

Veritas Storage Foundation™ and High Availability Installation Guide

Solaris

6.0 Platform Release 1

Veritas Storage Foundation™ and High Availability Installation Guide

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

Product version: 6.0 PR1

Document version: 6.0PR1.0

Legal Notice

Copyright © 2012 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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

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

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

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

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

Support agreement resources

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

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

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec Web site.

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

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

doc_feedback@symantec.com

About Symantec Connect

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

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

Contents

Technical Support	4
Section 1 Installation overview and planning	17
Chapter 1 Introducing Storage Foundation and High Availability Solutions	19
About Veritas products	19
About Storage Foundation and High Availability	19
About Veritas Cluster Server	20
About Veritas high availability agents	20
About Veritas Volume Replicator	20
About Veritas Storage Foundation Cluster File System	21
About Veritas Storage Foundation for Sybase ASE CE	21
About Veritas Storage Foundation for Oracle® RAC	21
About Veritas graphical user interfaces	21
About Veritas Operations Manager	21
About Storage Foundation and High Availability features	22
About LLT and GAB	22
About configuring SFHA clusters for data integrity	23
About I/O fencing for SFHA in virtual machines that do not support SCSI-3 PR	24
About I/O fencing components	24
About global clusters	26
Chapter 2 Planning to install Storage Foundation and High Availability	29
About installation and configuration methods	29
Downloading the Storage Foundation and High Availability software	30
Chapter 3 System requirements	33
Release notes	33
Hardware compatibility list (HCL)	34
Supported operating systems	34

	Veritas File System requirements	34
	Disk space requirements	35
	Discovering product versions and various requirement information	35
	Database requirements	36
	I/O fencing requirements	36
	Coordinator disk requirements for I/O fencing	36
	CP server requirements	37
	Non-SCSI-3 I/O fencing requirements	40
	Number of nodes supported	40
Chapter 4	Licensing Veritas products	41
	About Veritas product licensing	41
	Setting or changing the product level for keyless licensing	42
	Installing Veritas product license keys	44
Section 2	Installing Veritas Storage Foundation and High Availability	45
Chapter 5	Preparing to install Storage Foundation High Availability	47
	Installation preparation overview	47
	About using ssh or rsh with the Veritas installer	48
	Setting up shared storage	49
	Setting up shared storage: SCSI disks	49
	Setting up shared storage: Fibre Channel	52
	Creating the /opt directory	53
	Creating root user	53
	Setting environment variables	54
	Mounting the product disc	54
	Assessing the system for installation readiness	55
	Symantec Operations Readiness Tools	55
Chapter 6	Installing Storage Foundation and High Availability using the script-based installer	57
	About the Veritas installer	57
	Installing Storage Foundation and High Availability using the installer	58
	Installing language packages	61
	Manually installing packages on solaris brand non-global zones	61

	Manually installing packages on solaris10 brand zones	62
Section 3	Configuring Veritas Storage Foundation and High Availability	63
Chapter 7	Preparing to configure Storage Foundation and High Availability	65
	About planning to configure I/O fencing	65
	Setting up the CP server	69
	Planning your CP server setup	70
	Installing the CP server using the installer	71
	Configuring the CP server cluster in secure mode	71
	Setting up shared storage for the CP server database	72
	Configuring the CP server using the configuration utility	73
	Configuring the CP server manually	82
	Verifying the CP server configuration	84
Chapter 8	Configuring Storage Foundation and High Availability	85
	Configuring Storage Foundation and High Availability Solutions	85
	Configuring Storage Foundation High Availability using the installer	85
Chapter 9	Configuring Storage Foundation High Availability for data integrity	107
	Setting up disk-based I/O fencing using installsfa	107
	Initializing disks as VxVM disks	107
	Checking shared disks for I/O fencing	108
	Configuring disk-based I/O fencing using installsfa	112
	Setting up disk-based I/O fencing manually	115
	Removing permissions for communication	115
	Identifying disks to use as coordinator disks	116
	Setting up coordinator disk groups	116
	Creating I/O fencing configuration files	117
	Modifying VCS configuration to use I/O fencing	118
	Verifying I/O fencing configuration	120
	Setting up server-based I/O fencing using installsfa	120
	Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsfa	120
	Setting up server-based I/O fencing manually	122

Preparing the CP servers manually for use by the SF HA cluster	123
Configuring CoordPoint agent to monitor coordination points	126
Verifying server-based I/O fencing configuration	128
Setting up non-SCSI-3 fencing in virtual environments manually	129
Sample /etc/vxfenmode file for non-SCSI-3 fencing	131
Enabling or disabling the preferred fencing policy	134

Section 4 Verifying the Veritas Storage Foundation and High Availability installation 137

Chapter 10 Verifying the installation	139
Performing a postcheck on a node	139
Verifying that the products were installed	140
Installation log files	140
Using the installation log file	140
Using the summary file	140
Starting and stopping processes for the Veritas products	141
Checking Veritas Volume Manager processes	141
Checking Veritas File System installation	141
Verifying Veritas File System kernel installation	142
Verifying command installation	142
Verifying the LLT, GAB, and VCS configuration files	142
Verifying LLT, GAB, and cluster operation	143
Verifying LLT	143
Verifying the cluster	146
Verifying the cluster nodes	146

Section 5 Adding and removing nodes 151

Chapter 11 Adding a node to a cluster	153
About adding a node to a cluster	153
Before adding a node to a cluster	154
Adding a node to a cluster	156
Adding a node to a cluster using the SFHA installer	156
Adding the node to a cluster manually	159
Configuring server-based fencing on the new node	166
Adding the new node to the vxfen service group	167

	Starting VCS after adding the new node	168
	Updating the Storage Foundation for Databases (SFDB) repository after adding a node	168
Chapter 12	Removing a node from a cluster	171
	Removing a node from a cluster	171
	Verifying the status of nodes and service groups	172
	Deleting the departing node from SFHA configuration	173
	Modifying configuration files on each remaining node	176
	Removing the node configuration from the CP server	177
	Removing security credentials from the leaving node	178
	Unloading LLT and GAB and removing VCS on the departing node	178
	Updating the Storage Foundation for Databases (SFDB) repository after removing a node	180
Section 6	Uninstalling Veritas Storage Foundation and High Availability	181
Chapter 13	Uninstalling Storage Foundation and High Availability	183
	About removing Storage Foundation and High Availability	183
	Preparing to uninstall	184
	Preparing to remove Veritas Volume Manager	184
	Preparing to remove Veritas File System	190
	Disabling VCS agents for VVR the agents on a system	191
	Removing the Replicated Data Set	192
	Uninstalling SFHA packages using the script-based installer	194
	Removing the CP server configuration using the removal script	195
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product	197
	Manually uninstalling Storage Foundation and High Availability packages on non-global zones	198
Section 7	Installation reference	199
Appendix A	Installation scripts	201
	About installation scripts	201
	Installation script options	202
	About using the postcheck option	206

Appendix B	Response files	209
	About response files	209
	Installing SFHA using response files	210
	Configuring SFHA using response files	210
	Uninstalling SFHA using response files	211
	Syntax in the response file	211
	Response file variables to install, upgrade, or uninstall Storage Foundation and High Availability	212
	Response file variables to configure Storage Foundation and High Availability	214
	Sample response file for SFHA configuration	223
	Sample response file for SFHA install	223
Appendix C	Tunable files for installation	225
	About setting tunable parameters using the installer or a response file	225
	Setting tunables for an installation, configuration, or upgrade	226
	Setting tunables with no other installer-related operations	227
	Setting tunables with an un-integrated response file	228
	Preparing the tunables file	229
	Setting parameters for the tunables file	229
	Tunables value parameter definitions	230
Appendix D	Configuring I/O fencing using a response file	237
	Configuring I/O fencing using response files	237
	Response file variables to configure disk-based I/O fencing	238
	Sample response file for configuring disk-based I/O fencing	239
	Response file variables to configure server-based I/O fencing	240
	Sample response file for configuring server-based I/O fencing	242
	Sample response file for configuring non-SCSI-3 server-based I/O fencing	242
	Response file variables to configure non-SCSI-3 server-based I/O fencing	243
Appendix E	Configuration files	245
	About the LLT and GAB configuration files	245
	About the AMF configuration files	248
	About the VCS configuration files	249
	Sample main.cf file for VCS clusters	251
	Sample main.cf file for global clusters	252

	About I/O fencing configuration files	254
	Sample configuration files for CP server	256
	Sample main.cf file for CP server hosted on a single node that runs VCS	257
	Sample main.cf file for CP server hosted on a two-node SFHA cluster	259
	Sample CP server configuration (/etc/vxcps.conf) file output	262
Appendix F	Configuring the secure shell or the remote shell for communications	263
	About configuring secure shell or remote shell communication modes before installing products	263
	Manually configuring and passwordless ssh	264
	Restarting the ssh session	268
	Enabling and disabling rsh for Solaris	269
Appendix G	Storage Foundation and High Availability components	271
	Storage Foundation and High Availability installation packages	271
	Veritas Cluster Server installation packages	274
	Chinese language packages	274
	Japanese language packages	275
	Veritas Storage Foundation obsolete and reorganized installation packages	275
Appendix H	Troubleshooting installation issues	279
	Restarting the installer after a failed connection	279
	What to do if you see a licensing reminder	279
	Troubleshooting information	280
	Incorrect permissions for root on remote system	280
	Inaccessible system	281
Appendix I	Troubleshooting cluster installation	283
	Unmount failures	283
	Command failures	283
	Installer cannot create UUID for the cluster	284
	The vxfsntsthdw utility fails when SCSI TEST UNIT READY command fails	284
	Troubleshooting CP server	285

	Troubleshooting issues related to the CP server service group	286
	Checking the connectivity of CP server	286
	Troubleshooting server-based fencing on the SFHA cluster nodes	286
	Issues during fencing startup on SF HA cluster nodes set up for server-based fencing	287
	Issues during online migration of coordination points	287
	Vxfen service group activity after issuing the vxfsnwap command	288
Appendix J	Sample SF HA cluster setup diagrams for CP server-based I/O fencing	289
	Configuration diagrams for setting up server-based I/O fencing	289
	Two unique client clusters served by 3 CP servers	289
	Client cluster served by highly available CPS and 2 SCSI-3 disks	290
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	291
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	293
Appendix K	Reconciling major/minor numbers for NFS shared disks	295
	Reconciling major/minor numbers for NFS shared disks	295
	Checking major and minor numbers for disk partitions	296
	Checking the major and minor number for VxVM volumes	299
Appendix L	Configuring LLT over UDP using IPv4	303
	Using the UDP layer for LLT	303
	When to use LLT over UDP	303
	Manually configuring LLT over UDP using IPv4	303
	Broadcast address in the /etc/llttab file	304
	The link command in the /etc/llttab file	304
	The set-addr command in the /etc/llttab file	305
	Selecting UDP ports	306
	Configuring the netmask for LLT	307
	Configuring the broadcast address for LLT	307
	Sample configuration: direct-attached links	308
	Sample configuration: links crossing IP routers	309

Appendix M	Configuring LLT over UDP using IPv6	313
	Using the UDP layer of IPv6 for LLT	313
	When to use LLT over UDP	313
	Manually configuring LLT over UDP using IPv6	313
	Sample configuration: direct-attached links	314
	Sample configuration: links crossing IP routers	316
Appendix N	Compatability issues when installing Storage Foundation High Availability with other products	319
	Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present	319
	Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present	320
	Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present	320
Appendix O	Changing root user into role	321
	Changing root user into root role	321
Index		323

Installation overview and planning

- [Chapter 1. Introducing Storage Foundation and High Availability Solutions](#)
- [Chapter 2. Planning to install Storage Foundation and High Availability](#)
- [Chapter 3. System requirements](#)
- [Chapter 4. Licensing Veritas products](#)

Introducing Storage Foundation and High Availability Solutions

This chapter includes the following topics:

- [About Veritas products](#)
- [About Veritas graphical user interfaces](#)
- [About Storage Foundation and High Availability features](#)

About Veritas products

The following products are available for this release.

About Storage Foundation and High Availability

Veritas Storage Foundation by Symantec includes Veritas File System by Symantec (VxFS) and Veritas Volume Manager by Symantec (VxVM) with various feature levels.

Veritas File System is a high-performance journaling file system that provides easy management and quick-recovery for applications. Veritas File System delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity.

Veritas Volume Manager removes the physical limitations of disk storage. You can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Veritas Volume Manager also provides easy-to-use, online storage management tools to reduce downtime.

You add high availability functionality to Storage Foundation HA by installing Veritas Cluster Server software.

VxFS and VxVM are a part of all Veritas Storage Foundation products. Do not install or update VxFS or VxVM as individual components.

Veritas Storage Foundation has the following products:

- Storage Foundation Standard
- Storage Foundation Standard HA
- Storage Foundation Enterprise
- Storage Foundation Enterprise HA

About Veritas Storage Foundation Basic

Storage Foundation Basic supports all Storage Foundation Standard features, however, there are deployment and technical support limitations.

About Veritas Cluster Server

Veritas Cluster Server by Symantec (VCS) is a clustering solution that provides the following benefits:

- Reduces application downtime
- Facilitates the consolidation and the failover of servers
- Manages a range of applications in heterogeneous environments

About Veritas high availability agents

Veritas agents provide high availability for specific resources and applications. Each agent manages resources of a particular type. For example, the Oracle agent manages Oracle databases. Agents typically start, stop, and monitor resources and report state changes.

About Veritas Volume Replicator

Veritas Volume Replicator by Symantec is an optional, separately-licensable feature that is fully integrated with Veritas Volume Manager. This component replicates data to remote locations over any standard IP network to provide continuous data availability.

Volume Replicator is available with Veritas Storage Foundation Standard and Enterprise products.

About Veritas Storage Foundation Cluster File System

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

Storage Foundation Cluster File System HA adds the failover functionality of Veritas Cluster Server. This functionality can protect everything from a single critical database instance to very large multiple-application clusters in networked environments. Veritas Storage Foundation Cluster File System also provides increased automation and intelligent management of availability and performance.

You can license Veritas Volume Replicator with this product.

About Veritas Storage Foundation for Sybase ASE CE

The solution stack comprises the Veritas Cluster Server (VCS), Veritas Cluster Volume Manager (CVM), Veritas Cluster File System (CFS), and Veritas Storage Foundation, which includes the base Veritas Volume Manager (VxVM) and Veritas File System (VxFS).

About Veritas Storage Foundation for Oracle® RAC

Veritas Storage Foundation for Oracle® RAC by Symantec is an integrated suite of Veritas storage management and high-availability software. The software is engineered to improve performance, availability, and manageability of Real Application Cluster (RAC) environments. Certified by Oracle Corporation, Veritas Storage Foundation for Oracle RAC delivers a flexible solution that makes it easy to deploy and manage RAC.

You can license Veritas Volume Replicator with this product.

About Veritas graphical user interfaces

The following are descriptions of Veritas GUIs.

About Veritas Operations Manager

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

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

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

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

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

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

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

About Storage Foundation and High Availability features

The following section describes different features in the Storage Foundation and High Availability product.

About LLT and GAB

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

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

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

Optimizing LLT media speed settings on private NICs

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

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

Guidelines for setting the media speed of the LLT interconnects

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

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

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

About configuring SFHA clusters for data integrity

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

See [“About planning to configure I/O fencing”](#) on page 65.

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

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

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

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

See [“Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsfa”](#) on page 120.

See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 129.

About I/O fencing components

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

- Data disks—Store shared data
See [“About data disks”](#) on page 25.
- Coordination points—Act as a global lock during membership changes
See [“About coordination points”](#) on page 25.

About data disks

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

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

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. SFHA prevents split-brain when vxfen races for control of the coordination points and the winner partition fences the ejected nodes from accessing the data disks.

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

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

■ Coordinator disks

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

You can configure coordinator disks to use Veritas Volume Manager Dynamic Multi-pathing (DMP) feature. Dynamic Multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. So, you can configure I/O fencing to use either DMP devices or the underlying raw character devices. I/O fencing uses SCSI-3 disk policy that is either raw or dmp based on the disk device that you use. The disk policy is dmp by default.

See the *Veritas Storage Foundation Administrator's Guide*.

■ Coordination point servers

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

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

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

About preferred fencing

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

You can configure the preferred fencing policy using the cluster-level attribute PreferredFencingPolicy as follows:

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

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

See [“Enabling or disabling the preferred fencing policy”](#) on page 134.

About global clusters

Global clusters provide the ability to fail over applications between geographically distributed clusters when disaster occurs. You require a separate license to

configure global clusters. You must add this license during the installation. The installer only asks about configuring global clusters if you have used the global cluster license.

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

Planning to install Storage Foundation and High Availability

This chapter includes the following topics:

- [About installation and configuration methods](#)
- [Downloading the Storage Foundation and High Availability software](#)

About installation and configuration methods

You can install and configure SFHA using Veritas installation programs or using native operating system methods.

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

- The Veritas product installer
The installer displays a menu that simplifies the selection of installation options.
See [“About the Veritas installer”](#) on page 57.
- The product-specific installation scripts
The installation scripts provide a command-line interface to install a specific product. The product-specific scripts enable you to specify some additional command-line options. Installing with the installation script is also the same as specifying SFHA from the installer menu.
- 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 systems.

See “[About response files](#)” on page 209.

Downloading the Storage Foundation and High Availability software

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

To download the trialware version of the software

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

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

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

See “[About installation scripts](#)” on page 201.

To download the software

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

The estimated space for download, gunzip, and tar extract is 2 GB for SPARC and 1.5 GB for Opteron.

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 35.

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

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

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

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

System requirements

This chapter includes the following topics:

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

Release notes

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

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

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

Hardware compatibility list (HCL)

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

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

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

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

Supported operating systems

For information on supported operating systems, see the *Storage Foundation and High Availability Release Notes*.

Veritas File System requirements

Veritas File System requires that the values of the Solaris variables `lwp_default_stksize` and `svc_default_stksize` are at least 0x8000. When you install the Veritas File System package, `VRTSvxfs`, the `VRTSvxfs` packaging scripts check the values of these variables in the kernel. If the values are less than the required values, the installer gives a warning. You need to set the values and reboot the system.

To avoid an unexpected need for a reboot, verify the values of the variables before installing Veritas File System.

Use the following commands to check the values of the variables:

```
# echo "lwp_default_stksize/X" | mdb -k
lwp_default_stksize:
lwp_default_stksize:          8000

# echo "svc_default_stksize/X" | mdb -k
svc_default_stksize:
svc_default_stksize:          8000
```

If the values shown are less than 8000, you can expect a reboot after installation.

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

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

```
set lwp_default_stksize=0x8000
set rpcmod:svc_default_stksize=0x8000
```

Disk space requirements

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

Use the script-based installer `-precheck` option to determine if there is sufficient space.

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

If you have downloaded SFHA, use the following command:

```
# ./installsfha -precheck
```

Discovering product versions and various requirement information

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

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

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

- The installed version of all released Storage Foundation and High Availability Suite of products
- The required packages or patches (if applicable) that are missing

- The available updates (including patches or hotfixes) from Symantec Operations Readiness Tools (SORT) for the installed products

To run the version checker

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

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

Database requirements

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

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

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

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 36.
- CP servers

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

Coordinator disk requirements for I/O fencing

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

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

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

CP server requirements

SFHA 6.0 PR1 clusters (application clusters) support coordination point servers (CP servers) which are hosted on the following VCS and SFHA versions:

- VCS 6.0PR1, VCS 6.0, 5.1SP1, or 5.1 single-node cluster
Single-node VCS clusters with VCS 5.1 SP1 RP1 and later or VCS 6.0 and later that hosts CP server does not require LLT and GAB to be configured.
- SFHA 6.0PR1, SFHA 6.0, 5.1SP1, or 5.1 cluster

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

See the *Veritas Cluster Server Installation Guide*.

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

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

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

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

Table 3-1 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: ■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB) ■ 300 MB in /usr ■ 20 MB in /var ■ 10 MB in /etc (for the CP server database) See “Disk space requirements” on page 35.
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the CP servers.
RAM	Each CP server requires at least 512 MB.
Network	Network hardware capable of providing TCP/IP connection between CP servers and SFHA clusters (application clusters).

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

Table 3-2 CP server supported operating systems and versions

CP server	Operating system and version
CP server hosted on a VCS single-node cluster or on an SFHA cluster	CP server supports any of the following operating systems: ■ AIX 6.1 and 7.1 ■ HP-UX 11i v3 ■ Linux: ■ RHEL 5 ■ RHEL 6 ■ SLES 10 ■ SLES 11 ■ Solaris 10 ■ Oracle Solaris 11 Review other details such as supported operating system levels and architecture for the supported operating systems. See the <i>Veritas Cluster Server Release Notes</i> or the <i>Veritas Storage Foundation High Availability Release Notes</i> for that platform.

Following are the CP server networking requirements and recommendations:

- Symantec recommends that network access from the application clusters to the CP servers should be made highly-available and redundant. The network connections require either a secure LAN or VPN.
- The CP server uses the TCP/IP protocol to connect to and communicate with the application clusters by these network paths. The CP server listens for messages from the application clusters using TCP port 14250. This is the default port that can be changed during a CP server configuration.
Symantec recommends that you configure multiple network paths to access a CP server. If a network path fails, CP server does not require a restart and continues to listen on one of the other available virtual IP addresses.
- The CP server supports either Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with the application clusters. If the CP server is configured to use an IPv6 virtual IP address, then the application clusters should also be on the IPv6 network where the CP server is being hosted.
- When placing the CP servers within a specific network configuration, you must take into consideration the number of hops from the different application cluster nodes to the CP servers. As a best practice, Symantec recommends that the number of hops and network latency from the different application cluster nodes to the CP servers should be equal. This ensures that if an event occurs that results in an I/O fencing scenario, there is no bias in the race due to the number of hops between the nodes.

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

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

For secure communications between the SF HA cluster and CP server, consider the following requirements and suggestions:

- In a secure communication environment, all CP servers that are used by the application cluster must be configured with security enabled. A configuration where the application cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.

- For non-secure communication between CP server and application clusters, there is no need to configure Symantec Product Authentication Service. In non-secure mode, authorization is still provided by CP server for the application cluster users. The authorization that is performed only ensures that authorized users can perform appropriate actions as per their user privileges on the CP server.

For information about establishing secure communications between the application cluster and CP server, see the *Veritas Cluster Server Administrator's Guide*.

Non-SCSI-3 I/O fencing requirements

Supported virtual environment for non-SCSI-3 fencing:

- Oracle Solaris 11
Oracle VM Server for SPARC 2.1
Guest operating system: Solaris 10, Oracle Solaris 11

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

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

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

Licensing Veritas products

This chapter includes the following topics:

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

About Veritas product licensing

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

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

www.symantec.com/techsupp/

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, such as Veritas Operations Manager (VOM). If you do not comply with the above terms, continuing to use the Symantec product is a violation of your end user license agreement, and results in warning messages.

For more information about keyless licensing, see the following URL:

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

If you 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 42.
See the `vxkeyless (1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.
See “Installing Veritas product license keys” on page 44.
See the `vxlicinst (1m)` manual page.

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

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

Setting or changing the product level for keyless licensing

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

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

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

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

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

To set or change the product level

- 1 Show your current working directory:

```
# pwd
```

Output resembles:

```
/opt/VRTSvlic/bin
```

- 2 View the current setting for the product level.

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

- 3 View the possible settings for the product level.

```
# ./vxkeyless displayall
```

- 4 Set the desired product level.

```
# ./vxkeyless set prod_levels
```

where *prod_levels* is a comma-separated list of keywords. The keywords are the product levels as shown by the output of step 3.

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas 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
```

Installing Veritas Storage Foundation and High Availability

- [Chapter 5. Preparing to install Storage Foundation High Availability](#)
- [Chapter 6. Installing Storage Foundation and High Availability using the script-based installer](#)

Preparing to install Storage Foundation High Availability

This chapter includes the following topics:

- [Installation preparation overview](#)
- [About using ssh or rsh with the Veritas installer](#)
- [Setting up shared storage](#)
- [Creating the /opt directory](#)
- [Creating root user](#)
- [Setting environment variables](#)
- [Mounting the product disc](#)
- [Assessing the system for installation readiness](#)

Installation preparation overview

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

Table 5-1 Installation overview

Installation task	Section
Obtain product licenses.	See “About Veritas product licensing” on page 41.

Table 5-1 Installation overview (*continued*)

Installation task	Section
Download the software, or insert the product DVD.	See “Mounting the product disc” on page 54.
Create root user.	See “Creating root user” on page 53.
Set environment variables.	See “Setting environment variables” on page 54.
Create the <code>/opt</code> directory, if it does not exist.	See “Creating the /opt directory” on page 53.
Configure the secure shell (ssh) on all nodes.	
Verify that hardware, software, and operating system requirements are met.	See “Release notes” on page 33.
Check that sufficient disk space is available.	See “Disk space requirements” on page 35.
Use the installer to install the products.	See “About the Veritas installer” on page 57.

About using ssh or rsh with the Veritas installer

The installer uses passwordless secure shell (ssh) or remote shell (rsh) communications among systems. The installer uses the ssh or rsh daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. You then provide the installer with the superuser passwords for the systems where you plan to install. The ssh or rsh communication among the systems is removed when the installation process completes, unless the installation abruptly terminates. If installation terminated abruptly, use the installation script's `-comcleanup` option to remove the ssh or rsh configuration from the systems.

See [“Installation script options”](#) on page 202.

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

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

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 263.

Setting up shared storage

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

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

See [“About planning to configure I/O fencing”](#) on page 65.

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

Setting up shared storage: SCSI disks

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

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

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

To set up shared storage

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

Refer to the documentation that is shipped with the host adapters, the storage, and the systems.
- 2 With both nodes powered off, power on the storage devices.
- 3 Power on one system, but do not allow it to boot. If necessary, halt the system so that you can use the ok prompt.

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

4 Find the paths to the host adapters:

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

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

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

```
{0} ok nvedit
```

As you edit the script, note the following points:

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

In this example, edit the `nvrarc` script as follows:

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

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

```
{0} ok nvstore
```

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

```
{0} ok printenv nvrामrc
```

You can re-edit the file to make corrections:

```
{0} ok nvedit
```

Or, discard the changes if necessary by entering:

```
{0} ok nvquit
```

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

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

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

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

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

Permit the system to continue booting.

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

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

Permit the system to continue booting.

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up shared storage

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

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

- 3 Boot each system with the `reconfigure devices` option:

```
ok boot -r
```

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

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

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

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

Creating the /opt directory

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

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

Creating root user

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

To change root role into a user

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

```
% su - root
```

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

```
# roles admin

root

# usermod -R " " admin
```

- 3 Change the root role into a user.

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

- 4 Verify the change.

```
■ # getent user_attr root

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

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

```
■ # userattr type root
```

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

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

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

See [“Changing root user into root role”](#) on page 321.

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, SFHA commands are in `/opt/VRTS/bin`. SFHA 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
$ MANPATH=/usr/share/man:/opt/VRTS/man:$MANPATH
$ export PATH MANPATH
```

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

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

Mounting the product disc

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

To mount the product disc

- 1 Log in as superuser on a system where you want to install SFHA.
 The system from which you install SFHA need not be part of the cluster. The systems must be in the same subnet.
- 2 Insert the product disc into a DVD drive that is connected to your system.
- 3 If Solaris volume management software is running on your system, the software disc automatically mounts as /cdrom/cdrom0.
- 4 If Solaris volume management software is not available to mount the DVD, you must mount it manually. After you insert the software disc, enter:

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

Where c0t6d0s2 is the default address for the disc drive.

Assessing the system for installation readiness

Symantec provides the following tool for assessing your system, to ensure that the system meets the requirements for installing Storage Foundation 6.0 PR1.

Symantec Operations Readiness Tools

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

See [“Symantec Operations Readiness Tools”](#) on page 55.

Symantec Operations Readiness Tools

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

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

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

- Create automatic email notifications for changes in patches, documentation, and array-specific modules.

To access SORT, go to:

<https://sort.symantec.com>

Installing Storage Foundation and High Availability using the script-based installer

This chapter includes the following topics:

- [About the Veritas installer](#)
- [Installing Storage Foundation and High Availability using the installer](#)
- [Installing language packages](#)
- [Manually installing packages on solaris brand non-global zones](#)
- [Manually installing packages on solaris10 brand zones](#)

About the Veritas installer

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

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

See [“About installation scripts”](#) on page 201.

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

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

See [“Installation script options”](#) on page 202.

Installing Storage Foundation and High Availability using the installer

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

The following sample procedure is based on the installation of Storage Foundation on a single system.

To install Storage Foundation and High Availability

- 1 Set up the systems so that the commands execute on remote machines without prompting for passwords or confirmations with remote shell or secure shell communication utilities.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 263.

- 2 Load and mount the software disc. If you downloaded the software, navigate to the top level of the download directory and skip the next step.

See [“Mounting the product disc”](#) on page 54.

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

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

- 4 From this directory, type the following command to start the installation on the local system. Use this command to install on remote systems if secure shell or remote shell communication modes are configured:

```
# ./installer
```

- 5 Enter **Y** to install and press Return.

- 6 When the list of available products is displayed, select Storage Foundation and High Availability, enter the corresponding number, and press Return.
- 7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).
- 8 Select from one of the following installation options:
 - Minimal packages: installs only the basic functionality for the selected product.
 - Recommended packages: installs the full feature set without optional packages.
 - All packages: installs all available packages.

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

- 9 You are prompted to enter the system names where you want to install the software. Enter the system name or names and then press Enter.

```
Enter the system names separated by spaces:  
[q,?] host1
```

- 10 After the system checks complete, the installer displays a list of the packages to be installed. Press Enter to continue with the installation.
- 11 You need to synchronize the system clocks of your application servers or have them point to an NTP server. After the system check, if the nodes have time difference, the installer prompts:

```
Do you want to synchronize system clock with NTP server(s)?  
[y,n,q] (y)
```

- 12 The installer can configure remote shell or secure shell communications for you among systems, however each system needs to have RSH or SSH servers installed. You also need to provide the superuser passwords for the systems. Note that for security reasons, the installation program neither stores nor caches these passwords.
- 13 The installer may prompt for previous Veritas Volume Manager configurations.

- 14** Choose the licensing method. Answer the licensing questions and follow the prompts.

Note: The keyless license option enables you to install without entering a key. However, you still need a valid license to install and use Veritas products. Keyless licensing requires that you manage the systems with a Management Server.

- 15** The installer prompts you to configure SFHA. You can continue with configuration if you answer **y**.

See "Configuration for Storage Foundation High Availability" for more information.

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

```
1) SF Standard HA
2) SF Enterprise HA
b) Back to previous menu
```

```
Select product mode to license: [1-2,b,q,?] (2) 1
```

- 17** When prompted, decide to enable replication or not:

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

When prompted, decide to enable the Global Cluster option or not:

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

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

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

Check the log file, if needed, to confirm the installation and configuration. Follow the prompts and reboot as necessary.

Installing language packages

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

To install the language packages on the server

- 1 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.
- 2 Install the language packages using the `install_lp` command.

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

Manually installing packages on solaris brand non-global zones

With Oracle Solaris 11, you need to manually install SFHA packages inside non-global zones. The native non-global zones are called solaris brand zones.

- 1 Ensure that the SMF service, `svc:/application/pkg/system-repository:default` is online on the global zone.

```
# svcctl svc:/application/pkg/system-repository
```

- 2 Log on to the non-global zone as a superuser.
- 3 Copy the `VRTSpkgs.p5p` package from the `pkgs` directory from the installation media to the non-global zone.
- 4 Add a file-based repository in the non-global zone.

```
# pkg set-publisher -P -g /<packagelocationpath>/VRTSpkgs.p5p Symantec
```

- 5 Install the required packages.

```
# pkg install VRTSperl VRTSvlic VRTSvcsc VRTSvcscag VRTSvcsea
```

- 6 Remove the publisher on the non-global zone.

```
# pkg unset-publisher Symantec
```

Note: Perform Steps 2 through 6 on each non-global zone.

Manually installing packages on solaris10 brand zones

You need to manually install SFHA 6.0 packages inside the solaris10 brand zones.

1. Boot the zone.
2. Log on to the solaris10 brand zone as a super user.
3. Install the following SFHA packages on the brand zone.
 - VRT\$perl
 - VRT\$vc\$cs
 - VRT\$vc\$csag
 - VRT\$vc\$sea

Note: Perform Steps 1 through 3 on each solaris10 brand zone.

Configuring Veritas Storage Foundation and High Availability

- [Chapter 7. Preparing to configure Storage Foundation and High Availability](#)
- [Chapter 8. Configuring Storage Foundation and High Availability](#)
- [Chapter 9. Configuring Storage Foundation High Availability for data integrity](#)

Preparing to configure Storage Foundation and High Availability

This chapter includes the following topics:

- [About planning to configure I/O fencing](#)
- [Setting up the CP server](#)

About planning to configure I/O fencing

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

You can configure disk-based I/O fencing or server-based I/O fencing. If your enterprise setup has multiple clusters that use VCS for clustering, Symantec recommends you to configure server-based I/O fencing.

The coordination points in server-based fencing can include only CP servers or a mix of CP servers and coordinator disks. Symantec also supports server-based fencing with a single coordination point which is a single highly available CP server that is hosted on an SFHA cluster.

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

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

[Figure 7-1](#) illustrates a high-level flowchart to configure I/O fencing for the SFHA cluster.

Figure 7-1 Workflow to configure I/O fencing





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

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



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

Using the installsfa	See “Setting up disk-based I/O fencing using installsfa” on page 107. See “Setting up server-based I/O fencing using installsfa” on page 120. See “Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsfa” on page 120.
Using response files	See “Response file variables to configure disk-based I/O fencing” on page 238. See “Response file variables to configure server-based I/O fencing” on page 240. See “Response file variables to configure non-SCSI-3 server-based I/O fencing” on page 243. See “Configuring I/O fencing using response files” on page 237.

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

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

Setting up the CP server

[Table 7-1](#) lists the tasks to set up the CP server for server-based I/O fencing.

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

Task	Reference
Plan your CP server setup	See “Planning your CP server setup” on page 70.
Install the CP server	See “Installing the CP server using the installer” on page 71.
Configure the CP server cluster in secure mode	See “Configuring the CP server cluster in secure mode” on page 71.
Set up shared storage for the CP server database	See “Setting up shared storage for the CP server database” on page 72.
Configure the CP server	See “Configuring the CP server using the configuration utility” on page 73. See “Configuring the CP server manually” on page 82.
Verify the CP server configuration	See “Verifying the CP server configuration” on page 84.

Planning your CP server setup

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

To plan your CP server setup

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

Symantec recommends hosting the CP server on an SFHA cluster to make the CP server highly available.
- 2 If you host the CP server on an SFHA cluster, review the following information. Make sure you make the decisions and meet these prerequisites when you set up the CP server:
 - You must configure disk-based fencing during the SFHA configuration.
 - You must set up shared storage for the CP server database during your CP server setup.
 - Decide whether you want to configure server-based fencing for the SF HA cluster (application cluster) with a single CP server as coordination point or with at least three coordination points.
Symantec recommends using at least three coordination points.
- 3 Decide whether you want to configure the CP server cluster in secure mode.

Symantec recommends configuring the CP server cluster in secure mode to secure the communication between the CP server and its clients (SFHA clusters). It also secures the HAD communication on the CP server cluster.
- 4 Set up the hardware and network for your CP server.
- 5 Have the following information handy for CP server configuration:
 - Name for the CP server
The CP server name should not contain any special characters. CP server name can include alphanumeric characters, underscore, and hyphen.
 - Port number for the CP server
Allocate a TCP/IP port for use by the CP server.
Valid port range is between 49152 and 65535. The default port number is 14250.
 - Virtual IP address, network interface, netmask, and networkhosts for the CP server
You can configure multiple virtual IP addresses for the CP server.

Installing the CP server using the installer

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

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

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

CP server setup uses a single system

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

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

See the *Veritas Cluster Server Installation Guide* for instructions on installing and configuring VCS.

Proceed to configure the CP server.

See “[Configuring the CP server using the configuration utility](#)” on page 73.

See “[Configuring the CP server manually](#)” on page 82.

CP server setup uses multiple systems

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

Meet the following requirements for CP server:

- During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.
- During configuration, configure disk-based fencing (scsi3 mode).

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

Configuring the CP server cluster in secure mode

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

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

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

To configure the CP server cluster in secure mode

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

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

```
# installvcs -security
```

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

```
# installsfha -security
```

Setting up shared storage for the CP server database

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

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

To set up shared storage for the CP server database

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

For example:

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

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

For example:

```
# vxdg import cps_dg
```


3 Create a mirrored volume over the disk group.

For example:

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

4 Create a file system over the volume.

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

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

```
AIX # mkfs -V vxfs /dev/vx/rdisk/cps_dg/cps_volume
```

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

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

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

Configuring the CP server using the configuration utility

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

Perform one of the following procedures:

For CP servers on single-node VCS cluster:	See “To configure the CP server on a single-node VCS cluster” on page 73.

For CP servers on an SFHA cluster:	See “To configure the CP server on an SFHA cluster” on page 77.

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

- 1 Verify that the VRTScps package is installed on the node.
- 2 Run the CP server configuration script on the node where you want to configure the CP server:

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

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

The configuration utility then runs the following preconfiguration checks:

- Checks to see if a single-node VCS cluster is running with the supported platform.

The CP server requires VCS to be installed and configured before its configuration.

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

- 4 Enter the name of the CP server.

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

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

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

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

- Enter valid virtual IP addresses:

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

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

You can also use IPv6 address.

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

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

```
Using default port: 14250
```

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

```
Using default port: 14250
```

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

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

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

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

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

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

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

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

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

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

Following is the CP Server configuration information:

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

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

- 10** The configuration utility proceeds with the configuration process, and creates a `vxcps.conf` configuration file.

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

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

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

```
Enter how many NIC resources you want to configure [1 to 2]: 2
```

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

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

```
Enter a valid network interface for virtual IP 10.209.83.85  
on mycps1.symantecexample.com: net0  
Enter a valid network interface for virtual IP 10.209.83.87  
on mycps1.symantecexample.com: net0
```

- 13** Enter the NIC resource you want to associate with the virtual IP addresses.

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

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

```
Symantec recommends configuring NetworkHosts attribute to ensure  
NIC resource to be online always.  
Do you want to add NetworkHosts attribute for the NIC device  
net0 on system mycps1? [y/n] : y  
Enter a valid IP address to configure NetworkHosts for  
NIC net0 on system mycps1 : 10.209.83.86  
Do you want to add another Network Host ?[y/n] : n
```

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

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

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

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

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

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

17 Run the `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 `vxcperv` process and other resources to the VCS configuration in the CP server service group (CPSSG).

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

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

See [“Sample configuration files for CP server”](#) on page 256.

To configure the CP server on an SFHA cluster

- 1** Verify that the `VRTScps` package is installed on each node.
- 2** Make sure that you have configured passwordless `ssh` or `rsh` on the CP server cluster nodes.

- 3 Run the CP server configuration script on any node in the cluster:

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

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

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

The configuration utility then runs the following preconfiguration checks:

- Checks to see if an SFHA cluster is running with the supported platform.
The CP server requires SFHA to be installed and configured before its configuration.
- Checks to see if the CP server is already configured on the system.
If the CP server is already configured, then the configuration utility informs the user and requests that the user unconfigure the CP server before trying to configure it.

- 5 Enter the name of the CP server.

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

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

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

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

- Enter valid virtual IP addresses:

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

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

You can also use IPv6 address.

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

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

Using default port: 14250

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

Using default port: 14250

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

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

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

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

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

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

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

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

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

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

10 Verify and confirm the CP server configuration information.

Following is the CP Server configuration information:

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

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

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

The following output is for one node:

```
Successfully generated the /etc/vxcps.conf  
configuration file.  
Successfully created directory /etc/VRTScps/db.  
Creating mount point /etc/VRTScps/db on  
mycps1.symantecexample.com.  
Copying configuration file /etc/vxcps.conf to  
mycps1.symantecexample.com  
  
Configuring CP Server Service Group (CPSSG) for this cluster  
-----
```

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

Enter how many NIC resources you want to configure [1 to 2]: 2

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

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

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

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

```
Enter a valid interface for virtual IP 10.209.83.85
on all the systems : net0
```

- 15** Enter the NIC resource you want to associate with the virtual IP addresses.

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

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

```
Symantec recommends configuring NetworkHosts attribute to ensure
NIC resource to be online always.
Do you want to add NetworkHosts attribute for the NIC device
net0 on system mycps1? [y/n] : y
Enter a valid IP address to configure NetworkHosts for
NIC net0 on system mycps1 : 10.209.83.86
Do you want to add another Network Host ?[y/n] : n
```

- 17** Enter the netmask for each virtual IP address.

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

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

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

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

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

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

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

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

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

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

For example:

```
# hagrp -state CPSSG
```

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

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

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

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

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

See [“Sample configuration files for CP server”](#) on page 256.

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

To manually configure the CP server

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

```
# hstop -local
```

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

See [“Sample configuration files for CP server”](#) on page 256.

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

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

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

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

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

Based on whether you have configured the CP server cluster in secure mode or not, do the following:

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

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

- 5 Start VCS on all the cluster nodes.

```
# hstart
```

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

```
# hagrps -state CPSSG
```

Output similar to the following appears:

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

Verifying the CP server configuration

Perform the following steps to verify the CP server configuration.

To verify the CP server configuration

- 1 Verify that the following configuration files are updated with the information you provided during the CP server configuration process:
 - `/etc/vxcps.conf` (CP server configuration file)
 - `/etc/VRTSvcs/conf/config/main.cf` (VCS configuration file)
 - `/etc/VRTScps/db` (default location for CP server database)
- 2 Run the `cpsadm` command to check if the `vxcpsserv` process is listening on the configured Virtual IP.

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

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

Configuring Storage Foundation and High Availability

This chapter includes the following topics:

- [Configuring Storage Foundation and High Availability Solutions](#)

Configuring Storage Foundation and High Availability Solutions

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

Use the following procedures to configure Storage Foundation High Availability and clusters using the installer.

Configuring Storage Foundation High Availability using the installer

Storage Foundation HA configuration requires configuring the HA (VCS) cluster. Perform the following tasks to configure the cluster.

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

[Table 8-1](#) lists the tasks that are involved in configuring SFHA using the script-based installer.

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

Task	Reference
Start the software configuration	See “Starting the software configuration” on page 87.
Specify the systems where you want to configure SFHA	See “Specifying systems for configuration” on page 88.
Configure the basic cluster	See “Configuring the cluster name” on page 89. See “Configuring private heartbeat links” on page 89.
Configure virtual IP address of the cluster (optional)	See “Configuring the virtual IP of the cluster” on page 92.
Configure the cluster in secure mode (optional)	See “Configuring the cluster in secure mode” on page 93.
Add VCS users (required if you did not configure the cluster in secure mode)	See “Adding VCS users” on page 96.
Configure SMTP email notification (optional)	See “Configuring SMTP email notification” on page 97.
Configure SNMP email notification (optional)	See “Configuring SNMP trap notification” on page 99.
Configure global clusters (optional) Note: You must have enabled Global Cluster Option when you installed SFHA.	See “Configuring global clusters” on page 101.
Complete the software configuration	See “Completing the SFHA configuration” on page 102.

Required information for configuring Storage Foundation and High Availability Solutions

To configure Storage Foundation High Availability, the following information is required:

See also the *Veritas Cluster Server Installation Guide*.

- A unique Cluster name
- A unique Cluster ID number between 0-65535
- Two or more NIC cards per system used for heartbeat links

One or more heartbeat links are configured as private links and one heartbeat link may be configured as a low priority link.

You can configure Storage Foundation High Availability in secure mode.

Running SFHA in Secure Mode guarantees that all inter-system communication is encrypted and that users are verified with security credentials. When running in Secure Mode, NIS and system usernames and passwords are used to verify identity. SFHA usernames and passwords are no longer used when a cluster is running in Secure Mode.

The following information is required to configure SMTP notification:

- The domain-based hostname of the SMTP server
- The email address of each SMTP recipient
- A minimum severity level of messages to be sent to each recipient

The following information is required to configure SNMP notification:

- System names of SNMP consoles to receive VCS trap messages
- SNMP trap daemon port numbers for each console
- A minimum severity level of messages to be sent to each console

Starting the software configuration

You can configure SFHA using the Veritas product installer or the `installsfha` command.

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

To configure SFHA using the product installer

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

```
# ./installer
```

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

- 3 From the opening Selection Menu, choose: c for "Configure an Installed Product."
- 4 From the displayed list of products to configure, choose the corresponding number for your product:
Storage Foundation and High Availability

To configure SFHA using the `installsfha` program

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

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

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

Specifying systems for configuration

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

To specify system names for configuration

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

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

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

The installer does the following tasks:

- Checks that the local node running the installer can communicate with remote nodes
If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases.
- Makes sure that the systems are running with the supported operating system
- Makes sure the installer started from the global zone
- Checks whether SFHA is installed

- Exits if Storage Foundation 6.0 PR1 is not installed

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

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

See [“About planning to configure I/O fencing”](#) on page 65.

Configuring the cluster name

Enter the cluster information when the installer prompts you.

To configure the cluster

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

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

Configuring private heartbeat links

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

See [“Using the UDP layer for LLT”](#) on page 303.

The following procedure helps you configure LLT over Ethernet.

To configure private heartbeat links

- 1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or UDP.
 - Option 1: LLT over Ethernet (answer installer questions)

Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.

Skip to step 2.
 - Option 2: LLT over UDP (answer installer questions)

Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.

Skip to step 3.

- Option 3: Automatically detect configuration for LLT over Ethernet
 Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.

Skip to step 5.

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

The installer discovers and lists the network interface cards.

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

- For Solaris SPARC:

You must not enter the network interface card that is used for the public network (typically net0.)

```
Enter the NIC for the first private heartbeat link on galaxy:
[b,q,?] net0
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on galaxy:
[b,q,?] net1
Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n)

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

- For Solaris x64:

You must not enter the network interface card that is used for the public network (typically net0.)

```
Enter the NIC for the first private heartbeat link on galaxy:
[b,q,?] e1000g1
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on galaxy:
[b,q,?] e1000g2
Would you like to configure a third private heartbeat link?
[y,n,q,b,?] (n)
```

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

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

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

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

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

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

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

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

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

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

- 6 Enter a unique cluster ID:

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

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

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

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

Configuring the virtual IP of the cluster

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

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

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

To configure the virtual IP of the cluster

- 1 Review the required information to configure the virtual IP of the cluster.
- 2 When the system prompts whether you want to configure the virtual IP, enter `y`.
- 3 Confirm whether you want to use the discovered public NIC on the first system.

Do one of the following:

- If the discovered NIC is the one to use, press `Enter`.
- If you want to use a different NIC, type the name of a NIC to use and press `Enter`.

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

4 Confirm whether you want to use the same public NIC on all nodes.

Do one of the following:

- If all nodes use the same public NIC, enter **y**.
- If unique NICs are used, enter **n** and enter a NIC for each node.

```
Is net0 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

Configuring the cluster in secure mode

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

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

To configure a secure cluster, enter **y**.

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

```
# haclus -value SecureClus

1
```

Configuring a secure cluster node by node

For environments that do not support passwordless ssh or passwordless rsh, you cannot use the `-security` option to enable secure mode for your cluster. Instead, you can use the `-securityonnode` option to configure a secure cluster node by node.

[Table 8-2](#) lists the tasks that you must perform to configure a secure cluster.

Table 8-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See “Configuring the first node” on page 94.
Configure security on the remaining nodes	See “Configuring the remaining nodes” on page 95.

Table 8-2 Configuring a secure cluster node by node *(continued)*

Task	Reference
Complete the manual configuration steps	See “ Completing the secure cluster configuration ” on page 95.

Configuring the first node

Perform the following steps on one node in your cluster.

To configure security on the first node

- 1 Ensure that you are logged in as superuser.
- 2 Enter the following command:

```
# /opt/VRTS/install/installsfha -securityonnode
```

The installer lists information about the cluster, nodes, and service groups. If VCS is not configured or if VCS is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure.

```
VCS is not running on all systems in this cluster. All VCS systems
must be in RUNNING state. Do you want to continue? [y,n,q] (n) y
```

```
1) Perform security configuration on first node and export
security configuration files.
```

```
2) Perform security configuration on remaining nodes with
security configuration files.
```

```
Select the option you would like to perform [1-2,q.?] 1
```

Warning: All configurations about cluster users are deleted when you configure the first node. You can use the `/opt/VRTSvcs/bin/hauser` command to create cluster users manually.

- 3 The installer completes the secure configuration on the node. It specifies the location of the security configuration files and prompts you to copy these files to the other nodes in the cluster. The installer also specifies the location of log files, summary file, and response file.
- 4 Copy the security configuration files from the `/var/VRTSvcs/vcsauth/bkup` directory to temporary directories on the other nodes in the cluster.

Configuring the remaining nodes

On each of the remaining nodes in the cluster, perform the following steps.

To configure security on each remaining node

- 1 Ensure that you are logged in as superuser.
- 2 Enter the following command:

```
# /opt/VRTS/install/installsfha -securityonnode
```

The installer lists information about the cluster, nodes, and service groups. If VCS is not configured or if VCS is not running on all nodes of the cluster, the installer prompts whether you want to continue configuring security. It then prompts you for the node that you want to configure. Enter **2**.

```
VCS is not running on all systems in this cluster. All VCS systems
must be in RUNNING state. Do you want to continue? [y,n,q] (n) y
```

```
1) Perform security configuration on first node and export
security configuration files.
```

```
2) Perform security configuration on remaining nodes with
security configuration files.
```

```
Select the option you would like to perform [1-2,q,?] 2
```

The installer completes the secure configuration on the node. It specifies the location of log files, summary file, and response file.

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

- 1 On the first node, freeze all service groups except the ClusterService service group.

```
# /opt/VRTSvcs/bin/haconf -makerw
```

```
# /opt/VRTSvcs/bin/hagrp -list Frozen=0
```

```
# /opt/VRTSvcs/bin/hagrp -freeze groupname -persistent
```

```
# /opt/VRTSvcs/bin/haconf -dump -makero
```

- 2 On the first node, stop the VCS engine.

```
# /opt/VRTSvcs/bin/CmdServer/hastop -all -force
```

- 3 On all nodes, stop the CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

- 4 On the first node, edit the `/etc/VRTSvcs/conf/config/main.cf` file to resemble the following:

```
cluster clus1 (  
  SecureClus = 1  
)
```

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

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

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

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

- 7 On all nodes, start CmdServer.

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

- 8 On the first node, unfreeze the service groups.

```
# /opt/VRTSvcs/bin/haconf -makerw  
  
# /opt/VRTSvcs/bin/hagrp -list Frozen=1  
  
# /opt/VRTSvcs/bin/hagrp -unfreeze groupname -persistent  
  
# /opt/VRTSvcs/bin/haconf -dump -makero
```

Adding VCS users

If you have enabled a secure VCS cluster, you do not need to add VCS users now. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

To add VCS users

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

```
Do you wish to accept the default cluster credentials of
'admin/password'? [y,n,q] (y) n
Enter the user name: [b,q,?] (admin)
Enter the password:
Enter again:
```

- 3 To add a user, enter **y** at the prompt.

```
Do you want to add another user to the cluster? [y,n,q] (y)
```

- 4 Enter the user's name, password, and level of privileges.

```
Enter the user name: [b,q,?] smith
Enter New Password:*****

Enter Again:*****
Enter the privilege for user smith (A=Administrator, O=Operator,
G=Guest): [b,q,?] a
```

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

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

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

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation. Refer to the *Veritas Cluster Server Administrator's Guide* for more information.

To configure SMTP email notification

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

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

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

See [“Configuring SNMP trap notification”](#) on page 99.

- 3 Provide information to configure SMTP notification.

Provide the following information:

■ Enter the NIC information.

```
Active NIC devices discovered on galaxy: net0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (net0)
Is net0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

■ Enter the SMTP server’s host name.

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

■ Enter the email address of each recipient.

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

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

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

- 4 Add more SMTP recipients, if necessary.

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

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

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

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

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

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

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

5 Verify and confirm the SMTP notification information.

```
NIC: net0
```

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

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

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

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

Configuring SNMP trap notification

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

Note that you can also configure the notification after installation.

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

To configure the SNMP trap notification

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

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

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

See [“Configuring global clusters”](#) on page 101.

3 Provide information to configure SNMP trap notification.

Provide the following information:

■ Enter the NIC information.

```
Active NIC devices discovered on galaxy: net0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (net0)
Is net0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

■ Enter the SNMP trap daemon port.

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

■ Enter the SNMP console system name.

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

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

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

4 Add more SNMP consoles, if necessary.

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

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

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

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

5 Verify and confirm the SNMP notification information.

NIC: net0

SNMP Port: 162

Console: saturn receives SNMP traps for Error or higher events

Console: jupiter receives SNMP traps for SevereError or higher events

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

Configuring global clusters

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

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

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

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

To configure the global cluster option

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

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

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

3 Provide information to configure this cluster as global cluster.

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

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

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

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

For IPv4: Global Cluster Option configuration verification:

```
NIC: net0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

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

On Solaris x64, an example for the NIC's port is net0.

For IPv6 Global Cluster Option configuration verification:

```
NIC: net0
IP: 2001:454e:205a:110:203:baff:ffee:10
Prefix: 64
```

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

On Solaris x64, an example for the NIC's port is net0.

Completing the SFHA configuration

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

To complete the SFHA configuration

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

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

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

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

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

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

summary file	Describes the cluster and its configured resources.
log file	Details the entire configuration.
response file	Contains the configuration information that can be used to perform secure or unattended installations on other systems. See “Configuring SFHA using response files” on page 210.

Verifying and updating licenses on the system

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

See [“Checking licensing information on the system”](#) on page 103.

See [“Updating product licenses using vxlicinst”](#) on page 104.

Checking licensing information on the system

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

To check licensing information

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

```
# vxlicrep
```

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

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

```
License Key           = xxx-xxx-xxx-xxx-xxx
Product Name          = Storage Foundation and High Availability
Serial Number         = xxxxx
License Type          = PERMANENT
OEM ID                = xxxxx

Features :=
Platform              = Solaris
Version               = 6.0 PR1
Tier                  = 0
Reserved              = 0
Mode                  = VCS
```

Updating product licenses using vxlicinst

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

See [“Replacing a SFHA demo license with a permanent license”](#) on page 104.

To update product licenses

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

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

Replacing a SFHA demo license with a permanent license

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

To replace a demo key

- 1 Make sure you have permissions to log in as root on each of the nodes in the cluster.

- 2 Shut down SFHA on all nodes in the cluster:

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

This command does not shut down any running applications.

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

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

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

```
# vxlicrep
```

- 5 Start SFHA on each node:

```
# hstart
```


Configuring Storage Foundation High Availability for data integrity

This chapter includes the following topics:

- [Setting up disk-based I/O fencing using `installsfha`](#)
- [Setting up disk-based I/O fencing manually](#)
- [Setting up server-based I/O fencing using `installsfha`](#)
- [Setting up non-SCSI-3 server-based I/O fencing in virtual environments using `installsfha`](#)
- [Setting up server-based I/O fencing manually](#)
- [Setting up non-SCSI-3 fencing in virtual environments manually](#)
- [Enabling or disabling the preferred fencing policy](#)

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

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

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

- 1 List the new external disks or the LUNs as recognized by the operating system.
On each node, enter:

```
# devfsadm
```

- 2 To initialize the disks as VxVM disks, use one of the following methods:

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

```
# vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i c2t13d0
```

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

Checking shared disks for I/O fencing

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

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

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

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

Checking that disks support SCSI-3 involves the following tasks:

- Verifying the Array Support Library (ASL)
See [“Verifying Array Support Library \(ASL\)”](#) on page 109.

- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 110.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfcntl utility”](#) on page 110.

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
libvxCLARion.so	DGC	All
libvxFJTSYe6k.so	FUJITSU	E6000
libvxFJTSYe8k.so	FUJITSU	All
libvxap.so	SUN	All
libvxatf.so	VERITAS	ATFNODES
libvxcompellent.so	COMPELNT	Compellent Vol
libvxcopan.so	COPANSYS	8814, 8818

- 3 Scan all disk drives and their attributes, update the VxVM device list, and reconfigure DMP with the new devices. Type:

```
# vxdisk scandisks
```

See the Veritas Volume Manager documentation for details on how to add and configure disks.

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the `vxfcntlshdw` 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 SFHA.
- 2 Ensure that both nodes are connected to the same disk during the testing. Use the `vxfenadm` command to verify the disk serial number.

```
# vxfenadm -i diskpath
```

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

For example, an EMC disk is accessible by the `/dev/rdisk/c1t1d0s2` path on node A and the `/dev/rdisk/c2t1d0s2` path on node B.

From node A, enter:

```
# vxfenadm -i /dev/rdisk/c1t1d0s2
```

```
Vendor id : EMC
Product id : SYMMETRIX
Revision : 5567
Serial Number : 42031000a
```

The same serial number information should appear when you enter the equivalent command on node B using the `/dev/rdisk/c2t1d0s2` path.

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

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

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

Testing the disks using vxfcntlshdw utility

This procedure uses the `/dev/rdisk/c1t1d0s2` disk in the steps.

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

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

```
The disk /dev/rdisk/clt1d0s2 is ready to be configured for I/O Fencing on
node galaxy
```

For more information on how to replace coordinator disks, refer to the *Veritas Cluster Server Administrator's Guide*.

To test the disks using `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/cxtxdxsx
for raw: /dev/rdisk/cxtxdxsx
Make sure it's the same disk as seen by nodes
IP_adrs_ofgalaxy and IP_adrs_of_nebula
/dev/rdisk/c2t13d0s2
```

```
Enter the disk name to be checked for SCSI-3 PGR on node
IP_adrs_of_nebula in the format:
for dmp: /dev/vx/rdmp/cxtxdxsx
for raw: /dev/rdisk/cxtxdxsx
Make sure it's the same disk as seen by nodes
IP_adrs_ofgalaxy and IP_adrs_of_nebula
/dev/rdisk/c2t13d0s2
```

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

- 5 Review the output as the utility performs the checks and reports its activities.
- 6 If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node galaxy.

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

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

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

Configuring disk-based I/O fencing using installsfa

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

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

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

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

The installsfa starts with a copyright message and verifies the cluster information.

Note the location of log files which you can access in the event of any problem with the configuration process.

- 2 Confirm that you want to proceed with the I/O fencing configuration at the prompt.

The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 6.0 PR1 is configured properly.

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

```
Select the fencing mechanism to be configured in this  
Application Cluster [1-4,b,q] 2
```

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

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

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

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

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

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

If the available VxVM CDS disks are less than the required, installer asks whether you want to initialize more disks as VxVM disks. Choose the disks you want to initialize as VxVM disks and then use them to create new disk group.

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
 - Enter the disk group name.
- 6 Verify that the coordinator disks you chose meet the I/O fencing requirements. You must verify that the disks are SCSI-3 PR compatible using the `vxfcntlshdw` utility and then return to this configuration program.
- See [“Checking shared disks for I/O fencing”](#) on page 108.
- 7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.
- 8 Enter the I/O fencing disk policy that you chose to use. For example:

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

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
 - Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information
- 9 Verify and confirm the I/O fencing configuration information that the installer summarizes.
- 10 Review the output as the configuration program does the following:
- Stops VCS and I/O fencing on each node.
 - Configures disk-based I/O fencing and starts the I/O fencing process.
 - Updates the VCS configuration file `main.cf` if necessary.
 - Copies the `/etc/vxfenmode` file to a date and time suffixed file `/etc/vxfenmode-date-time`. This backup file is useful if any future fencing configuration fails.
 - Starts VCS on each node to make sure that the SFHA is cleanly configured to use the I/O fencing feature.

- 11 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.
- 12 Configure the Coordination Point agent to monitor the coordinator disks.
See [“Configuring CoordPoint agent to monitor coordination points”](#) on page 126.

Setting up disk-based I/O fencing manually

[Table 9-1](#) lists the tasks that are involved in setting up I/O fencing.

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

Task	Reference
Initializing disks as VxVM disks	See “Initializing disks as VxVM disks” on page 107.
Identifying disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 116.
Checking shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 108.
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 116.
Creating I/O fencing configuration files	See “Creating I/O fencing configuration files” on page 117.
Modifying SFHA configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 118.
Configuring CoordPoint agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 126.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 120.

Removing permissions for communication

Make sure you completed the installation of SFHA and the verification of disk support for I/O fencing. If you used `rsh`, remove the temporary `rsh` access permissions that you set for the nodes and restore the connections to the public network.

If the nodes use `ssh` for secure communications, and you temporarily removed the connections to the public network, restore the connections.

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 107.

Review the following procedure to identify disks to use as coordinator disks.

To identify the coordinator disks

- 1 List the disks on each node.

For example, execute the following commands to list the disks:

```
# vxdisk -o alldgs list
```

- 2 Pick three SCSI-3 PR compliant shared disks as coordinator disks.

See [“Checking shared disks for I/O fencing”](#) on page 108.

Setting up coordinator disk groups

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

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

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

The following example procedure assumes that the disks have the device names `c1t1d0s2`, `c2t1d0s2`, and `c3t1d0s2`.

To create the `vxfencoorddg` disk group

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

```
# vxvg init vxfencoorddg c1t1d0s2 c2t1d0s2 c3t1d0s2
```

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

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

- 1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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

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

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

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

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

Modifying VCS configuration to use I/O fencing

After you add coordination points and configure I/O fencing, add the `UseFence = SCSI3` cluster attribute to the VCS configuration file `/etc/VRTSvcs/conf/config/main.cf`.

If you reset this attribute to `UseFence = None`, VCS does not make use of I/O fencing abilities while failing over service groups. However, I/O fencing needs to be disabled separately.

To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

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

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

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

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

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

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

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

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

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

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

For example, on each remaining node, enter:

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

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

- Start the I/O fencing driver.

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

```
# svcadm enable vxfen
```

- Start VCS.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

- 1 On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is SCSI3 and the SCSI3 disk policy is dmp:

```
I/O Fencing Cluster Information:
```

```
=====
```

```
Fencing Protocol Version: 201
```

```
Fencing Mode: SCSI3
```

```
Fencing SCSI3 Disk Policy: dmp
```

```
Cluster Members:
```

```
* 0 (galaxy)
```

```
1 (nebula)
```

```
RFSM State Information:
```

```
node 0 in state 8 (running)
```

```
node 1 in state 8 (running)
```

- 2 Verify that the disk-based I/O fencing is using the specified disks.

```
# vxfenconfig -l
```

Setting up server-based I/O fencing using installsfha

See [“Setting up server-based I/O fencing using installsfha”](#) on page 120.

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

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

To configure I/O fencing using the installsfa in a non-SCSI-3 PR-compliant setup

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

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

The installsfa 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 SFHA 6.0 PR1 is configured properly.

- 3 Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster
[1-4,b,q] 1
```

- 4 Enter **n** to confirm that your storage environment does not support SCSI-3 PR.

```
Does your storage environment support SCSI3 PR?
[y,n,q] (y) n
```

- 5 Confirm that you want to proceed with the non-SCSI-3 I/O fencing configuration at the prompt.
- 6 Enter the number of CP server coordination points you want to use in your setup.
- 7 Enter the following details for each CP server:
 - Enter the virtual IP address or the fully qualified host name.
 - Enter the port address on which the CP server listens for connections. The default value is 14250. You can enter a different port address. Valid values are between 49152 and 65535.

The installer assumes that these values are identical from the view of the SF HA cluster nodes that host the applications for high availability.

- 8 Verify and confirm the CP server information that you provided.
- 9 Verify and confirm the SF HA cluster configuration information.

Review the output as the installer performs the following tasks:

- Updates the CP server configuration files on each CP server with the following details:
 - Registers each node of the SF HA cluster with the CP server.
 - Adds CP server user to the CP server.
 - Adds SF HA cluster to the CP server user.
 - Updates the following configuration files on each node of the SF HA cluster
 - `/etc/vxfenmode` file
 - `/etc/vxenvIRON` file
 - `/etc/llttab` file
- 10 Review the output as the installer stops SFHA on each node, starts I/O fencing on each node, updates the VCS configuration file `main.cf`, and restarts SFHA with non-SCSI-3 server-based fencing.
- Confirm to configure the CP agent on the SF HA cluster.
- 11 Confirm whether you want to send the installation information to Symantec.
- 12 After the installer configures I/O fencing successfully, note the location of summary, log, and response files that installer creates.
- The files provide useful information which can assist you with the configuration, and can also assist future configurations.

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

Task	Reference
Preparing the CP servers for use by the SFHA cluster	See “ Preparing the CP servers manually for use by the SF HA cluster ” on page 123.
Modifying I/O fencing configuration files to configure server-based I/O fencing	
Modifying SFHA configuration to use I/O fencing	See “ Modifying VCS configuration to use I/O fencing ” on page 118.

Table 9-2 Tasks to set up server-based I/O fencing manually (*continued*)

Task	Reference
Configuring Coordination Point agent to monitor coordination points	See “ Configuring CoordPoint agent to monitor coordination points ” on page 126.
Verifying the server-based I/O fencing configuration	See “ Verifying server-based I/O fencing configuration ” on page 128.

Preparing the CP servers manually for use by the SF HA cluster

Use this procedure to manually prepare the CP server for use by the SF HA 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
Node #1 - SF HA cluster	galaxy
Node #2 - SF HA cluster	nebula
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the SF HA cluster

- 1 Determine the cluster name and uuid on the SF HA cluster.

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

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster clus1

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2-bb31-00306eea460a}
```

- 2 Use the `cpsadm` command to check whether the SF HA cluster and nodes are present in the CP server.

For example:

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

ClusName	UUID	Hostname (Node ID)	Registered
clus1	{f0735332-1dd1-11b2-bb31-00306eea460a}	galaxy(0)	0
clus1	{f0735332-1dd1-11b2-bb31-00306eea460a}	nebula(1)	0

If the output does not show the cluster and nodes, then add them as described in the next step.

For detailed information about the `cpsadm` command, see the *Veritas Cluster Server Administrator's Guide*.

3 Add the SF HA cluster and nodes to each CP server.

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

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

Cluster clus1 added successfully

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

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

Node 0 (galaxy) successfully added

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

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

Node 1 (nebula) successfully added

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

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

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

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

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

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

5 Add the users to the CP server.

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

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

```
User CPSADM@VCS_SERVICES@cluster_uuid  
successfully added
```

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

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

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

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

Configuring CoordPoint agent to monitor coordination points

The following procedure describes how to manually configure the CoordPoint agent to monitor coordination points.

The CoordPoint agent can monitor CP servers and SCSI-3 disks.

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

To configure CoordPoint agent to monitor coordination points

- 1** Ensure that your SF HA cluster has been properly installed and configured with fencing enabled.
- 2** Create a parallel service group vxfen and add a coordpoint resource to the vxfen service group using the following commands:

```
# haconf -makerw
# hagrps -add vxfen
# hagrps -modify vxfen SystemList galaxy 0 nebula 1
# hagrps -modify vxfen AutoFailOver 0
# hagrps -modify vxfen Parallel 1
# hagrps -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 SF HA cluster using the `hares` commands. For example:

```
# hares -state coordpoint
```

The following is an example of the command and output::

```
# hares -state coordpoint

# Resource      Attribute    System      Value
coordpoint      State        galaxy      ONLINE
coordpoint      State        nebula      ONLINE
```

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

The agent log contains detailed CoordPoint agent monitoring information; including information about whether the CoordPoint agent is able to access all the coordination points, information to check on which coordination points the CoordPoint agent is reporting missing keys, etc.

To view all such information in the engine log, change the dbg level for that node using the following commands:

```
# haconf -makerw

# hatype -modify Coordpoint LogDbg 10

# haconf -dump -makero
```

The agent log can now be viewed at the following location:

```
/var/VRTSvcs/log/engine_A.log
```

Verifying server-based I/O fencing configuration

Follow the procedure described below to verify your server-based I/O fencing configuration.

To verify the server-based I/O fencing configuration

- 1 Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

Note: For troubleshooting any server-based I/O fencing configuration issues, refer to the *Veritas Cluster Server Administrator's Guide*.

- 2 Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```

If the output displays `single_cp=1`, it indicates that the application cluster uses a CP server as the single coordination point for server-based fencing.

Setting up non-SCSI-3 fencing in virtual environments manually

To manually set up I/O fencing in a non-SCSI-3 PR compliant setup

- 1 Configure I/O fencing in customized mode with only CP servers as coordination points.

See [“Setting up server-based I/O fencing manually”](#) on page 122.

- 2 Make sure that the SFHA cluster is online and check that the fencing mode is customized.

```
# vxfenadm -d
```

- 3 Make sure that the cluster attribute `UseFence` is set to `SCSI3`.

```
# haclus -value UseFence
```

- 4 On each node, edit the `/etc/vxenviro` file as follows:

```
data_disk_fencing=off
```

- 5 On each node, edit the `/kernel/drv/vxfen.conf` file as follows:

```
vxfen_vxfnd_tmt=25
```

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

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

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

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

- Run the following command:

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

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

```
set-timer sendhbcap:3000
```

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

- Make the VCS configuration file writable:

```
# haconf -makerw
```

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

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

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

- Run the following command to verify the value:

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

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

The command should not list any resources.

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

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

- Make the VCS configuration file read-only

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

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


```
# svcadm disable -t vcs
```
 - After VCS takes all services offline, run the following command to stop VxFEN:


```
# svcadm disable -t vxfen
```
 - On each node, run the following commands to restart VxFEN and VCS:


```
# svcadm enable vxfen
```

Sample /etc/vxfenmode file for non-SCSI-3 fencing

```
=====
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps        - use a coordination point server with optional script
#              controlled scsi3 disks
#
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is required
# only if customized coordinator disks are being used.
```

```
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
# scsi3_disk_policy=dmp

#
# Seconds for which the winning sub cluster waits to allow for the
# losing subcluster to panic & drain I/Os. Useful in the absence of
# SCSI3 based data disk fencing
loser_exit_delay=55

#
# Seconds for which vxfsend process wait for a customized fencing
# script to complete. Only used with vxfsen_mode=customized
vxfsen_script_timeout=25

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

#
# Specify 3 or more odd number of coordination points in this file,
# one in its own line. They can be all-CP servers, all-SCSI-3
# compliant coordinator disks, or a combination of CP servers and
# SCSI-3 compliant coordinator disks. Please ensure that the CP
# server coordination points are numbered sequentially and in the
# same order on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as:
#
# cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames over
# different subnets, all of the IPs/names can be specified in a
# comma separated list as follows:
```

```
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,...,
#   [<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#   is the serial number of the CPS as a coordination point; must
#   start with 1.
# <vip>
#   is the virtual IP address of the CPS, must be specified in
#   square brackets ("[]").
# <vhn>
#   is the virtual hostname of the CPS, must be specified in square
#   brackets ("[]").
# <port>
#   is the port number bound to a particular <vip/vhn> of the CPS.
#   It is optional to specify a <port>. However, if specified, it
#   must follow a colon (":") after <vip/vhn>. If not specified, the
#   colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>, a
# default port can be specified as follows:
#
# port=<default_port>
#
# Where <default_port> is applicable to all the <vip/vhn>s for
# which a <port> is not specified. In other words, specifying <port>
# with a <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[mycps1.company.com]
# cps2=[192.168.0.25]
# cps3=[mycps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [mycps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#   [192.168.0.23]
```

```
# [mycps1.company.com]
# [192.168.0.25]
# - if default port 57777 were not specified, port 14250 would be used
# for all remaining <vip/vhn>s:
# [192.168.0.23]
# [mycps1.company.com]
# [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoordg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
cps1=[mycps1.company.com]
cps2=[mycps2.company.com]
cps3=[mycps3.company.com]
port=14250
=====
```

Enabling or disabling the preferred fencing policy

You can enable or disable the preferred fencing feature for your I/O fencing configuration.

You can enable preferred fencing to use system-based race policy or group-based race policy. If you disable preferred fencing, the I/O fencing configuration uses the default count-based race policy.

See [“About preferred fencing”](#) on page 26.

To enable preferred fencing for the I/O fencing configuration

- 1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

- 2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

- 3 To enable system-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as System.

```
# haclus -modify PreferredFencingPolicy System
```

- Set the value of the system-level attribute FencingWeight for each node in the cluster.

For example, in a two-node cluster, where you want to assign galaxy five times more weight compared to nebula, run the following commands:

```
# hasys -modify galaxy FencingWeight 50
```

```
# hasys -modify nebula FencingWeight 10
```

- Save the VCS configuration.

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

- 4 To enable group-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Group.

```
# haclus -modify PreferredFencingPolicy Group
```

- Set the value of the group-level attribute Priority for each service group. For example, run the following command:

```
# hagrps -modify service_group Priority 1
```

Make sure that you assign a parent service group an equal or lower priority than its child service group. In case the parent and the child service groups are hosted in different subclusters, then the subcluster that hosts the child service group gets higher preference.

- Save the VCS configuration.

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

- 5 To view the fencing node weights that are currently set in the fencing driver, run the following command:

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

- 1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

- 2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

- 3 To disable preferred fencing and use the default race policy, set the value of the cluster-level attribute PreferredFencingPolicy as Disabled.

```
# haconf -makerw
```

```
# haclus -modify PreferredFencingPolicy Disabled
```

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


Verifying the Veritas Storage Foundation and High Availability installation

- [Chapter 10. Verifying the installation](#)

Verifying the installation

This chapter includes the following topics:

- [Performing a postcheck on a node](#)
- [Verifying that the products were installed](#)
- [Installation log files](#)
- [Starting and stopping processes for the Veritas products](#)
- [Checking Veritas Volume Manager processes](#)
- [Checking Veritas File System installation](#)
- [Verifying the LLT, GAB, and VCS configuration files](#)
- [Verifying LLT, GAB, and cluster operation](#)

Performing a postcheck on a node

The installer's `postcheck` command can help you to determine installation-related problems and provide troubleshooting information.

See [“About using the postcheck option”](#) on page 206.

To run the `postcheck` command on a node

- 1 Run the installer with the `-postcheck` option.

```
# ./installer -postcheck system_name
```

- 2 Review the output for installation-related information.

Verifying that the products were installed

Verify that the SFHA products are installed.

Use the `pkg info` command to check which packages have been installed.

```
# pkg info -l VRTSvlic package_name package_name ...
```

You can verify the version of the installed product. Use the following command:

```
# /opt/VRTS/install/installsfha -version
```

Use the following sections to further verify the product installation.

Installation log files

After every product installation, the installer creates three text files:

- Installation log file
- Response file
- Summary file

The name and location of each file is displayed at the end of a product installation, and are always located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

Using the installation log file

The installation log file contains all commands executed during the procedure, their output, and errors generated by the commands. This file is for debugging installation problems and can be used for analysis by Veritas Support.

Using the summary file

The summary file contains the results of the installation by the installer or product installation scripts. The summary includes the list of the packages, and the status (success or failure) of each package. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

To stop the processes

- ◆ Use the `-stop` option to stop the product installation script.

For example, to stop the product's processes, enter the following command:

```
# ./installer -stop
```

To start the processes

- ◆ Use the `-start` option to start the product installation script.

For example, to start the product's processes, enter the following command:

```
# ./installer -start
```

Checking Veritas Volume Manager processes

Use the following procedure to verify that Volume Manager processes are running.

To confirm that key Volume Manager processes are running

- ◆ Type the following command:

```
# ps -ef | grep vx
```

Entries for the `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxcached`, and `vxconfigbackupd` processes should appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

Checking Veritas File System installation

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

Verifying Veritas File System kernel installation

To ensure that the file system driver is loaded, enter:

```
# modinfo | grep vxfs
```

The `modinfo` command displays information about all modules loaded on the system. If the `vxfs` module is loaded, you will see an entry corresponding to `vxfs`. If not, follow the instructions load and then unload the file system module to complete the process.

Verifying command installation

[Table 10-1](#) lists the directories with Veritas File System commands.

Table 10-1 VxFS command locations

Location	Contents
/etc/fs/vxfs	Contains the Veritas <code>mount</code> command and QuickLog commands required to mount file systems.
/usr/lib/fs/vxfs/bin	Contains the VxFS type-specific switch-out commands.
/opt/VRTSvxfs/sbin	Contains the Veritas-specific commands.
/opt/VRTS/bin	Contains symbolic links to all Veritas-specific commands installed in the directories listed above.

Determine whether these subdirectories are present:

```
# ls /etc/fs/vxfs
# ls /usr/lib/fs/vxfs/bin
# ls /opt/VRTSvxfs/sbin
# ls /opt/VRTS/bin
```

Make sure you have adjusted the environment variables accordingly.

See [“Setting environment variables”](#) on page 54.

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 VCS configuration files”](#) on page 249.

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 143.
- 4 Verify GAB operation.
- 5 Verify the cluster operation.
 See [“Verifying the cluster”](#) on page 146.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

- 1 Log in as superuser on the node galaxy.
- 2 Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node           State           Links
*0 galaxy      OPEN              2
 1 nebula      OPEN              2
```

Each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
Node           State           Links
* 0 galaxy      OPEN              2
 1 nebula      OPEN              2
 2 saturn      OPEN              1
```

- 3
- Log in as superuser on the node nebula.
- 4
- Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
Node           State           Links
 0 galaxy      OPEN              2
*1 nebula      OPEN              2
```

- 5
- To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv active
```

The output on galaxy resembles the following:

■ For Solaris SPARC:

Node	State	Link	Status	Address
*0 galaxy	OPEN			
		net1	UP	08:00:20:93:0E:34


```
1 nebula      OPEN      net2 UP      08:00:20:93:0E:38
                                net1 UP      08:00:20:8F:D1:F2
                                net2 DOWN
```

■ For Solaris x64:

Node	State	Link	Status	Address
*0 galaxy	OPEN			
		e1000g:1	UP	08:00:20:93:0E:34
		e1000g:2	UP	08:00:20:93:0E:38
1 nebula	OPEN			
		e1000g:1	UP	08:00:20:8F:D1:F2
		e1000g:2	DOWN	

The command reports the status on the two active nodes in the cluster, galaxy and nebula.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- An address for each link

However, the output in the example shows different details for the node nebula. The private network connection is possibly broken or the information in the /etc/llttab file may be incorrect.

- 6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node galaxy in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
Port  Usage      Cookie
0     gab       0x0
      opens:   0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects: 0 1
7     gab       0x7
      opens:   0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects: 0 1
31    gab       0x1F
```

```
opens:      0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
connects:   0 1
```

Verifying 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.

Note: The example in the following procedure is for SPARC. x64 clusters have different command output.

To verify the cluster nodes

- ◆ On one of the nodes, type the `hasys -display` command:

```
# hasys -display
```

The example shows the output when the command is run on the node galaxy. The list continues with similar information for nebula (not shown) and any other nodes in the cluster.

#System	Attribute	Value
galaxy	AgentsStopped	0
galaxy	AvailableCapacity	100
galaxy	CPUBinding	BindTo None CPUNumber 0
galaxy	CPUThresholdLevel	Critical 90 Warning 80 Note 70 Info 60
galaxy	CPUUsage	0
galaxy	CPUUsageMonitoring	Enabled 0 ActionThreshold 0 ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
galaxy	Capacity	100
galaxy	ConfigBlockCount	130
galaxy	ConfigChecksum	46688
galaxy	ConfigDiskState	CURRENT
galaxy	ConfigFile	/etc/VRTSvcscs/conf/config
galaxy	ConfigInfoCnt	0
galaxy	ConfigModDate	Thu Sep 22 07:14:23 CDT 2011
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	

galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	6.0.00.0
galaxy	FencingWeight	0
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	HostUtilization	CPU 7 Swap 0
galaxy	LLTNodeId	0
galaxy	LicenseType	PERMANENT_SITE
galaxy	Limits	
galaxy	LinkHbStatus	net1 UP net2 UP
galaxy	LoadTimeCounter	0
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	PhysicalServer	
galaxy	ShutdownTimeout	600
galaxy	SourceFile	./main.cf
galaxy	SwapThresholdLevel	Critical 90 Warning 80 Note 70 Info 60
galaxy	SysInfo	Solaris:galaxy,Generic_ 118558-11,5.9,sun4u
galaxy	SysName	galaxy
galaxy	SysState	RUNNING
galaxy	SystemLocation	

galaxy	SystemOwner	
galaxy	SystemRecipients	
galaxy	TFrozen	0
galaxy	TRSE	0
galaxy	UpDownState	Up
galaxy	UserInt	0
galaxy	UserStr	
galaxy	VCSFeatures	NONE
galaxy	VCSMode	

Adding and removing nodes

- [Chapter 11. Adding a node to a cluster](#)
- [Chapter 12. Removing a node from a cluster](#)

Adding a node to a cluster

This chapter includes the following topics:

- [About adding a node to a cluster](#)
- [Before adding a node to a cluster](#)
- [Adding a node to a cluster](#)
- [Configuring server-based fencing on the new node](#)
- [Starting VCS after adding the new node](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node](#)

About adding a node to a cluster

After you install SFHA and create a cluster, you can add and remove nodes from the cluster. You can create clusters of up to 64 nodes.

You can add a node:

- Using the product installer
- Manually

The following procedure provides a summary of the tasks required to add a node to an existing SFHA cluster.

To add a node to a cluster

- 1 Complete the prerequisites and preparatory tasks before adding a node to the cluster.
See [“Before adding a node to a cluster”](#) on page 154.
- 2 Add a new node to the cluster.
See [“Adding a node to a cluster using the SFHA installer”](#) on page 156.
See [“Adding the node to a cluster manually”](#) on page 159.
- 3 If you are using the Storage Foundation for Databases (SFDB) tools, you must update the repository database.
See [“Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node”](#) on page 168.

The example procedures describe how to add a node to an existing cluster with two nodes.

Before adding a node to a cluster

Before preparing to add the node to an existing SFHA cluster, perform the required preparations.


- Verify hardware and software requirements are met.
- Set up the hardware.
- Prepare the new node.

To verify hardware and software requirements are met

- 1 Review hardware and software requirements for SFHA.
See [“Hardware compatibility list \(HCL\)”](#) on page 34.
- 2 Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
- 3 Verify the existing cluster is a SFHA cluster and that SFHA is running on the cluster.

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in [Figure 11-1](#).

Figure 11-1 Adding a node to a two-node cluster using two switches



To set up the hardware

1 Connect the SFHA private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a two-node cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

[Figure 11-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.

2 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.

Prepare the new node before you add it to an existing SFHA cluster.

To prepare the new node

- 1 Verify that the new node meets installation requirements.

```
# ./installsfha -precheck
```
- 2 Install SFHA on the new system using the `-install` option to install SFHA. Do not configure SFHA when prompted.
- 3 You can restart the new node after installation is complete. Configure the new node using the configuration from the existing cluster nodes.

Adding a node to a cluster

You can use one of the following methods to add a node to an existing SFHA cluster:

SFHA installer	See “Adding a node to a cluster using the SFHA installer” on page 156.
Manual	See “Adding the node to a cluster manually” on page 159.

Note: Before you add the node, make sure that SFHA is not configured on the node.

Adding a node to a cluster using the SFHA installer

You can add a node using the `-addnode` option with the SFHA installer.

The SFHA installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:
`/etc/llttab`

```
/etc/VRTSvcs/conf/sysname
```

- Copies the following files on the new node:

```
/etc/llthosts
```

```
/etc/gabtab
```

```
/etc/VRTSvcs/conf/config/main.cf
```

- Copies the following files from the existing cluster to the new node

```
/etc/vxfenmode
```

```
/etc/vxfendg
```

```
/etc/vx/.uuids/clusuuid
```

```
/etc/default/llt
```

```
/etc/default/gab
```

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

- Generate security credentials on the new node if the CPS server of existing cluster is secure
- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.

At the end of the process, the new node joins the SFHA cluster.

Note: If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the SFHA installer with the -addnode option.

```
# cd /opt/VRTS/install
```

```
# ./installsfha -addnode
```

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

- 3 Enter the name of a node in the existing SFHA cluster.

The installer uses the node information to identify the existing cluster.

```
Enter one node of the SFHA cluster to which
you would like to add one or more new nodes: 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,?] **net1**

Enter the NIC for the second private heartbeat
link on saturn: [b,q,?] **net2**

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

- 7 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 8 Review and confirm the information.
- 9 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

Enter the NIC for VCS to use on saturn: **net3**

SFHA is configured on the cluster. Do you want to
configure it on the new node(s)? [y,n,q] (y) n

- 10 If the existing cluster uses server-based fencing in secure mode, the installer will configure server-based fencing in secure mode on the new nodes.

The installer then starts all the required Veritas processes and joins the new node to cluster.
- 11 Confirm that the new node has joined the SFHA cluster using `lltstat -n` and `gabconfig -a` commands.

Adding the node to a cluster manually

Perform this procedure after you install SFHA only if you plan to add the node to the cluster manually.

To add the node manually to the cluster

- 1 Start the Volume Manager.

See [“Starting Volume Manager on the new node”](#) on page 159.
- 2 Configure LLT and GAB.

See [“Configuring LLT and GAB on the new node”](#) on page 160.
- 3 If the CPS server of existing cluster is secure, generate security credentials on the new node.

See [“Setting up the node to run in secure mode”](#) on page 162.
- 4 Configure fencing for the new node to match the fencing configuration on the existing cluster.

If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node.

See [“Starting fencing on the new node”](#) on page 165.
- 5 Start VCS.

See [“Starting VCS after adding the new node”](#) on page 168.
- 6 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 165.

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 `installsfha` 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

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 net1 /dev/net:1 - ether - -
link net2 /dev/net:2 - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 4 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

```
/sbin/gabconfig -c -nN
```

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

- 5 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 6 Use vi or another text editor to create the file `/etc/VRTSvcS/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 7 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

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

For Solaris 10:

```
# svcadm enable llt  
  
# svcadm enable gab  
  
# svcadm restart vxodm
```

- 9 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a  
GAB Port Memberships  
=====
```

Port a	gen	df204	membership	012
Port d	gen	df207	membership	012

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

[Table 11-1](#) uses the following information for the following command examples.

Table 11-1 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.

Configuring the authentication broker on node saturn

To configure the authentication broker on node saturn

- 1 Extract the embedded authentication files and copy them to temporary directory:

```
# mkdir -p /var/VRTSvcs/vcsauth/bkup  
  
# cd /tmp; gunzip -c /opt/VRTSvcs/bin/VxAT.tar.gz | tar xvf -
```

- 2 Edit the setup file manually:

```
# /cat /etc/vx/.uuids/clusuuid 2>&1
```

The output is a string denoting the UUID. This UUID (without { and }) is used as the ClusterName for the setup file.

```
{UUID}
```

```
# cat /tmp/eat_setup 2>&1
```

The file content must resemble the following example:

```
AcceptorMode=IP_ONLY  
  
BrokerExeName=vcsauthserver  
  
ClusterName=UUID  
  
DataDir=/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER  
  
DestDir=/opt/VRTSvcs/bin/vcsauth/vcsauthserver  
  
FipsMode=0  
  
IPPort=14149  
  
RootBrokerName=vcsroot_uuid  
  
SetToRBPlusABorNot=0  
  
SetupPDRs=1  
  
SourceDir=/tmp/VxAT/version
```

- 3 Set up the embedded authentication file:

```
# cd /tmp/VxAT/version/bin/edition_number; \
./broker_setup.sh/tmp/eat_setup

/opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssregctl -s -f
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER/root/.VRTSat/profile \
/VRTSatlocal.conf -b 'Security\Authentication \
\Authentication Broker' -k UpdatedDebugLogFileName \
-v /var/VRTSvcs/log/vcsauthserver.log -t string
```

- 4 Copy the broker credentials from one node in the cluster to saturn by copying the entire bkup directory.

The bkup directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/

# ls

CMDSERVER  CPSADM  CPSEVER  HAD  VCS_SERVICES  WAC
```

- 5 Import the VCS_SERVICES domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER -f /var/VRTSvcs/vcsauth/bkup \
/VCS_SERVICES -p password
```

- 6 Import the credentials for HAD, CMDSERVER, CPSADM, CPSEVER, and WAC.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCS_SERVICES -f /var/VRTSvcs/vcsauth/bkup \
/HAD -p password
```

- 7 Start the vcsauthserver process on saturn.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh
```

8 Perform the following tasks:

```
# mkdir /var/VRTSvcs/vcsauth/data/CLIENT

# mkdir /var/VRTSvcs/vcsauth/data/TRUST

# export EAT_DATA_DIR='/var/VRTSvcs/vcsauth/data/TRUST'

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat setuptrust -b \
localhost:14149 -s high
```

9 Create the /etc/VRTSvcs/conf/config/.secure file:

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

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/default/vxfen
/etc/vxfendg
/etc/vxfenmode
```

If you are using pure CP server-based fencing on the existing cluster, then only the /etc/vxfenmode file needs to be copied on the new node.

- 2 Start fencing on the new node:

```
# svcadm enable vxfen
```

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 net0 -sys saturn
```

```
# hares -modify gconic Device net0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

Configuring server-based fencing on the new node

This section describes the procedures to configure server-based fencing on a new node. Depending on whether server-based fencing is configured in secure or non-secure mode on the existing cluster, perform the tasks in one of the following procedures:

- Server-based fencing in non-secure mode:
[To configure server-based fencing in non-secure mode on the new node](#)
- Server-based fencing in secure mode:
[To configure server-based fencing with security on the new node](#)

To configure server-based fencing in non-secure mode on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

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

The new node must be listed in the command output.

- 4 Add the VCS user cpsclient@saturn to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e cpsclient@saturn \  
-f cps_operator -g vx
```

User cpsclient@saturn successfully added

To configure server-based fencing with security on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

Node 2 (saturn) successfully added

- 3 Verify that the new node is added to the CP server configuration:

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

The new node must be listed in the output.

Adding the new node to the vxfen service group

Perform the steps in the following procedure to add the new node to the vxfen service group.

To add the new node to the vxfen group using the CLI

- 1 On one of the nodes in the existing SF HA 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 SF HA cluster:

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

Starting VCS after adding the new node

Start VCS on the new node.

To start VCS on the new node

- ◆ Start VCS on the new node:

```
# hstart
```

Updating the Storage Foundation for Databases (SFDB) repository after adding a node

If you are using Database Storage Checkpoints, Database FlashSnap, or SmartTier for Oracle in your configuration, update the SFDB repository to enable access for the new node after it is added to the cluster.

To update the SFDB repository after adding a node

- 1 Copy the `/var/vx/vxdba/rep_loc` file from one of the nodes in the cluster to the new node.
- 2 If the `/var/vx/vxdba/auth/user-authorizations` file exists on the existing cluster nodes, copy it to the new node.

If the `/var/vx/vxdba/auth/user-authorizations` file does not exist on any of the existing cluster nodes, no action is required.

This completes the addition of the new node to the SFDB repository.

For information on using SFDB tools features, refer to *Veritas Storage Foundation: Storage and Availability Management for Oracle Databases*.

Removing a node from a cluster

This chapter includes the following topics:

- [Removing a node from a cluster](#)

Removing a node from a cluster

[Table 12-1](#) specifies the tasks that are involved in removing a node from a cluster. In the example procedure, the cluster consists of nodes galaxy, nebula, and saturn; node saturn is to leave the cluster.

Table 12-1 Tasks that are involved in removing a node

Task	Reference
■ Back up the configuration file. ■ Check the status of the nodes and the service groups.	See “Verifying the status of nodes and service groups” on page 172.
■ Switch or remove any SFHA service groups on the node departing the cluster. ■ Delete the node from SFHA configuration.	See “Deleting the departing node from SFHA configuration” on page 173.
Modify the llthosts(4) and gabtab(4) files to reflect the change.	See “Modifying configuration files on each remaining node” on page 176.
If the existing cluster is configured to use server-based I/O fencing, remove the node configuration from the CP server.	See “Removing the node configuration from the CP server” on page 177.

Table 12-1 Tasks that are involved in removing a node *(continued)*

Task	Reference
For a cluster that is running in a secure mode, remove the security credentials from the leaving node.	See “Removing security credentials from the leaving node ” on page 178.
On the node departing the cluster: ■ Modify startup scripts for LLT, GAB, and SFHA to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. ■ Remove the SFHA packages.	See “Unloading LLT and GAB and removing VCS on the departing node” on page 178.

Verifying the status of nodes and service groups

Start by issuing the following commands from one of the nodes to remain in the cluster node galaxy or node nebula in our example.

To verify the status of the nodes and the service groups

- 1 Make a backup copy of the current configuration file, main.cf.

```
# cp -p /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.goodcopy
```

- 2 Check the status of the systems and the service groups.

```
# hastatus -summary

-- SYSTEM STATE
-- System      State      Frozen
A galaxy      RUNNING    0
A nebula      RUNNING    0
A saturn      RUNNING    0

-- GROUP STATE
-- Group      System      Probed    AutoDisabled    State
B grp1       galaxy      Y         N                ONLINE
B grp1       nebula      Y         N                OFFLINE
B grp2       galaxy      Y         N                ONLINE
B grp3       nebula      Y         N                OFFLINE
B grp3       saturn      Y         N                ONLINE
B grp4       saturn      Y         N                ONLINE
```

The example output from the `hastatus` command shows that nodes galaxy, nebula, and saturn are the nodes in the cluster. Also, service group grp3 is configured to run on node nebula and node saturn, the departing node. Service group grp4 runs only on node saturn. Service groups grp1 and grp2 do not run on node saturn.

Deleting the departing node from SFHA configuration

Before you remove a node from the cluster you need to identify the service groups that run on the node.

You then need to perform the following actions:

- Remove the service groups that other service groups depend on, or
- Switch the service groups to another node that other service groups depend on.

To remove or switch service groups from the departing node

- 1** Switch failover service groups from the departing node. You can switch grp3 from node saturn to node nebula.

```
# hagrps -switch grp3 -to nebula
```

- 2** Check for any dependencies involving any service groups that run on the departing node; for example, grp4 runs only on the departing node.

```
# hagrps -dep
```

- 3** If the service group on the departing node requires other service groups—if it is a parent to service groups on other nodes—unlink the service groups.

```
# haconf -makerw
```

```
# hagrps -unlink grp4 grp1
```

These commands enable you to edit the configuration and to remove the requirement grp4 has for grp1.

- 4** Stop SFHA on the departing node:

```
# hastop -sys saturn
```

To stop VCS using SMF, run the following command:

```
# svcadm disable vcs
```

- 5 Check the status again. The state of the departing node should be EXITED. Make sure that any service group that you want to fail over is online on other nodes.

```
# hastatus -summary

-- SYSTEM STATE
-- System      State          Frozen
A galaxy      RUNNING      0
A nebula      RUNNING      0
A saturn      EXITED       0

-- GROUP STATE
-- Group      System      Probed   AutoDisabled   State
B grp1       galaxy      Y        N              ONLINE
B grp1       nebula      Y        N              OFFLINE
B grp2       galaxy      Y        N              ONLINE
B grp3       nebula      Y        N              ONLINE
B grp3       saturn      Y        Y              OFFLINE
B grp4       saturn      Y        N              OFFLINE
```

- 6 Delete the departing node from the SystemList of service groups grp3 and grp4.

```
# hagrps -modify grp3 SystemList -delete saturn
# hagrps -modify grp4 SystemList -delete saturn
```

- 7 For the service groups that run only on the departing node, delete the resources from the group before you delete the group.

```
# hagrps -resources grp4
    processx_grp4
    processy_grp4
# hares -delete processx_grp4
# hares -delete processy_grp4
```

- 8 Delete the service group that is configured to run on the departing node.

```
# hagrps -delete grp4
```

9 Check the status.

```
# hastatus -summary
-- SYSTEM STATE
-- System      State      Frozen
A galaxy       RUNNING   0
A nebula       RUNNING   0
A saturn       EXITED    0

-- GROUP STATE
-- Group      System      Probed   AutoDisabled   State
B grp1       galaxy      Y        N               ONLINE
B grp1       nebula      Y        N               OFFLINE
B grp2       galaxy      Y        N               ONLINE
B grp3       nebula      Y        N               ONLINE
```

10 Delete the node from the cluster.

```
# hasys -delete saturn
```

11 Save the configuration, making it read only.

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

Modifying configuration files on each remaining node

Perform the following tasks on each of the remaining nodes of the cluster.

To modify the configuration files on a remaining node

- 1 If necessary, modify the `/etc/gabtab` file.

No change is required to this file if the `/sbin/gabconfig` command has only the argument `-c`. Symantec recommends using the `-nN` option, where *N* is the number of cluster systems.

If the command has the form `/sbin/gabconfig -c -nN`, where *N* is the number of cluster systems, make sure that *N* is not greater than the actual number of nodes in the cluster. When *N* is greater than the number of nodes, GAB does not automatically seed.

Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`.

- 2 Modify `/etc/llthosts` file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 galaxy
1 nebula
2 saturn
```

To:

```
0 galaxy
1 nebula
```

Removing the node configuration from the CP server

After removing a node from a SFHA cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

Note: The `cpsadm` command is used to perform the steps in this procedure. For detailed information about the `cpsadm` command, see the *Veritas Cluster Server Administrator's Guide*.

To remove the node configuration from the CP server

- 1 Log into the CP server as the root user.
- 2 View the list of VCS users on the CP server, using the following command:

```
# cpsadm -s cp_server -a list_users
```

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

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
# cpsadm -s cp_server -a rm_node -h saturn -c clus1 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
# cpsadm -s cp_server -a list_nodes
```

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node saturn. Perform the following steps.

To remove the security credentials

- 1 Stop the AT process.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh \  
stop
```

- 2 Remove the credentials.

```
# rm -rf /var/VRTSvcs/vcsauth/data/
```

Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node that is departing the cluster.

If you have configured Storage Foundation and High Availability as part of the Storage Foundation and High Availability products, you may have to delete other dependent packages before you can delete all of the following ones.

To unconfigure and unload LLT and GAB and remove SFHA

- 1 If you had configured I/O fencing in enabled mode, then stop I/O fencing.

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

- 2 Unconfigure GAB and LLT:

```
# /sbin/gabconfig -U
# /sbin/lltconfig -U
```

- 3 Unload the GAB and LLT modules from the kernel.

- Determine the kernel module IDs:

```
# modinfo | grep gab
# modinfo | grep llt
```

The module IDs are in the left-hand column of the output.

- Unload the module from the kernel:

```
# modunload -i gab_id
# modunload -i llt_id
```

- 4 Disable the startup files to prevent LLT, GAB, or SFHA from starting up:

```
# /usr/sbin/svcadm disable -t llt
# /usr/sbin/svcadm disable -t gab
# /usr/sbin/svcadm disable -t vcs
```

- 5 To determine the packages to remove, enter:

```
# pkginfo | grep VRTS
```

- 6 To permanently remove the SFHA packages from the system, use the `pkgrm` command. Start by removing the following packages, which may have been optionally installed, in the order shown:

```
# pkgrm VRTSvcsea
# pkgrm VRTSat
# pkgrm VRTSvcSag
# pkgrm VRTScps
# pkgrm VRTSvcS
# pkgrm VRTSamf
# pkgrm VRTSvxfen
# pkgrm VRTSgab
# pkgrm VRTSllt
# pkgrm VRTSspt
# rpm -e VRTSsfcp160
# pkgrm VRTSperl
# pkgrm VRTSvlic
```

- 7 Remove the LLT and GAB configuration files.

```
# rm /etc/llttab
# rm /etc/gabtab
# rm /etc/llthosts
```

- 8 Remove the language packages and patches.

Updating the Storage Foundation for Databases (SFDB) repository after removing a node

After removing a node from a cluster, you do not need to perform any steps to update the SFDB repository.

For information on updating the SFDB repository after adding a node to the cluster:

See [“Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node”](#) on page 168.

For information on removing the SFDB repository after removing the product:

See [“Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product”](#) on page 197.

Uninstalling Veritas Storage Foundation and High Availability

- [Chapter 13. Uninstalling Storage Foundation and High Availability](#)

Uninstalling Storage Foundation and High Availability

This chapter includes the following topics:

- [About removing Storage Foundation and High Availability](#)
- [Preparing to uninstall](#)
- [Disabling VCS agents for VVR the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling SFHA packages using the script-based installer](#)
- [Removing the CP server configuration using the removal script](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product](#)
- [Manually uninstalling Storage Foundation and High Availability packages on non-global zones](#)

About removing Storage Foundation and High Availability

This section covers uninstallation requirements and steps to uninstall the Veritas software.

Only users with superuser privileges can uninstall Storage Foundation and High Availability.

Warning: Failure to follow the instructions in the following sections may result in unexpected behavior.

Preparing to uninstall

Review the following removing the Veritas software.

Preparing to remove Veritas Volume Manager

This section describes the steps you need to take before removing Veritas Volume Manager (VxVM) to preserve the contents of the volumes.

Warning: Failure to follow the preparations in this section might result in unexpected behavior.

Moving volumes to disk partitions

Use the following procedure to move volumes incrementally to disk partitions.

To move volumes incrementally to disk partitions

- 1 Evacuate disks using `vxdiskadm`, the VOM GUI, or the `vxevac` utility.

Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.
- 2 Remove the evacuated disks from VxVM control by entering:


```
# vxdg rmdisk diskname  
# vxdisk rm devname
```
- 3 Decide which volume to move first, and if the volume is mounted, unmount it.
- 4 If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that you have applied the `sync` command to the data on the volume.
- 5 Create a partition on free disk space of the same size as the volume using the `format` command.

If there is not enough free space for the partition, add a new disk to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this first volume.

- 6 Copy the data on the volume onto the newly created disk partition using a command such as `dd`.

```
# dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/dsk/c2t2d2s7
```

where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly created partition.

- 7 Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.
- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Stop and remove the volume from VxVM using the commands.

```
# vxvol -g diskgroup stop volume_name
# vxedit -rf rm volume_name
```

- 10 Remove any free disks (those having no subdisks defined on them) by removing the volumes from VxVM control.

To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.

```
# vxprint -g diskgroup -F '%snum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VxVM control.

```
# vxdg rmdisk diskname
# vxdisk rm devname
```

Use the free space created for adding the data from the next volume you want to remove.

- 11 After you successfully convert all volumes into disk partitions, reboot the system.
- 12 After the reboot, make sure none of the volumes are open by using the `vxprint` command.

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

- 13 If any volumes remain open, repeat the steps listed above.

Example of moving volumes to disk partitions on Solaris

This example shows how to move the data on a volume to a disk partition. In the example, there are three disks: `disk1` and `disk2` are subdisks on volume `vol01` and `disk3` is a free disk. The data on `vol01` is copied to `disk3` using `vxevac`.

These are the contents of the disk group `voldg` before the data on `vol01` is copied to `disk3`.

```
# vxprint -g voldg -ht
DG NAME  NCONFIG  NLOG     MINORS   GROUP-ID
DM NAME  DEVICE   TYPE     PRIVLEN  PUBLLEN  STATE
RV NAME  RLINK_CNT KSTATE   STATE    PRIMARY  DATAVOL  SRL
RL NAME  RVG      KSTATE   STATE    REM_HOST  REM_DG    REM_RLNK
V  NAME  RVG      KSTATE   STATE    LENGTH   READPOL   PREFPLEX  UTYPE
PL NAME  VOLUME   KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID  MODE
SD NAME  PLEX     DISK     DISKOFFS LENGTH    [COL/]OFF DEVICE    MODE
SV NAME  PLEX     VOLNAME  NVOLLAYR LENGTH    [COL/]OFF AM/NM     MODE
DC NAME  PARENTVOL LOGVOL
SP NAME  SNAPVOL  DCO

dg voldg default      default 115000
1017856044.1141.hostname.veritas.com

dm disk1 c1t12d0s2 sliced  2591    17900352 -
dm disk2 c1t14d0s2 sliced  2591    17899056 -
dm disk3 c1t3d0s2 sliced  2591    17899056 -

v  vol1  -          ENABLED ACTIVE  4196448 ROUND    -          fsgen
pl pl1   vol1    ENABLED ACTIVE  4196448 CONCAT   -          RW
sd sd1   pl1     disk1    0        2098224  0        c1t12d0  ENA
sd sd2   pl1     disk2    0        2098224  2098224  c1t14d0  ENA
```

Evacuate `disk1` to `disk3`.

```
# /etc/vx/bin/vxevac -g voldg disk1 disk3
# vxprint -g voldg -ht

DG NAME  NCONFIG  NLOG     MINORS   GROUP-ID
DM NAME  DEVICE   TYPE     PRIVLEN  PUBLLEN  STATE
RV NAME  RLINK_CNT KSTATE   STATE    PRIMARY  DATAVOL  SRL
RL NAME  RVG      KSTATE   STATE    REM_HOST  REM_DG    REM_RLNK
V  NAME  RVG      KSTATE   STATE    LENGTH   READPOL   PREFPLEX  UTYPE
PL NAME  VOLUME   KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID  MODE
SD NAME  PLEX     DISK     DISKOFFS LENGTH    [COL/]OFF DEVICE    MODE
```

```
SV NAME    PLEX          VOLNAME  NVOLLAYR  LENGTH    [COL/]OFF AM/NM    MODE
DC NAME    PARENTVOL  LOGVOL
SP NAME    SNAPVOL    DCO
```

```
dg voldg default      default  115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2  sliced  2591    17900352 -
dm disk2 c1t14d0s2  sliced  2591    17899056 -
dm disk3 c1t3d0s2   sliced  2591    17899056 -
```

```
v  vol1  -          ENABLED  ACTIVE  4196448  ROUND    -          fsgen
pl pl1   vol1      ENABLED  ACTIVE  4196448  CONCAT   -          RW
sd disk3-01l1    disk3    0        2098224  0        c1t3d0   ENA
sd sd2   pl1       disk2    0        2098224  2098224  c1t14d0  ENA
```

Evacuate disk2 to disk3.

```
# /etc/vx/bin/vxevac -g voldg disk2 disk3
# vxprint -g voldg -ht
```

```
DG NAME      NCONFIG  NLOG    MINORS  GROUP-ID
DM NAME      DEVICE   TYPE    PRIVLEN PUBLEN  STATE
RV NAME      RLINK_CNT KSTATE  STATE   PRIMARY DATAVOLS SRL
RL NAME      RVG      KSTATE  STATE   REM_HOST REM_DG   REM_RLNK
V  NAME      RVG      KSTATE  STATE   LENGTH  READPOL  PREFPLEX UTYPE
PL NAME      VOLUME   KSTATE  STATE   LENGTH  LAYOUT   NCOL/WID MODE
SD NAME      PLEX     DISK    DISKOFFS LENGTH  [COL/]OFF DEVICE  MODE
SV NAME      PLEX     VOLNAME NVOLLAYR LENGTH  [COL/]OFF AM/NM   MODE
DC NAME      PARENTVOL LOGVOL
SP NAME      SNAPVOL   DCO
```

```
dg voldg default      default  115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2  sliced  2591    17900352 -
dm disk2 c1t14d0s2  sliced  2591    17899056 -
dm disk3 c1t3d0s2  sliced  2591    17899056 -
```

```
v  vol1  -          ENABLED  ACTIVE  4196448  ROUND    -          fsgen
pl pl1   vol1      ENABLED  ACTIVE  4196448  CONCAT   -          RW
sd disk3-01 pl1    disk3    0        2098224  0        c1t3d0   ENA
sd disk3-02 pl1    disk3    2098224  2098224  2098224  c1t3d0   ENA
```

Remove the evacuated disks from VxVM control.

```
# vxdisk -g voldg list
DEVICE      TYPE      DISK      GROUP      STATUS
c1t3d0s2    sliced    disk3     voldg      online
c1t12d0s2    sliced    disk1     voldg      online
c1t14d0s2    sliced    disk2     voldg      online

# vxvg rmdisk disk1
# vxvg rmdisk disk2
# vxdisk rm c1t12d0
# vxdisk rm c1t14d0
```

Verify that the evacuated disks have been removed from VxVM control.

```
# vxdisk -g voldg list
DEVICE      TYPE      DISK      GROUP      STATUS
c1t3d0s2    sliced    disk3     voldg      online
```

Check to see whether the volume you want to move first is mounted.

```
# mount | grep vol1
/vol1 on /dev/vx/dsk/voldg/vol1
read/write/setuid/log/nolargefiles/dev=12dc138 on Wed Apr
3 10:13:11 2002
```

Create a partition on free disk space of the same size as the volume. In this example, a 2G partition is created on disk1 (c1t12d0s1).

```
# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
 0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
    /sbus@1f,0/SUNW,fas@e,8800000/sd@0,0
 1. c1t3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
    /sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
 2. c1t9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
    /sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
 3. c1t10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
    /sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
 4. c1t11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
    /sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
 5. c1t12d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
    /sbus@1f,0/SUNW,fas@2,8800000/sd@c,0
```

6. clt14d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@e,0
7. clt15d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
/sbus@1f,0/SUNW,fas@2,8800000/sd@f,0

Specify disk (enter its number): **5**

selecting clt12d0

[disk formatted]

FORMAT MENU:

- disk - select a disk
- type - select (define) a disk type
- partition - select (define) a partition table
- current - describe the current disk
- format - format and analyze the disk
- repair - repair a defective sector
- label - write label to the disk
- analyze - surface analysis
- defect - defect list management
- backup - search for backup labels
- verify - read and display labels
- save - save new disk/partition definitions
- inquiry - show vendor, product and revision
- volname - set 8-character volume name
- !**<cmd>** - execute **<cmd>**, then return
- quit

format> **p**

PARTITION MENU:

- 0 - change '0' partition
- 1 - change '1' partition
- 2 - change '2' partition
- 3 - change '3' partition
- 4 - change '4' partition
- 5 - change '5' partition
- 6 - change '6' partition
- 7 - change '7' partition
- select - select a predefined table
- modify - modify a predefined partition table
- name - name the current table
- print - display the current table
- label - write partition map and label to the disk
- !**<cmd>** - execute **<cmd>**, then return
- quit

```

partition> 1
Part      Tag      Flag      Cylinders      Size      Blocks
  1 unassigned    wm         0           0      (0/0/0)         0
Enter partition id tag[unassigned]:
Enter partition permission flags[wm]:
Enter new starting cyl[0]:
Enter partition size[0b, 0c, 0.00mb, 0.00gb]: 2.00gb
partition> 1
Ready to label disk, continue? y

partition> p
Current partition table (unnamed):
Total disk cylinders available: 13814 + 2 (reserved cylinders)
Part      Tag      Flag      Cylinders      Size      Blocks
  0 unassigned    wm         0           0      (0/0/0)         0
  1 unassigned    wm         0 - 3236      2.00GB      (3237/0/0)    4195152
partition> q

```

Copy the data on `vol01` to the newly created disk partition.

```
# dd if=/dev/vx/dsk/voldg/vol01 of=/dev/dsk/c1t12d0s1
```

In the `/etc/vfstab` file, remove the following entry.

```
/dev/vx/dsk/voldg/vol1 /dev/vx/rdsk/voldg/vol1 /vol1 vxfs 4 yes rw
```

Replace it with an entry for the newly created partition.

```
/dev/dsk/c1t12d0s1 /dev/rdsk/c1t12d0s1 /vol01 vxfs 4 yes rw
```

Mount the disk partition.

```
# mount -F vxfs /dev/dsk/c1t12d0s1 /vol01
```

Remove `vol01` from VxVM.

```
# vxdedit -rf rm /dev/vx/dsk/voldg/vol01
```

To complete the procedure, follow the remaining steps.

Preparing to remove Veritas File System

The `VRTSvxfs` package cannot be removed if there are any mounted VxFS file systems or Storage Checkpoints. Unmount the VxFS file systems and Storage Checkpoints before uninstalling Veritas Storage Foundation. After you remove

the `VRTSvxfs` package, VxFS file systems are not mountable or accessible until another `VRTSvxfs` package is installed.

To unmount a file system

- 1 Check if any VxFS file systems are mounted.

```
# cat /etc/mnttab | grep vxfs
```

- 2 Unmount any file systems.

```
# umount special | mount_point
```

Specify the file system to be unmounted as a *mount_point* or *special* (the device on which the file system resides). See the `umount_vxfs(1M)` manual page for more information about this command and its available options.

You can use the `-a` option to unmount all file systems except `/`, `/usr`, `/usr/kvm`, `/var`, `/proc`, `/dev/fd`, and `/tmp`.

To unmount a Storage Checkpoint

- 1 Check if any Storage Checkpoints are mounted.

```
# cat /etc/mnttab | grep vxfs
```

- 2 Unmount any Storage Checkpoints.

```
# umount /checkpoint_name
```

Disabling VCS agents for VVR the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

To disable the agents

- 1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
# hagrps -state service_group -sys system_name
```

If none of the service groups is online, skip to 3.

- 2 If the service group is online, take it offline.

To take the service group offline without bringing it online on any other system in the cluster, enter:

```
# hagrps -offline service_group -sys system_name
```

- 3 Stop the agent on the system by entering:

```
# haagent -stop agent_name -sys system_name
```

When you get the message `Please look for messages in the log file, check the file /var/VRTSvcs/log/engine_A.log for a message confirming that each agent has stopped.`

You can also use the `ps` command to confirm that the agent is stopped.

- 4 Remove the system from the `SystemList` of the service group. If you disable the agent on all the systems in the `SystemList`, you can also remove the service groups and resource types from the VCS configuration.

Read information on administering VCS from the command line.

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

Removing the Replicated Data Set

If you use VVR, you need to perform the following steps. This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

To remove the Replicated Data Set

- 1 Verify that all RLINKs are up-to-date:

```
# vxrlink -g diskgroup status rlink_name
```

If the Secondary is not required to be up-to-date, proceed to [2](#) and stop replication using the `-f` option with the `vradmin stoprep` command.

- 2 Stop replication to the Secondary by issuing the following command on any host in the RDS:

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

```
# vradmin -g diskgroup stoprep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 3 Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

```
# vradmin -g diskgroup delsec local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 4 Remove the Primary from the RDS by issuing the following command on the Primary:

```
# vradmin -g diskgroup delpri local_rvgname
```

When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

The RDS is removed.

- 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 SFHA packages using the script-based installer

Use the following procedure to remove SFHA products.

Not all packages may be installed on your system depending on the choices that you made when you installed the software.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in SFHA 6.0 PR1 with a previous version of SFHA.

Language packages are uninstalled when you uninstall the English language packages.

To shut down and remove the installed SFHA packages

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems.

```
# umount /mount_point
```

- 3 If the VxVM package (`VRTSvxxvm`) is installed, read and follow the uninstallation procedures for VxVM.

See [“Preparing to remove Veritas Volume Manager”](#) on page 184.

- 4 Make sure you have performed all of the prerequisite steps.

- 5 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 6 Move to the `/opt/VRTS/install` directory and run the uninstall script.

```
# cd /opt/VRTS/install
```

For Veritas Storage Foundation High Availability

```
# ./uninstallsfha
```

- 7 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall SFHA, for example,

```
host1:
```

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

- 8 The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the packages are uninstalled.

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

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

- 10 To verify the removal of the packages, use the `pkginfo` command.

```
# pkginfo | grep VRTS
```

Removing the CP server configuration using the removal script

This section describes how to remove the CP server configuration from a node or a cluster that hosts the CP server.

Warning: Ensure that no SFHA cluster (application cluster) uses the CP server that you want to unconfigure.

You can use the CP server configuration utility (`configure_cps.pl`) to remove the CP server configuration. This utility performs the following tasks when you choose to unconfigure the CP server:

- Removes all CP server configuration files
- Removes the VCS configuration for CP server

After you run this utility, you can uninstall VCS from the node or the cluster.

Note: You must run the configuration utility only once per CP server (which can be on a single-node VCS cluster or an SFHA cluster), when you want to remove the CP server configuration.

To remove the CP server configuration

- 1 To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2 Select option 3 from the menu to unconfigure the CP server.

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

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

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

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

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

- 3 Review the warning message and confirm that you want to unconfigure the CP server.

```
WARNING: Unconfiguring Coordination Point Server stops the
vxcpserv process. VCS clusters using this server for
coordination purpose will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)
(Default:n) :y
```

- 4 Review the screen output as the script performs the following steps to remove the CP server configuration:

- Stops the CP server
- Removes the CP server from VCS configuration
- Removes resource dependencies
- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration

- 5 Answer **y** to delete the CP server database.

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

- 6 Answer **y** at the prompt to confirm the deletion of the CP server database.

```
Warning: This database won't be available if CP server
is reconfigured on the cluster. Are you sure you want to
proceed with the deletion of database? (y/n) (Default:n) :
```

- 7 Answer **y** to delete the CP server configuration file and log files.

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 8 Run the `hagrp -state` command to ensure that the CPSSG service group has been removed from the node. For example:

```
root@mycps1.symantecexample.com # hagrp -state CPSSG
```

```
VCS WARNING V-16-1-40131 Group CPSSG does not exist
in the local cluster
```

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file disables the SFDB tools.

To remove the SFDB repository

- 1 Identify the SFDB repositories created on the host.

```
# cat /var/vx/vxdba/rep_loc

{
  "sfae_rept_version" : 1,
  "oracle" : {
    "SFAEDB" : {
      "location" : "/data/sfaedb/.sfae",
      "old_location" : "",
      "alias" : [
        "sfaedb"
      ]
    }
  }
}
```

- 2 Remove the directory identified by the `location` key.

```
# rm -rf /data/sfaedb/.sfae
```

- 3 Remove the repository location file.

```
# rm -rf /var/vx/vxdba/rep_loc
```

This completes the removal of the SFDB repository.

Manually uninstalling Storage Foundation and High Availability packages on non-global zones

1. Log on to the non-global zone as a super user.
2. Uninstall SFHA packages from Solaris brand zones.

```
# pkg uninstall VRTSperl VRTSvlic VRTSvcS VRTSvcSag VRTSvcsea
```

3. Uninstall SFHA packages from solaris10 brand zones.

```
# pkgrm VRTSperl VRTSvlic VRTSvcS VRTSvcSag VRTSvcsea
```

Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Response files](#)
- [Appendix C. Tunable files for installation](#)
- [Appendix D. Configuring I/O fencing using a response file](#)
- [Appendix E. Configuration files](#)
- [Appendix F. Configuring the secure shell or the remote shell for communications](#)
- [Appendix G. Storage Foundation and High Availability components](#)
- [Appendix H. Troubleshooting installation issues](#)
- [Appendix I. Troubleshooting cluster installation](#)
- [Appendix J. Sample SF HA cluster setup diagrams for CP server-based I/O fencing](#)
- [Appendix K. Reconciling major/minor numbers for NFS shared disks](#)
- [Appendix L. Configuring LLT over UDP using IPv4](#)
- [Appendix M. Configuring LLT over UDP using IPv6](#)

- [Appendix N. Compatability issues when installing Storage Foundation High Availability with other products](#)
- [Appendix O. Changing root user into role](#)

Installation scripts

This appendix includes the following topics:

- [About installation scripts](#)
- [Installation script options](#)
- [About using the postcheck option](#)

About installation scripts

Veritas Storage Foundation and High Availability Solutions products 6.0 PR1 provides several installation scripts. You can find these scripts at the root of the product media in the scripts directory.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from the Symantec download site, which does not include the installer, use the appropriate product installation script.

The following product installation scripts are available:

Veritas Cluster Server (VCS)	<code>installvcs</code>
Veritas Storage Foundation (SF)	<code>installsf</code>
Veritas Storage Foundation and High Availability (SFHA)	<code>installsfha</code>
Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)	<code>installsfcfsha</code>
Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)	<code>installsfprac</code>
Veritas Storage Foundation for Sybase ASE CE (SF Sybase CE)	<code>installsfbasece</code>

Veritas Volume Manager	<code>installvm</code>
Veritas File System	<code>installfs</code>
Veritas Dynamic Multi-pathing	<code>installdmp</code>
Symantec VirtualStore	<code>installsvs</code>

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 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 201.

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 required for the specified product. The packages are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
<code>-comcleanup</code>	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
<code>-configure</code>	Configures the product after installation.
<code>-fencing</code>	Configures I/O fencing in a running cluster.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-hostfile <i>full_path_to_file</i>	Specifies the location of a file that contains a list of hostnames on which to install.
-install	The -install option is used to install products on systems.
-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.
-jumpstart <i>dir_path</i>	Produces a sample finish file for Solaris JumpStart installation. The <i>dir_path</i> indicates the path to the directory in which to create the finish file.
-keyfile <i>ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes -i <i>ssh_key_file</i> to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
-logpath <i>log_path</i>	Specifies a directory other than /opt/VRTS/install/logs as the location where installer log files, summary files, and response files are saved.
-makeresponsefile	Use the -makeresponsefile option only to generate response files. No actual software installation occurs when you use this option.
-minpkgs	Displays the minimal packages required for the specified product. The packages are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See allpkgs option.
-nolic	Allows installation of product packages without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-pkg info	Displays a list of packages and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkg info option with the installvcs script to display VCS packages.
-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 displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.
-pkgtable	Displays product's packages in correct installation order by group.
-postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
-recpkgs	Displays the recommended packages required for the specified product. The packages are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See allpkgs option.
-redirect	Displays progress details without showing the progress bar.
-requirements	The -requirements option displays required OS version, required packages, file system space, and other system requirements in order to install the product.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
<code>-responsefile <i>response_file</i></code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<code>-rolling_upgrade</code>	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
<code>-rollingupgrade_phase1</code>	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel packages get upgraded to the latest version
<code>-rollingupgrade_phase2</code>	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."
<code>-rootpath <i>root_path</i></code>	Specifies an alternative root directory on which to install packages. On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code> command.
<code>-rsh</code>	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.
<code>-serial</code>	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.

Table A-1 Available command line options (continued)

Command Line Option	Function
-settnables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the -tunablesfile option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-tmppath <i>tmp_path</i>	Specifies a directory other than /var/tmp as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
-uninstall	The -uninstall option is used to uninstall products from systems.
-tunablesfile	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing packages where applicable for the product. Provides a summary that includes the count of the installed and any missing packages where applicable. Lists the installed hotfixes and available updates for the installed product if an Internet connection is available.

About using the postcheck option

You can use the installer's post-check to determine installation-related problems and to aid in troubleshooting.

Note: This command option requires downtime for the node.

When you use the `postcheck` option, it can help you troubleshoot the following VCS-related issues:

- The heartbeat link does not exist.
- The heartbeat link cannot communicate.
- The heartbeat link is a part of a bonded or aggregated NIC.
- A duplicated cluster ID exists.
- The `VRTSllt` pkg version is not consistent on the nodes.
- The `llt-linkinstall` value is incorrect.
- The `llthosts(4)` or `llttab(4)` configuration is incorrect.
- the `/etc/gabtab` file is incorrect.
- The incorrect GAB `linkinstall` value exists.
- The `VRTSgab` pkg version is not consistent on the nodes.
- The `main.cf` file or the `types.cf` file is invalid.
- The `/etc/VRTSvcs/conf/sysname` file is not consistent with the hostname.
- The cluster UUID does not exist.
- The `uuidconfig.pl` file is missing.
- The `VRTSvcs` pkg version is not consistent on the nodes.
- The `/etc/vxfenmode` file is missing or incorrect.
- The `/etc/vxfendg` file is invalid.
- The `vxfen` link-install value is incorrect.
- The `VRTSvxfen` pkg version is not consistent.

The `postcheck` option can help you troubleshoot the following SFHA or SFCFSHA issues:

- Volume Manager cannot start because the `/etc/vx/reconfig.d/state.d/install-db` file has not been removed.
- Volume Manager cannot start because the `Volboot` file is not loaded.
- Volume Manager cannot start because no license exists.

- Cluster Volume Manager cannot start because the CVM configuration is incorrect in the `main.cf` file. For example, the `Autostartlist` value is missing on the nodes.
- Cluster Volume Manager cannot come online because the node ID in the `/etc/llthosts` file is not consistent.
- Cluster Volume Manager cannot come online because `Vxfen` is not started.
- Cluster Volume Manager cannot start because `gab` is not configured.
- Cluster Volume Manager cannot come online because of a CVM protocol mismatch.
- Cluster Volume Manager group name has changed from "cvm", which causes CVM to go offline.

See [“Performing a postcheck on a node”](#) on page 139.

Response files

This appendix includes the following topics:

- [About response files](#)
- [Installing SFHA using response files](#)
- [Configuring SFHA using response files](#)
- [Uninstalling SFHA using response files](#)
- [Syntax in the response file](#)
- [Response file variables to install, upgrade, or uninstall Storage Foundation and High Availability](#)
- [Response file variables to configure Storage Foundation and High Availability](#)
- [Sample response file for SFHA configuration](#)
- [Sample response file for SFHA install](#)

About response files

The installer or product installation script generates a response file during any installation, configuration, upgrade, or uninstall procedure. The response file contains the configuration information that you entered during the procedure. When the procedure completes, the installation script displays the location of the response files.

You can use the response file for future installation procedures by invoking an installation script with the `-responsefile` option. The response file passes arguments to the script to automate the installation of that product. You can edit the file to automate installation and configuration of additional systems.

You can generate a response file using the `-makeresponsefile` option.

Installing SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA installation on one cluster to install SFHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To install SFHA using response files

- 1 Make sure the systems where you want to install SFHA 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 SFHA.
- 4 Edit the values of the response file variables as necessary.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installation from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file  
  
# ./installsfha -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Configuring SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA configuration on one cluster to configure SFHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure SFHA using response files

- 1 Make sure the SFHA packages are installed on the systems where you want to configure SFHA.
- 2 Copy the response file to one of the cluster systems where you want to configure SFHA.

- 3 Edit the values of the response file variables as necessary.

To configure optional features, you must define appropriate values for all the response file variables that are related to the optional feature.

See [“Response file variables to configure Storage Foundation and High Availability”](#) on page 214.

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

```
# /opt/VRTS/install/installsfha -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Uninstalling SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA uninstallation on one cluster to uninstall SFHA on other clusters.

To perform an automated uninstallation

- 1 Make sure that you meet the prerequisites to uninstall SFHA.
- 2 Copy the response file to the system where you want to uninstall SFHA.
- 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/uninstallsfha -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Syntax in the response file

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

For example, in the case of a string value:

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value", "value", "value"];
```

Response file variables to install, upgrade, or uninstall Storage Foundation and High Availability

Table B-1 lists the response file variables that you can define to configure SFHA.

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SFHA

Variable	Description
CFG{opt}{install}	Installs SFHA packages. Configuration can be performed at a later time using the <code>-configure</code> option. List or scalar: scalar Optional or required: optional
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required
\$CFG{opt}{vxkeyless}	Installs the product with keyless license. List or scalar: scalar Optional or required: optional
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SFHA (*continued*)

Variable	Description
CFG{opt}{pkgpath}	Defines a location, typically an NFS mount, from which all remote systems can install product packages. The location must be accessible from all target systems. List or scalar: scalar Optional or required: optional
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the packages that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{rsh}	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. List or scalar: scalar Optional or required: optional
CFG{donotinstall} {package}	Instructs the installation to not install the optional packages in the list. List or scalar: list Optional or required: optional
CFG{donotremove} {package}	Instructs the uninstallation to not remove the optional packages in the list. List or scalar: list Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
\$CFG{opt}{prodmode}	List of modes for product List or scalar: list Optional or required: optional

Table B-1

Response file variables specific to installing, upgrading, or uninstalling SFHA *(continued)*

Variable	Description
CFG{opt}{upgrade}	Upgrades all packages installed, without configuration. List or scalar: list Optional or required: optional
CFG{opt}{uninstall}	Uninstalls SFHA packages. List or scalar: scalar Optional or required: optional
CFG{mirrordgname}{system}	Splits the target disk group name for a system. List or scalar: scalar Optional or required: optional
CFG{splitmirror}{system}	Indicates the system where you want a split mirror backup disk group created. List or scalar: scalar Optional or required: optional

Response file variables to configure Storage Foundation and High Availability

Table B-2 lists the response file variables that you can define to configure SFHA.

Table B-2

Response file variables specific to configuring Storage Foundation and High Availability

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the packages are already installed. (Required) Set the value to 1 to configure SFHA.
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)

Table B-2 Response file variables specific to configuring Storage Foundation and High Availability (*continued*)

Variable	List or Scalar	Description
CFG{systems}	List	List of systems on which the product is to be configured. (Required)
CFG{prod}	Scalar	Defines the product to be configured. The value is VCS60 for VCS. (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of ssh as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtprecp, and smtpprsev), the SNMP trap notification (snmppport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Table B-3 lists the response file variables that specify the required information to configure a basic SFHA cluster.

Table B-3 Response file variables specific to configuring a basic SFHA cluster

Variable	List or Scalar	Description
CFG{vcs_clusterid}	Scalar	An integer between 0 and 65535 that uniquely identifies the cluster. (Required)
CFG{vcs_clustername}	Scalar	Defines the name of the cluster. (Required)
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)
CFG{fencingenabled}	Scalar	In a SFHA configuration, defines if fencing is enabled. Valid values are 0 or 1. (Required)

Table B-4 lists the response file variables that specify the required information to configure LLT over Ethernet.

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

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

[Table B-5](#) lists the response file variables that specify the required information to configure LLT over UDP.

Table B-5 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltoverudp}=1	Scalar	Indicates whether to configure heartbeat link using LLT over UDP. (Required)

Table B-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG {vcs_udplinklowpri<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table B-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table B-6 lists the response file variables that specify the required information to configure virtual IP for SFHA cluster.

Table B-6 Response file variables specific to configuring virtual IP for SFHA cluster

Variable	List or Scalar	Description
CFG{vcs_csgnic} {system}	Scalar	Defines the NIC device to use on a system. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_csgvip}	Scalar	Defines the virtual IP address for the cluster. (Optional)
CFG{vcs_csgnetmask}	Scalar	Defines the Netmask of the virtual IP address for the cluster. (Optional)

[Table B-7](#) lists the response file variables that specify the required information to configure the SFHA cluster in secure mode.

Table B-7 Response file variables specific to configuring SFHA cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonnode}	Scalar	Specifies that the securityonnode option is being used.
CFG{securityonnode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. ■ 1—Configure the first node ■ 2—Configure the other node
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.

[Table B-8](#) lists the response file variables that specify the required information to configure VCS users.

Table B-8 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	List of encoded passwords for VCS users The value in the list can be "Administrators Operators Guests" Note: The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list. (Optional)
CFG{vcs_username}	List	List of names of VCS users (Optional)

Table B-8 Response file variables specific to configuring VCS users (*continued*)

Variable	List or Scalar	Description
CFG{vcs_userpriv}	List	List of privileges for VCS users Note: The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list. (Optional)

[Table B-9](#) lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table B-9 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)
CFG{vcs_smtprev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

[Table B-10](#) lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table B-10 Response file variables specific to configuring VCS notifications using SNMP

Variable	List or Scalar	Description
CFG{vcs_snmpport}	Scalar	Defines the SNMP trap daemon port (default=162). (Optional)
CFG{vcs_snmpcons}	List	List of SNMP console system names (Optional)
CFG{vcs_snmpcsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

[Table B-11](#) lists the response file variables that specify the required information to configure SFHA global clusters.

Table B-11 Response file variables specific to configuring SFHA global clusters

Variable	List or Scalar	Description
CFG{vcs_gconic} {system}	Scalar	Defines the NIC for the Virtual IP that the Global Cluster Option uses. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_gcovip}	Scalar	Defines the virtual IP address to that the Global Cluster Option uses. (Optional)
CFG{vcs_gconetmask}	Scalar	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. (Optional)

Sample response file for SFHA configuration

The following example shows a response file for configuring Storage Foundation High Availability.

```
#####
#Auto generated sfha responsefile #
#####

our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{rsh}=1;
$CFG{vcs_allowcomms}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
$CFG{opt}{prodmode}="SF Enterprise HA";
$CFG{opt}{configure}=1;
$CFG{prod}="SFHA51";
$CFG{systems}=[ qw( system01 system02 ) ];
$CFG{vm_restore_cfg}{system01}=0;
$CFG{vm_restore_cfg}{system02}=0;
$CFG{vcs_clusterid}=127;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_username}=[ qw(admin operator) ];
$CFG{vcs_userenpw}=[ qw(JlmElgLimHmmKumGlj bQOsOUUnVQoOUUnTQsOSnUQuOUUnPQtOS) ];
$CFG{vcs_userpriv}=[ qw(Administrators Operators) ];
$CFG{vcs_lltlink1}{"system01"}="net1";
$CFG{vcs_lltlink2}{"system01"}="net2";
$CFG{vcs_lltlink1}{"system02"}="net1";
$CFG{vcs_lltlink2}{"system02"}="net2";
$CFG{opt}{uuid}=normC;
$CFG{opt}{logpath}="/opt/VRTS/install/logs/installsf-xxxxxx/installsf-xxxxxx.response";

1;
```

Sample response file for SFHA install

The following example shows a response file for installing Storage Foundation High Availability.

```
#####
#Auto generated sfha responsefile #
#####
```

Sample response file for SFHA install

```
our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
$CFG{opt}{prodmode}="SF Enterprise HA";
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{prod}="SFHA51";
$CFG{systems}=[ qw( system01 system02 ) ];
$CFG{keys}{system01}=["XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX"];
$CFG{keys}{system02}=["XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX"];
$CFG{opt}{uuid}=normI;
$CFG{opt}{logpath}="/opt/VRTS/install/logs/SxRT-5.1-2009-03-10a";

1;
```


Tunable files for installation

This appendix includes the following topics:

- [About setting tunable parameters using the installer or a response file](#)
- [Setting tunables for an installation, configuration, or upgrade](#)
- [Setting tunables with no other installer-related operations](#)
- [Setting tunables with an un-integrated response file](#)
- [Preparing the tunables file](#)
- [Setting parameters for the tunables file](#)
- [Tunables value parameter definitions](#)

About setting tunable parameters using the installer or a response file

You can set non-default product and system tunable parameters using a tunables file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing. The tunables file passes arguments to the installer script to set tunables. With the file, you can set the tunables for the following operations:

- When you install, configure, or upgrade systems.

```
# ./installer -tunablesfile tunables_file_name
```

See “[Setting tunables for an installation, configuration, or upgrade](#)” on page 226.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -settunables [  
  system1 system2 ...]
```

See [“Setting tunables with no other installer-related operations”](#) on page 227.

- When you apply the tunables file with an un-integrated response file.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name
```

See [“Setting tunables with an un-integrated response file”](#) on page 228.

See [“About response files”](#) on page 209.

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

See [“Tunables value parameter definitions”](#) on page 230.

Setting tunables for an installation, configuration, or upgrade

You can use a tunables file for installation procedures to set non-default tunables. You invoke the installation script with the `tunablesfile` option. The tunables file passes arguments to the script to set the selected tunables. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 230.

Note: Certain tunables only take effect after a system reboot.

To set the non-default tunables for an installation, configuration, or upgrade

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 229.
- 2 Make sure the systems where you want to install SFHA meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 230.

Note: Certain tunables only take effect after a system reboot.

To set tunables with no other installer-related operations

- 1 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 229.
- 2 Make sure the systems where you want to install SFHA meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-setttunables` option.

```
# ./installer -tunablesfile tunables_file_name -setttunables [  
  sys123 sys234 ...]
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 230.

Note: Certain tunables only take effect after a system reboot.

To set tunables with an un-integrated response file

- 1 Make sure the systems where you want to install SFHA meet the installation requirements.
- 2 Complete any preinstallation tasks.
- 3 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 229.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-settunables` option.

```
# ./installer -responsefile response_file_name -tunablesfile  
tunables_file_name -settunables
```

Where *response_file_name* is the full path name for the response file and *tunables_file_name* is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Preparing the tunables file

A tunables file is a Perl module and consists of an opening and closing statement, with the tunables defined between. Use the hash symbol at the beginning of the line to comment out the line. The tunables file opens with the line "our %TUN;" and ends with the return true "1;" line. The final return true line only needs to appear once at the end of the file. Define each tunable parameter on its own line.

Format the tunable parameter as follows:

```
$TUN{"tunable_name"}{ "system_name" | "*" }=value_of_tunable;
```

For the *system_name*, use the name of the system, its IP address, or a wildcard symbol. The *value_of_tunable* depends on the type of tunable you are setting. End the line with a semicolon.

The following is an example of a tunables file.

```
#  
# Tunable Parameter Values:  
#  
our %TUN;  
  
$TUN{"tunable1"}{ "*" }=1024;  
$TUN{"tunable3"}{ "sys123" }="SHA256";  
  
1;
```

Setting parameters for the tunables file

Each tunables file defines different tunable parameters. The values that you can use are listed in the description of each parameter. Select the tunables that you want to add to the tunables file and then configure each parameter.

See [“Tunables value parameter definitions”](#) on page 230.

Each line for the parameter value starts with \$TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{ "node123" }=16;
```

In this example, you are changing the *dmp_daemon_count* value from its default of 10 to 16. You can use the wildcard symbol "*" for all systems. For example:

```
$TUN{"dmp_daemon_count"}{"*"}=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Veritas Storage Foundation and High Availability Solutions Tuning Guide* for detailed information on product tunable ranges and recommendations .

[Table C-1](#) describes the supported tunable parameters that can be specified in a tunables file.

Table C-1 Supported tunable parameters

Tunable	Description
dmp_cache_open	(Veritas Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_daemon_count	(Veritas Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_delayq_interval	(Veritas Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_fast_recovery	(Veritas Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_health_time	(Veritas Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_log_level	(Veritas Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

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

Tunable	Description
dmp_low_impact_probe	(Veritas Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_lun_retry_timeout	(Veritas Dynamic Multi-Pathing) The retry period for handling transient errors. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_fabric	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Veritas Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_multipathing	(Veritas Dynamic Multi-Pathing) Whether DMP will intercept the I/Os directly on the raw OS paths or not. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_support	(Veritas Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_path_age	(Veritas Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_pathswitch_blks_shift	(Veritas Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

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

Tunable	Description
dmp_probe_idle_lun	(Veritas Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_threshold	(Veritas Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_cycles	(Veritas Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_interval	(Veritas Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_policy	(Veritas Dynamic Multi-Pathing) The policy used by DMP path restoration thread. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_state	(Veritas Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_retry_count	(Veritas Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_scsi_timeout	(Veritas Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_sfg_threshold	(Veritas Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_stat_interval	(Veritas Dynamic Multi-Pathing) The time interval between gathering DMP statistics. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

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

Tunable	Description
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer sets only the system default value of max_diskq. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer sets only the system default value of read_ahead. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer sets only the system default value of read_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer sets only the system default value of read_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
vol_checkpoint_default	(Veritas File System) Size of VxVM checkpoints (sectors). This tunable requires system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for VERITAS Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for VERITAS Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires system reboot to take effect.

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

Tunable	Description
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of VERITAS Volume Replicator (bytes).
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (sectors). This tunable requires system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (sectors). This tunable requires system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Veritas Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by VERITAS Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions that can exist on a non-sequential DRL. This tunable requires system reboot to take effect.

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

Tunable	Description
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (sectors). This tunable requires system reboot to take effect.
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires system reboot to take effect.
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires system reboot to take effect.

Table C-1 Supported tunable parameters (continued)

Tunable	Description
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires system reboot to take effect.
vx_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. This tunable requires system reboot to take effect.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. This tunable requires system reboot to take effect.
vxfs_mbuf	(Veritas File System) Maximum memory used for VxFS buffer cache. This tunable requires system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer sets only the system default value of write_nstream. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer sets only the system default value of write_pref_io. Refer to the tuneftab(4) manual page for setting this tunable for a specified block device.

Configuring I/O fencing using a response file

This appendix includes the following topics:

- [Configuring I/O fencing using response files](#)
- [Response file variables to configure disk-based I/O fencing](#)
- [Sample response file for configuring disk-based I/O fencing](#)
- [Response file variables to configure server-based I/O fencing](#)
- [Sample response file for configuring non-SCSI-3 server-based I/O fencing](#)
- [Response file variables to configure non-SCSI-3 server-based I/O fencing](#)

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for SFHA.

To configure I/O fencing using response files

- 1 Make sure that SFHA 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 65.

- 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 239.
- See “Sample response file for configuring server-based I/O fencing” on page 242.
- 4
- Edit the values of the response file variables as necessary.
- See “Response file variables to configure disk-based I/O fencing” on page 238.
- See “Response file variables to configure server-based I/O fencing” on page 240.
- 5
- Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfha -responsefile /tmp/response_file
```

Where /tmp/response_file is the response file’s full path name.

Response file variables to configure disk-based I/O fencing

Table D-1 lists the response file variables that specify the required information to configure disk-based I/O fencing for SFHA.

Table D-1 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode ■ 4—Fencing migration when the cluster is online (Required)

Table D-1

Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG {fencing_scsi3_disk_policy}	Scalar	Specifies the I/O fencing mechanism. This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the fencing_scsi3_disk_policy variable and either the fencing_dgname variable or the fencing_newdg_disks variable. (Optional)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) Note: You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.
CFG{fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) Note: You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.

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 238.

```
#  
# Configuration Values:
```

```
#
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;

$CFG{prod}="SFHA60";

$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_dgname}="fendg";
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{fencing_newdg_disks}=
    [ qw(c1t1d0s2 c2t1d0s2 c3t1d0s2) ];
$CFG{fencing_option}=2;
```

Response file variables to configure server-based I/O fencing

You can use a coordination point server-based fencing response file to configure server-based customized I/O fencing.

[Table D-2](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table D-2 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.

Table D-2 Coordination point server (CP server) based fencing response file definitions (*continued*)

Response file field	Definition
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. Note: This field is obsolete if the <code>fencing_config_cpagent</code> field is given a value of '0'.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_reusedg}	This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks). Enter either a "1" or "0". Entering a "1" indicates reuse, and entering a "0" indicates do not reuse. When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text, such as " <code>\$CFG{fencing_reusedg}=0</code> " or " <code>\$CFG{fencing_reusedg}=1</code> " before proceeding with a silent installation.
CFG {fencing_dgname}	The name of the disk group to be used in the customized fencing, where at least one disk is being used.
CFG {fencing_disks}	The disks being used as coordination points if any.
CFG {fencing_ncp}	Total number of coordination points being used, including both CP servers and disks.
CFG {fencing_ndisks}	The number of disks being used.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_ports}	The port that the virtual IP address or the fully qualified host name of the CP server listens on.
CFG {fencing_scsi3_disk_policy}	The disk policy that the customized fencing uses. The value for this field is either "raw" or "dmp"

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing:

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cps_vips}{"10.200.117.145"}=[ qw(10.200.117.145) ];
$CFG{fencing_dgname}="vxvfencoorddg";
$CFG{fencing_disks}=[ qw(emc_clariion0_37 emc_clariion0_13) ];
$CFG{fencing_scsi3_disk_policy}="raw";
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=2;
$CFG{fencing_ports}{"10.200.117.145"}=14250;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFHA60";
$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Sample response file for configuring non-SCSI-3 server-based I/O fencing

The following is a sample response file used for non-SCSI-3 server-based I/O fencing :

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.198.89.251 10.198.89.252 10.198.89.253) ];
$CFG{fencing_cps_vips}{"10.198.89.251"}=[ qw(10.198.89.251) ];
$CFG{fencing_cps_vips}{"10.198.89.252"}=[ qw(10.198.89.252) ];
$CFG{fencing_cps_vips}{"10.198.89.253"}=[ qw(10.198.89.253) ];
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=0;
$CFG{fencing_ports}{"10.198.89.251"}=14250;
$CFG{fencing_ports}{"10.198.89.252"}=14250;
$CFG{fencing_ports}{"10.198.89.253"}=14250;
$CFG{non_scsi3_fencing}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFHA60";
$CFG{systems}=[ qw(galaxy nebula) ];
```

```
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Response file variables to configure non-SCSI-3 server-based I/O fencing

[Table D-3](#) lists the fields in the response file that are relevant for non-SCSI-3 server-based customized I/O fencing.

See [“About I/O fencing for SFHA in virtual machines that do not support SCSI-3 PR”](#) on page 24.

Table D-3 Non-SCSI-3 server-based I/O fencing response file definitions

Response file field	Definition
CFG{non_scsi3_fencing}	Defines whether to configure non-SCSI-3 server-based I/O fencing. Valid values are 1 or 0. Enter 1 to configure non-SCSI-3 server-based I/O fencing.
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. Note: This field is obsolete if the <code>fencing_config_cpagent</code> field is given a value of '0'.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_ncp}	Total number of coordination points (CP servers only) being used.

Table D-3

Non-SCSI-3 server-based I/O fencing response file definitions

(continued)

Response file field	Definition
CFG {fencing_ports}	The port of the CP server that is denoted by <i>cps</i> .

Configuration files

This appendix includes the following topics:

- [About the LLT and GAB configuration files](#)
- [About the AMF configuration files](#)
- [About the VCS configuration files](#)
- [About I/O fencing configuration files](#)
- [Sample configuration files for CP server](#)

About the LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires `/etc/llthosts` and `/etc/llttab` files. GAB requires `/etc/gabtab` file.

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

Table E-1 LLT configuration files

File	Description
/etc/default/llt	This file stores the start and stop environment variables for LLT: ■ LLT_START—Defines the startup behavior for the LLT module after a system reboot. Valid values include: 1—Indicates that LLT is enabled to start up. 0—Indicates that LLT is disabled to start up. ■ LLT_STOP—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include: 1—Indicates that LLT is enabled to shut down. 0—Indicates that LLT is disabled to shut down. The installer sets the value of these variables to 1 at the end of SFHA configuration.
/etc/llthosts	The file <code>llthosts</code> is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster. For example, the file <code>/etc/llthosts</code> contains the entries that resemble: <pre>0 galaxy 1 nebula</pre>

Table E-1 LLT configuration files (continued)

File	Description
/etc/llttab	The file <code>llttab</code> contains the information that is derived during installation and used by the utility <code>lltconfig(1M)</code>. After installation, this file lists the private network links that correspond to the specific system. For example, the file <code>/etc/llttab</code> contains the entries that resemble the following: ■ For Solaris SPARC: <pre>set-node galaxy set-cluster 2 link net1 /dev/net/net1 - ether - - link net2 /dev/net/net2 - ether - -</pre> ■ For Solaris x64: <pre>set-node galaxy set-cluster 2 link net1 /dev/net/net1 - ether - - link net2 /dev/net/net2 - ether - -</pre> The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the <code>link</code> command. These lines identify the two network cards that the LLT protocol uses. If you configured a low priority link under LLT, the file also includes a "link-lowpri" line. Refer to the <code>llttab(4)</code> manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the <code>llttab</code> file.

Table E-2 lists the GAB configuration files and the information that these files contain.

Table E-2 GAB configuration files

File	Description
/etc/default/gab	This file stores the start and stop environment variables for GAB: ■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: 1—Indicates that GAB is enabled to start up. 0—Indicates that GAB is disabled to start up. ■ GAB_STOP—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include: 1—Indicates that GAB is enabled to shut down. 0—Indicates that GAB is disabled to shut down. The installer sets the value of these variables to 1 at the end of SFHA configuration.
/etc/gabtab	After you install SFHA, the file /etc/gabtab contains a <code>gabconfig (1)</code> command that configures the GAB driver for use. The file /etc/gabtab contains a line that resembles: <pre>/sbin/gabconfig -c -nN</pre> The <code>-c</code> option configures the driver for use. The <code>-nN</code> specifies that the cluster is not formed until at least <i>N</i> nodes are ready to form the cluster. Symantec recommends that you set <i>N</i> to be the total number of nodes in the cluster. Note: Symantec does not recommend the use of the <code>-c -x</code> option for <code>/sbin/gabconfig</code>. Using <code>-c -x</code> can lead to a split-brain condition.

About the AMF configuration files

Asynchronous Monitoring Framework (AMF) kernel driver provides asynchronous event notifications to the VCS agents that are enabled for intelligent resource monitoring.

Table E-3 lists the AMF configuration files.

Table E-3 AMF configuration files

File	Description
/etc/default/amf	This file stores the start and stop environment variables for AMF: ■ AMF_START—Defines the startup behavior for the AMF module after a system reboot or when AMF is attempted to start using the init script. Valid values include: 1—Indicates that AMF is enabled to start up. (default) 0—Indicates that AMF is disabled to start up. ■ AMF_STOP—Defines the shutdown behavior for the AMF module during a system shutdown or when AMF is attempted to stop using the init script. Valid values include: 1—Indicates that AMF is enabled to shut down. (default) 0—Indicates that AMF is disabled to shut down.
/etc/amftab	After you install VCS, the file <code>/etc/amftab</code> contains a <code>amfconfig(1)</code> command that configures the AMF driver for use. The AMF init script uses this <code>/etc/amftab</code> file to configure the AMF driver. The <code>/etc/amftab</code> file contains the following line by default: <pre>/opt/VRTSamf/bin/amfconfig -c</pre>

About the VCS configuration files

VCS configuration files include the following:

- **main.cf**
The installer creates the VCS configuration file in the `/etc/VRTSvcs/conf/config` folder by default during the SFHA configuration. The `main.cf` file contains the minimum information that defines the cluster and its nodes.
See [“Sample main.cf file for global clusters”](#) on page 252.
- **types.cf**
The file `types.cf`, which is listed in the include statement in the `main.cf` file, defines the VCS bundled types for VCS resources. The file `types.cf` is also located in the folder `/etc/VRTSvcs/conf/config`.
Additional files similar to `types.cf` may be present if agents have been added, such as `OracleTypes.cf`.
- **/etc/default/vcs**
This file stores the start and stop environment variables for VCS engine:

- **VCS_START**—Defines the startup behavior for VCS engine after a system reboot. Valid values include:
 - 1—Indicates that VCS engine is enabled to start up.
 - 0—Indicates that VCS engine is disabled to start up.
- **VCS_STOP**—Defines the shutdown behavior for VCS engine during a system shutdown. Valid values include:
 - 1—Indicates that VCS engine is enabled to shut down.
 - 0—Indicates that VCS engine is disabled to shut down.

The installer sets the value of these variables to 1 at the end of SFHA configuration.

Note the following information about the VCS configuration file after installing and configuring VCS:

- The cluster definition includes the cluster information that you provided during the configuration. This definition includes the cluster name, cluster address, and the names of users and administrators of the cluster.
Notice that the cluster has an attribute `UserNames`. The `installsfha` creates a user "admin" whose password is encrypted; the word "password" is the default password.
- If you set up the optional I/O fencing feature for VCS, then the `UseFence = SCSI3` attribute is present.
- If you configured the cluster in secure mode, the `main.cf` includes "SecureClus = 1" cluster attribute.
- The `installsfha` creates the `ClusterService` service group if you configured the virtual IP, SMTP, SNMP, or global cluster options.

The service group also has the following characteristics:

- The group includes the IP and NIC resources.
- The service group also includes the notifier resource configuration, which is based on your input to `installsfha` prompts about notification.
- The `installsfha` also creates a resource dependency tree.
- If you set up global clusters, the `ClusterService` service group contains an Application resource, `wac` (wide-area connector). This resource's attributes contain definitions for controlling the cluster in a global cluster environment.

Refer to the *Veritas Cluster Server Administrator's Guide* for information about managing VCS global clusters.

Refer to the *Veritas Cluster Server Administrator's Guide* to review the configuration concepts, and descriptions of main.cf and types.cf files for Solaris systems.

Sample main.cf file for VCS clusters

The following sample main.cf file is for a three-node cluster in secure mode.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"


cluster vcs02 (
    SecureClus = 1
)


system sysA (
)


system sysB (
)


system sysC (
)


group ClusterService (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)


NIC csgnic (
    Device = net0
    NetworkHosts = { "10.182.13.1" }
)


NotifierMngr ntfr (
    SnmpConsoles = { jupiter" = SevereError }
    SmtServer = "smtp.example.com"
    SmtRecipients = { "ozzie@example.com" = SevereError }
```

```

    )

ntfr requires csgnic

// resource dependency tree
//
//     group ClusterService
//     {
//         NotifierMngr ntfr
//         {
//             NIC csgnic
//         }
//     }

```

Sample main.cf file for global clusters

If you installed SFHA with the Global Cluster option, note that the ClusterService group also contains the Application resource, wac. The wac resource is required to control the cluster in a global cluster environment.

In the following main.cf file example, bold text highlights global cluster specific entries.

```

include "types.cf"

cluster vcs03 (
    ClusterAddress = "10.182.13.50"
    SecureClus = 1
)

system sysA (

)

system sysB (

)

system sysC (

)

group ClusterService (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
)

```

```
OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
)

IP gcoip (
    Device = net0
    Address = "10.182.13.50"
    NetMask = "255.255.240.0"
)

NIC csgnic (
    Device = net0
    NetworkHosts = { "10.182.13.1" }
)

NotifierMngr ntfr (
    SnmpConsoles = { jupiter = SevereError }
    Smtperver = "smtp.example.com"
    Smtpercipients = { "ozzie@example.com" = SevereError }
)

gcoip requires csgnic
ntfr requires csgnic
wac requires gcoip

// resource dependency tree
//
//     group ClusterService
//     {
//         NotifierMngr ntfr
//         {
//             NIC csgnic
//         }
//     }
//     Application wac
//     {
//         IP gcoip
```

```
//      {
//      NIC csgnic
//      }
//      }
//      }
```

About I/O fencing configuration files

Table E-4 lists the I/O fencing configuration files.

Table E-4 I/O fencing configuration files

File	Description
/etc/default/vxfen	This file stores the start and stop environment variables for I/O fencing: ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. The installer sets the value of these variables to 1 at the end of SFHA configuration.
/etc/vxfendg	This file includes the coordinator disk group information. This file is not applicable for server-based fencing.

Table E-4 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfenmode	This file contains the following parameters: ■ vxfen_mode ■ scsi3—For disk-based fencing ■ customized—For server-based fencing ■ disabled—To run the I/O fencing driver but not do any fencing operations. ■ vxfen_mechanism This parameter is applicable only for server-based fencing. Set the value as cps. ■ scsi3_disk_policy ■ dmp—Configure the vxfen module to use DMP devices The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp. ■ raw—Configure the vxfen module to use the underlying raw character devices Note: You must use the same SCSI-3 disk policy on all the nodes. ■ security This parameter is applicable only for server-based fencing. 1—Indicates that communication with the CP server is in secure mode. This setting is the default. 0—Indicates that communication with the CP server is in non-secure mode. ■ List of coordination points This list is required only for server-based fencing configuration. Coordination points in 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. Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points and multiple IP addresses for each CP server. ■ single_cp This parameter is applicable for server-based fencing which uses a single highly available CP server as its coordination point. Also applicable for when you use a coordinator disk group with single disk. ■ autoseed_gab_timeout This parameter enables GAB automatic seeding of the cluster even when some cluster nodes are unavailable. This feature requires that I/O fencing is enabled. 0—Turns the GAB auto-seed feature on. Any value greater than 0 indicates the number of seconds that GAB must delay before it automatically seeds the cluster. -1—Turns the GAB auto-seed feature off. This setting is the default.

Table E-4 I/O fencing configuration files (continued)

File	Description
/etc/vxfentab	When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points. Note: The /etc/vxfentab file is a generated file; do not modify this file. For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows: ■ Raw disk: <pre>/dev/rdisk/c1t1d0s2 /dev/rdisk/c2t1d0s2 /dev/rdisk/c3t1d2s2</pre> ■ DMP disk: <pre>/dev/vx/rdmp/c1t1d0s2 /dev/vx/rdmp/c2t1d0s2 /dev/vx/rdmp/c3t1d0s2</pre> For server-based fencing, the /etc/vxfentab file also includes the security settings information. For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.

Sample configuration files for CP server

The /etc/vxcps.conf file determines the configuration of the coordination point server (CP server.)

See “[Sample CP server configuration \(/etc/vxcps.conf\) file output](#)” on page 262.

The following are example main.cf files for a CP server that is hosted on a single node, and a CP server that is hosted on an SFHA cluster.

- The main.cf file for a CP server that is hosted on a single node:
See “[Sample main.cf file for CP server hosted on a single node that runs VCS](#)” on page 257.
- The main.cf file for a CP server that is hosted on an SFHA cluster:
See “[Sample main.cf file for CP server hosted on a two-node SFHA cluster](#)” on page 259.

Note: The CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with SFHA clusters (application clusters). The example main.cf files use IPv4 addresses.

Sample main.cf file for CP server hosted on a single node that runs VCS

The following is an example of a single CP server node main.cf.

For this CP server single node main.cf, note the following values:

- Cluster name: cps1
- Node name: mycps1

```
include "types.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"

// cluster name: cps1
// CP server: mycps1

cluster cps1 (
    UserNames = { admin = bMNfMHmJNiNNlVNhMK, haris = fopKojNvpHouNn,
                  "mycps1.symantecexample.com@root@vx" = aj,
                  "root@mycps1.symantecexample.com" = hq }
    Administrators = { admin, haris,
                       "mycps1.symantecexample.com@root@vx",
                       "root@mycps1.symantecexample.com" }
    SecureClus = 1
    HacliUserLevel = COMMANDROOT
)

system mycps1 (
)

group CPSSG (
    SystemList = { mycps1 = 0 }
    AutoStartList = { mycps1 }
)

IP cpsvip1 (
    Critical = 0
    Device @mycps1 = net0
    Address = "10.209.3.1"
```

```
        NetMask = "255.255.252.0"
    )

    IP cpsvip2 (
        Critical = 0
        Device @mycps1 = net1
        Address = "10.209.3.2"
        NetMask = "255.255.252.0"
    )

    NIC cpsnic1 (
        Critical = 0
        Device @mycps1 = net0
        PingOptimize = 0
        NetworkHosts @mycps1 = { "10.209.3.10" }
    )

    NIC cpsnic2 (
        Critical = 0
        Device @mycps1 = net1
        PingOptimize = 0
    )

    Process vxcpserve (
        PathName = "/opt/VRTScps/bin/vxcpserve"
        ConfInterval = 30
        RestartLimit = 3
    )

    Quorum quorum (
        QuorumResources = { cpsvip1, cpsvip2 }
    )

cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
vxcpserve requires quorum

// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1
```

```
//      {
//      NIC cpsnic1
//      }
// IP cpsvip2
//      {
//      NIC cpsnic2
//      }
// Process vxcpserve
//      {
//      Quorum quorum
//      }
// }
```

Sample main.cf file for CP server hosted on a two-node SFHA cluster

The following is an example of a main.cf, where the CP server is hosted on an SFHA cluster.

For this CP server hosted on an SFHA cluster main.cf, note the following values:

- Cluster name: cps1
- Nodes in the cluster: mycps1, mycps2

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTTypes.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"

// cluster: cps1
// CP servers:
// mycps1
// mycps2

cluster cps1 (
    UserNames = { admin = ajkCjeJgkFkkIskEjh,
                  "mycps1.symantecexample.com@root@vx" = JK,
                  "mycps2.symantecexample.com@root@vx" = dl }
    Administrators = { admin, "mycps1.symantecexample.com@root@vx",
                       "mycps2.symantecexample.com@root@vx" }
    SecureClus = 1
)

system mycps1 (
```

```
    )

system mycps2 (
)

group CPSSG (
    SystemList = { mycps1 = 0, mycps2 = 1 }
    AutoStartList = { mycps1, mycps2 } )

DiskGroup cpsdg (
    DiskGroup = cps_dg
)

IP cpsvip1 (
    Critical = 0
    Device @mycps1 = net0
    Device @mycps2 = net0
    Address = "10.209.81.88"
    NetMask = "255.255.252.0"
)

IP cpsvip2 (
    Critical = 0
    Device @mycps1 = net1
    Device @mycps2 = net1
    Address = "10.209.81.89"
    NetMask = "255.255.252.0"
)

Mount cpsmount (
    MountPoint = "/etc/VRTScps/db"
    BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
    FSType = vxfs
    FsckOpt = "-y"
)

NIC cpsnic1 (
    Critical = 0
    Device @mycps1 = net0
    Device @mycps2 = net0
    PingOptimize = 0
    NetworkHosts @mycps1 = { "10.209.81.10" }
)
```

```

NIC cpsnic2 (
    Critical = 0
    Device @mycps1 = net1
    Device @mycps2 = net1
    PingOptimize = 0
)

Process vxcperv (
    PathName = "/opt/VRTScps/bin/vxcperv"
)

Quorum quorum (
    QuorumResources = { cpsvip1, cpsvip2 }
)

Volume cpsvol (
    Volume = cps_volume
    DiskGroup = cps_dg
)

cpismount requires cpsvol
cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
cpsvol requires cpsdg
vxcperv requires cpismount
vxcperv requires quorum

// resource dependency tree
//
// group CPSSG
// {
//   IP cpsvip1
//   {
//     NIC cpsnic1
//   }
//   IP cpsvip2
//   {
//     NIC cpsnic2
//   }
//   Process vxcperv
//   {

```

```
//      Quorum quorum
//      Mount cpsmount
//      {
//          Volume cpsvol
//          {
//              DiskGroup cpsdg
//          }
//      }
//  }
// }
```

Sample CP server configuration (/etc/vxcps.conf) file output

The following is an example of a coordination point server (CP server) configuration file `/etc/vxcps.conf` output.

```
## The vxcps.conf file determines the
## configuration for Veritas CP Server.
cps_name=mycps1
vip=[10.209.81.88]
vip=[10.209.81.89]:56789
port=14250
security=1
db=/etc/VRTScps/db
```

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

- [About configuring secure shell or remote shell communication modes before installing products](#)
- [Manually configuring and passwordless ssh](#)
- [Restarting the ssh session](#)
- [Enabling and disabling rsh for Solaris](#)

About configuring secure shell or remote shell communication modes before installing products

Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installer is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install Veritas software.

You can install products to remote systems using either secure shell (`ssh`) or remote shell (`rsh`). Symantec recommends that you use `ssh` as it is more secure than `rsh`.

This section contains an example of how to set up `ssh` password free communication. The example sets up `ssh` between a source system (`system1`) that

contains the installation directories, and a target system (system2). This procedure also applies to multiple target systems.


Manually configuring and passwordless ssh

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

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

Figure F-1 illustrates this procedure.

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



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

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

To create the DSA key pair

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

```
system1 # cd /
```

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

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

System output similar to the following is displayed:

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

- 3 Press Enter to accept the default location of `/.ssh/id_dsa`.
- 4 When the program asks you to enter the passphrase, press the Enter key twice.

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

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

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

```
system2 # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

```
system2 # chmod go-w /.ssh
```

To append the public key from the source system to the `authorized_keys` file on the target system, using secure file transfer

- 1** Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems (system2 in this example).

To enable SFTP, the `/etc/ssh/sshd_config` file must contain the following two lines:

```
PermitRootLogin          yes
Subsystem                sftp          /usr/lib/ssh/sftp-server
```

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

To restart ssh on Solaris 10, type the following command:

```
system1 # svcadm restart ssh
```

- 3** From the source system (system1), move the public key to a temporary file on the target system (system2).

Use the secure file transfer program.

In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

Use the following command for secure file transfer:

```
system1 # sftp system2
```

If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

```
Connecting to system2 ...
The authenticity of host 'system2 (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
Are you sure you want to continue connecting (yes/no)?
```

- 4** Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'system2,10.182.00.00'
(DSA) to the list of known hosts.
root@system2 password:
```

- 5** Enter the root password of system2.

- 6** At the `sftp` prompt, type the following command:

```
sftp> put /.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /.ssh/id_dsa.pub to /id_dsa.pub
```

- 7** To quit the SFTP session, type the following command:

```
sftp> quit
```

- 8** To begin the `ssh` session on the target system (system2 in this example), type the following command on system1:

```
system1 # ssh system2
```

Enter the root password of system2 at the prompt:

```
password:
```

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

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

- 10** After the `id_dsa.pub` public key file is copied to the target system (system2), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, enter the following command on system2:

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

- 11** To log out of the `ssh` session, enter the following command:

```
system2 # exit
```

- 12 When you install from a source system that is also an installation target, also add the local system `id_dsa.pub` key to the local `authorized_keys` file. The installation can fail if the installation source system is not authenticated.

To add the local system `id_dsa.pub` key to the local `authorized_keys` file, enter the following command:

```
system1 # cat /.ssh/id_dsa.pub >> /.ssh/authorized_keys
```

- 13 Run the following commands on the source installation system. If your ssh session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and make the key globally available to the user `root`:

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

```
Identity added: //ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

To verify that you can connect to a target system

- 1 On the source system (system1), enter the following command:

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

where `system2` is the name of the target system.

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

Restarting the ssh session

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

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

To restart ssh

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

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

- 2 Make the key globally available for the user `root`

```
system1 # ssh-add
```

Enabling and disabling rsh for Solaris

The following section describes how to enable remote shell on Solaris system.

Veritas recommends configuring a secure shell environment for Veritas product installations.

See [“Manually configuring and passwordless ssh”](#) on page 264.

See the operating system documentation for more information on configuring remote shell.

To enable rsh

- 1 To determine the current status of `rsh` and `rlogin`, type the following command:

```
# inetadm | grep -i login
```

If the service is enabled, the following line is displayed:

```
enabled online svc:/network/login:rlogin
```

If the service is not enabled, the following line is displayed:

```
disabled disabled svc:/network/login:rlogin
```

- 2 To enable a disabled `rsh/rlogin` service, type the following command:

```
# inetadm -e rlogin
```

- 3 To disable an enabled `rsh/rlogin` service, type the following command:

```
# inetadm -d rlogin
```

- 4 Modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. This file must be modified for each user who remotely accesses the system using rsh. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system having access to the local system. For example, if the root user must remotely access `system1` from `system2`, you must add an entry for `system2.companyname.com` in the `.rhosts` file on `system1`.

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

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

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

Storage Foundation and High Availability components

This appendix includes the following topics:

- [Storage Foundation and High Availability installation packages](#)
- [Veritas Cluster Server installation packages](#)
- [Chinese language packages](#)
- [Japanese language packages](#)
- [Veritas Storage Foundation obsolete and reorganized installation packages](#)

Storage Foundation and High Availability installation packages

[Table G-1](#) shows the package name and contents for each English language package for Storage Foundation and High Availability. The table also gives you guidelines for which packages to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and High Availability and Veritas Cluster Server (VCS) packages, the combined functionality is called Storage Foundation and High Availability and High Availability.

See [“Veritas Cluster Server installation packages”](#) on page 274.

Table G-1 Storage Foundation and High Availability packages

packages	Contents	Configuration
VRTSaslapm	Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries Required for the support and compatibility of various storage arrays.	Minimum
VRTSperl	Perl 5.12.2 for Veritas	Minimum
VRTSvlic	Veritas License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Minimum
VRTSvxfs	Veritas File System binaries Required for VxFS file system support.	Minimum
VRTSvxvm	Veritas Volume Manager binaries, scripts, and utilities. Required for VxVM volume manager support.	Minimum
VRTSdbed	Veritas Storage Foundation for Databases	Recommended
VRTSodm	Veritas ODM Driver for VxFS Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager allows Oracle 9i and 10g to improve performance and manage system bandwidth.	Recommended

Table G-1 Storage Foundation and High Availability packages (*continued*)

packages	Contents	Configuration
VRTSsfcp60	Veritas Storage Foundation Common Product Installer The Storage Foundation Common Product installer package contains the installer libraries and product scripts that perform the following: ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ removing nodes ■ etc. You can use these script to simplify the native operating system installations, configurations, and upgrades.	Minimum
VRTSsfmh	Veritas Storage Foundation Managed Host 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:	Recommended
VRTSspt	Veritas Software Support Tools	Recommended
VRTSfsadv	Minimum Veritas File System Advanced Solutions by Symantec (Solaris SPARC only).	Minimum
VRTSfssdk	Veritas File System Software Developer Kit 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 packages

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

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

Table G-2 VCS installation packages

package	Contents	Configuration
VRTSgab	Veritas Cluster Server group membership and atomic broadcast services	Minimum
VRTSllt	Veritas Cluster Server low-latency transport	Minimum
VRTSamf	Veritas Cluster Server Asynchronous Monitoring Framework	Minimum
VRTSvcs	Veritas Cluster Server	Minimum
VRTSvcsag	Veritas Cluster Server Bundled Agents	Minimum
VRTSvxfen	Veritas I/O Fencing	Minimum
VRTSvcsea	Consolidated database and enterprise agent packages	Recommended
VRTScps	Veritas Coordination Point Server The Coordination Point Server is an alternate mechanism for I/O fencing. It implements I/O fencing through a client/server architecture and can provide I/O fencing for multiple VCS clusters.	All

Chinese language packages

The following table shows the package name and contents for each Chinese language package.

Table G-3 Chinese language packages

package	Contents
VRTSzhvm	Chinese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages

Japanese language packages

The following table show the package name and contents for each Japanese language package.

Table G-4 Japanese language packages

package	Contents
VRTSjacav	Japanese Veritas Cluster Server Agents for Storage Foundation Cluster File System – Manual Pages and Message Catalogs by Symantec
VRTSjacs	Veritas Cluster Server Japanese Message Catalogs by Symantec
VRTSjacse	Japanese Veritas High Availability Enterprise Agents by Symantec
VRTSjadba	Japanese Veritas Oracle Real Application Cluster Support package by Symantec
VRTSjadbe	Japanese Veritas Storage Foundation for Oracle from Symantec – Message Catalogs
VRTSjafs	Japanese Veritas File System – Message Catalog and Manual Pages
VRTSjaodm	Veritas Oracle Disk Manager Japanese Message Catalog and Manual Pages by Symantec
VRTSjavm	Japanese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages
VRTSmulic	Multi-language Symantec License Utilities

Veritas Storage Foundation obsolete and reorganized installation packages

Table G-5 lists the packages that are obsolete or reorganized for Storage Foundation and High Availability.

Table G-5 Veritas Storage Foundation obsolete and reorganized packages

package	Description
Obsolete and reorganized for 6.0	
VRTSat	Obsolete
VRTSatZH	Obsolete
VRTSatJA	Obsolete
Obsolete and reorganized for 5.1	
Infrastructure	
SYMClma	Obsolete
VRTSaa	Included in VRTSsfmh
VRTSccg	Included in VRTSsfmh
VRTSdbms3	Obsolete
VRTSicsco	Obsolete
VRTSjre	Obsolete
VRTSjre15	Obsolete
VRTSmh	Included in VRTSsfmh
VRTSpxb	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product packages	
VRTSacclib	Obsolete The following information is for installations, upgrades, and uninstallations using the script-based installer. ■ For fresh installations VRTSacclib is not installed. ■ For upgrades, the existing VRTSacclib is uninstalled and a new VRTSacclib is installed. ■ For uninstallation, VRTSacclib is not uninstalled.

Table G-5 Veritas Storage Foundation obsolete and reorganized packages
(continued)

package	Description
VRTSalloc	Obsolete
VRTScmccc	Obsolete
VRTScmcm	Obsolete
VRTScmcs	Obsolete
VRTScscm	Obsolete
VRTScscw	Obsolete
VRTScsocw	Obsolete
VRTScssim	Obsolete
VRTScutil	Obsolete
VRTSd2gui	Included in VRTSdbed
VRTSdb2ed	Included in VRTSdbed
VRTSdbcom	Included in VRTSdbed
VRTSdbed	Included in VRTSdbed
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSdsa	Obsolete
VRTSfas	Obsolete
VRTSfasag	Obsolete
VRTSfsman	Included in the product's main package.
VRTSfsmnd	Included in the product's main package.
VRTSfspro	Included in VRTSsfmh
VRTSgapms	Obsolete
VRTSmapro	Included in VRTSsfmh
VRTSorgui	Obsolete

Table G-5

Veritas Storage Foundation obsolete and reorganized packages

(continued)

package	Description
VRTSsybed	Included in VRTSdbed
VRTSvail	Obsolete
VRTSvcldb	Included in VRTSvcsea
VRTSvcsmn	Included in VRTSvcsc
VRTSvcscr	Included in VRTSvcsea
VRTSvcssy	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvcsc
VRTSvdid	Obsolete
VRTSvmman	Included in the product's main package.
VRTSvmpro	Included in VRTSsfmh
VRTSvrw	Obsolete
VRTSvxmsa	Obsolete
Documentation	All Documentation packages obsolete

Troubleshooting installation issues

This appendix includes the following topics:

- [Restarting the installer after a failed connection](#)
- [What to do if you see a licensing reminder](#)
- [Troubleshooting information](#)
- [Incorrect permissions for root on remote system](#)
- [Inaccessible system](#)

Restarting the installer after a failed connection

If an installation is killed because of a failed connection, you can restart the installer to resume the installation. The installer detects the existing installation. The installer prompts you whether you want to resume the installation. If you resume the installation, the installation proceeds from the point where the installation failed.

What to do if you see a licensing reminder

In this release, you can install without a license key. In order to comply with the End User License Agreement, you must either install a license key or make the host managed by a Management Server. If you do not comply with these terms within 60 days, the following warning messages result:

```
WARNING V-365-1-1 This host is not entitled to run Veritas Storage  
Foundation/Veritas Cluster Server.As set forth in the End User
```

License Agreement (EULA) you must complete one of the two options set forth below. To comply with this condition of the EULA and stop logging of this message, you have <nn> days to either:

- make this host managed by a Management Server (see <http://go.symantec.com/sfhakeyless> for details and free download), or
- add a valid license key matching the functionality in use on this host using the command 'vxlicinst'

To comply with the terms of the EULA, and remove these messages, you must do one of the following within 60 days:

- Install a valid license key corresponding to the functionality in use on the host. After you install the license key, you must validate the license key using the following command:

```
# /opt/VRTS/bin/vxkeyless
```

- Continue with keyless licensing by managing the server or cluster with a management server.

For more information about keyless licensing, see the following URL:
<http://go.symantec.com/sfhakeyless>

Troubleshooting information

The VRTSpt package provides a group of tools for troubleshooting a system and collecting information on its configuration. The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. Although the tools are not required for the operation of any Veritas product, Symantec recommends installing them should a support case be needed to be opened with Symantec Support. If you are unfamiliar with their use and purpose, use caution when using them or use them in concert with Symantec Support.

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Failed to setup rsh communication on 10.198.89.241:
'rsh 10.198.89.241 <command>' failed
Trying to setup ssh communication on 10.198.89.241.
```



```
Failed to setup ssh communication on 10.198.89.241:
Login denied
```

```
Failed to login to remote system(s) 10.198.89.241.
Please make sure the password(s) are correct and superuser(root)
can login to the remote system(s) with the password(s).
If you want to setup rsh on remote system(s), please make sure
rsh with command argument ('rsh <host> <command>') is not
denied by remote system(s).
```

```
Either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

```
Would you like the installer to setup ssh/rsh communication
automatically between the nodes?
Superuser passwords for the systems will be asked. [y,n,q] (y) n
```

```
System verification did not complete successfully
```

```
The following errors were discovered on the systems:
```

```
The ssh permission denied on 10.198.89.241
rsh exited 1 on 10.198.89.241
either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

Suggested solution: You need to set up the systems to allow remote access using ssh **OR** rsh.

Note: Remove remote shell permissions after completing the SFHA installation and configuration.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
Verifying systems: 12% .....
Estimated time remaining: 0:10 1 of 8
Checking system communication ..... Done
System verification did not complete successfully
```

```
The following errors were discovered on the systems:  
cannot resolve hostname host1  
Enter the system names separated by spaces: q,? (host1)
```

Suggested solution: Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

Troubleshooting cluster installation

This appendix includes the following topics:

- [Unmount failures](#)
- [Command failures](#)
- [Installer cannot create UUID for the cluster](#)
- [The vxfsentsthdw utility fails when SCSI TEST UNIT READY command fails](#)
- [Troubleshooting CP server](#)
- [Troubleshooting server-based fencing on the SFHA cluster nodes](#)
- [Issues during online migration of coordination points](#)

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.

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 54.

- 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.

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 SFHA configuration, upgrade, or when you add a node to the cluster using the installer.

Workaround: To start SFHA, you must run the `uuidconfig.pl` script manually to configure the UUID on each cluster node.

To configure the cluster UUID when you create a cluster manually

- ◆ On one node in the cluster, perform the following command to populate the cluster UUID on each node in the cluster.

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -configure nodeA  
nodeB ... nodeN
```

Where `nodeA`, `nodeB`, through `nodeN` are the names of the cluster nodes.

The `vxfcntlsthdw` utility fails when SCSI TEST UNIT READY command fails

While running the `vxfcntlsthdw` utility, you may see a message that resembles as follows:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node  
FAILED.
```

Contact the storage provider to have the hardware configuration fixed.

The disk array does not support returning success for a `SCSI TEST UNIT READY` command when another host has the disk reserved using SCSI-3 persistent reservations. This happens with the Hitachi Data Systems 99XX arrays if bit 186 of the system mode option is not enabled.

Troubleshooting CP server

All CP server operations and messages are logged in the `/var/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 system that hosts the CP server.

The following files contain logs and text files that may be useful in understanding and troubleshooting a CP server:

- `/var/VRTScps/log/cpsrvr_[ABC].log`
- `/var/VRTSvcs/log/vcsauthserver.log` (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.

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs that may be useful in understanding and troubleshooting fencing-related issues on a SF HA cluster (client cluster) node.

See [“Troubleshooting issues related to the CP server service group”](#) on page 286.

See [“Checking the connectivity of CP server”](#) on page 286.

See [“Issues during fencing startup on SF HA cluster nodes set up for server-based fencing”](#) on page 287.

See [“Issues during online migration of coordination points”](#) on page 287.

Troubleshooting issues related to the CP server service group

If you cannot bring up the CPSSG service group after the CP server configuration, perform the following steps:

- Verify that the CPSSG service group and its resources are valid and properly configured in the VCS configuration.
- Check the VCS engine log (`/var/VRTSvcs/log/engine_[ABC].log`) to see if any of the CPSSG service group resources are **FAULTED**.
- Review the sample dependency graphs to make sure the required resources are configured correctly.

Checking the connectivity of CP server

You can test the connectivity of CP server using the `cpsadm` command.

You must have set the environment variables `CPS_USERNAME` and `CPS_DOMAINTYPE` to run the `cpsadm` command on the SF HA cluster (client cluster) nodes.

To check the connectivity of CP server

- ◆ Run the following command to check whether a CP server is up and running at a process level:

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

where `cp_server` is the virtual IP address or virtual hostname on which the CP server is listening.

Troubleshooting server-based fencing on the SFHA cluster nodes

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs files that may be useful in understanding and troubleshooting fencing-related issues on a SFHA cluster (application cluster) node.

Issues during fencing startup on SF HA cluster nodes set up for server-based fencing

Table I-1 Fencing startup issues on SF HA cluster (client cluster) nodes

Issue	Description and resolution
<code>cpsadm</code> command on the SF HA cluster gives connection error	If you receive a connection error message after issuing the <code>cpsadm</code> command on the SF HA cluster, perform the following actions: ■ Ensure that the CP server is reachable from all the SF HA cluster nodes. ■ Check that the SF HA cluster nodes use the correct CP server virtual IP or virtual hostname and the correct port number. Check the <code>/etc/vxfenmode</code> file. ■ Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.
Authorization failure	Authorization failure occurs when the CP server's nodes or users are not added in the CP server configuration. Therefore, fencing on the SF HA cluster (client cluster) node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points. To resolve this issue, add the CP server node and user in the CP server configuration and restart fencing. See “Preparing the CP servers manually for use by the SF HA cluster” on page 123.
Authentication failure	If you had configured secure communication between the CP server and the SF HA cluster (client cluster) nodes, authentication failure can occur due to the following causes: ■ Symantec Product Authentication Services (AT) is not properly configured on the CP server and/or the SF HA cluster. ■ The CP server and the SF HA cluster nodes use different root brokers, and trust is not established between the authentication brokers:

Issues during online migration of coordination points

During online migration of coordination points using the `vxfenswap` utility, the operation is automatically rolled back if a failure is encountered during validation of coordination points from any of the cluster nodes.

Validation failure of the new set of coordination points can occur in the following circumstances:

- The `/etc/vxfenmode` file is not updated on all the SF HA 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 SF HA 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 SF HA cluster nodes to the CP server(s).
- Cluster, nodes, or users for the SF HA cluster nodes have not been added on the new CP servers, thereby causing authorization failure.

Vxfen service group activity after issuing the `vxfenswap` command

The Coordination Point agent reads the details of coordination points from the `vxfenconfig -l` output and starts monitoring the registrations on them.

Thus, during `vxfenswap`, when the `vxfenmode` file is being changed by the user, the Coordination Point agent does not move to `FAULTED` state but continues monitoring the old set of coordination points.

As long as the changes to `vxfenmode` file are not committed or the new set of coordination points are not reflected in `vxfenconfig -l` output, the Coordination Point agent continues monitoring the old set of coordination points it read from `vxfenconfig -l` output in every monitor cycle.

The status of the Coordination Point agent (either `ONLINE` or `FAULTED`) depends upon the accessibility of the coordination points, the registrations on these coordination points, and the fault tolerance value.

When the changes to `vxfenmode` file are committed and reflected in the `vxfenconfig -l` output, then the Coordination Point agent reads the new set of coordination points and proceeds to monitor them in its new monitor cycle.

Sample SF HA 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:
[Figure J-1](#)
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
[Figure J-2](#)
- 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

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.


Client cluster served by highly available CPS and 2 SCSI-3 disks

[Figure J-1](#) displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to customized with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfencoorddg`. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure J-1 Client cluster served by highly available CP server and 2 SCSI-3 disks




Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure J-2 displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks (one from each site) are part of disk group vxencoorddg.
The third coordination point is a CP server on a single node VCS cluster.

Figure J-2 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

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfencoorddg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Reconciling major/minor numbers for NFS shared disks

This appendix includes the following topics:

- [Reconciling major/minor numbers for NFS shared disks](#)

Reconciling major/minor numbers for NFS shared disks

Your configuration may include disks on the shared bus that support NFS. You can configure the NFS file systems that you export on disk partitions or on Veritas Volume Manager volumes.

An example disk partition name is `/dev/dsk/c1t1d0s2`.

An example volume name is `/dev/vx/dsk/shareddg/vol3`. Each name represents the block device on which the file system is to be mounted.

In a VCS cluster, block devices providing NFS service must have the same major and minor numbers on each cluster node. Major numbers identify required device drivers (such as a Solaris partition or a VxVM volume). Minor numbers identify the specific devices themselves. NFS also uses major and minor numbers to identify the exported file system.

Major and minor numbers must be verified to ensure that the NFS identity for the file system is the same when exported from each node.

Checking major and minor numbers for disk partitions

The following sections describe checking and changing, if necessary, the major and minor numbers for disk partitions used by cluster nodes.

To check major and minor numbers on disk partitions

- ◆ Use the following command on all nodes exporting an NFS file system. This command displays the major and minor numbers for the block device.

```
# ls -lL block_device
```

The variable *block_device* refers to a partition where a file system is mounted for export by NFS. Use this command on each NFS file system. For example, type:

```
# ls -lL /dev/dsk/c1t1d0s2
```

Output on Node A resembles:

```
crw-r----- 1 root sys 32,1 Dec 3 11:50 /dev/dsk/c1t1d0s2
```

Output on Node B resembles:

```
crw-r----- 1 root sys 32,1 Dec 3 11:55 /dev/dsk/c1t1d0s2
```

Note that the major numbers (32) and the minor numbers (1) match, satisfactorily meeting the requirement for NFS file systems.

To reconcile the major numbers that do not match on disk partitions

- 1 Reconcile the major and minor numbers, if required. For example, if the output in the previous section resembles the following, perform the instructions beginning step 2:

Output on Node A:

```
crw-r----- 1 root sys 32,1 Dec 3 11:50 /dev/dsk/c1t1d0s2
```

Output on Node B:

```
crw-r----- 1 root sys 36,1 Dec 3 11:55 /dev/dsk/c1t1d0s2
```

- 2 Place the VCS command directory in your path.

```
# export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin
```


- 3 Attempt to change the major number on System B (now 36) to match that of System A (32). Use the command:

```
# haremajor -sd major_number
```

For example, on Node B, enter:

```
# haremajor -sd 32
```

- 4 If the command succeeds, go to step 8.
- 5 If the command fails, you may see a message resembling:

```
Error: Preexisting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

- 6 Notice that the number 36 (the major number on Node A) is not available on Node B. Run the `haremajor` command on Node B and change it to 128,

```
# haremajor -sd 128
```

- 7 Run the same command on Node A. If the command fails on Node A, the output lists the available numbers. Rerun the command on both nodes, setting the major number to one available to both.
- 8 Reboot each system on which the command succeeds.
- 9 Proceed to reconcile the major numbers for your next partition.

To reconcile the minor numbers that do not match on disk partitions

- 1 In the example, the minor numbers are 1 and 3 and are reconciled by setting to 30 on each node.
- 2 Type the following command on both nodes using the name of the block device:

```
# ls -l /dev/dsk/c1t1d0s2
```

Output from this command resembles the following on Node A:

```
lrwxrwxrwx 1 root  root  83 Dec 3 11:50
/dev/dsk/c1t1d0s2      -> ../../
devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:d,raw
```

The device `name` (in bold) includes the slash following the word `devices`, and continues to, but does not include, the colon.

- 3 Type the following command on both nodes to determine the instance numbers that the SCSI driver uses:

```
# grep sd /etc/path_to_inst | sort -n -k 2,2
```

Output from this command resembles the following on Node A:

```
"/sbus@1f,0/QLGC,isp@0,10000/sd@0,0" 0 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@1,0" 1 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@2,0" 2 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@3,0" 3 "sd"
.
.
"/sbus@1f,0/SUNW,fas@e,8800000/sd@d,0" 27 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@e,0" 28 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@f,0" 29 "sd"
```

In the output, the instance numbers are in the second field.

The instance number that is associated with the device name that matches the name for Node A displayed in step 2, is "1."

- 4 Compare instance numbers for the device in the output on each node.

After you review the instance numbers, perform one of the following tasks:

- If the instance number from one node is unused on the other— it does not appear in the output of step 3—edit `/etc/path_to_inst`.
You edit this file to make the second node's instance number similar to the number of the first node.
- If the instance numbers in use on both nodes, edit `/etc/path_to_inst` on both nodes. Change the instance number that is associated with the device name to an unused number. The number needs to be greater than the highest number that other devices use. For example, the output of step 3 shows the instance numbers that all devices use (from 0 to 29). You edit the file `/etc/path_to_inst` on each node and reset the instance numbers to 30.

- 5 Type the following command to reboot each node on which `/etc/path_to_inst` was modified:

```
# reboot -- -rv
```

Checking the major and minor number for VxVM volumes

The following sections describe checking and changing, if necessary, the major and minor numbers for the VxVM volumes that cluster systems use.

To check major and minor numbers on VxVM volumes

- 1 Place the VCS command directory in your path. For example:

```
# export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin
```

- 2 To list the devices, use the `ls -lL block_device` command on each node:

```
# ls -lL /dev/vx/dsk/shareddg/vol3
```

On Node A, the output may resemble:

```
brw----- 1 root root 32,43000 Mar 22 16:4 1
/dev/vx/dsk/shareddg/vol3
```

On Node B, the output may resemble:

```
brw----- 1 root root 36,43000 Mar 22 16:4 1
/dev/vx/dsk/shareddg/vol3
```

- 3 Import the associated shared disk group on each node.

- 4 Use the following command on each node exporting an NFS file system. The command displays the major numbers for `vxio` and `vxspec` that Veritas Volume Manager uses. Note that other major numbers are also displayed, but only `vxio` and `vxspec` are of concern for reconciliation:

```
# grep vx /etc/name_to_major
```

Output on Node A:

```
vxdump 30
vxio 32
vxspec 33
vxfen 87
vxglm 91
```

Output on Node B:

```
vxdump 30
vxio 36
vxspec 37
vxfen 87
vxglm 91
```

- 5 To change Node B's major numbers for `vxio` and `vxspec` to match those of Node A, use the command:

```
haremajor -vx major_number_vxio major_number_vxspec
```

For example, enter:

```
# haremajor -vx 32 33
```

If the command succeeds, proceed to step 8. If this command fails, you receive a report similar to the following:

```
Error: Preexisting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

- 6 If you receive this report, use the `haremajor` command on Node A to change the major number (32/33) to match that of Node B (36/37). For example, enter:

```
# haremajor -vx 36 37
```

If the command fails again, you receive a report similar to the following:

```
Error: Preexisting major number 36
These are available numbers on this node: 126...
Check /etc/name_to_major on all systems for
available numbers.
```

- 7 If you receive the second report, choose the larger of the two available numbers (in this example, 128). Use this number in the `haremajor` command to reconcile the major numbers. Type the following command on both nodes:

```
# haremajor -vx 128 129
```

- 8 Reboot each node on which `haremajor` was successful.
- 9 If the minor numbers match, proceed to reconcile the major and minor numbers of your next NFS block device.
- 10 If the block device on which the minor number does not match is a volume, consult the `vxvg(1M)` manual page. The manual page provides instructions on reconciling the Veritas Volume Manager minor numbers, and gives specific reference to the `reminor` option.

Node where the vxio driver number have been changed require rebooting.

Configuring LLT over UDP using IPv4

This appendix includes the following topics:

- [Using the UDP layer for LLT](#)
- [Manually configuring LLT over UDP using IPv4](#)

Using the UDP layer for LLT

SFHA provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

LLT over UDP is slower than LLT over Ethernet. Use LLT over UDP only when the hardware configuration makes it necessary.

Manually configuring LLT over UDP using IPv4

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on separate subnets. Set the broadcast address in `/etc/llttab` explicitly depending on the subnet for each link. See [“Broadcast address in the `/etc/llttab` file”](#) on page 304.

- 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 306.
- Set the broadcast address correctly for direct-attached (non-routed) links.
See [“Sample configuration: direct-attached links”](#) on page 308.
- 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 309.

Broadcast address in the `/etc/llttab` file

The broadcast address is set explicitly for each link in the following example.

- Display the content of the `/etc/llttab` file on the first node galaxy:

```
galaxy # cat /etc/llttab

set-node galaxy
set-cluster 1
link link1 /dev/udp - udp 50000 - 192.168.9.1 192.168.9.255
link link2 /dev/udp - udp 50001 - 192.168.10.1 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

- Display the content of the `/etc/llttab` file on the second node nebula:

```
nebula # cat /etc/llttab

set-node nebula
set-cluster 1
link link1 /dev/udp - udp 50000 - 192.168.9.2 192.168.9.255
link link2 /dev/udp - udp 50001 - 192.168.10.2 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

The link command in the `/etc/llttab` file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See “[Sample configuration: direct-attached links](#)” on page 308.
- See “[Sample configuration: links crossing IP routers](#)” on page 309.

[Table L-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 L-1 Field description for link command in /etc/llttab

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example /dev/udp.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “ Selecting UDP ports ” on page 306.
<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>	■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the /etc/llttab file

The `set-addr` command in the /etc/llttab file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 309.

[Table L-2](#) describes the fields of the set-addr command.

Table L-2 Field description for set-addr command in /etc/llttab

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.

Table L-2 Field description for set-addr command in /etc/llttab (continued)

Field	Description
<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
-----
      *.sunrpc            Idle
      *. *               Unbound
      *.32771             Idle
      *.32776             Idle
      *.32777             Idle
      *.name              Idle
      *.biff              Idle
      *.talk              Idle
      *.32779             Idle
      .
      .
      .
      *.55098             Idle
      *.syslog            Idle
      *.58702             Idle
      *. *               Unbound
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# ifconfig interface_name netmask netmask
```

For example:

- For the first network interface on the node galaxy:

```
IP address=192.168.9.1, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

For the first network interface on the node nebula:

```
IP address=192.168.9.2, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

- For the second network interface on the node galaxy:

```
IP address=192.168.10.1, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

For the second network interface on the node nebula:

```
IP address=192.168.10.2, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.

An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.


```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100
```

```
link link1 /dev/udp - udp 50000 - 192.168.30.1 192.168.30.255
link link2 /dev/udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: direct-attached links

Figure L-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure L-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /dev/udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 /dev/udp - udp 50001 - 192.1.3.1 192.1.3.255
```


The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /dev/udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 /dev/udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

[Figure L-2](#) depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure L-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1
```

```

link link1 /dev/udp - udp 50000 - 192.1.3.1 -
link link2 /dev/udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr      0 link1 192.1.1.1
set-addr      0 link2 192.1.2.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb   0
set-arp       0
    
```

The /etc/llttab file on Node 0 resembles:

```

set-node Node0
set-cluster 1

link link1 /dev/udp - udp 50000 - 192.1.1.1 -
link link2 /dev/udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr      1 link1 192.1.3.1
set-addr      1 link2 192.1.4.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb   0
set-arp       0
    
```


Configuring LLT over UDP using IPv6

This appendix includes the following topics:

- [Using the UDP layer of IPv6 for LLT](#)
- [Manually configuring LLT over UDP using IPv6](#)

Using the UDP layer of IPv6 for LLT

Storage Foundation 6.0 PR1 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

Manually configuring LLT over UDP using IPv6

The following checklist is to configure LLT over UDP:


- For UDP6, the multicast address is set to "-".
- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.

- Make sure the IPv6 addresses in the /etc/llttab files are consistent with the IPv6 addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
- 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 316.

Sample configuration: direct-attached links

[Figure M-1](#) depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure M-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
```

```
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```


The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

Sample configuration: links crossing IP routers

[Figure M-2](#) depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure M-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

#set address of each link for all peer nodes in the cluster
```

```
#format: set-addr node-id link tag-name address
set-addr 0 link1 fe80::21a:64ff:fe92:1b46
set-addr 0 link2 fe80::21a:64ff:fe92:1b47
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

The `/etc/llttab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 1 link1 fe80::21a:64ff:fe92:1a92
set-addr 1 link2 fe80::21a:64ff:fe92:1a93
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

Compatibility issues when installing Storage Foundation High Availability with other products

This appendix includes the following topics:

- [Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present](#)

Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present

Installing Storage Foundation when other Veritas products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where SFM or VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the SFM Central Server and Managed Host packages as is.
- When uninstalling Storage Foundation products where SFM or VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where SFM or VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspb and VRTSsico. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspb, VRTSsico, and VRTSat.

Changing root user into role

This appendix includes the following topics:

- [Changing root user into root role](#)

Changing root user into root role

On Oracle Solaris 11, to perform installation, you need to create root user. This means that a local user cannot assume the root role. After installation, you may want to turn root user into root role for a local user, who can log in as root.

1. Log in as root user.
2. Change the root account into role.

```
# rolemod -K type=role root

# getent user_attr root

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

3. Assign the root role to a local user who was unassigned the role.

```
# usermod -R root admin
```

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

Index

A

- about
 - global clusters 26
- adding
 - users 96
- agents
 - disabling 191
- attributes
 - UseFence 118

B

- block device
 - partitions
 - example file name 295
 - volumes
 - example file name 295

C

- cables
 - cross-over Ethernet 155
- cluster
 - removing a node from 171
 - verifying operation 146
- command failures 284
- commands
 - format 52
 - hastatus 146
 - hasys 146
 - lltconfig 245
 - lltstat 143
 - vxdisksetup (initializing disks) 107
 - vxlicinst 104
 - vxlicrep 103
- configuring
 - rsh 48
 - ssh 48
- configuring SFHA
 - script-based installer 85
- configuring VCS
 - adding users 96

- configuring VCS (*continued*)
 - event notification 97, 99
 - global clusters 101
 - secure mode 93
 - starting 87
- controllers
 - SCSI 49
- coordinator disks
 - DMP devices 25
 - for I/O fencing 25
 - setting up 116

D

- data disks
 - for I/O fencing 25
- disabling the agents 191
- disks
 - adding and initializing 107
 - coordinator 116
 - testing with vxfststhdw 108
 - verifying node access 110

E

- Ethernet controllers 155

F

- FC-AL controllers 52

G

- GAB
 - description 22
- gabtab file
 - verifying after installation 245
- global clusters 26
 - configuration 101

H

- hastatus -summary command 146
- hasys -display command 146

hubs
 independent 155

I

I/O fencing
 checking disks 108
 setting up 115
 shared storage 108
 I/O fencing requirements
 non-SCSI-3 40
 installing
 post 102

L

license keys
 adding with vxlicinst 104
 replacing demo key 104
 licenses
 information about 103
 links
 private network 245
 LLT
 description 22
 interconnects 23
 verifying 143
 lltconfig command 245
 llthosts file
 verifying after installation 245
 lltstat command 143
 llttab file
 verifying after installation 245
 log files 285

M

main.cf file
 contents after installation 251
 main.cf files 256
 major and minor numbers
 checking 296, 299
 shared devices 295
 manual pages
 potential problems 283
 troubleshooting 283
 media speed 23
 optimizing 22
 mounting
 software disc 54

N

NFS services
 shared storage 295
 nodes
 adding application nodes
 configuring GAB 160
 configuring LLT 160
 configuring VXFEN 160
 starting Volume Manager 159
 non-SCSI-3 fencing
 manual configuration 129
 setting up 129
 non-SCSI-3 I/O fencing
 requirements 40
 non-SCSI3 fencing
 setting up 120
 using installsfha 120

O

optimizing
 media speed 22

P

PATH variable
 VCS commands 143
 persistent reservations
 SCSI-3 49
 problems
 accessing manual pages 283
 executing file system commands 284

R

removing
 the Replicated Data Set 192
 removing a system from a cluster 171
 Replicated Data Set
 removing the 192
 rsh 88
 configuration 48

S

script-based installer
 SFHA configuration overview 85
 SCSI driver
 determining instance numbers 297
 SCSI-3
 persistent reservations 49

- SCSI-3 persistent reservations
 - verifying 115
- SFHA
 - about 21
 - configuring 85
 - coordinator disks 116
- SFHA installation
 - verifying
 - cluster operations 143
 - GAB operations 143
 - LLT operations 143
- shared storage
 - Fibre Channel
 - setting up 52
 - NFS services 295
- SMTP email notification 97
- SNMP trap notification 99
- ssh 88
 - configuration 48
- starting configuration
 - installvcs program 88
 - Veritas product installer 87
- storage
 - setting up shared fibre 52
- Symantec Product Authentication Service 93
- system state attribute value 146

T

- troubleshooting
 - accessing manual pages 283
 - executing file system commands 284
- tunables file
 - about setting parameters 225
 - parameter definitions 230
 - preparing 229
 - setting for configuration 226
 - setting for installation 226
 - setting for upgrade 226
 - setting parameters 229
 - setting with no other operations 227
 - setting with un-integrated response file 228

V

- VCS
 - command directory path variable 143
 - configuration files
 - main.cf 249

- verifying installation
 - kernel component 141
- Veritas Operations Manager 21
- Volume Manager
 - Fibre Channel 52
- vradmin
 - delpri 193
 - stoprep 193
- vxdisksetup command 107
- vxlicinst command 104
- vxlicrep command 103