setting up replication using VVR on the primary site](#)
- [Setting up replication using VVR on the secondary site](#)
- [Starting replication of application database volume](#)
- [Configuring VCS to replicate the database volume using VVR](#)
- [Using VCS commands on SFCFS global clusters](#)
- [Using VVR commands on SFCFS global clusters](#)

## About configuring global clustering using VVR

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 product licensing”](#) on page 21.

After setting up two clusters running SFCFS, you can configure a global cluster environment with VVR. You must modify both cluster configurations to support replication in the global cluster environment.

Configuring SFCFS for global clusters requires:

- Setting up both clusters as part of a global cluster environment.  
See [“About setting up a global cluster in an SFCFS environment”](#) on page 250.
- Setting up replication for clusters at both sites.  
See [“Setting up replication using VVR on the primary site”](#) on page 262.  
See [“Setting up replication using VVR on the secondary site”](#) on page 265.
- Starting replication of the database.  
See [“Starting replication of application database volume”](#) on page 269.
- Configuring VCS for replication on clusters at both sites.  
See [“Configuring VCS to replicate the database volume using VVR”](#) on page 272.

## Setting up replication using VVR on 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 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 SRL volume on the primary site”](#) on page 262.
- Creating the Replicated Volume Group (RVG) on the primary site.  
See [“Setting up the Replicated Volume Group \(RVG\) on the primary site”](#) on page 263.

### Creating the SRL volume on the primary site

Create the SRL. The SRL is a volume in the RVG. The RVG also holds the data volumes for replication.

- The data volume on the secondary site has the same name and the same size as the data volume on the primary site.
- The SRL on the secondary site has the same name and the same size as the SRL on the primary site.
- The data volume and SRL volume should exist in the same disk group.
- If possible, create SRLs on disks without other volumes.
- Mirror SRLs and data volumes 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 master or the slave:

```
# vxctl -c mode
```

- 3 From the master node, issue the following command:

```
# vxassist -g oradatadg make rac1_srl 1500M nmirror=2 disk4 disk5
```

- 4 Using the following command, 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 `hagrps -state cvm` command.

- 3 Check that the CVM group is online.
- 4 On each site, verify vradadmin is running:

```
# ps -ef |grep vradadmin
    root  536594  598036   0 12:31:25      0  0:00 grep vradmin
```

If vradadmin is not running start it:

```
# vxstart_vvr
VxVM VVR INFO V-5-2-3935 Using following ports:
heartbeat: 4145
vradmind: 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/vradmind
    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          *.*
```

To create the RVG

The command to create the primary RVG takes the form:

```
vradadmin -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

For example, to create the *rac1\_rvg* RVG, enter:

```
# vradadmin -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*.

# Setting up replication using VVR on 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:

- If you have not already done so, create a disk group on the storage on the primary site. For example:
- Create volumes for the database and SRL on the secondary site.  
See [“Creating the data and SRL volumes on the secondary site”](#) on page 265.
- Edit the `/etc/vx/vras/.rdg` file on the secondary site.  
See [“Editing the /etc/vx/vras/.rdg files”](#) on page 266.
- 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 266.
- Create the replication objects on the secondary site.  
See [“Setting up the disk group on secondary site for replication”](#) on page 267.

## 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 reflect the sizes and names of the corresponding volumes in 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 volume 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
```

- 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 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 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 en0:1
- The virtual IP address is 10.10.9.101

- The net mask is 255.255.255.0
  - ```
# ifconfig en0 10.10.9.101 netmask 255.255.255.0 \  
broadcast 10.180.95.255 alias  
# ifconfig en0 up
```
- 2 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 the network interface is en0:1
  - virtual IP address is 10.11.9.102
  - net mask is 255.255.255.0
- 3 Define the virtual IP addresses to correspond to a virtual cluster host name on the primary site and a virtual cluster host name 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 10.10.9.101
- 4 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 on 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: `oradatadg`
- `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`
- `sec_host` is the virtual IP address or resolvable virtual host name of the cluster on the secondary site.  
For example: `clus2`

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 \  
clus1 clus2
```

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
- 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: oradatadg
datavol_cnt: 1
vset_cnt: 0
srl: rac1_srl
RLinks:
name=rlk_10.11.9.102_ 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_10.10.9.101_ rac1_rvg, detached=on,
synchronous=off
```

---

**Note:** Once the replication is started the value off detached flag will change the status from OFF to ON.

---

## Starting replication of application database volume

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: `srlprot_autodcm`
- Packet size: `packet_size=8400`
- Network protocol: `protocol=UDP`

Method of initial synchronization:

- Automatic synchronization
- Full synchronization with 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 Checkpoint

Use the `vradmin` command with the Checkpoint option to start replication using full synchronization with Checkpoint.

### To start replication using full synchronization with 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 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 racl_ckpt syncrvg racl_rvg  
clus2
```

- 2 To start replication after full synchronization, enter the following command:

```
# vradmin -g oradatadg -c racl_ckpt startrep racl_rvg  
clus2
```

## Verifying replication status

Verify that replication is properly functioning.

### To verify replication status

- 1 Use the `vxprint` command on the primary site:

```
vxprint -g diskgroup -l rlink_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_oradatadg  
Rlink: rlk_clus2_oradatadg  
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: racl_ckpt
```

```
flags: write enabled attached consistent connected  
asynchronous
```

## Configuring VCS to replicate the database volume using VVR

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.

### About modifying the VCS configuration for replication

The following resources must be configured or modified for replication:

- Log owner group
- RVG group
- CVMVolDg resource
- RVGSharedPri resource
- application database service group

For detailed examples of service group modification:

For more information on service replication resources:

See the *Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide*.

#### Log owner group

Create a log owner group including the RVGLogowner resources. The RVGLogowner resources are used by:

- RLINKs for the RVG
- RVGLogowner resource. The RVG and its associated disk group are defined as attributes for the RVGLogowner resource.

The RVG log owner service group has an online local firm dependency on the service group containing the RVG.

The VCS uses the following agents to control the following resources:

- RVGLogowner agent to control the RVGLogowner resource

- RVGShared agent to control the RVGShared resource

## RVG group

Create an RVG group that includes the RVGShared resource replication objects. Define the RVGShared resource and CVMVolDg 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.

## CVMVolDg resource

The CVMVolDg resource does not have volumes specified for the CVMVolume attribute; the volumes are contained in the RVG resource. The CVMVolume attribute for the CVMVolDg 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 CVM service group.

For a detailed description of the CVMVolDg agent in this guide:

See “[CVMVolDg agent](#)” on page 342.

## RVGSharedPri resource

Add the RVGSharedPri resource to the existing application database service group. The CVMVolDg resource must be removed from the existing application database service group.

## application database service group

The existing application database service group is a parallel group consisting of the application database resource, CVMVolDg resource, and CFSSMount 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.

## 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 CVMVolDg resource from the existing application database service group to the newly created RVGShared service group.

#### 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 SFCFS 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 = en0  
    Address = "10.10.9.101"  
    NetMask = "255.255.255.0"  
)
```

```

NIC nic (
    Device = en0
    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 bolded attribute in the example that follows.
- Add the ClusterFailOverPolicy cluster attribute. Symantec recommends using the Manual value. See the bolded attribute 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.
- 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 }  
)  
  
CFSSMount oradata_mnt (  
    MountPoint = "/oradata"  
    BlockDevice = "/dev/vx/dsk/oradatadg/racdb_vol"  
)  
  
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.

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

Wait for port h to stop on all nodes, and then restart VCS with the new configuration on all primary nodes:

```
# hstart
```

## 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.

### 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 SFCFS installation to see the CVM configuration.

In our example, the secondary site has clus2 consisting of the nodes and . 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 = { = 0, = 1 }
    AutoStartList = { , }
)

IP logowner_ip (
    Device = en0
    Address = "10.11.9.102"
    NetMask = "255.255.255.0"
)

NIC nic (
    Device = en0
    NetworkHosts = { "10.10.8.1" }
    NetworkType = ether
)

RVGLogowner logowner (
    RVG = racl_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 = { = 0, = 1 }
    Parallel = 1
    AutoStartList = { , }
)

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.
- 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 = { = 0, = 1 }  
    ClusterList = { clus2 = 0, clus1 = 1 }  
    Parallel = 1  
    OnlineRetryInterval = 300  
    ClusterFailOverPolicy = Manual  
    Authority = 1  
    AutoStartList = { , }  
)  
  
CFMount oradata_mnt (  
    MountPoint = "/oradata"  
    BlockDevice = "/dev/vx/dsk/oradatadg/racdb_vol"  
)
```

**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:

```
# hstart
```

- 11 Verify that VCS brings all resources online. On one node, enter the following command:

```
# hagr -display
```

The application, RVG, and CVM groups are online on both nodes of the primary site. The RVGLogOwner group is 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 `hagr -online` command. This information applies to the secondary site, except for the application group which must be offline.

On the primary site, enter the following commands:

```
# hagr -online rlogowner -sys galaxy
```

```
# hagr -online database_grp -sys galaxy
```

```
VCS WARNING V-16-1-50817 Please use hagr -online -force to online a group for the first time
```

```
# hagr -online -force database_grp -sys galaxy
```

On the secondary site, enter the following command:

```
# hagr -online rlogowner -sys
```

- 12 Verify the service groups and their resources that are brought online. On one node, enter the following command:

```
# hagr -display
```

The application service group is offline on the secondary site, but the CVM, RVG log owner, and RVG groups are online.

This completes the setup for an SFCFS global cluster using VVR for replication. Symantec recommends testing a global cluster before putting it into production.

## Using VCScommands on SFCFS global clusters

For information on the VCS commands for global clusters:

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

## Using VVR commands on SFCFS global clusters

If you have two SFCFS clusters configured to use VVR for replication, the following administrative functions are available:

- Migration of the role of the primary site to the remote site
- Takeover of the primary site role by the secondary site

### About migration and takeover of the primary site role

Migration 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 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.

### 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 `hagrps -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.

```
# hagrps -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:

```
# vxrlink -g data_disk_group status rlk_clus1_priv_rac1_rvg
```

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:

```
# hagrps -online database_grp -any
```

## Migrating the role of new primary site back to the original primary site

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 `hagrps -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:

```
# hagrps -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 primary cluster.

For example:

```
# vxrlink -g data_disk_group status rlk_rac_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:

```
# hagrps -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 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.

Takeover options are:

- [Disaster](#)
- [Outage](#)

- [Disconnect](#)
- [Replica](#)

## Disaster

When the cluster on the primary site is inaccessible and appears dead, the administrator declares the failure type as "disaster." For example, fire may destroy a data center, including the primary site and all data in the volumes. After making this declaration, the administrator can bring the service group online on the secondary site, which now has the role as "primary" site.

## Outage

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.

After the declaration for an outage occurs, the RVGSharedPri 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.

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.

See "[Examples for takeover and resynchronization](#)" on page 286.

See "[Replica](#)" on page 286.

## Disconnect

When both clusters are functioning properly and the heartbeat link between the clusters fails, a split-brain condition exists. In this case, the administrator can declare the failure as "disconnect," which means no attempt will occur to take over the role of the primary site at the secondary site. This declaration is merely advisory, generating a message in the VCS log indicating the failure results from a network outage rather than a server outage.

## Replica

In the rare case where the current primary site becomes inaccessible while data is resynchronized from that site to the original primary site using the fast fail back method, the administrator at the original primary site may resort to using a data snapshot (if it exists) taken before the start of the fast fail back operation. In this case, the failure type is designated as "replica".

## Examples for takeover and resynchronization

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 rac_cluster101
```

- 2 After declaring the state of the remote cluster, bring the Oracle 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:

```
# vxrvg -g data_disk_group -F snapshot rac1_rvg
```

See the *Veritas Volume Replicator Administrator's Guide* for details on RVG snapshots.

- 2 Resynchronize the RVG. From the CVM master 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
```

To determine which node is the CVM master node, type:

```
# vxdctl -c mode
```

- 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 `vxrvg` command with the `snapback` option to reattach the snapshot volumes on the original primary site to the original volumes in the specified RVG:

```
# vxrvg -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 CVR

The following topic headings indicate either probable symptoms of a problem, or the procedures required for a solution.

## Updating the rlink

If the rlink is not up to date, use the `hares -action` command with the `resync` action token to synchronize the RVG.

The following command example is issued on any node (`galaxy`, in this case) in the primary cluster, specifying the `RVGSharedPri` resource, `ora_vvr_shpri`:

```
# hares -action ora_vvr_shpri resync -sys galaxy
```

## VCS agents to manage wide-area failover

VCS agents now 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.

VCS provides agents for other array-based or application-based solutions. This section covers the replication agents that is bundled with VVR. See the VCS replication agent documentation for more details.

---

**Note:** See the Veritas Cluster Server Agents for Veritas Volume Replicator Configuration Guide for more information about the RVG and RVGPrimary agents.

---

---

**Note:** The RVGSnapshot agent is not supported for SFCFS.

---

### 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.

# Uninstalling Storage Foundation Cluster File System

This chapter includes the following topics:

- [Preparing to uninstall a Storage Foundation product](#)
- [About removing Veritas Storage Foundation](#)
- [Shutting down cluster operations](#)
- [Moving volumes to physical disks](#)
- [Disabling the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling SFCFS with the Veritas Web-based installer](#)
- [Uninstalling SFCFS filesets using the script-based installer](#)
- [Removing Storage Foundation products using SMIT](#)
- [Uninstalling Storage Foundation Cluster File System](#)
- [Removing the CP server configuration using the removal script](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product](#)

# Preparing to uninstall a Storage Foundation product

Complete the following preparations to uninstall a Storage Foundation product.

---

**Warning:** Failure to follow the preparations that are outlined in this chapter can result in loss of data.

---

To remove Veritas Storage Foundation, complete the following preparations before the uninstallation:

- Back up all VxFS file systems in full and move the files in all VxFS file systems to native file systems backed with LVM logical volumes. Raw application data stored in VxVM logical volumes must be moved to LVM logical volumes.
- Remove all but one copy of file systems and databases.
- Remove all but one plex from volumes that contain multiple plexes (mirrors). To display a list of all volumes, use the command:

```
# vxprint -Ath
```

To remove a plex, use the command:

```
# vxplex -g diskgroup -o rm dis plex
```

- If a remaining plex contains multiple subdisks, consolidate the subdisks into a single subdisk using the commands:

```
# vxassist -g diskgroup mirror volume layout=contig  
# vxplex -g diskgroup -o rm dis plex
```

Sufficient space on another disk is required for this operation to complete.

- Modify `/etc/filesystems` to remove or change entries for VxFS file systems that were moved to native file systems.
- Move all data from volumes created from multiple regions of storage, including striped or spanned volumes, onto a single disk or appropriate LVM logical volume. This can be done using one of the following three methods:
  - Back up the system to tape or other media and recover the system from this.
  - Move volumes incrementally (evacuate) onto logical volumes. Evacuation moves subdisks from the source disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to LVM volumes. See [“Moving volumes to physical disks”](#) on page 291.

## About removing Veritas Storage Foundation

This section covers uninstallation requirements and steps to uninstall the Veritas software.

Only users with superuser privileges can uninstall Veritas Storage Foundation.

---

**Warning:** Failure to follow the instructions in the following sections may result in unexpected behavior.

---

## 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.

---

## Moving volumes to physical disks

You can use the following steps to move data off of VxVM volumes.

**To move data off of VxVM volumes**

- 1 Evacuate as many disks as possible by using one of the following methods:
  - the "Remove a disk" option in `vxdiskadm`
  - the Veritas Enterprise Administrator

- the `vxevac` script from the command line.

- 2 Remove the evacuated disks from Veritas Volume Manager control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name  
# /usr/lib/vxvm/bin/vxdiskunsetup -C disk_access_name  
# vxdisk rm disk_access_name
```

For example:

```
# vxdg -g mydg rmdisk mydg01  
# /usr/lib/vxvm/bin/vxdiskunsetup -C hdisk1  
# vxdisk rm hdisk01
```

- 3 Decide which volume to move first. If the volume to be moved is mounted, unmount it. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume has been synchronized.
- 4 On the free disk space, create an LVM logical volume that is the same size as the VxVM volume. If there is not enough free space for the logical volume, add a new disk to the system for the first volume to be removed. For subsequent volumes, you can use the free space generated by the removal of the first volume.
- 5 Copy the data on the volume onto the newly created LVM logical volume using the following command:

```
# dd if=/dev/vx/dsk/diskgroup/volume of=/dev/vgvol
```

where *diskgroup* is the name of a VxVM disk group, *volume* is the old volume in that disk group, and *vgvol* is a newly created LVM volume.

If the volume contains a VxFS file system, the user data managed by VxFS in the volume must be backed up or copied to a native AIX file system in an LVM logical volume.

- 6 The entries in `/etc/filesystems` for volumes holding VxFS file systems, that were copied to native file systems in step 5, must be modified according to the change in step 5.
- 7 Mount the disk if the corresponding volume was previously mounted.
- 8 Remove the volume from VxVM using the following command:

```
# vxedit -g diskgroup -rf rm volume
```

- 9 Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

```
# vxprint -g diskgroup -F "%sdnum" disk_media_name
```

- 10 If the return code is not 0, there are still some subdisks on this disk that must be subsequently removed. If the return code is 0, remove the disk from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name  
# vxdisk rm disk_access_name
```

- 11 Copy the data in the next volume to be removed to the newly created free space.
- 12 Reboot the system after all volumes have been converted successfully. Verify that no open volumes remain after the system reboot using the following command:

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

- 13 If any volumes remain open, repeat the steps listed above.

## 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:

```
# hagr -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:

```
# hagr -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

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.

- 5 If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

```
# vxedit -r -g diskgroup rm srl_name
```

## Uninstalling SFCFS with the Veritas Web-based installer

This section describes uninstalling Storage Foundation Cluster File System or Storage Foundation Cluster File System High Availability with the Veritas Web-based installer.

### To uninstall SFCFS

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Start the Web-based installer.  
See [“Starting the Veritas Web-based installer”](#) on page 57.
- 4 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 5 Select Storage Foundation Cluster File System or Storage Foundation Cluster File System High Availability from the Product drop-down list, and click **Next**.
- 6 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Validate**.
- 7 After the validation completes successfully, click **Uninstall** to uninstall SFCFS on the selected system.
- 8 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 9 After the installer stops the processes, the installer removes the products from the specified system.  
Click **Next**.
- 10 After the uninstall completes, the installer displays the location of the log and summary files. If required, view the files to confirm the status of the removal.
- 11 Click **Finish**. The webinstaller prompts you for another task.

## Uninstalling SFCFS filesets using the script-based installer

Use the following procedure to remove SFCFS products.

Not all filesets may be installed on your system depending on the choices that you made when you installed the software.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 29.

### To shut down and remove the installed SFCFS filesets

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.
- 2 Unmount all mount points for VxFS file systems.

```
# umount /mount_point
```

**3** If the VxVM package (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

**4** Stop the VEA Service.

```
# /opt/VRTS/bin/vxsvcctl stop
```

**5** Make sure you have performed all of the prerequisite steps.

**6** In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hactop -local
```

To stop VCS processes on all systems:

```
# hactop -all
```

**7** Move to the `/opt/VRTS/install` directory and run the uninstall script.

```
# cd /opt/VRTS/install
```

For Veritas Storage Foundation Cluster File System

```
# ./uninstallsfcfs
```

**8** The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall SFCFS, for example, `host1`:

```
Enter the system names separated by spaces from which to  
uninstall Storage Foundation: host1
```

**9** The uninstall script prompts you to select Storage Foundation Cluster File System or Storage Foundation Cluster File System High Availability.

**10** The uninstall script prompts you to confirm the uninstall. If you respond yes, the processes are stopped and the filesets are uninstalled.

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

**11** Most filesets have kernel components. In order to ensure complete removal, a system reboot is recommended after all filesets have been removed.

## Removing Storage Foundation products using SMIT

Use the following procedure to remove Storage Foundation products using SMIT.

To remove the packages using SMIT:

- 1 Enter this command to invoke SMIT:

```
# smit
```

- 2 In SMIT, select **Software Installation and Maintenance**> **Software Maintenance and Utilities**> **Remove Installed Software**.
- 3 Under the "SOFTWARE name" menu, press F4 or Esc-4 to list all software installed on the system.
- 4 Enter "/" for Find, type "VRTS" to find all Veritas packages, and select the packages that you want to remove.
- 5 Reboot the system after removing all Storage Foundation packages.
- 6 Depending on the choices that were made when Storage Foundation was originally installed, you may find that not all of the listed Storage Foundation packages are installed on the system. You may also choose to remove the `VRTSvlic` licensing package unless this is required by other Veritas software.

## Uninstalling Storage Foundation Cluster File System

If you need to uninstall SFCFS software. Use the `uninstallsfcfs` script.

To uninstall SFCFS HA

- 1 Log in as superuser.
- 2 Stop the cluster:

```
# hastop -all
```

Do not use the `hastop -force` command to stop VCS.

- 3 Change directory to `/opt/VRTS/install`:

```
# cd /opt/VRTS/install
```

- 4 Run the `uninstallsfcfs` command to uninstall SFCFS. The `uninstallsfcfs` script uses `ssh` to communicate with remote nodes as default:

```
# ./uninstallsfcfs
```

If you want to use `rsh` you must specify on the command line:

```
# ./uninstallsfcfs -rsh
```

- 5 Enter the system names to uninstall SFCFS.

```
Enter the system names separated by spaces on which to  
uninstall SFCFS: system01 system02
```

- 6 Enter `y` to uninstall SFCFS.

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

## Removing the CP server configuration using the removal script

This section describes how to remove the CP server configuration from a node or cluster hosting the CP server.

---

**Warning:** Ensure that no SFCFS cluster is using the CP server that will have its CP server configuration removed.

---

A configuration utility that is part of VRTScps package is used to remove the CP server configuration. When using the configuration utility, a configuration removal script is run and the following tasks are performed:

- All CP server configuration files are removed
- The VCS configuration for CP server is removed

After running the utility and script, you can then uninstall VCS from the node or cluster.

---

**Note:** The configuration script has to run only once per CP server (which can be on a single node or SFHA cluster), when removing the CP server configuration.

---

The configuration utility performs the following steps to remove the CP server configuration:

- Offlines the CP server service group (CPSSG), if it is online
- Removes the CPSSG service group from the VCS configuration

The following procedure describes how to remove the CP server configuration.

**To remove the CP server configuration**

- 1 To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2 The Veritas Coordination Point Server Configuration utility appears with an option menu.

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

Select one of the following:

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

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

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

- 3 Select option 3 to unconfigure the Coordination Point Server.
- 4 A warning appears and prompts you to confirm the action to unconfigure the Coordination Point Server.

Enter "y" to proceed.

```
Unconfiguring Coordination Point Server stops the vxcperv process.  
VCS clusters using this server for coordination purpose  
will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)[Default:n] :y
```

- 5 After entering "y" to proceed, messages appear informing you of the progress in removing the CP server configuration.

When the CP server configuration has been unconfigured, a success message appears.

For an example of the messages from a single node VCS cluster:

```
A single node VCS cluster is currently configured.  
Stopping the CP server ...  
  
Removing the CP Server from VCS configuration..  
  
Removing resource dependencies...  
Deleting the resources configured under CPSSG service group...  
Deleting the CPSSG service group...  
  
Successfully unconfigured the Veritas Coordination Point Server.
```

For an example of the messages from a CP server on an SFHA cluster:

```
A multinode CP Server cluster is currently configured.  
Stopping the CP server ...  
  
Removing the CP Server from VCS configuration..  
  
Removing resource dependencies...  
Deleting the resources configured under CPSSG service group...  
Deleting the CPSSG service group...  
  
Successfully unconfigured the Veritas Coordination Point Server.
```

- 6 You are then prompted to delete the CP server database. Enter "y" to delete the database.

For example:

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

- 7 You are then prompted to delete the CP server configuration file and log files. Enter "y" to delete these files.

For example:

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 8 Run the following `hagrp -state` command to ensure that the CPSSG resource 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 will disable the SFDB tools.

### To remove the SFDB repository

- 1 Change directories to the location of the local lookup information for the Oracle SID.

For example:

```
# cd /var/vx/vxdba/$ORACLE_SID
```

- 2 Identify the SFDB repository file and any associated links:

For example:

```
# ls -al
```

```
lrwxrwxrwx 1 oracle oinstall 26 Jul 21 13:58 .sfdb_rept -> \
/ora_data1/TEST/.sfdb_rept
```

```
# cd /ora_data1/TEST
```

Follow the symlink of .sfdb\_rept.

- 3 Remove the repository directory containing the repository file and all backups.

For example:

```
# rm -rf .sfdb_rept
```

- 4 Remove the local lookup directory for the Oracle SID:

```
# cd /var/vx/vxdba
```

```
# rm -rf $ORACLE_SID
```

This completes the removal of the SFDB repository.



# Installation scripts

This appendix includes the following topics:

- [About installation scripts](#)
- [Installation script options](#)

## About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.1 provides several installation scripts.

To install a fresh installation on a system, or to upgrade from Veritas Storage Foundation and High Availability Solutions version prior to 5.1, the recommended installation method is to use the common product installer. To use the common product installer, run the `installer` command.

See [“About the common product installer”](#) on page 49.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from an electronic download site, which does not include the common product installer, use the appropriate product installation script.

The following product installation scripts are available:

|                                                        |                           |
|--------------------------------------------------------|---------------------------|
| Veritas Cluster Server (VCS)                           | <code>installvcs</code>   |
| Veritas Storage Foundation (SF)                        | <code>installsf</code>    |
| Veritas Storage Foundation Cluster File System (SFCFS) | <code>installsfcfs</code> |
| Veritas Storage Foundation for Oracle RAC (SFRAC)      | <code>installsfrac</code> |

Symantec Product Authentication Service (AT)      `installat`

Veritas Volume Manager                              `installvm`

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-1** shows command line options for the product 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 installation scripts”](#) on page 305.

**Table A-1**            Available command line options

| Command Line Option       | Function                                                                                                                                                                                                                                          |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>system1 system2...</i> | 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 packages and patches required for the specified product. The packages and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network. |
| <code>-configure</code>   | Configures the product after installation.                                                                                                                                                                                                        |
| <code>-fencing</code>     | Configures I/O fencing in a running cluster.                                                                                                                                                                                                      |

**Table A-1** Available command line options (*continued*)

| Command Line Option                | Function                                                                                                                                                                                                                                                                                                                                                                                                                         |
|------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -ha                                | <p>Specifies that the Storage Foundation High Availability software is installed or displayed. Otherwise, the list of Storage Foundation packages excludes the Veritas Cluster Server packages.</p> <p>This option only applies to the <code>installsfcfs</code> script when one of the following options is specified:</p> <ul style="list-style-type: none"> <li>■ -allpkgs</li> <li>■ -recpkgs</li> <li>■ -minpkgs</li> </ul> |
| -hostfile <i>full_path_to_file</i> | Specifies the location of a file that contains a list of hostnames on which to install.                                                                                                                                                                                                                                                                                                                                          |
| -installallpkgs                    | Specifies that all packages are installed.                                                                                                                                                                                                                                                                                                                                                                                       |
| -installminpkgs                    | Specifies that the minimum package set is installed.                                                                                                                                                                                                                                                                                                                                                                             |
| -installrecpkgs                    | Specifies that the required package set is installed.                                                                                                                                                                                                                                                                                                                                                                            |
| -keyfile <i>ssh_key_file</i>       | Specifies a key file for secure shell (SSH) installs. This option passes <code>-i ssh_key_file</code> to every SSH invocation.                                                                                                                                                                                                                                                                                                   |
| -license                           | Registers or updates product licenses on the specified systems.                                                                                                                                                                                                                                                                                                                                                                  |
| -logpath <i>log_path</i>           | 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.                                                                                                                                                                                                                                                                     |
| -makeresponsefile                  | Generates a response file without performing an actual installation, configuration, or uninstallation.                                                                                                                                                                                                                                                                                                                           |
| -minpkgs                           | Displays the minimal packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.                                                                                                     |

**Table A-1** Available command line options (*continued*)

| Command Line Option                | Function                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -pkginfo                           | Displays a list of filesets and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS filesets.                                                                                    |
| -pkgpath <i>package_path</i>       | Designates the path of a directory that contains all packages to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.                                                                                                                                                 |
| -pkgset                            | Discovers and lists the 5.1 packages installed on the systems that you specify.                                                                                                                                                                                                                                                            |
| -pkgtable                          | Displays the SFCFS 5.1 packages in the correct installation order.                                                                                                                                                                                                                                                                         |
| -precheck                          | Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.                                                                                                                                                                             |
| -recpkgs                           | Displays the recommended packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.           |
| -redirect                          | Displays progress details without showing the progress bar.                                                                                                                                                                                                                                                                                |
| -responsefile <i>response_file</i> | 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. |

**Table A-1** Available command line options (*continued*)

| Command Line Option      | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -rsh                     | Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.<br><br>See “ <a href="#">Configuring secure shell (ssh) or remote shell before installing products</a> ” on page 29.                                                                                                                                                                                                                      |
| -security                | Enable or disable Symantec Product Authentication Service in a VCS cluster that is running. Install and configure Root Broker for Symantec Product Authentication Service.<br><br>You can specify this option with the <code>installvcs</code> , <code>installsf</code> or <code>installsfefs</code> scripts.<br><br>For more information about Symantec Product Authentication Service in a VCS cluster, see the <i>Veritas Cluster Server Installation Guide</i> . |
| -serial                  | Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.                                                                                                                                                                                                                                    |
| -start                   | Starts the daemons and processes for the specified product.                                                                                                                                                                                                                                                                                                                                                                                                          |
| -stop                    | Stops the daemons and processes for the specified product.                                                                                                                                                                                                                                                                                                                                                                                                           |
| -tmppath <i>tmp_path</i> | Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.                                                                                                                                                                                                                            |
| -upgrade                 | Specifies that an existing version of the product exists and you plan to upgrade it.                                                                                                                                                                                                                                                                                                                                                                                 |



# Response files

This appendix includes the following topics:

- [About response files](#)
- [Installing SFCFS using response files](#)
- [Configuring SFCFS using response files](#)
- [Upgrading SFCFS using response files](#)
- [Uninstalling SFCFS using response files](#)
- [Syntax in the response file](#)
- [Response file variable definitions](#)
- [Sample response file for SFCFS install](#)
- [Sample response file for SFCFS configure](#)

## 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 response file option.

## Installing SFCFS using response files

Typically, you can use the response file that the installer generates after you perform SFCFS installation on one cluster to install SFCFS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

### To install SFCFS using response files

- 1 Make sure the systems where you want to install SFCFS 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 SFCFS.
- 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 installation from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file  
  
# ./installsfcfs -responsefile /tmp/response_file
```

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

## Configuring SFCFS using response files

Typically, you can use the response file that the installer generates after you perform SFCFS configuration on one cluster to configure SFCFS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

### To configure SFCFS using response files

- 1 Make sure the SFCFS filesets are installed on the systems where you want to configure SFCFS.
- 2 Copy the response file to one of the cluster systems where you want to configure SFCFS.

- 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.

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

```
# /opt/VRTS/install/installsfdfs -responsefile /tmp/response_file
```

Where */tmp/response\_file* is the response file's full path name.

## Upgrading SFCFS using response files

Typically, you can use the response file that the installer generates after you perform SFCFS upgrade on one cluster to upgrade SFCFS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

### To perform automated SFCFS upgrade

- 1 Make sure the systems where you want to upgrade SFCFS meet the upgrade requirements.
- 2 Make sure the pre-upgrade tasks are completed.
- 3 Copy the response file to one of the cluster systems where you want to upgrade SFCFS.
- 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
```

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

Where */tmp/response\_file* is the response file's full path name.

## Uninstalling SFCFS using response files

Typically, you can use the response file that the installer generates after you perform SFCFS uninstallation on one cluster to uninstall SFCFS on other clusters.

### To perform automated SFCFS uninstallation

- 1 Make sure that you meet the pre-requisites to uninstall SFCFS.
- 2 Copy the response file to one of the cluster systems where you want to uninstall SFCFS.
- 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/installsfcs -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 variable definitions

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (CSGNIC, CSGVIP, and CSGNETMASK) must be defined if any are defined. The same is true for the SMTP notification (SMTPSERVER, SMTPRECP, and SMTPRSEV), the SNMP trap notification (SNMPPORT, SNMPCONS, and SNMPCSEV), and the Global Cluster Option (CGONIC, GCOVIP, and GCONETMASK).

[Table B-1](#) lists the variables that are used in the response file and their definitions.

**Table B-1** Response file variables

| Variable              | Description                                                                                                                                                                           |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{install}     | Installs SFCFS filesets. Configuration can be performed at a later time using the <code>-configure</code> option.<br><br>List or scalar: scalar<br>Optional or required: optional     |
| CFG{accepteula}       | Specifies whether you agree with the EULA.pdf file on the media.<br><br>List or scalar: scalar<br>Optional or required: required                                                      |
| \$CFG{opt}{vxkeyless} | Installs the product with keyless license.<br><br>List or scalar: scalar<br>Optional or required: optional                                                                            |
| CFG{systems}          | List of systems on which the product is to be installed, uninstalled, or configured.<br><br>List or scalar: list<br>Optional or required: required                                    |
| CFG{systemscfs}       | List of systems for configuration if secure environment prevents the installer to install SFCFS on all systems at once.<br><br>List or scalar: list<br>Optional or required: required |
| CFG{product}          | Defines the product to be installed, uninstalled, or configured.<br><br>List or scalar: scalar<br>Optional or required: required                                                      |
| CFG{opt}{keyfile}     | Defines the location of an ssh keyfile that is used to communicate with all remote systems.<br><br>List or scalar: scalar<br>Optional or required: optional                           |

**Table B-1** Response file variables (*continued*)

| Variable                    | Description                                                                                                                                                                                                                                        |
|-----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{at_rootdomain}          | <p>Defines the name of the system where the root broker is installed.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                                                      |
| CFG{opt}{patchpath}         | <p>Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>         |
| CFG{opt}{pkgpath}           | <p>Defines a location, typically an NFS mount, from which all remote systems can install product filesets. 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} {fileset} | <p>Instructs the installation to not install the optional filesets in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>                                                                                              |
| CFG{donotremove} {fileset}  | <p>Instructs the uninstallation to not remove the optional filesets in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>                                                                                             |

**Table B-1** Response file variables (*continued*)

| Variable                        | Description                                                                                                                                                                                                                                |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{sfcs_fencingenabled}      | When SFCFS is configured, defines if fencing is enabled.<br>Scalar<br>Required<br>0 or 1                                                                                                                                                   |
| CFG{vcs_clustername}            | Defines the name of the cluster.<br>List or scalar: scalar<br>Optional or required: required                                                                                                                                               |
| CFG{vcs_clusterid}              | An integer between 0 and 65535 that uniquely identifies the cluster.<br>List or scalar: scalar<br>Optional or required: required                                                                                                           |
| CFG{opt}{logpath}               | Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.<br>List or scalar: scalar<br>Optional or required: optional                                                                    |
| CFG{opt}{configure}             | Performs the configuration after the filesets are installed using the <code>-install</code> option.<br>List or scalar: scalar<br>Optional or required: optional                                                                            |
| CFG{vcs_lltlink#} {system}      | Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (LLTLINK1 and LLTLINK2). Up to four LLT links can be configured.<br>List or scalar: scalar<br>Optional or required: required |
| CFG{vcs_lltlinklowpri} {system} | Defines a low priority heartbeat link. Typically, LLTLINKLOWPRI is used on a public network link to provide an additional layer of communication.<br>List or scalar: scalar<br>Optional or required: optional                              |

**Table B-1** Response file variables (*continued*)

| Variable            | Description                                                                                                                                                                                                                                                                                                   |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_csgnic}     | <p>Defines the NIC for Cluster Management Console to use on a system. 'ALL' can be entered as a system value if the same NIC is used on all systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                                  |
| CFG{csgvip}         | <p>Defines the virtual IP address that the Cluster Management Console uses.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                                                                                                           |
| CFG{vcs_csgnetmask} | <p>Defines the Netmask of the virtual IP address that the Cluster Management Console uses.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                                                                                            |
| CFG{vcs_smtpserver} | <p>Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                                                          |
| CFG{vcs_smtprecp}   | <p>List of full email addresses (example: user@symantecexample.com) of SMTP recipients.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>                                                                                                                                                 |
| CFG{vcs_smtpsev}    | <p>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.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p> |
| CFG{vcs_snmpport}   | <p>Defines the SNMP trap daemon port (default=162).</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                                                                                                                                   |

**Table B-1** Response file variables (*continued*)

| Variable                 | Description                                                                                                                                                                                                                                                                               |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_snmpcons}        | List of SNMP console system names<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                                                               |
| CFG{vcs_snmpcsev}        | 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.<br>List or scalar: list<br>Optional or required: optional |
| CFG{vcs_gconic} {system} | Defines the NIC for the Virtual IP that the Global Cluster Option uses. 'ALL' can be entered as a system value if the same NIC is used on all systems.<br>List or scalar: scalar<br>Optional or required: optional                                                                        |
| CFG{vcs_gcovip}          | Defines the virtual IP address to that the Global Cluster Option uses.<br>List or scalar: scalar<br>Optional or required: optional                                                                                                                                                        |
| CFG{vcs_gconetmask}      | Defines the Netmask of the virtual IP address that the Global Cluster Option uses.<br>List or scalar: scalar<br>Optional or required: optional                                                                                                                                            |
| CFG{vcs_userenpw}        | List of encoded passwords for users<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                                                             |
| CFG{vcs_username}        | List of names of users<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                                                                          |

**Table B-1** Response file variables (*continued*)

| Variable                   | Description                                                                                                                                                                                                                                              |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| \$CFG{vcs_securitymenuopt} | Specifies the menu option to choose to configure the cluster in secure mode.<br>List or scalar: scalar<br><ul style="list-style-type: none"> <li>■ 1–Automatic</li> <li>■ 2–Semi-automatic</li> <li>■ 3–Manual</li> </ul> Optional or required: optional |
| \$CFS{vcs_clustername}     | Defines the name of the cluster.<br>Optional or required: optional                                                                                                                                                                                       |
| CFG{vcs_userpriv}          | List of privileges for users<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                                   |
| \$CFG{opt}{prodmode}       | List of modes for product<br>List or scalar: list<br>Optional or required: optional                                                                                                                                                                      |
| CFG{opt}{upgrade}          | Upgrades all filesets installed, without configuration.<br>List or scalar: list<br>Optional or required: optional                                                                                                                                        |
| CFG{opt}{uninstall}        | Uninstalls SFCFS filesets.<br>List or scalar: scalar<br>Optional or required: optional                                                                                                                                                                   |

## Sample response file for SFCFS install

The following example shows a response file for installing Storage Foundation Cluster File System.

```
#####
#Auto generated sfcfs responsefile #
#####
```

```

our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{rsh}=1;
$CFG{opt}{trace}=0;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{upi}="SFCFS";
$CFG{prod}="SFCFS51";
$CFG{systems}=[ qw( system01 system02 ) ];
$CFG{opt}{uuid}=normI;
$CFG{opt}{logpath}="/opt/VRTS/install/logs/installsfdfs-xxxxxx/installsfdfs

1;

```

## Sample response file for SFCFS configure

The following example shows a response file for configuring Storage Foundation Cluster File System.

```

#####
#Auto generated sfdfs responsefile #
#####

```

```

our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{rsh}=1;
$CFG{opt}{trace}=0;
$CFG{vcs_allowcomms}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
$CFG{opt}{configure}=1;
$CFG{upi}="SFCFS";
0016: $CFG{prod}="SFCFS51";
0017: $CFG{systems}=[ qw( system01 system02 ) ];
0018: $CFG{sfdfs_cvmtimeout}=200;
0019: $CFG{sfdfs_fencingenabled}=0;
0020: $CFG{vm_newnames_file}{system01}=0;
0021: $CFG{vm_restore_cfg}{system01}=0;

```

**Sample response file for SFCFS configure**

```
0022: $CFG{obc_agentfqhn}{system01}="system01.example.com";
0023: $CFG{vm_newnames_file}{system02}=0;
0024: $CFG{vm_restore_cfg}{system02}=0;
0025: $CFG{obc_agentfqhn}{system02}="system02.example.com";
0026: $CFG{obc_mode}="STANDALONE";
0027: $CFG{opt}{noextrapkgs}=1;
0028: $CFG{vcs_clusterid}=127;
0029: $CFG{vcs_clustername}="uxrt5_sol";
0030: $CFG{vcs_username}=[ qw(admin operator) ];
0031: $CFG{vcs_userenpw}=[ qw(JlmElgLimHmmKumGlj bQOsOUUnVQoOUUnTQsOSnUQuOUUnPQtOS) ];
0032: $CFG{vcs_userpriv}=[ qw(Administrators Operators) ];
0033: $CFG{vcs_lltlink1}{system01}="bge1";
0034: $CFG{vcs_lltlink2}{system01}="bge2";
0035: $CFG{vcs_lltlink1}{system02}="bge1";
0036: $CFG{vcs_lltlink2}{system02}="bge2";
0037: $CFG{vcs_enabled}=1;
0038: $CFG{opt}{uuid}=normC;
0039: $CFG{opt}{logpath}="/opt/VRTS/install/logs/installsfcfs-xxxxxx/installsfcfs-xxxxxx.response";
0040:
0041: 1;
```

# Configuring I/O fencing using a response file

This appendix includes the following topics:

- [Response file variables to configure disk-based I/O fencing](#)
- [Sample response file for configuring disk-based I/O fencing](#)
- [Configuring I/O fencing using response files](#)
- [Response file variables to configure server-based I/O fencing](#)

## Response file variables to configure disk-based I/O fencing

[Table C-1](#) lists the response file variables that specify the required information to configure disk-based I/O fencing for SFCFS.

**Table C-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.<br>(Required) |

**Table C-1** Response file variables specific to configuring disk-based I/O fencing  
*(continued)*

| Variable                              | List or Scalar | Description                                                                                                                                                                                                                                  |
|---------------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vxfen_config_fencing_option}      | Scalar         | Specifies the I/O fencing configuration mode.<br><br><ul style="list-style-type: none"> <li>■ 1—Coordination Point Server-based I/O fencing</li> <li>■ 2—Coordinator disk-based I/O fencing</li> <li>■ 3—Disabled mode</li> </ul> (Required) |
| CFG {vxfen_config_fencing_mechanism}  | Scalar         | Specifies the I/O fencing mechanism.<br>(Optional)                                                                                                                                                                                           |
| CFG{vxfen_config_fencing_dg}          | Scalar         | Specifies the disk group for I/O fencing.<br>(Optional)<br><br><b>Note:</b> You must define either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.                                                    |
| CFG{vxfen_config_fencing_newdg_disks} | List           | Specifies the disks to use to create a new disk group for I/O fencing.<br>(Optional)<br><br><b>Note:</b> You must define either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.                       |

## Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions. See [“Response file variables to configure disk-based I/O fencing”](#) on page 323.

```
#
# Configuration Values:
#
```

```
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;

$CFG{prod}="SFCFS51";

$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{vxfen_config_fencing_dg}="fendg";
$CFG{vxfen_config_fencing_mechanism}="dmp";
$CFG{vxfen_config_fencing_newdg_disks}=
  [ qw(rhdisk75 rhdisk76 rhdisk77) ];
$CFG{vxfen_config_fencing_option}=2;
```

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

### To configure I/O fencing using response files

- 1 Make sure that Storage Foundation Cluster File System 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 83.
- 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 324.  
See [“Sample response file for configuring server-based I/O fencing”](#) on page 328.

- 4 Edit the values of the response file variables as necessary.  
See “[Response file variables to configure disk-based I/O fencing](#)” on page 323.  
See “[Response file variables to configure server-based I/O fencing](#)” on page 326.
- 5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfdfs -responsefile /tmp/response_file
```

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

## Response file variables to configure server-based I/O fencing

You can use a CP server response file to configure server-based customized I/O fencing. The installer uses the CP server response file for the following types of I/O fencing configurations:

- Client cluster fencing (server-based I/O fencing configuration itself)  
The installer configures server-based customized I/O fencing on the SFCFS cluster without prompting for user input.
- Disk-based fencing with the disk group already created  
The installer configures fencing in disk-based mode on the SFCFS cluster without prompting for user input.  
Disk-based fencing configuration is one in which SCSI-3 disks are used as the only coordination points.  
Disk-based fencing with the disk group already created means that the disk group consisting of the coordinating disks already exists on the SFCFS cluster nodes.
- Disk-based fencing with the disk group to be created  
The installer creates the disk group and configures fencing properly on all the nodes in the SFCFS cluster without user intervention.  
Disk-based fencing with the disk group to be created means that the disk group does not exist yet, but will be created with the disks mentioned as coordination point.

[Table C-2](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

**Table C-2** CP server response file definitions

| Response file field        | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencing_cpc_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>                                                                                                                                                                                                                                                    |
| fencing_cpc_cpagentgrp     | <p>Name of the service group which will have the Coordination Point agent resource as part of it.</p> <p><b>Note:</b> This field is obsolete if the <code>fencing_cpc_config_cpagent</code> field is given a value of '0'.</p>                                                                                                                                                                                                                                                                                                                                                   |
| fencing_cpc_cps            | <p>Virtual IP address or Virtual hostname of the CP servers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| fencing_cpc_reusedg        | <p>This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks).</p> <p>Enter either a "1" or "0".</p> <p>Entering a "1" indicates reuse, and entering a "0" indicates do not reuse.</p> <p>When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text , such as<br/> <code>"\$CFG{fencing_cpc_reusedg}=0"</code> or<br/> <code>"\$CFG{fencing_cpc_reusedg}=1"</code> before proceeding with a silent installation.</p> |
| fencing_cpc_dgname         | <p>The name of the disk group to be used in the customized fencing, where at least one disk is being used.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| fencing_cpc_diffab         | <p>This response field indicates whether the CP servers and SFCFS clusters use different root brokers.</p> <p>Entering a "1" indicates that they are using different root brokers.</p> <p>Entering a "0" indicates that they are not using different root brokers.</p>                                                                                                                                                                                                                                                                                                           |

**Table C-2** CP server response file definitions (*continued*)

| Response file field   | Definition                                                                                                                                                                        |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencing_cpc_disks     | The disks being used as coordination points if any.                                                                                                                               |
| fencing_cpc_ncps      | Total number of coordination points being used, including both CP servers and disks.                                                                                              |
| fencing_cpc_ndisks    | The number of disks being used.                                                                                                                                                   |
| fencing_cpc_ports     | The port of the CP server that is denoted by <i>cps</i> .                                                                                                                         |
| fencing_cpc_ccab      | The name of the authentication broker (AB) for any one of the SFCFS cluster nodes.                                                                                                |
| fencing_cpc_cpsabport | The port at which the authentication broker (AB) mentioned above listens for authentication..                                                                                     |
| fencing_cpc_ccabport  | The port at which the authentication broker (AB) mentioned above listens for authentication.                                                                                      |
| fencing_cpc_mechanism | The disk mechanism that is used by customized fencing.<br><br>The value for this field is either "raw" or "dmp"                                                                   |
| fencing_cpc_cpsab     | The name of the authentication broker (AB) for any one of the CP servers.                                                                                                         |
| fencing_cpc_security  | This field indicates whether security is enabled or not<br><br>Entering a "1" indicates that security is enabled.<br>Entering a "0" indicates that security has not been enabled. |

## Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing :

# Storage Foundation Cluster File System components

This appendix includes the following topics:

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

## Veritas Storage Foundation installation filesets

[Table D-1](#) shows the fileset name and contents for each English language fileset for Veritas Storage Foundation. The table also gives you guidelines for which filesets to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and Veritas Cluster Server (VCS) filesets, the combined functionality is called Storage Foundation and High Availability.

See [“Veritas Cluster Server installation filesets”](#) on page 331.

**Table D-1** Veritas Storage Foundation filesets

| fileset    | Contents                                                                                                                                                | Configuration |
|------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| VRTSaslapm | Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries<br><br>Required for the support and compatibility of various storage arrays. | Minimum       |

**Table D-1** Veritas Storage Foundation filesets (*continued*)

| <b>fileset</b> | <b>Contents</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>Configuration</b> |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| VRTSat         | <p>Symantec Product Authentication Service</p> <p>Installs the Symantec Product Authentication Service, which provides authentication services to other Symantec products.</p> <p>This package contains a server and client component. The server provides services for a root broker, authentication broker, or both.</p> <p>The client allows Symantec products to communicate with the brokers.</p> <p>Required to use Symantec Product Authentication Service.</p> | All                  |
| VRTSperl       | Perl 5.10.0 for Veritas                                                                                                                                                                                                                                                                                                                                                                                                                                                | Minimum              |
| VRTSveki       | <p>Veritas Kernel Interface</p> <p>Contains a common set of modules that other Veritas drivers use.</p>                                                                                                                                                                                                                                                                                                                                                                | Minimum              |
| VRTSvlic       | <p>Veritas License Utilities</p> <p>Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.</p>                                                                                                                                                                                                                                                  | Minimum              |
| VRTSvxfs       | <p>Veritas File System binaries</p> <p>Required for VxFS file system support.</p>                                                                                                                                                                                                                                                                                                                                                                                      | Minimum              |
| VRTSvxvm       | Veritas Volume Manager binaries                                                                                                                                                                                                                                                                                                                                                                                                                                        | Minimum              |
| VRTSdbed       | Veritas Storage Foundation for Oracle                                                                                                                                                                                                                                                                                                                                                                                                                                  | Recommended          |
| VRTSob         | Veritas Enterprise Administrator                                                                                                                                                                                                                                                                                                                                                                                                                                       | Recommended          |

**Table D-1** Veritas Storage Foundation filesets (*continued*)

| fileset   | Contents                                                                                                                                                                                                                                                                                                                                                                                                 | Configuration |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| VRTSodm   | ODM Driver for VxFS<br><br>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   |
| VRTSsfmh  | Veritas Storage Foundation Managed Host<br><br>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:<br><br><a href="http://www.symantec.com/business/storage-foundation-manager">http://www.symantec.com/business/storage-foundation-manager</a> | Recommended   |
| VRTSspt   | Veritas Software Support Tools                                                                                                                                                                                                                                                                                                                                                                           | Recommended   |
| VRTSfssdk | Veritas File System Software Developer Kit<br><br>For VxFS APIs, the package contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.                                                                                                                                                                                       | All           |

## Veritas Cluster Server installation filesets

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

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

See “[Veritas Storage Foundation installation filesets](#)” on page 329.

**Table D-2** VCS installation filesets

| fileset   | Contents                                                                                                                                                                                                                                | Configuration |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| VRTSgab   | Veritas Cluster Server group membership and atomic broadcast services                                                                                                                                                                   | Minimum       |
| VRTSllt   | Veritas Cluster Server low-latency transport                                                                                                                                                                                            | Minimum       |
| VRTSvc    | Veritas Cluster Server                                                                                                                                                                                                                  | Minimum       |
| VRTSvcsg  | Veritas Cluster Server Bundled Agents                                                                                                                                                                                                   | Minimum       |
| VRTSvxfen | Veritas I/O Fencing                                                                                                                                                                                                                     | Minimum       |
| VRTScutil | Veritas Cluster Server Utilities                                                                                                                                                                                                        | Recommended   |
| VRTSvcsea | Consolidated database and enterprise agent packages                                                                                                                                                                                     | Recommended   |
| VRTScps   | Veritas Coordination Point Server<br><br>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. | Advanced      |

## Veritas Cluster File System installation filesets

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

When you install all CFS filesets and all the packages that comprise Storage Foundation and Veritas Cluster Server, the resulting functionality is called Storage Foundation Cluster File System.

See [“Veritas Storage Foundation installation filesets”](#) on page 329.

See [“Veritas Cluster Server installation filesets”](#) on page 331.

**Table D-3** CFS installation filesets

| fileset  | 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 filesets

[Table D-4](#) lists the filesets that are obsolete or reorganized for Storage Foundation and Storage Foundation High Availability.

**Table D-4** Veritas Storage Foundation obsolete and reorganized filesets

| fileset        | Description          |
|----------------|----------------------|
| Infrastructure |                      |
| SYMClma        | Obsolete             |
| VRTSaa         | Included in VRTSsmfh |
| VRTSccg        | Included in VRTSsmfh |
| VRTSdbms3      | Obsolete             |
| VRTSicsco      | Obsolete             |
| VRTSjre        | Obsolete             |
| VRTSjre15      | Obsolete             |
| VRTSmh         | Included in VRTSsmfh |
| VRTSobc33      | Obsolete             |
| VRTSobgui      | Obsolete             |
| VRTSpbx        | Obsolete             |
| VRTSsfm        | Obsolete             |

**Table D-4** Veritas Storage Foundation obsolete and reorganized filesets  
*(continued)*

| <b>fileset</b>   | <b>Description</b>    |
|------------------|-----------------------|
| VRTSweb          | Obsolete              |
| Product packages |                       |
| VRTSacclib       | Obsolete              |
| VRTSalloc        | Obsolete              |
| VRTScmccc        | Obsolete              |
| VRTScmcs         | Obsolete              |
| VRTScscm         | Included in VRTScutil |
| VRTScscw         | Included in VRTScutil |
| VRTScsocw        | Included in VRTScutil |
| VRTScssim        | Included in VRTScutil |
| VRTSd2gui        | Included in VRTSdbed  |
| VRTSdb2ed        | Included in VRTSdbed  |
| VRTSdbcom        | Included in VRTSdbed  |
| VRTSdbed         | Included in VRTSdbed  |
| VRTSdcli         | Obsolete              |
| VRTSddlpr        | Obsolete              |
| VRTSdsa          | Obsolete              |
| VRTSfsman        | Included in mainpkg   |
| VRTSfsmnd        | Included in mainpkg   |
| VRTSfspro        | Included in VRTSsmfh  |
| VRTSgapms        | Obsolete              |
| VRTSmapro        | Included in VRTSsmfh  |
| VRTSorgui        | Obsolete              |
| VRTSvail         | Obsolete              |

**Table D-4** Veritas Storage Foundation obsolete and reorganized filesets  
(continued)

| fileset       | Description                         |
|---------------|-------------------------------------|
| VRTSvcldb     | Included in VRTSvcsea               |
| VRTSvcsor     | Included in VRTSvcsea               |
| VRTSvcsvr     | Included in VRTSvc                  |
| VRTSvdid      | Obsolete                            |
| VRTSvmman     | Included in mainpkg                 |
| VRTSvmpro     | Included in VRTSsmfh                |
| VRTSvrpro     | Included in VRTSob                  |
| VRTSvrw       | Obsolete                            |
| VRTSvxmsa     | Obsolete                            |
| Documentation | All Documentation packages obsolete |

**Veritas Storage Foundation obsolete and reorganized installation filesets**

# High availability agent information

This appendix includes the following topics:

- [About agents](#)
- [CVMCluster agent](#)
- [CVMVxconfigd agent](#)
- [CVMVolDg agent](#)
- [CFSMount 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 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 SFCFS agent are described in this appendix.

## VCS agents included within SFCFS

SFCFS includes the following VCS agents:

- CVMCluster agent
- CVMVxconfigd agent
- CVMVolDG agent
- CFSSMount agent
- Coordination Point agent

An SFCFS 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 User's Guide*

## CVMCluster agent

The CVMCluster agent controls system membership on the cluster port that is associated with Veritas Volume Manager (VxVM).

### Entry points for CVMCluster agent

[Table E-1](#) describes the entry points used by the CVMCluster agent.

**Table E-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 E-2](#) describes the user-modifiable attributes of the CVMCluster resource type.

**Table E-2** CVMCluster agent attributes

| Attribute    | Dimension          | Description                                                                                                   |
|--------------|--------------------|---------------------------------------------------------------------------------------------------------------|
| CVMClustName | string-scalar      | Name of the cluster.                                                                                          |
| CVMNodeAddr  | string-association | List of host names and IP addresses.                                                                          |
| CVMNodeId    | string-association | Associative list. The first part names the system; the second part contains the LLT ID number for the system. |
| CVMTransport | string-scalar      | Specifies the cluster messaging mechanism.<br>Default = gab<br><b>Note:</b> Do not change this value.         |
| PortConfigd  | integer-scalar     | The port number that is used by CVM for vxconfigd-level communication.                                        |
| PortKmsgd    | integer-scalar     | The port number that is used by CVM for kernel-level communication.                                           |
| CVMTimeout   | integer-scalar     | Timeout in seconds used for CVM cluster reconfiguration. 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 SFCFS 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
)
```

For a more extensive `main.cf` example that includes the CVMCluster resource:

## 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 SFCFS installation configures the `syslog` option for the CVMVxconfigd agent.

## Entry points for CVMVxconfigd agent

[Table E-3](#) describes the entry points for the CVMVxconfigd agent.

**Table E-3** Vxconfigd entry points

| Entry Point | Description                                  |
|-------------|----------------------------------------------|
| Online      | Starts the vxconfigd daemon                  |
| Offline     | N/A                                          |
| Monitor     | Monitors whether vxconfigd daemon is running |

## Attribute definition for CVMVxconfigd agent

[Table E-4](#) describes the modifiable attributes of the CVMVxconfigd resource type.

**Table E-4** CVMVxconfigd agent attribute

| Attribute        | Dimension | Description                                                                                                                                 |
|------------------|-----------|---------------------------------------------------------------------------------------------------------------------------------------------|
| CVMVxconfigdArgs | keylist   | List of the arguments that are sent to the online entry point.<br><br>Symantec recommends always specifying the <code>syslog</code> option. |

## 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 }
)
```

For a more extensive `main.cf` that includes the `CVMVxconfigd` resource:

## CVMVolDg agent

The CVMVolDg agent represents and controls CVM diskgroups and CVM volumes within the diskgroups. The global nature of CVM diskgroups and volumes requires importing them only once on the CVM master node.

Configure the CVMVolDg agent for each disk group used by an service group. A disk group must be configured to only one service group.If cluster file systems are used for the database, configure the CFSMount agent for each volume in the disk group.

## Entry points for CVMVolDg agent

[Table E-5](#) describes the entry points used by the CVMVolDg agent.

**Table E-5** CVMVolDg agent entry points

| Entry Point | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Online      | <p>Starts all volumes in the shared disk group specified by the CVMVolume attribute.</p> <p>Makes sure that the shared disk groups are imported. Note that the CVMVolDg agent does not import shared disk groups. If the disk group is not imported, the online script returns an error.</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>Sets the diskgroup activation mode to off so that all the volumes in diskgroup are invalid.</p>                                                                                                                                                                                                                                                                                                                                                                           |

**Table E-5** CVMVolDg agent entry points (*continued*)

| Entry Point | Description                                                                                                                                                          |
|-------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Monitor     | Monitors specified critical volumes in the diskgroup. The CVMVolume attribute specifies these volumes. SFCFS requires specifying at least one volume in a diskgroup. |
| Clean       | Cleans up the temporary files created by the online entry point.                                                                                                     |

## Attribute definition for CVMVolDg agent

[Table E-6](#) describes the user-modifiable attributes of the CVMVolDg resource type.

**Table E-6** CVMVolDg agent attributes

| Attribute     | Dimension      | Description                                                                                              |
|---------------|----------------|----------------------------------------------------------------------------------------------------------|
| CVMDiskGroup  | string-scalar  | Names the diskgroup.                                                                                     |
| CVMVolume     | string-keylist | Lists critical volumes in the diskgroup. SFCFS requires specifying at least one volume in the diskgroup. |
| CVMActivation | string-scalar  | Sets the activation mode for the diskgroup.<br>Default = sw (shared-write)                               |

## CVMVolDg agent type definition

The CVMTypes.cf file includes the CVMVolDg type definition.

```

type CVMVolDg (
    static keylist RegList = { CVMActivation }
    static str ArgList[] = { CVMDiskGroup, CVMVolume,
        CVMActivation }
    str CVMDiskGroup
    keylist CVMVolume[]
    str CVMActivation
    temp int voldg_stat
)

```

## CVMVolDg agent sample configuration

Each service group requires a CVMVolDg resource type to be defined. The following is a sample configuration.

```
CVMVolDg ora_voldg (
    CVMDiskGroup = oradatadg
    CVMVolume = { oradata1, oradata2 }
    CVMActivation = sw
)
```

To see CVMVolDg defined in a more extensive example:

## 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.

### Entry points for CFSMount agent

[Table E-7](#) provides the entry points for the CFSMount agent.

**Table E-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.                                              |

### Attribute definition for CFSMount agent

[Table E-8](#) lists user-modifiable attributes of the CFSMount Agent resource type.

**Table E-8** CFSMount Agent attributes

| Attribute   | Dimension     | Description                       |
|-------------|---------------|-----------------------------------|
| MountPoint  | string-scalar | Directory for the mount point.    |
| BlockDevice | string-scalar | Block device for the mount point. |

**Table E-8** CFSMount Agent attributes (*continued*)

| Attribute              | Dimension      | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| NodeList               | string-keylist | List of nodes on which to mount. If NodeList is NULL, the agent uses the service group system list.                                                                                                                                                                                                                                                                                                                                                                           |
| MountOpt<br>(optional) | string-scalar  | Options for the mount command. To create a valid MountOpt attribute string: <ul style="list-style-type: none"> <li>■ Use the VxFS type-specific options only.</li> <li>■ Do not use the -o flag to specify the VxFS-specific options.</li> <li>■ Do not use the -v vxfs file system type option.</li> <li>■ Be aware the cluster option is not required.</li> <li>■ Specify options in comma-separated list: <pre>ro ro,cluster blkclear,mincache=closesync</pre> </li> </ul> |
| Policy (optional)      | string-scalar  | 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.                                                                                                                                                                                                       |

## CFSMount agent type definition

The `CFSTypes.cf` file includes the CFSMount agent type definition.

```
type CFSMount (
    static keylist RegList = { MountOpt, Policy, NodeList }
    static int FaultOnMonitorTimeouts = 1

    static int InfoTimeout = 0
    static int OnlineRetryLimit = 16

    static int OnlineWaitLimit = 0
    static str ArgList[] = { MountPoint, BlockDevice,
        MountOpt }
    NameRule = resource.MountPoint
    str MountPoint
    str MountType
```

```
str BlockDevice
str MountOpt

keylist NodeList
keylist Policy

temp str Primary
str SetPrimary
str RemountRes
str ForceOff
```

## CFSMount agent sample configuration

Each service group requires a CFSMount resource type to be defined.

```
CFSMount ora_mount (
    MountPoint = "/oradata"
    BlockDevice = "/dev/vx/dsk/oradatadg/oradatavol1"
    Primary = nebula;
)
```

To see CFSMount defined in a more extensive example:

# Troubleshooting information

This appendix includes the following topics:

- [Restarting the installer after a failed connection](#)
- [What to do if you see a licensing reminder](#)
- [Troubleshooting an installation on AIX](#)
- [Storage Foundation Cluster File System installation issues](#)
- [Storage Foundation Cluster File System 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' and validate using the command 'vxkeyless set NONE'

To comply with the terms of the EULA, and remove these messages, you must do one of the following within 60 days:

- Install a valid license key corresponding to the functionality in use on the host. See “[Installing Veritas product license keys](#)” on page 108. After you install the license key, you must validate the license key using the following command:

```
# vxkeyless set NONE
```

- Continue with keyless licensing by managing the server or cluster with a management server. For more information about keyless licensing, see the following URL: <http://go.symantec.com/sfhakeyless>

## Troubleshooting an installation on AIX

Save a copy of `/var/adm/ras/errtmpl` and `/etc/trcfmt` files before you install the product. If the packages fail to install due to the template file is corrupted error message, replace `/var/adm/ras/errtmpl` file and `/etc/trcfmt` file with the ones that you had saved, uninstall all the packages installed.

See “[Preparing to uninstall a Storage Foundation product](#)” on page 290.

Then reinstall.

See “[About the common product installer](#)” on page 49.

## Storage Foundation Cluster File System installation issues

If you encounter any issues installing SFCFS, 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.

Suggested solution: You need to set up the systems to allow remote access using `ssh` or `rsh`.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 29.

---

**Note:** Remove remote shell permissions after completing the SFCFS 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 maximum number of processes allowed per user may not be large enough. This kernel attribute is a tunable and can be changed on any node of the cluster.

To determine the current value of "Maximum number of PROCESSES allowed per user", enter:

```
# lsattr -H -E -l sys0 -a maxuproc
```

To see the default value of this tunable and its valid range of values, enter:

```
# odmget -q "attribute=maxuproc" PdAt
```

If necessary, you can change the value of the tunable using the `smitty` interface:

```
# smitty chgsys
```

You can also directly change the `CuAt` class using the following command:

```
# chdev -l sys0 -a maxuproc=600
```

Increasing the value of the parameter takes effect immediately; otherwise the change takes effect after a reboot.

See the `smitty` and `chdev` 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.

Suggested solution: Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

## Storage Foundation Cluster File System 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 6 or 7.
- 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: 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 -V 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 28.
- 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/jeopardy

Network partition (or split brain) is a condition where a network failure can be misinterpreted as a failure of one or more nodes in a cluster. If one system in the cluster incorrectly assumes that another system failed, it may restart applications already running on the other system, thereby corrupting data. CFS tries to prevent this by having redundant heartbeat links.

At least one link must be active to maintain the integrity of the cluster. If all the links go down, after the last network link is broken, the node can no longer communicate with other nodes in the cluster. Thus the cluster is in one of two possible states. Either the last network link is broken (called a network partition condition), or the last network link is okay, but the node crashed, in which case it is not a network partition problem. It is not possible to identify whether it is the first or second state, so a kernel message is issued to indicate that a network partition may exist and there is a possibility of data corruption.

Jeopardy is a condition where a node in the cluster has a problem connecting to other nodes. In this situation, the link or disk heartbeat may be down, so a jeopardy warning may be displayed. Specifically, this message appears when a node has only one remaining link to the cluster and that link is a network link. This is considered a critical event because the node may lose its only remaining connection to the network.

---

**Warning:** Do not remove the communication links while shared storage is still connected.

---

## Low memory

Under heavy loads, software that manages heartbeat communication links may not be able to allocate kernel memory. If this occurs, a node halts to avoid any chance of network partitioning. Reduce the load on the node if this happens frequently.

A similar situation may occur if the values in the `/etc/llttab` files on all cluster nodes are not correct or identical.



# Troubleshooting cluster installation

This appendix includes the following topics:

- [Installer cannot create UUID for the cluster](#)
- [The vxfcntl utility fails when SCSI TEST UNIT READY command fails](#)
- [Troubleshooting on the CP server](#)
- [Troubleshooting server-based I/O fencing on the SFCFS cluster](#)
- [Troubleshooting server-based I/O fencing in mixed mode](#)

## 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 SFCFS configuration, upgrade, or when you add a node to the cluster using the installer.

**Workaround:** To start SFCFS, you must run the `uuidconfig.pl` script manually to configure the UUID on each cluster node.

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

## 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 on the CP server

All the CP server operations and messages are logged in the `/var/VRTScps/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 single node VCS or SFHA cluster hosting the CP server.

The following files contain logs and text files that may be useful in understanding and troubleshooting a CP server:

- `/var/VRTScps/log/cpsrvr_[ABC].log`
- `/var/VRTSat/vrtsat_broker.txt` (Security related)

If the `vxcperv` process fails on the CP server, then review the following diagnostic files:

- `/var/VRTScps/diag/FFDC_CPS_<pid>_vxcperv.log`
- `/var/VRTScps/diag/stack_<pid>_vxcperv.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.

---

## CP server service group issues

If you cannot bring up the CPSSG service group after the CP server configuration, verify that the CPSSG service group and its resources are valid and properly configured in the VCS configuration.

Check the VCS engine log to see if any of the CPSSG service group resources are FAULTED. The engine log is located in the following directory:

```
/var/VRTSvcs/log/engine_[ABC].log
```

The resources that are configured under the CPSSG service groups are displayed in the following figures:

- CPSSG group and dependency figure for CP server hosted on a single node VCS cluster:
- CPSSG group and dependency figure for CP server hosted on an SFHA cluster:

---

**Note:** For information about general VCS troubleshooting procedures, refer to the Veritas™ Cluster Server User's Guide, Version 5.1.

---

## Testing the connectivity of the CP server

The connectivity of the CP server can be tested using the `cpsadm` command. The following `cpsadm` command tests 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.

Issuing the command on the SFCFS cluster nodes requires the environment variables `CPS_USERNAME` and `CPS_DOMAINTYPE` to be set.

## Troubleshooting server-based I/O fencing on the SFCFS cluster

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs and text files that may be useful in understanding and/or troubleshooting fencing-related issues on a SFCFS cluster node.

## Issues during server-based fencing start up on SFCFS cluster node

The following issues may occur during fencing start up on the SFCFS cluster node:

- `cpsadm` command on the SFCFS cluster gives connection error
- Authentication failure
- Authorization failure
- Preexisting split-brain

### **cpsadm** command on the SFCFS cluster node gives connection error

If you receive a connection error message after issuing the `cpsadm` command on the SFCFS cluster, perform the following actions:

- Ensure that the CP server is reachable from all the SFCFS cluster nodes.
- Check that the correct CP server virtual IP/virtual hostname and port number are being used by the SFCFS cluster nodes.  
Check the `/etc/vxfenmode` file.
- Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.

### **Authorization failure**

Authorization failure occurs when the CP server's SFCFS cluster nodes or users are not added in the CP server configuration. Therefore, fencing on the SFCFS cluster node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points. To resolve this issue, add the SFCFS cluster node and user in the CP server configuration and restart fencing. Refer to the following section:

See [“Preparing the CP servers manually for use by the SFCFS cluster”](#) on page 133.

### **Preexisting split-brain**

To illustrate preexisting split-brain, assume there are three CP servers acting as coordination points. One of the three CP servers then becomes inaccessible. While in this state, also one client node leaves the cluster. When the inaccessible CP server restarts, it has a stale registration from the node which left the SFCFS cluster. In this case, no new nodes can join the cluster. Each node that attempts to join the cluster gets a list of registrations from the CP server. One CP server includes an extra registration (of the node which left earlier). This makes the joiner node conclude that there exists a preexisting split-brain between the joiner

node and the node which is represented by the stale registration. The situation is similar to that of preexisting split-brain, with coordinator disks, where the problem is solved by the administrator running the `vxfcntlclearpre` command. A similar solution is required using the `cpsadm` command.

The following `cpsadm` command can be used to clear a registration on a CP server:

```
# cpsadm -s cp_server -a unreg_node -c cluster_name -n nodeid
```

where `cp_server` is the virtual IP address or virtual hostname on which the CP server is listening, `cluster_name` is the VCS name for the SFCFS cluster, and `nodeid` specifies the node id of SFCFS cluster node.

After removing all stale registrations, the joiner node will be able to join the cluster.

## 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 all 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 SFCFS 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 SFCFS 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 SFCFS cluster nodes to the CP server(s).
- The cluster or nodes or users for the SFCFS cluster nodes have not been added on the new CP servers, thereby causing authorization failure.

### Vxfen service group activity after issuing the `vxfenswap` command

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.

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 re-elected 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.

## Troubleshooting server-based I/O fencing in mixed mode

The following procedure can be used to troubleshoot a mixed I/O fencing configuration (configuration using both coordinator disks and CP server for I/O fencing). This procedure involves using the following commands to obtain I/O fencing information:

- To obtain I/O fencing cluster information on the coordinator disks, run the following command on one of the cluster nodes:

```
# vxfenadm -s diskname
```

Any keys other than the valid keys used by the cluster nodes that appear in the command output are spurious keys.

- To obtain I/O fencing cluster information on the CP server, run the following command on one of the cluster nodes:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp\_server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster name* is the VCS name for the SFCFS cluster. Nodes which are not in GAB membership, but registered with CP server indicate a pre-existing network partition.

Note that when running this command on the SFCFS cluster nodes, you need to first export the `CPS_USERNAME` and `CPS_DOMAINTYPE` variables.

The `CPS_USERNAME` value is the user name which is added for this node on the CP server.

- To obtain the user name, run the following command on the CP server:

```
# cpsadm -s cp_server -a list_users
```

where *cp\_server* is the virtual IP address or virtual hostname on which the CP server is listening.

The CPS\_DOMAINTYPE value is vx.

The following are export variable command examples:

```
# export CPS_USERNAME=_HA_VCS_test-system@HA_SERVICES@test-system.symantec.com
```

```
# export CPS_DOMAINTYPE=vx
```

Once a pre-existing network partition is detected using the above commands, all spurious keys on the coordinator disks or CP server must be removed by the administrator.

### Troubleshooting mixed I/O fencing configuration (coordinator disks and CP server)

- 1 Review the current I/O fencing configuration by accessing and viewing the information in the `vxfenmode` file.

Enter the following command on one of the SFCFS cluster nodes:

```
# cat /etc/vxfenmode

vxfen_mode=customized
vxfen_mechanism=cps
scsi3_disk_policy=dmp
security=0
cps1=[10.140.94.101]:14250
vxfendg=vxfencoordg
```

- 2 Review the I/O fencing cluster information.

Enter the `vxfenadm -d` command on one of the cluster nodes:

```
# vxfenadm -d

I/O Fencing Cluster Information:
=====

Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:

    * 0 (galaxy)
      1 (nebula)

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

**3 Review the SCSI registration keys for the coordinator disks used in the I/O fencing configuration.**

Enter the `vxfenadm -s` command on each of the SFCFS cluster nodes.

```
# vxfenadm -s /dev/vx/rdmp/3pardata0_190
```

```
Device Name: /dev/vx/rdmp/3pardata0_190
Total Number Of Keys: 2
key[0]:
    [Numeric Format]: 86,70,66,69,65,68,48,48
    [Character Format]: VFBEAD00
    [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
    [Numeric Format]: 86,70,66,69,65,68,48,49
    [Character Format]: VFBEAD01
*    [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

```
# vxfenadm -s /dev/vx/rdmp/3pardata0_191
```

```
Device Name: /dev/vx/rdmp/3pardata0_191
Total Number Of Keys: 2
key[0]:
    [Numeric Format]: 86,70,66,69,65,68,48,48
    [Character Format]: VFBEAD00
    [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
    [Numeric Format]: 86,70,66,69,65,68,48,49
    [Character Format]: VFBEAD01
*    [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

**4 Review the CP server information about the cluster nodes.**

On the CPS server, run the `cpsadm list nodes` command to review a list of nodes in the cluster.

The command syntax is as follows:

```
# cpsadm -s cp_server -a list_nodes
```

where *cp server* is the virtual IP address or virtual hostname on which the CP server is listening.

For example:

```
# /opt/VRTS/bin/cpsadm -s 10.140.94.101 -a list_nodes
```

| ClusName | UUID                                   | Hostname(Node ID)  | Registered |
|----------|----------------------------------------|--------------------|------------|
| gl-rh2   | {25aeb8c6-1dd2-11b2-95b5-a82227078d73} | node_101(0)        | 0          |
| gl-rh2   | {25aeb8c6-1dd2-11b2-95b5-a82227078d73} | node_102(1)        | 0          |
| cpstest  | {a0cf10e8-1dd1-11b2-87dc-080020c8fa36} | node_220(0)        | 0          |
| cpstest  | {a0cf10e8-1dd1-11b2-87dc-080020c8fa36} | node_240(1)        | 0          |
| ictwo    | {f766448a-1dd1-11b2-be46-5d1da09d0bb6} | node_330(0)        | 0          |
| ictwo    | {f766448a-1dd1-11b2-be46-5d1da09d0bb6} | sassette(1)        | 0          |
| fencing  | {e5288862-1dd1-11b2-bc59-0021281194de} | CDC-SFLAB-CD-01(0) | 0          |
| fencing  | {e5288862-1dd1-11b2-bc59-0021281194de} | CDC-SFLAB-CD-02(1) | 0          |
| gl-su2   | {8f0a63f4-1dd2-11b2-8258-d1bcc1356043} | gl-win03(0)        | 0          |
| gl-su2   | {8f0a63f4-1dd2-11b2-8258-d1bcc1356043} | gl-win04(1)        | 0          |
| gl-su1   | {2d2d172e-1dd2-11b2-bc31-045b4f6a9562} | gl-win01(0)        | 0          |
| gl-su1   | {2d2d172e-1dd2-11b2-bc31-045b4f6a9562} | gl-win02(1)        | 0          |
| gl-ax4   | {c17cf9fa-1dd1-11b2-a6f5-6dbd1c4b5676} | gl-ax06(0)         | 0          |
| gl-ax4   | {c17cf9fa-1dd1-11b2-a6f5-6dbd1c4b5676} | gl-ax07(1)         | 0          |
| gl-ss2   | {da2be862-1dd1-11b2-9fb9-0003bac43ced} | galaxy(0)          | 1          |
| gl-ss2   | {da2be862-1dd1-11b2-9fb9-0003bac43ced} | nebula(1)          | 1          |

**5** Review the CP server list membership.

On the CP server, run the following command to review the list membership. The command syntax is as follows:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp\_server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster\_name* is the VCS name for the SFCFS cluster.

For example:

```
# cpsadm -s 10.140.94.101 -a list_membership -c gl-ss2
```

```
List of registered nodes: 0 1
```

## Checking keys on coordination points when vxfen\_mechanism value is set to cps

When I/O fencing is configured in customized mode and the vxfen\_mechanism value is set to cps, the recommended way of reading keys from the coordination points (coordinator disks and CP servers) is as follows:

- For coordinator disks, the disks can be put in a file and then information about them supplied to the vxfenadm command.

For example:

```
# vxfenadm -s all -f file_name
```

- For CP servers, the cpsadm command can be used to obtain the membership of the SFCFS cluster.

For example:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

Where *cp\_server* is the virtual IP address or virtual hostname on which CP server is configured, and *cluster\_name* is the VCS name for the SFCFS cluster.



# Sample SFCFS 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:

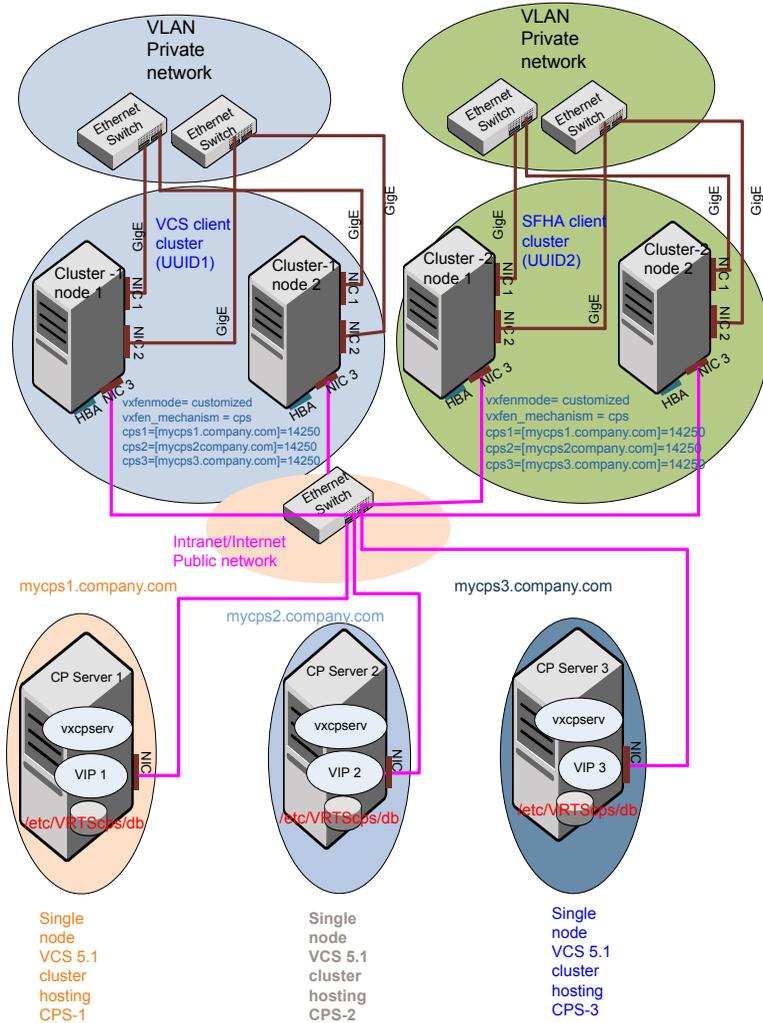
- Two unique client clusters that are served by 3 CP servers:
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:  
See [Figure H-2](#) on page 369.
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:  
See [Figure H-3](#) on page 371.
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:

### Two unique client clusters served by 3 CP servers

[\[Unresolved xref\]](#) 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 H-1** Two unique client clusters served by 3 CP servers



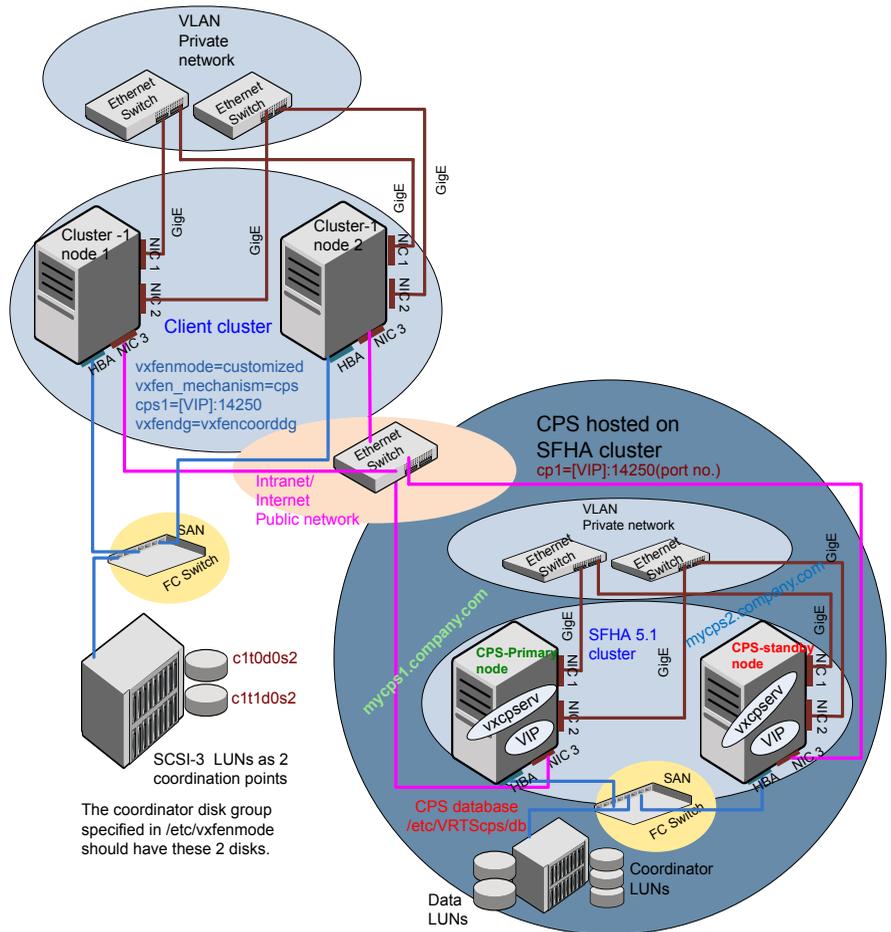
## Client cluster served by highly available CPS and 2 SCSI-3 disks

Figure H-2 displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to customized with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

**Figure H-2** Client cluster served by highly available CP server and 2 SCSI-3 disks



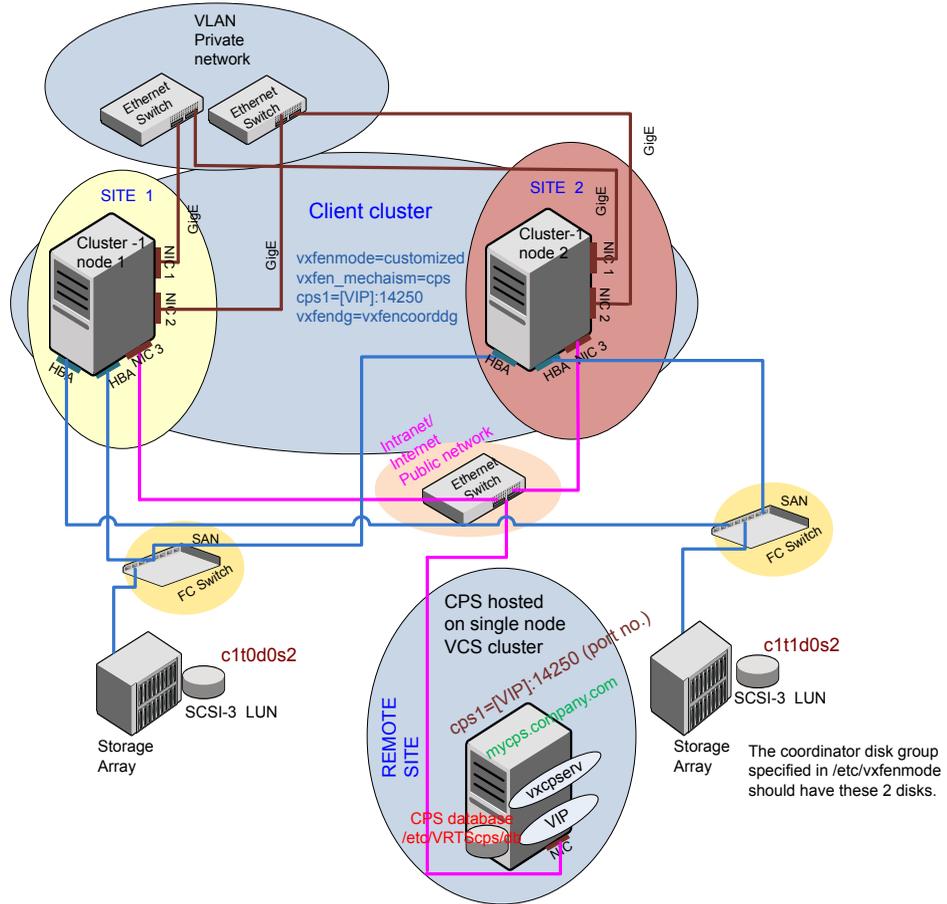
## Two node campus cluster served by remote CP server and 2 SCSI-3 disks

[Figure H-3](#) displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoorddg`. The third coordination point is a CP server on a single node VCS cluster.

**Figure H-3** Two node campus cluster served by remote CP server and 2 SCSI-3



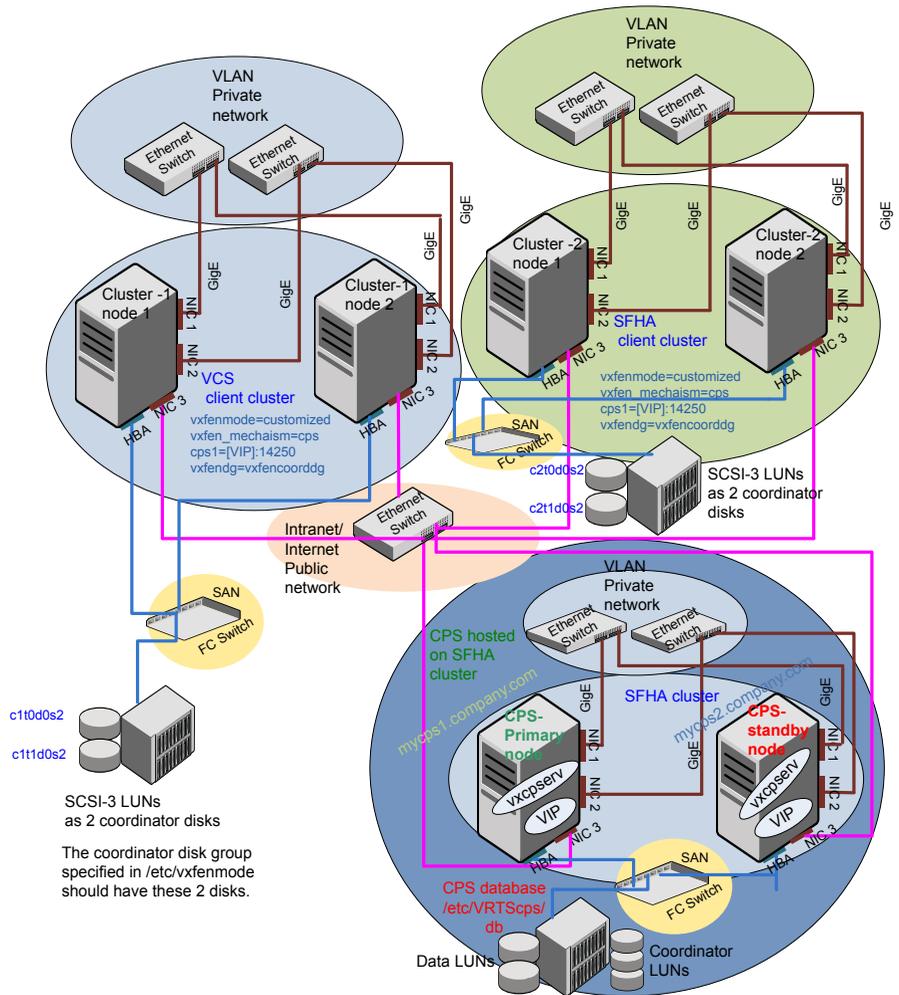
## Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

[Unresolved xref] displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

**Figure H-4** Multiple client clusters served by highly available CP server and 2 SCSI-3 disks



## Changing NFS server major numbers for VxVM volumes

This appendix includes the following topics:

- [Changing NFS server major numbers for VxVM volumes](#)

### Changing NFS server major numbers for VxVM volumes

Use the `haremajor` command to determine and reassign the major number that a system uses for shared VxVM volume block devices. For Veritas Volume Manager, the major number is set to the `vxio` driver number. To be highly available, each NFS server in a VCS cluster must have the same `vxio` driver number, or major number.

**To list the major number currently in use on a system**

- ◆ Use the command:

```
# haremajor -v
55
```

Run this command on each cluster node. If major numbers are not the same on each node, you must change them on the nodes so that they are identical.

**To list the available major numbers for a system**

- ◆ Use the command:

```
# haremajor -a
54, 56..58, 60, 62..
```

The output shows the numbers that are not in use on the system where the command is issued.

**To reset the major number on a system**

- ◆ You can reset the major number to an available number on a system. For example, to set the major number to 75 type:

```
# haremajord -s 75
```

# Configuring LLT over UDP using IPv6

This appendix includes the following topics:

- [Using the UDP layer of IPv6 for LLT](#)
- [Configuring LLT over UDP using IPv6](#)

## Using the UDP layer of IPv6 for LLT

Veritas Storage Foundation Cluster File System 5.1 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

### When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

## Configuring LLT over UDP using IPv6

The following checklist is to configure LLT over UDP:

- For UDP6, the multicast address is set to "-".
- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.

- Make sure the IPv6 addresses in the `/etc/llttab` files are consistent with the IPv6 addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.  
See [“Selecting UDP ports”](#) on page 377.
- 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 379.

## 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 378.
- See [“Sample configuration: links crossing IP routers”](#) on page 379.

Note that some of the fields in [Table J-1](#) differ from the command for standard LLT links.

[Table J-1](#) describes the fields of the link command that are shown in the `/etc/llttab` file examples.

**Table J-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/xti/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.<br>See <a href="#">“Selecting UDP ports”</a> on page 377.                                                                          |
| <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 379.

[Table J-2](#) describes the fields of the `set-addr` command.

**Table J-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:

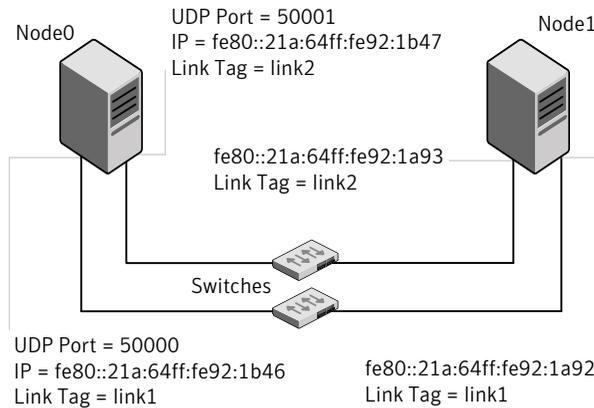
```
# netstat -a | more
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address Foreign Address (state)
tcp      0      0 *.32778      *.*          LISTEN
tcp      0      0 *.32781      *.*          LISTEN
udp4     0      0 *.daytime    *.*
udp4     0      0 *.time       *.*
udp4     0      0 *.sunrpc     *.*
udp      0      0 *.snmp       *.*
```

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 J-1** depicts a typical configuration of direct-attached links employing LLT over UDP.

**Figure J-1** A typical configuration of direct-attached links that use LLT over UDP



The configuration that the /etc/llttab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the /etc/llttab file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```

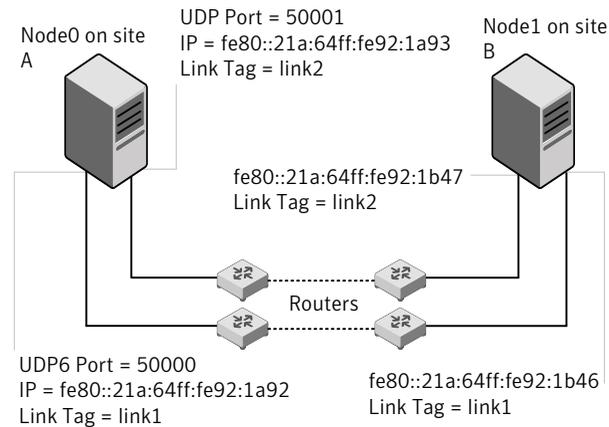
The file for Node 1 resembles:

```
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/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

## Sample configuration: links crossing IP routers

Figure J-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 J-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/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/xti/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/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 /dev/xti/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](#)
- [Configuring LLT over UDP](#)

## Using the UDP layer for LLT

Veritas Storage Foundation Cluster File System 5.1 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

### When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

LLT over UDP is slower than LLT over Ethernet. Use LLT over UDP only when the hardware configuration makes it necessary.

## Configuring LLT over UDP

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on different physical networks.

If the LLT private links are not on different physical networks, then make sure that the links are on separate subnets. Set the broadcast address in `/etc/llttab` explicitly depending on the subnet for each link.

See [“Broadcast address in the `/etc/llttab` file”](#) on page 382.

- 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 384.
- Set the broadcast address correctly for direct-attached (non-routed) links.  
See [“Sample configuration: direct-attached links”](#) on page 386.
- 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 387.

## 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/xti/udp - udp 50000 - 192.168.9.1 192.168.9.255
link link2 /dev/xti/udp - udp 50001 - 192.168.10.1 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

- Display the content of the `/etc/llttab` file on the second node nebula:

```
nebula # cat /etc/llttab

set-node nebula
set-cluster 1
link link1 /dev/xti/udp - udp 50000 - 192.168.9.2 192.168.9.255
link link2 /dev/xti/udp - udp 50001 - 192.168.10.2 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

## The link command in the /etc/llttab file

Review the link command information in this section for the /etc/llttab file. See the following information for sample configurations:

- See “[Sample configuration: direct-attached links](#)” on page 386.
- See “[Sample configuration: links crossing IP routers](#)” on page 387.

[Table K-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 K-1** Field description for link command in /etc/llttab

| Field                | Description                                                                                                                                                                                               |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>tag-name</i>      | A unique string that is used as a tag by LLT; for example link1, link2,....                                                                                                                               |
| <i>device</i>        | The device path of the UDP protocol; for example /dev/xti/udp.                                                                                                                                            |
| <i>node-range</i>    | Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.                                                                                                                 |
| <i>link-type</i>     | Type of link; must be "udp" for LLT over UDP.                                                                                                                                                             |
| <i>udp-port</i>      | Unique UDP port in the range of 49152-65535 for the link.<br>See “ <a href="#">Selecting UDP ports</a> ” on page 384.                                                                                     |
| <i>MTU</i>           | "-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.              |
| <i>IP address</i>    | IP address of the link on the local node.                                                                                                                                                                 |
| <i>bcast-address</i> | <ul style="list-style-type: none"> <li>■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address.</li> <li>■ "-" is the default for clusters spanning routers.</li> </ul> |

## 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 387.

[Table K-2](#) describes the fields of the set-addr command.

**Table K-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 | more
UDP
  Local Address          Remote Address         State
  -----
  *.*                   *.*                   Unbound
  *.32771                *.*                   Idle
  *.32776                *.*                   Idle
  *.32777                *.*                   Idle
  *.name                 *.*                   Idle
  *.biff                 *.*                   Idle
  *.talk                 *.*                   Idle
  *.32779                *.*                   Idle
  .
  .
  .
  *.55098                *.*                   Idle
  *.syslog               *.*                   Idle
  *.58702                *.*                   Idle
  *.*                   *.*                   Unbound
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

## Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# 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.

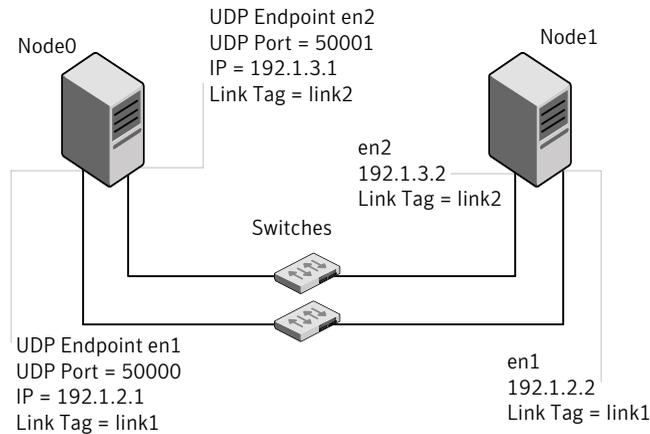
```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100
```

```
link link1 /dev/xti/udp - udp 50000 - 192.168.30.1  
192.168.30.255  
link link2 /dev/xti/udp - udp 50001 - 192.168.31.1  
192.168.31.255
```

## Sample configuration: direct-attached links

Figure K-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

**Figure K-1** A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0  
set-cluster 1  
#configure Links  
#link tag-name device node-range link-type udp port MTU \  
IP-address boast-address  
link link1 /dev/xti/udp - udp 50000 - 192.1.2.1 192.1.2.255  
link link2 /dev/xti/udp - udp 50001 - 192.1.3.1 192.1.3.255
```

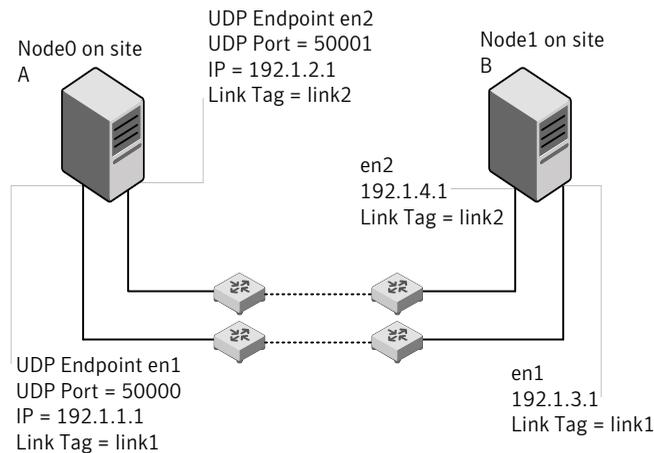
The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
# configure Links
# link tag-name device node-range link-type udp port MTU \
IP-address broadcast-address
link link1 /dev/xti/udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 /dev/xti/udp - udp 50001 - 192.1.3.2 192.1.3.255
```

## Sample configuration: links crossing IP routers

[Figure K-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 K-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/xti/udp - udp 50000 - 192.1.3.1 -
link link2 /dev/xti/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/xti/udp - udp 50000 - 192.1.1.1 -
link link2 /dev/xti/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
```

# Index

## A

### agents

- about 337
- CFSMount 344
- CVMCluster 338
- CVMVolDg 342
- CVMVxconfigd 340
- disabling 293
- of VCS 338

### application

- database replication 269

### applications, stopping 150

### attributes

- about agent attributes 337
- CFSMount agent 344
- CVMCluster agent 339
- CVMVolDg agent 339, 343
- CVMVxconfigd agent 341
- UseFence 120

## C

### cabling shared devices 36

### CFS

- mount and unmount failures 350
- synchronization 200
- troubleshooting 350

### CFSMount agent 344

- attributes 344
- entry points 344
- sample configuration 345–346
- type definition 345

### CFSTypes.cf 345

### cluster

- verifying operation 209

### command failures 351

### commands

- gabconfig 208
- gcoconfig 256
- hastatus 209
- hasys 210
- lltconfig 202

### commands (*continued*)

- lltstat 205
- vradmin 270
- vxassist 263, 265
- vxdisksetup (initializing disks) 109
- vxprint 271
- vxvol 263

### configuration file

- main.cf 201

### coordinator disks

- DMP devices 79
- for I/O fencing 79
- setting up 118

### CVM

- CVMTypes.cf file 339

### CVMCluster agent 338

- attributes 339
- entry points 338
- sample configuration 340
- type definition 339

### CVMTypes.cf

- definition, CVMCluster agent 339
- definition, CVMVolDg agent 343
- definition, CVMVxconfigd agent 341

### CVMVolDg agent 342

- attributes 343
- entry points 342
- sample configuration 343
- type definition 343

### CVMVxconfigd agent 340

- attributes 341
- CVMTypes.cf 341
- entry points 341
- sample configuration 342
- type definition 341

## D

### data disks

- for I/O fencing 79

### deinstalling the Volume Manager 290

### disabling the agents 293

**d**

- disks
  - adding and initializing 109
  - coordinator 118
  - testing with vxfcntlhdw 113
  - verifying node access 115

**F**

- Fibre Channel fabric 38
- files
  - main.cf 201
- freezing service groups 150

**G**

- GAB
  - port membership information 208
  - verifying 208
- gabconfig command 208
  - a (verifying GAB) 208
- gabtab file
  - verifying after installation 202
- gcoconfig command 256
- global clusters
  - configuration
    - using VVR 255
  - configuring VCS service groups 255
  - illustration of dependencies 272
  - migration 282–283
  - replicating database volumes 272
  - requirements 248
- groups
  - log owner 272
  - RVG 273

**H**

- hastatus -summary command 209
- hasys -display command 210
- high availability issues 352
  - low memory 353
  - network partition 352

**I**

- I/O fencing
  - checking disks 113
  - setting up 117
  - shared storage 113
- installing
  - Root Broker 73

**J**

- jeopardy 352

**L**

- links
  - private network 202
- LLT
  - verifying 205
- lltconfig command 202
- llthosts file
  - verifying after installation 202
- lltstat command 205
- llttab file
  - verifying after installation 202
- log files 356

**M**

- main.cf file 201
- manual pages
  - potential problems 351
  - troubleshooting 351
- membership information 208

**N**

- network partition 352
- nodes
  - adding application nodes
    - about 213
    - configuring GAB 220
    - configuring LLT 220
    - configuring VXFEN 220
    - enabling dlpi driver 214
    - starting Volume Manager 219
  - preparing application nodes
    - configuring CVM 228
  - removing allg nodes
    - editing VCS configuration files 238
    - GAB configuration 239
    - LLT configuration 238
    - modifying VCS configuration 239
  - removing application nodes
    - Oracle procedure 237
- NTP
  - network time protocol daemon 200

**P**

- PATH variable
  - VCS commands 205
- persistent reservations
  - SCSI-3 35
- planning to upgrade VVR 26
- port a
  - membership 208
- port h
  - membership 208
- port membership information 208
- preinstallation 26
- preparing to upgrade VVR 150
- primary site
  - creating SRL volume 262
  - setting up 251
  - setting up replication objects 263
  - VCS configuration 273, 277
- problems
  - accessing manual pages 351
  - executing file system commands 351

**Q**

- Quick I/O
  - performance on CFS 352

**R**

- removing
  - the Replicated Data Set 294
- Replicated Data Set
  - removing the 294
- replication
  - automatic synchronization 270
  - configuring on both sites 254
  - full synchronization with Checkpoint 270
  - modifying VCS configuration 272
  - options 249
  - setting up primary site 251
  - setting up secondary site 252
  - supported hardware 248
  - supported software 248
  - supported technologies 248
  - using VVR 262
  - verifying status 271
- resources
  - CVMVolDg 273
  - RVGSharedPri 273

- Root Broker
  - installing 73

**S**

- SAN
  - see Storage Area Network 38
- SCSI
  - changing initiator IDs 35
- SCSI ID
  - changing 36
  - verifying 36
- SCSI-3
  - persistent reservations 35
- SCSI-3 persistent reservations
  - verifying 117
- secondary site
  - configuring replication 265
  - creating SRL volume 265
  - setting up 252
  - setting up disk groups 267
  - setting up RLINKs 266
- service groups
  - freezing 150
  - VCS, for global clusters 255
- setup
  - cabling shared devices 36
  - SCSI Initiator ID 35
- SF Oracle RAC
  - takeover 284
- SFCFS
  - coordinator disks 118
  - illustration of global cluster dependencies 272
- SFCFS installation
  - verifying
    - cluster operations 205
    - GAB operations 205
    - LLT operations 205
- Shared storage
  - Fibre Channel 35
- shared storage
  - setting SCSI initiator ID 35
- split brain 352
- stopping
  - applications 150
- Storage Area Network 38
- Symantec Product Authentication Service 73
- system state attribute value 209

**T**

## troubleshooting

- accessing manual pages 351
- executing file system commands 351

**U**

## upgrading VVR

- from 4.0 27
- planning 26
- preparing 150

**V**

## VCS

- command directory path variable 205
- configuration, for database volume replication 272
- configuring service groups 255

## VCS configuration

- for replication 272

VCS Global cluster option. *See* GCOVeritas Volume Replicator. *See* VVR

## vradmin

- delpri 295
- stoprep 294

## vradmin command 270

## VVR

- application database replication 269
- configuring global clusters 255
- configuring on both sites 254
- database volume replication
  - configuring VCS 272
- defining heartbeat cluster objects 256
- defining remote clusters 256
- global cluster overview 261
- primary site
  - creating SRL volume 262
  - setting up replication objects 263
- replication
  - using automatic synchronization 270
  - using full synchronization with
    - Checkpoint 270
- replication agents 249
- secondary site
  - configuring replication 265
  - creating SRL volume 265
  - setting up disk groups 267
  - setting up RLINKs 266
- setting up primary site 251

VVR (*continued*)

- setting up replication 262
- setting up secondary site 252
- types of replication 249
- VCS configuration
  - application database service group 273
  - CVMoIDG resource 273
  - log owner group 272
  - primary site 273
  - RVG group 273
  - RVGSharedPri resource 273
  - secondary site 277
- verifying replication 271

## VVR 4.0

- planning an upgrade from 27

## vxassist command 263, 265

## vxdisksetup command 109

## vxprint command 271

## vxvol command 263