

Veritas Storage Foundation™ and High Availability Installation Guide

Linux

6.0

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

Document version: 6.0.5

Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

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

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Contents

Technical Support	4
Section 1 Installation overview and planning	19
Chapter 1 Introducing Storage Foundation and High Availability Solutions	20
About Veritas products	20
About Storage Foundation and High Availability	20
About Veritas Cluster Server	21
About Veritas high availability agents	21
About Veritas Volume Replicator	21
About Veritas Storage Foundation Cluster File System	22
About Veritas Storage Foundation for Sybase ASE CE	22
About Veritas Storage Foundation for Oracle® RAC	22
About Veritas graphical user interfaces	22
About Veritas Operations Manager	22
About Storage Foundation and High Availability features	23
About LLT and GAB	23
About configuring SFHA clusters for data integrity	24
About I/O fencing for SFHA in virtual machines that do not support SCSI-3 PR	25
About I/O fencing components	25
About global clusters	28
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	32
Release notes	32
Hardware compatibility list (HCL)	33

	Supported operating systems	33
	Storage Foundation memory requirements	33
	Disk space requirements	33
	Discovering product versions and various requirement information	33
	Database requirements	34
	VxVM licenses	34
	Cross-Platform Data Sharing licensing	35
	I/O fencing requirements	35
	Coordinator disk requirements for I/O fencing	36
	CP server requirements	36
	Non-SCSI-3 I/O fencing requirements	39
	Number of nodes supported	40
Chapter 4	Licensing Veritas products	41
	About Veritas SFHA Solutions product licensing	41
	Setting or changing the Veritas SFHA Solutions product level for keyless licensing	42
	Installing Veritas SFHA Solutions product license keys	44
Section 2	Installing Veritas Storage Foundation and High Availability	45
Chapter 5	Preparing to install Storage Foundation High Availability	46
	Installation preparation overview	46
	About using ssh or rsh with the Veritas installer	47
	Setting up shared storage	48
	Setting up shared storage: SCSI	48
	Setting up shared storage: Fibre Channel	49
	Setting environment variables	51
	Mounting the product disc	52
	Assessing the system for installation readiness	52
	Symantec Operations Readiness Tools	53
	Prechecking your systems using the Veritas installer	53

Chapter 6	Installing Storage Foundation and High Availability using the script-based installer	55
	About the Veritas installer	55
	Installing Storage Foundation and High Availability using the installer	56
Chapter 7	Installing Storage Foundation and High Availability using the Web-based installer	59
	About the Web-based installer	59
	Before using the Veritas Web-based installer	60
	Starting the Veritas Web-based installer	61
	Obtaining a security exception on Mozilla Firefox	61
	Performing a pre-installation check with the Veritas Web-based installer	62
	Setting installer options with the Web-based installer	62
	Installing SFHA with the Web-based installer	63
Chapter 8	Installing Storage Foundation and High Availability using other methods	65
	Installing SFHA using Kickstart	65
	Sample Kickstart configuration file	67
	Installing Storage Foundation and High Availability using yum	69
Section 3	Configuring Veritas Storage Foundation and High Availability	74
Chapter 9	Preparing to configure Storage Foundation and High Availability	75
	About planning to configure I/O fencing	75
	Typical Storage Foundation High Availability configuration with server-based I/O fencing	79
	Recommended CP server configurations	80
	Setting up the CP server	83
	Planning your CP server setup	83
	Installing the CP server using the installer	84
	Configuring the CP server cluster in secure mode	85
	Setting up shared storage for the CP server database	86
	Configuring the CP server using the configuration utility	87
	Configuring the CP server manually	96

	Verifying the CP server configuration	98
Chapter 10	Configuring Storage Foundation and High Availability	99
	Configuring Storage Foundation and High Availability Solutions	99
	Configuring Storage Foundation High Availability using the installer	99
	Configuring SFHA using the Web-based installer	119
Chapter 11	Configuring Storage Foundation High Availability for data integrity	125
	Setting up disk-based I/O fencing using installsfha	125
	Initializing disks as VxVM disks	126
	Checking shared disks for I/O fencing	126
	Configuring disk-based I/O fencing using installsfha	131
	Setting up disk-based I/O fencing manually	134
	Removing permissions for communication	134
	Identifying disks to use as coordinator disks	135
	Setting up coordinator disk groups	135
	Creating I/O fencing configuration files	136
	Modifying VCS configuration to use I/O fencing	137
	Verifying I/O fencing configuration	138
	Setting up server-based I/O fencing using installsfha	139
	Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsfha	139
	Configuring SFHA for data integrity using the Web-based installer	141
	Setting up server-based I/O fencing manually	144
	Preparing the CP servers manually for use by the Storage Foundation High Availability	145
	Configuring CoordPoint agent to monitor coordination points	148
	Verifying server-based I/O fencing configuration	150
	Setting up non-SCSI-3 fencing in virtual environments manually	151
	Sample /etc/vxfenmode file for non-SCSI-3 fencing	153
	Enabling or disabling the preferred fencing policy	156

Section 4	Upgrading Veritas Storage Foundation and High Availability	159
Chapter 12	Preparing to upgrade	160
	About upgrading	160
	About the different ways that you can upgrade	161
	Supported upgrade paths	161
	About using the installer to upgrade when the root disk is encapsulated	166
	Preparing to upgrade	166
	Getting ready for the upgrade	166
	Creating backups	168
	Tasks for upgrading the Storage Foundation for Databases (SFDB)	168
	Determining if the root disk is encapsulated	169
	Pre-upgrade tasks for migrating the SFDB repository database	169
	Pre-upgrade planning for Veritas Volume Replicator	169
	Preparing to upgrade VVR when VCS agents are configured	171
	Upgrading the array support	175
Chapter 13	Upgrading Storage Foundation and High Availability	177
	Upgrading Storage Foundation and High Availability from previous versions to 6.0	177
	Upgrading Storage Foundation and High Availability using the script-based installer	178
	Upgrading Storage Foundation and High Availability using the Veritas Web-based installer	182
	Upgrading Veritas Volume Replicator	183
	Upgrading VVR without disrupting replication	183
Chapter 14	Performing a rolling upgrade	186
	About rolling upgrades	186
	Performing a rolling upgrade using the script-based installer	188
	Performing a rolling upgrade of SFHA using the Web-based installer	190

Chapter 15	Performing a phased upgrade	193
	About phased upgrade	193
	Prerequisites for a phased upgrade	193
	Planning for a phased upgrade	193
	Phased upgrade limitations	194
	Phased upgrade example	194
	Phased upgrade example overview	195
	Performing a phased upgrade	196
	Moving the service groups to the second subcluster	196
	Upgrading the operating system on the first subcluster	200
	Upgrading the first subcluster	200
	Preparing the second subcluster	201
	Activating the first subcluster	205
	Upgrading the operating system on the second subcluster	206
	Upgrading the second subcluster	207
	Finishing the phased upgrade	208
Chapter 16	Performing post-upgrade tasks	211
	Optional configuration steps	211
	Re-joining the backup boot disk group into the current disk group	212
	Reverting to the backup boot disk group after an unsuccessful upgrade	212
	Post upgrade tasks for migrating the SFDB repository database	213
	Migrating from a 5.0 repository database to 6.0	213
	After upgrading from 5.0.x and before migrating SFDB	216
	Recovering VVR if automatic upgrade fails	216
	Post-upgrade tasks when VCS agents for VVR are configured	216
	Unfreezing the service groups	217
	Restoring the original configuration when VCS agents are configured	218
	Upgrading disk layout versions	219
	Upgrading VxVM disk group versions	220
	Updating variables	221
	Setting the default disk group	221
	About enabling LDAP authentication for clusters that run in secure mode	221
	Enabling LDAP authentication for clusters that run in secure mode	223
	Verifying the Storage Foundation and High Availability upgrade	229

Section 5	Verifying the Veritas Storage Foundation and High Availability installation	230
Chapter 17	Verifying the installation	231
	Performing a postcheck on a node	231
	Verifying that the products were installed	231
	Installation log files	232
	Using the installation log file	232
	Using the summary file	232
	Starting and stopping processes for the Veritas products	232
	Checking Veritas Volume Manager processes	233
	Verifying the LLT, GAB, and VCS configuration files	233
	Verifying LLT, GAB, and cluster operation	234
	Verifying LLT	234
	Verifying the cluster	236
	Verifying the cluster nodes	237
Section 6	Adding and removing nodes	240
Chapter 18	Adding a node to a cluster	241
	About adding a node to a cluster	241
	Before adding a node to a cluster	242
	Adding a node to a cluster using the product installer	244
	Adding a node using the Web-based installer	246
	Adding the node to a cluster manually	247
	Starting Volume Manager on the new node	255
	Configuring cluster files on the new node	255
	Configuring cluster files on the new node	250
	Starting fencing on the new node	260
	Starting VCS after adding the new node	263
	Configuring the ClusterService group for the new node	260
	Starting Volume Manager on the new node	255
	Configuring cluster files on the new node	255
	Setting up the node to run in secure mode	257
	Starting fencing on the new node	260
	Configuring the ClusterService group for the new node	260
	Configuring server-based fencing on the new node	261
	Adding the new node to the vxfen service group	262
	Starting VCS after adding the new node	263

	Updating the Storage Foundation for Databases (SFDB) repository after adding a node	263
Chapter 19	Removing a node from a cluster	265
	Removing a node from a cluster	265
	Verifying the status of nodes and service groups	266
	Deleting the departing node from SFHA configuration	267
	Modifying configuration files on each remaining node	270
	Removing the node configuration from the CP server	270
	Removing security credentials from the leaving node	271
	Unloading LLT and GAB and removing VCS on the departing node	271
	Updating the Storage Foundation for Databases (SFDB) repository after removing a node	273
Section 7	Uninstalling Veritas Storage Foundation and High Availability	274
Chapter 20	Uninstalling Storage Foundation and High Availability	275
	Removing VxFS file systems	275
	Removing rootability	276
	Moving volumes to disk partitions	277
	Moving volumes onto disk partitions using VxVM	277
	Disabling VCS agents for VVR the agents on a system	279
	Removing the Replicated Data Set	280
	Uninstalling SFHA RPMs using the script-based installer	281
	Uninstalling SFHA with the Veritas Web-based installer	283
	Removing license files (Optional)	284
	Removing the CP server configuration using the removal script	284
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product	286
Section 8	Installation reference	288
Appendix A	Installation scripts	289
	About Veritas Storage Foundation and High Availability Solutions installation scripts	289
	Installation script options	290

	About using the postcheck option	295
Appendix B	Response files	297
	About response files	297
	Installing SFHA using response files	298
	Configuring SFHA using response files	298
	Upgrading SFHA using response files	299
	Uninstalling SFHA using response files	300
	Syntax in the response file	300
	Response file variables to install, upgrade, or uninstall Storage Foundation and High Availability	301
	Response file variables to configure Storage Foundation and High Availability	303
	Sample response file for SFHA configuration	312
	Sample response file for SFHA install	312
	Sample response file for SF upgrade	313
	Sample response file for SFHA upgrade	313
Appendix C	Tunable files for installation	315
	About setting tunable parameters using the installer or a response file	315
	Setting tunables for an installation, configuration, or upgrade	316
	Setting tunables with no other installer-related operations	317
	Setting tunables with an un-integrated response file	318
	Preparing the tunables file	319
	Setting parameters for the tunables file	319
	Tunables value parameter definitions	320
Appendix D	Configuring I/O fencing using a response file	327
	Configuring I/O fencing using response files	327
	Response file variables to configure disk-based I/O fencing	328
	Sample response file for configuring disk-based I/O fencing	329
	Response file variables to configure server-based I/O fencing	330
	Sample response file for configuring server-based I/O fencing	332
	Sample response file for configuring non-SCSI-3 server-based I/O fencing	332
	Response file variables to configure non-SCSI-3 server-based I/O fencing	333

Appendix E	Configuration files	335
	About the LLT and GAB configuration files	335
	About the AMF configuration files	338
	About the VCS configuration files	339
	Sample main.cf file for VCS clusters	341
	Sample main.cf file for global clusters	342
	About I/O fencing configuration files	344
	Sample configuration files for CP server	347
	Sample main.cf file for CP server hosted on a single node that runs VCS	348
	Sample main.cf file for CP server hosted on a two-node SFHA cluster	350
	Sample CP server configuration (/etc/vxcps.conf) file output	353
Appendix F	Configuring the secure shell or the remote shell for communications	354
	About configuring secure shell or remote shell communication modes before installing products	354
	Manually configuring and passwordless ssh	355
	Restarting the ssh session	359
	Enabling rsh for Linux	359
Appendix G	Storage Foundation and High Availability components	361
	Storage Foundation and High Availability installation RPMs	361
	Veritas Cluster Server installation RPMs	364
	Veritas Storage Foundation obsolete and reorganized installation RPMs	365
Appendix H	Troubleshooting installation issues	369
	Restarting the installer after a failed connection	369
	What to do if you see a licensing reminder	369
	Troubleshooting information	370
	Incorrect permissions for root on remote system	370
	Inaccessible system	371
	Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 (2184482)	372
	Workaround	372

Appendix I	Troubleshooting cluster installation	373
	Unmount failures	373
	Command failures	373
	Installer cannot create UUID for the cluster	374
	The vxfsntsthdw utility fails when SCSI TEST UNIT READY command fails	374
	Troubleshooting CP server	375
	Troubleshooting issues related to the CP server service group	376
	Checking the connectivity of CP server	376
	Troubleshooting server-based fencing on the SFHA cluster nodes	376
	Issues during fencing startup on Storage Foundation High Availability nodes set up for server-based fencing	377
	Issues during online migration of coordination points	377
	Vxfsn service group activity after issuing the vxfsnwap command	378
Appendix J	Sample Storage Foundation High Availability setup diagrams for CP server-based I/O fencing	379
	Configuration diagrams for setting up server-based I/O fencing	379
	Two unique client clusters served by 3 CP servers	380
	Client cluster served by highly available CPS and 2 SCSI-3 disks	380
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	381
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	383
Appendix K	Configuring LLT over UDP using IPv4	384
	Using the UDP layer for LLT	384
	When to use LLT over UDP	384
	Manually configuring LLT over UDP using IPv4	384
	Broadcast address in the /etc/lfttab file	385
	The link command in the /etc/lfttab file	385
	The set-addr command in the /etc/lfttab file	386
	Selecting UDP ports	387
	Configuring the netmask for LLT	387
	Configuring the broadcast address for LLT	388
	Sample configuration: direct-attached links	388
	Sample configuration: links crossing IP routers	390

Appendix L	Configuring LLT over UDP using IPv6	392
	Using the UDP layer of IPv6 for LLT	392
	When to use LLT over UDP	392
	Manually configuring LLT over UDP using IPv6	392
	Sample configuration: direct-attached links	393
	Sample configuration: links crossing IP routers	394
Appendix M	Compatibility issues when installing Storage Foundation High Availability with other products	396
	Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present	397
	Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present	397
	Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present	397
	Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present	398
Index		399

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://www.symantec.com/operations-manager/support>. Veritas Storage Foundation Management Server is deprecated.

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

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 75.

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 installsfha”](#) on page 139.

See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 151.

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 26.
- Coordination points—Act as a global lock during membership changes

See [“About coordination points”](#) on page 26.

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 Storage Foundation High Availability 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 Storage Foundation High Availabilities running different operating systems can simultaneously access the CP server. TCP/IP based communication is used between the CP server and the Storage Foundation High Availabilities.

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 156.

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.
- 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 297.

- KickStart
You can use the Veritas product installer or the product-specific installation script to generate a Kickstart script file. Use the generated script to install Veritas RPMs from your Kickstart server.

Downloading the Storage Foundation and High Availability software

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

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

To download the trialware version of the software

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

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

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

See “[About Veritas Storage Foundation and High Availability Solutions installation scripts](#)” on page 289.

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:

1 GB

See [“Disk space requirements”](#) on page 33.

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

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

```
# df -k filesystem
```

Caution: When you select a location to download files, do not select a directory that contains Veritas products from a previous release or maintenance pack. 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](#)
- [Storage Foundation memory requirements](#)
- [Disk space requirements](#)
- [Discovering product versions and various requirement information](#)
- [Database requirements](#)
- [VxVM licenses](#)
- [Cross-Platform Data Sharing licensing](#)
- [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*.

Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Disk space requirements

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

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

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

If you have downloaded SFHA, you must 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 RPMs 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.

VxVM licenses

The following table shows the levels of licensing in Veritas Volume Manager and the features supported at each level.

[Table 3-1](#) describes the levels of licensing in Veritas Volume Manager and supported features.

Table 3-1 Levels of licensing in Veritas Volume Manager and supported features

VxVM License	Description of Supported Features
Full	Concatenation, spanning, rootability, volume resizing, multiple disk groups, co-existence with native volume manager, striping, mirroring, DRL logging for mirrors, striping plus mirroring, mirroring plus striping, RAID-5, RAID-5 logging, Smartsync, hot sparing, hot-relocation, online data migration, online relayout, volume snapshots, volume sets, Intelligent Storage Provisioning, FastResync with Instant Snapshots, Storage Expert, Device Discovery Layer (DDL), Dynamic Multipathing (DMP), and Veritas Operations Manager (VOM).
Add-on Licenses	Features that augment the Full VxVM license such as clustering functionality (cluster-shareable disk groups and shared volumes) and Veritas Volume Replicator.

Note: You need a Full VxVM license to make effective use of add-on licenses to VxVM.

To see the license features that are enabled in VxVM

- ◆ Enter the following command:

```
# vxdctl license
```

Cross-Platform Data Sharing licensing

The Cross-Platform Data Sharing (CDS) feature is also referred to as Portable Data Containers.

The ability to import a CDS disk group on a platform that is different from the platform on which the disk group was last imported is controlled by a CDS license. CDS licenses are included as part of the Veritas Storage Foundation license.

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

See “[CP server requirements](#)” on page 36.

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.

See “[Non-SCSI-3 I/O fencing requirements](#)” on page 39.

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

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

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

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

See the *Veritas Cluster Server Installation Guide*.

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 3-2](#) lists additional requirements for hosting the CP server.

Table 3-2 CP server hardware requirements

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

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

Table 3-3 CP server supported operating systems and versions

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

Following are the CP server networking requirements and recommendations:

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

For secure communication between the 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 Storage Foundation High Availability 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:

- VMware Server ESX 3.5, 4.0, and 5.0 on AMD Opteron or Intel Xeon EM64T (x86_64)
VMware vSphere 5.0
Guest operating system: See the *Storage Foundation and High Availability Release Notes* for the list of supported Linux operating systems.

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

About Veritas SFHA Solutions product licensing

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

www.symantec.com/techsupp/

During the installation, you can choose to either:

- Install a license key for the product and features that you want to install.
When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
- Continue to install without a license key.
The installer prompts for the product modes and options that you want to install, and then sets the required product level.
Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server, such as Veritas Operations Manager (VOM). If you do not comply with the above terms,

continuing to use the Symantec product is a violation of your end user license agreement, and results in warning messages.

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

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

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

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

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

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

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.

See [“Setting or changing the Veritas SFHA Solutions product level for keyless licensing”](#) on page 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 SFHA Solutions 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 Veritas SFHA Solutions product level for keyless licensing

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

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

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

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

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

To set or change the product level

- 1 Show your current working directory:

```
# pwd
```

Output resembles:

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

- 2 View the current setting for the product level.

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

- 3 View the possible settings for the product level.

```
# ./vxkeyless displayall
```

- 4 Set the desired product level.

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

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

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas SFHA Solutions product license keys

The `VRTSvlic` RPM 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](#)
- [Chapter 7. Installing Storage Foundation and High Availability using the Web-based installer](#)
- [Chapter 8. Installing Storage Foundation and High Availability using other methods](#)

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](#)
- [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 SFHA Solutions product licensing” on page 41.
Download the software, or insert the product DVD.	See “Downloading the Storage Foundation and High Availability software” on page 30. See “Mounting the product disc” on page 52.

Table 5-1 Installation overview (*continued*)

Installation task	Section
Set environment variables.	See “Setting environment variables” on page 51.
Configure the secure shell (ssh) on all nodes.	See “About configuring secure shell or remote shell communication modes before installing products” on page 354.
Verify that hardware, software, and operating system requirements are met.	See “Release notes” on page 32.
Check that sufficient disk space is available.	See “Disk space requirements” on page 33.
Use the installer to install the products.	See “About the Veritas installer” on page 55.

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 290.

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 354.

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.

Note: SFHA also supports non-SCSI3 server-based fencing configuration in virtual environments that do not support SCSI-3 PR-compliant storage.

See [“About planning to configure I/O fencing”](#) on page 75.

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

Setting up shared storage: SCSI

Perform the following steps to set up shared storage.

To set up shared storage

- 1 Connect the disk to the first cluster system.
- 2 Power on the disk.
- 3 Connect a terminator to the other port of the disk.
- 4 Boot the system. The disk is detected while the system boots.
- 5 Press CTRL+A to bring up the SCSI BIOS settings for that disk.

Set the following:

- Set Host adapter SCSI ID = 7, or to an appropriate value for your configuration.
- Set Host Adapter BIOS in Advanced Configuration Options to Disabled.

- 6 Format the shared disk and create required partitions on it.

Perform the following:

- Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is /dev/sdc. Identify whether the shared disk is sdc, sdb, and so on.
- Type the following command:

```
# fdisk /dev/shareddiskname
```


For example, if your shared disk is sdc, type:

```
# fdisk /dev/sdc
```

- Create disk groups and volumes using Volume Manager utilities.
- To apply a file system on the volumes, type:

```
# mkfs -t fs-type /dev/vx/dsk/disk-group/volume
```

For example, enter the following command:

```
# mkfs -t vxfs /dev/vx/dsk/dg/vol01
```

Where the name of the disk group is dg, the name of the volume is vol01, and the file system type is vxfs.

- 7 Power off the disk.
- 8 Remove the terminator from the disk and connect the disk to the other cluster system.
- 9 Power on the disk.
- 10 Boot the second system. The system can now detect the disk.
- 11 Press Ctrl+A to bring up the SCSI BIOS settings for the disk.
Set the following:
 - Set Host adapter SCSI ID = 6, or to an appropriate value for your configuration. Note that the SCSI ID should be different from the one configured on the first cluster system.
 - Set Host Adapter BIOS in Advanced Configuration Options to Disabled.
- 12 Verify that you can view the shared disk using the `fdisk` command.

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up shared storage for Fibre Channel

- 1 Connect the Fibre Channel disk to a cluster system.
- 2 Boot the system and change the settings of the Fibre Channel. Perform the following tasks for all QLogic adapters in the system:
 - Press Alt+Q to bring up the QLogic adapter settings menu.
 - Choose **Configuration Settings**.

- Click Enter.
 - Choose **Advanced Adapter Settings**.
 - Click Enter.
 - Set the Enable Target Reset option to **Yes** (the default value).
 - Save the configuration.
 - Reboot the system.
- 3 Verify that the system detects the Fibre Channel disks properly.
- 4 Create volumes. Format the shared disk and create required partitions on it and perform the following:
- Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is /dev/sdc. Identify whether the shared disk is sdc, sdb, and so on.
 - Type the following command:


```
# fdisk /dev/shareddiskname
```

For example, if your shared disk is sdc, type:

```
# fdisk /dev/sdc
```
 - Create disk groups and volumes using Volume Manager utilities.
 - To apply a file system on the volumes, type:


```
# mkfs -t fs-type /dev/vx/dsk/disk-group/volume
```

For example, enter the following command:

```
# mkfs -t vxfs /dev/vx/dsk/dg/vol01
```

Where the name of the disk group is dg, the name of the volume is vol01, and the file system type is vxfs.
- 5 Repeat step 2 and step 3 for all nodes in the clusters that require connections with Fibre Channel.
- 6 Power off this cluster system.
- 7 Connect the same disks to the next cluster system.

- 8 Turn on the power for the second system.
- 9 Verify that the second system can see the disk names correctly—the disk names should be the same.

See “[Verifying that the nodes have access to the same disk](#)” on page 128.

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`.

Specify `/opt/VRTS/bin` in your `PATH` after the path to the standard Linux commands.

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

To invoke the VxFS-specific `df`, `fsdb`, `ncheck`, or `umount` commands, type the full path name: `/opt/VRTS/bin/command`.

To set your `MANPATH` environment variable to include `/opt/VRTS/man` do the following:

- If you are using a shell such as `sh` or `bash`, enter the following:

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

- If you are using a shell such as `csh` or `tcsh`, enter the following:

```
% setenv MANPATH $(MANPATH) :/opt/VRTS/man
```

On a Red Hat system, also include the `1m` manual page section in the list defined by your `MANSECT` environment variable.

- If you are using a shell such as `sh` or `bash`, enter the following:

```
$ MANSECT=$MANSECT:1m; export MANSECT
```

- If you are using a shell such as `csh` or `tcsh`, enter the following:

```
% setenv MANSECT $(MANSECT) :1m
```

If you use the `man(1)` command to access manual pages, set `LC_ALL=C` in your shell to ensure that they display correctly.

Mounting the product disc

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 with the SFHA software into a drive that is connected to the system.

The disc is automatically mounted.

- 3 If the disc does not automatically mount, then enter:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

- 4 Navigate to the location of the RPMs.

```
# cd /mnt/cdrom/dist_arch/rpms
```

Where *dist* is rhel5, rhel6, sles10, or sles11, and *arch* is x86_64 for RHEL and SLES.

Assessing the system for installation readiness

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

Symantec Operations Readiness Tools

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

See [“Symantec Operations Readiness Tools”](#) on page 53.

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

See “Prechecking your systems using the Veritas installer” on page 53.

Symantec Operations Readiness Tools

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

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

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

To access SORT, go to:

<https://sort.symantec.com>

Prechecking your systems using the Veritas installer

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

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

To use the precheck option

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

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

```
# ./installer
```

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

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

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](#)

About the Veritas installer

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

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

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 289.

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 290.

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 354.

- 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 52.

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

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

Where *dist* is *rhel5*, *rhel6*, *sles10*, or *sles11*, and *arch* is *x86_64* for RHEL and SLES.

- 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 `␣` 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 RPMs: installs only the basic functionality for the selected product.
 - Recommended RPMs: installs the full feature set without optional RPMs.
 - All RPMs: installs all available RPMs.

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 RPMs 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 Storage Foundation and High Availability using the Web-based installer

This chapter includes the following topics:

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

About the Web-based installer

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

You use the `webinstaller` script to start and stop the Veritas XPortal Server `xprt1wid` process. The `webinstaller` script can also be used to check the status of the XPortal Server.

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

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

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

See [“Before using the Veritas Web-based installer”](#) on page 60.

See [“Starting the Veritas Web-based installer”](#) on page 61.

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 7-1 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Veritas products.	Must be a supported platform for Veritas Storage Foundation 6.0.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must use the same operating system as the target systems and must be at one of the supported operating system update levels.
Administrative system	The system where you run the Web browser to perform the installation.	Must have a Web browser. Supported browsers: <ul style="list-style-type: none"> ■ Internet Explorer 6, 7, and 8 ■ Firefox 3.x and later

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

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

```
# ./webinstaller start
```

The webinstaller script displays a URL. Note this URL.

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

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

- 2 On the administrative server, start the Web browser.
- 3 Navigate to the URL that the script displayed.
- 4 Certain browsers may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

- 5 Log in as superuser.

Obtaining a security exception on Mozilla Firefox

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

The following instructions are general. They may change because of the rapid release cycle of Mozilla browsers.

To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **Add Exception** button.
- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.

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

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

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

To perform a pre-installation check

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 61.
- 2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.
- 3 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Next**.
- 4 The installer performs the precheck and displays the results.
- 5 Click **Finish**. The installer prompts you for another task.

Setting installer options with the Web-based installer

You can use the Web-based installer for certain command-line installer options.

The supported options follow:

- `-serial`
- `-require path_to_hotfix_file`
- `-mediapath directory_path_to_install_media`
- `-logpath directory_path_to_save_logs`
- `-tmpopath directory_path_to_save_temp_files`

See [“Installation script options”](#) on page 290.

To use installer options

- 1 On the Web-installer's entry page, click the **Advanced Options** link.
- 2 In the Command line options field, enter the option that you want to use.
For example, if you want to use the serial option and the logpath option, enter:

```
-serial -logpath /opt/VRTS/install/advlogs
```

Where */opt/VRTS/install/advlogs* is the path that you want to use. Separate the command with a space.

- 3 Click the **OK** button and proceed.

Installing SFHA with the Web-based installer

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

To install SFHA using the Web-based installer

- 1 Perform preliminary steps.
See [“Performing a pre-installation check with the Veritas Web-based installer”](#) on page 62.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 61.
- 3 On the License agreement page, read the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 4 Indicate the systems where you want to install. Separate multiple system names with spaces. Click **Next**.
- 5 If you have not yet configured a communication mode among systems, you have the option to let the installer configure ssh or rsh. If you choose to allow this configuration, select the communication mode and provide the superuser passwords for the systems.
- 6 After the validation completes successfully, click **Next** to install SFHA on the selected system.
- 7 After the installation completes, you must choose your licensing method.
On the license page, select one of the following tabs:
 - Keyless licensing

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

For more information, go to the following website:

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

Click **Register**.

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

- 8 For Storage Foundation and High Availability, click **Next** to complete the configuration and start the product processes.

Note that you are prompted to configure only if the product is not yet configured.

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

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

- 9 The installer prompts you to configure the cluster. Select **Yes** to continue with configuring the product.

If you select **No**, you can exit the installer. You must configure the product before you can use SFHA.

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

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

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

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

Installing Storage Foundation and High Availability using other methods

This chapter includes the following topics:

- [Installing SFHA using Kickstart](#)
- [Sample Kickstart configuration file](#)
- [Installing Storage Foundation and High Availability using yum](#)

Installing SFHA using Kickstart

You can install SFHA using Kickstart. Kickstart is supported for Red Hat Enterprise Linux (RHEL5) and Red Hat Enterprise Linux 6 (RHEL6).

To install SFHA using Kickstart

- 1 Create a directory for the Kickstart configuration files.

```
# mkdir /kickstart_files/
```

- 2 Generate the Kickstart configuration files. The configuration files have the extension `.ks`. Do one of the following:

- To generate configuration files, enter the following command:

```
# ./installer -kickstart /kickstart_files/
```

The system lists the files.

- To generate only Storage Foundation and High Availability (SFHA) configuration files, use the `installsfha` script. Enter the following:

```
# ./installsfha -kickstart /kickstart_files/
```

The output includes the following:

```
The kickstart script for SFHA60 is generated at  
/kickstart_files/kickstart_sfha60.ks
```

- 3 Setup an NFS exported location which the Kickstart client can access. For example, if `/nfs_mount_kickstart` is the directory which has been NFS exported, the NFS exported location may look similar to the following:

```
# cat /etc/exports  
/nfs_mount_kickstart * (rw, sync, no_root_squash)
```

- 4 Copy the `rpms` directory from the installation media to the NFS location.
- 5 Verify the contents of the directory.

```
# ls /nfs_mount_kickstart/
```

- 6 In the SFHA Kickstart configuration file, modify the `BUILDSRC` variable to point to the actual NFS location. The variable has the following format:

```
BUILDSRC="hostname_or_ip:/nfs_mount_kickstart"
```

- 7 Append the entire modified contents of the Kickstart configuration file to the operating system `ks.cfg` file.
- 8 Launch the Kickstart installation for the operating system.
- 9 After the operating system installation is complete, check the file `/var/tmp/kickstart.log` for any errors related to the installation of Veritas RPMs and Veritas product installer scripts.
- 10 Verify that all the product RPMs have been installed. Enter the following command:

```
# rpm -qa | grep -i vrts
```

- 11 If you do not find any installation issues or errors, configure the product stack. Enter the following command:

For Storage Foundation High Availability, enter the following:

```
# /opt/VRTS/install/installsfha -configure node1 node2
```

- 12 Verify that all the configured llc links and gac ports have started.
- 13 Verify that the product is configured properly by entering commands such as `hastatus` and `lltstat -n`.

```
# hastatus -sum
```

```
-- SYSTEM STATE
```

```
-- System                State                Frozen
```

```
A swlx13                 RUNNING             0
```

```
A swlx14                 RUNNING             0
```

```
# lltstat -n
```

```
LLT node information:
```

```
Node                State    Links
```

```
0 swlx13            OPEN    2
```

```
* 1 swlx14          OPEN    2
```

- 14 Verify that the ODM RPM is installed and determine which mode it is running in.

For example, in the case of the SFHA stack, the ODM cluster status is shown as disabled.

- 15 If you configure the node in a secured mode, verify the VxSS service group status. For example:

```
# haclus -value SecureClus
```

```
1
```

Sample Kickstart configuration file

The following is a sample RedHat Enterprise Linux 5 (RHEL5) Kickstart configuration file.

```
# The packages that follow are required. These packages are automatically installed
# from the operating system's installation media.
```

```
%packages
libattr.i386
libacl.i386

%post --nochroot
# Add any necessary scripts or commands here.
# This generated kickstart file is only for
# the installation of products in the installation media.

PATH=$PATH:/sbin:/usr/sbin:/bin:/usr/bin
export PATH

#
# Notice:
# * Modify the BUILDSRC below according to your environment
# * The location that you specify in BUILDSRC should let the
#   Kickstart Server be NFS accessible.
# * Copy the rpms directory from installation media to the
#   BUILDSRC location.
#

BUILDSRC="<hostname_or_ipaddress_or_ipv6address>:/path/to/rpms"

#
# Notice:
# * You do not have to change the following scripts.
#

# Define the path variables.
ROOT=/mnt/sysimage
BUILDDIR="${ROOT}/build"
RPMDIR="${BUILDDIR}/rpms"

# Define the log path.
KSLOG="${ROOT}/var/tmp/kickstart.log"

echo "==== Executing kickstart post section: =====" >> ${KSLOG}

mkdir -p ${BUILDDIR}
mount -t nfs -o nolock,vers=3 <BUILDSRC> ${BUILDDIR} >> ${KSLOG} 2>&1
```

```
# Install the RPMs in the following order.
for RPM in VRTSvlic VRTSperl VRTSsfcp160 VRTSspt VRTSvxvm VRTSaslapm
VRTSob VRTSlnconv VRTSsfmh VRTSvxfs VRTSfssdk VRTSatClient
VRTSatServer VRTSllt VRTSgab VRTSvxfen VRTSsamf VRTSvcS VRTScps
VRTSvcSag VRTSvcSdr VRTSvcSea VRTSdbed VRTSodm

do
    echo "Installing package -- $RPM" >> ${KSLOG}
    rpm -U -v --root ${ROOT} ${RPMDIR}/${RPM}-* >> ${KSLOG} 2>&1
done

umount ${BUILDDIR}

${ROOT}/opt/VRTS/install/bin/UXRT60/add_install_scripts >> ${KSLOG} 2>&1

echo "==== Completed kickstart file ==== " >> ${KSLOG}

exit 0
```

Installing Storage Foundation and High Availability using yum

You can install SFHA using yum. yum is supported for Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6.

To install SFHA using yum

- 1 Run the `installsfha -pkginfo` command to get SFHA RPMs.

```
# ./installsfha -pkginfo
```

- 2 Add the SFHA RPMs into the yum repository. You can add SFHA RPMs into either a new repository or an existing repository with other RPMs. Use the `createrepo` command to create or update the repository. The operating system package `createrepo-ver-rel.noarch.rpm` provides the command.

- To create the new repository */path/to/new/repository/* for SFHA RPMs

1. Create an empty directory, for example: */path/to/new/repository*. The yum client systems should be able to access the directory with the HTTP, FTP, or file protocols.

```
# rm -rf /path/to/new/repository
# mkdir -p /path/to/new/repository
```

2. Copy all the SFHA RPMs into */path/to/new/repository/*.

```
# cp -f VRTSvlic-* VRTSperl-* ... VRTSsfcp60-* \
/path/to/new/repository
```

3. Use the `createrepo` command to create the repository.

```
# /usr/bin/createrepo /path/to/new/repository
```

Output resembles:

```
27/27 - VRTSsfcp60-6.0.000.000-GA_GENERIC.noarch.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
```

4. The metadata for this repository is created in */path/to/new/repository/repdata*.

■ To use an existing repository in */path/to/existing/repository/* for SFHA RPMs

1. Copy all the SFHA RPMs into */path/to/existing/repository/*. The yum client systems should be able to access the directory with the HTTP, FTP, or file protocols.

```
# cp -f VRTSvlic-* VRTSperl-* ... VRTSsfcp60-* \
/path/to/existing/repository
```

2. Use the `createrepo` command with the `--update` option to update the repository's metadata.

```
# createrepo --update /path/to/existing/repository
```

Output resembles:

```
27/27 * VRTSsfcp60-6.0.000.000-GA_GENERIC.noarch.rpm
Saving Primary metadata
Saving file lists metadata
Saving other metadata
```

3. The metadata in `/path/to/existing/repository/repodata` is updated for the newly added RPMs.

- **To create a package group for SFHA RPMs when the repository is created or updated (optional)**

1. Create an XML file, which you can name `SFHA_group.xml` in the repository directory. In the file specify the name, the id, the package list, and other information for the group. An example of this XML file for SFHA is:

```
# cat SFHA_group.xml
<comps>
  <group>
    <id>SFHA60</id>
    <name>SFHA60</name>
    <default>true</default>
    <description>RPMs of SFHA 6.0</description>
    <uservisible>true</uservisible>
    <packagelist>
      <packagereq type="default">VRTSvlic</packagereq>
      <packagereq type="default">VRTSperl</packagereq>
      ... [other RPMs for SFHA]
      <packagereq type="default">VRTSsfcp60</packagereq>
    </packagelist>
  </group>
</comps>
```

2. Create the group when the repository is created or updated. You can generate this XML file using the installer with the option `-yumgroupxml`.

```
# createrepo -g SFHA_group.xml /path/to/new/repository/
```

Or

```
# createrepo -g SFHA_group.xml --update /path/to/existing\
/repository/
```

Refer to the *Red Hat Enterprise Linux Deployment Guide* for more information on yum repository configuration.

3. Configure a yum repository on a client system.
 - Create a `.repo` file under `/etc/yum.repos.d/`. An example of this `.repo` file for SFHA is:

```
# cat /etc/yum.repos.d/SFHA.repo
[repo-SFHA]
name=Repository for SFHA
baseurl=file:///path/to/repository/
enabled=1
gpgcheck=0
```

The values for the `baseurl` attribute can start with `http://`, `ftp://`, or `file://`. The URL you choose needs to be able to access the `repodata` directory. It also needs to access all the SFHA RPMs in the repository that you create or update.

- Check the yum configuration. List SFHA RPMs.

```
# yum list 'VRTS*'
Available Packages
VRTSperl.x86_64          5.10.0.7-RHEL5.3      repo-SFHA
VRTSsfcp160.noarch      6.0.000.000-GA_GENERIC  repo-SFHA
VRTSvlic.x86_64         3.02.52.000-0         repo-SFHA
...
```

The SFHA RPMs may not be visible immediately if:

- the repository was visited before the SFHA RPMs were added, and
- the local cache of its metadata has not expired.

To eliminate the local cache of the repositories' metadata and get the latest information from the specified `baseurl`, run the following commands:

```
# yum clean expire-cache
# yum list 'VRTS*'
```

Refer to the *Red Hat Enterprise Linux Deployment Guide* for more information on yum repository configuration.

4 Install the RPMs on the target systems.

- **To install all the RPMs**

1. Specify each RPM name as its yum equivalent. For example:

```
# yum install VRTSvlic VRTSperl ... VRTSsfcp160
```

2. Specify all of SFHA's RPMs by its package glob. For example:

```
# yum install 'VRTS*'
```


3. Specify the group name if a group is configured for SFHA's RPMs. In this example, the group name is *SFHA60*:

```
# yum install @SFHA60
```

Or

```
# yum groupinstall SFHA60
```

■ To install one RPM at a time

1. Run the `installsfha -pkginfo` command to determine package installation order.
2. Use the same order as the output from the `installsfha -pkginfo` command:

```
# yum install VRTSvlic
# yum install VRTSperl
...
# yum install VRTSsfcp60
```

- 5 After you install all the RPMs, use the `/opt/VRTS/install/installsfha` script to license, configure, and start the product.

If the `VRTSsfcp60` package is installed before you use `yum` to install SFHA, this package is not reinstalled. The `/opt/VRTS/install/installsfha` script may not be created. If the script is not created, use the `/opt/VRTS/install/bin/UXRT60/add_install_scripts` script to create the `installsfha` or `uninstallsfha` scripts after all the other SFHA RPMs are installed.

```
# /opt/VRTS/install/bin/UXRT60/add_install_scripts
Creating install/uninstall scripts for installed products
Creating /opt/VRTS/install/installsfha for UXRT60
Creating /opt/VRTS/install/uninstallsfha for UXRT60
```

Configuring Veritas Storage Foundation and High Availability

- [Chapter 9. Preparing to configure Storage Foundation and High Availability](#)
- [Chapter 10. Configuring Storage Foundation and High Availability](#)
- [Chapter 11. 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.

See [Figure 9-2](#) on page 78.

[Figure 9-1](#) illustrates a high-level flowchart to configure I/O fencing for the SFHA cluster.

Figure 9-1 Workflow to configure I/O fencing

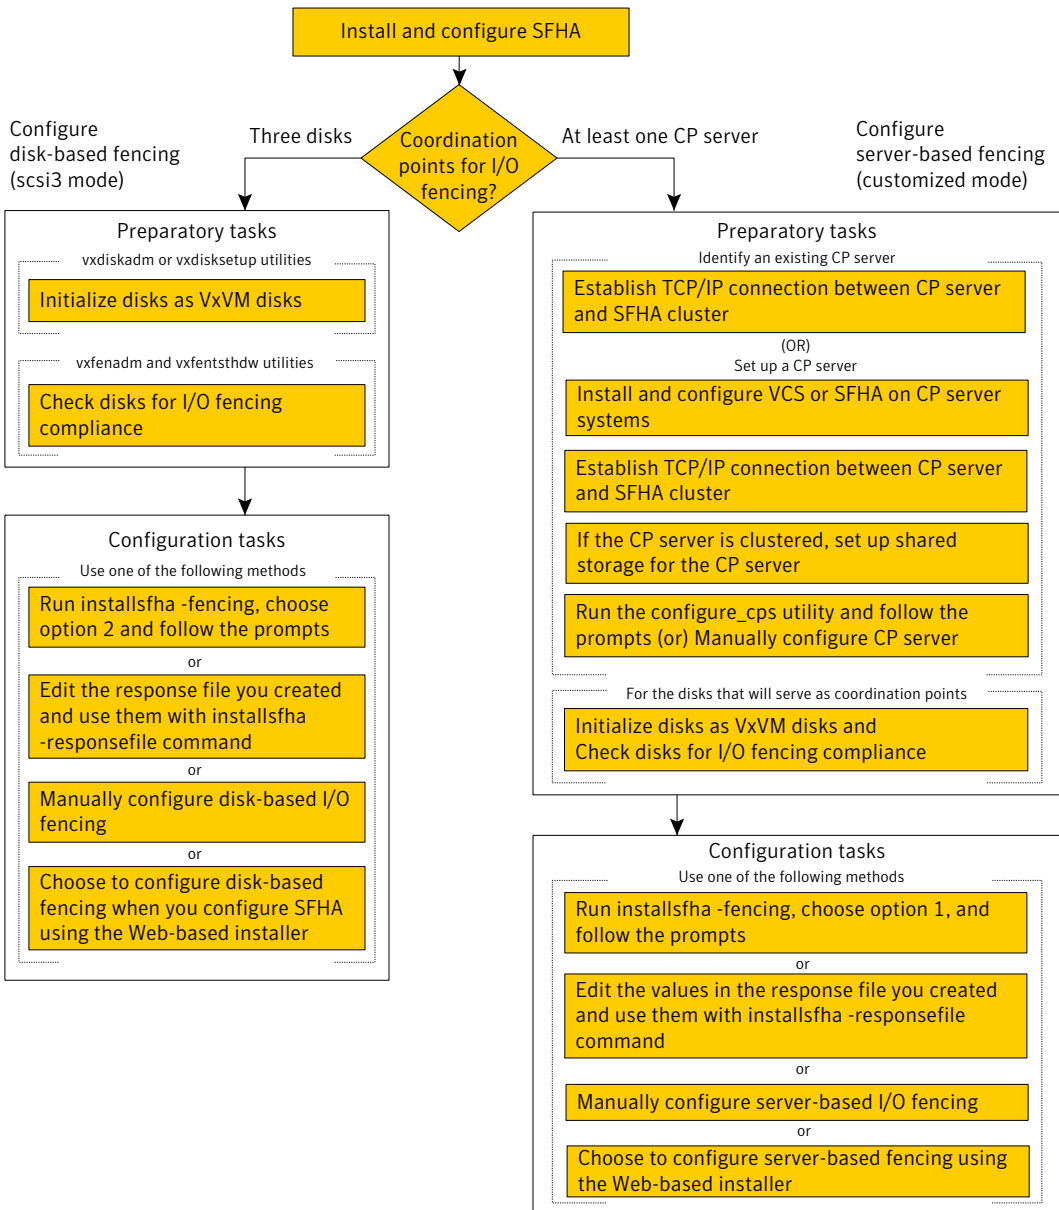
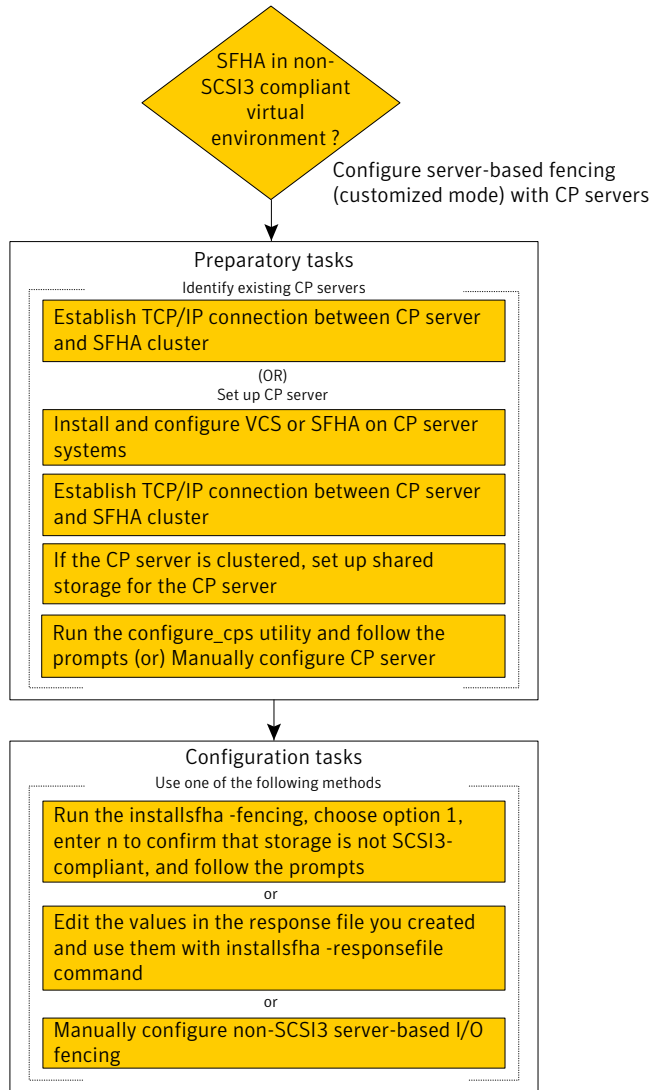


Figure 9-2 illustrates a high-level flowchart to configure non-SCSI-3 server-based I/O fencing for the SFHA cluster in virtual environments that do not support SCSI-3 PR.

Figure 9-2 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 `installsfha`
 - See [“Setting up disk-based I/O fencing using `installsfha`”](#) on page 125.
 - See [“Setting up server-based I/O fencing using `installsfha`”](#) on page 139.
 - See [“Setting up non-SCSI-3 server-based I/O fencing in virtual environments using `installsfha`”](#) on page 139.

- Using the Web-based installer
 - See [“Configuring SFHA for data integrity using the Web-based installer”](#) on page 141.

- Using response files
 - See [“Response file variables to configure disk-based I/O fencing”](#) on page 328.
 - See [“Response file variables to configure server-based I/O fencing”](#) on page 330.
 - See [“Response file variables to configure non-SCSI-3 server-based I/O fencing”](#) on page 333.
 - See [“Configuring I/O fencing using response files”](#) on page 327.

- Manually editing configuration files
 - See [“Setting up disk-based I/O fencing manually”](#) on page 134.
 - See [“Setting up server-based I/O fencing manually”](#) on page 144.
 - See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 151.


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

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

Typical Storage Foundation High Availability configuration with server-based I/O fencing

[Figure 9-3](#) displays a configuration using a Storage Foundation High Availability (with two nodes), a single CP server, and two coordinator disks. The nodes within the Storage Foundation High Availability are connected to and communicate with each other using LLT links.

Figure 9-3 CP server, Storage Foundation High Availability, and coordinator disks



Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points
 See [Figure 9-4](#) on page 81.
- Multiple application clusters use a single CP server and multiple pairs of coordinator disks (two) as their coordination points
 See [Figure 9-5](#) on page 82.
- Multiple application clusters use a single CP server as their coordination point
 This single coordination point fencing configuration must use a highly available CP server that is configured on an SFHA cluster as its coordination point.
 See [Figure 9-6](#) on page 82.

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

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

Figure 9-4 displays a configuration using three CP servers that are connected to multiple application clusters.

Figure 9-4 Three CP servers connecting to multiple application clusters




Figure 9-5 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.

Figure 9-5 Single CP server with two coordinator disks for each application cluster





Figure 9-6 displays a configuration using a single CP server that is connected to multiple application clusters.

Figure 9-6 Single CP server connecting to multiple application clusters



See [“Configuration diagrams for setting up server-based I/O fencing”](#) on page 379.

Setting up the CP server

Table 9-1 lists the tasks to set up the CP server for server-based I/O fencing.

Table 9-1 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 83.
Install the CP server	See “Installing the CP server using the installer” on page 84.
Configure the CP server cluster in secure mode	See “Configuring the CP server cluster in secure mode” on page 85.
Set up shared storage for the CP server database	See “Setting up shared storage for the CP server database” on page 86.
Configure the CP server	See “Configuring the CP server using the configuration utility” on page 87. See “Configuring the CP server manually” on page 96.
Verify the CP server configuration	See “Verifying the CP server configuration” on page 98.

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 Storage Foundation High Availability (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.
 See [“CP server requirements”](#) on page 36.
- 5 Have the following information handy for CP server configuration:
 - Name for the CP server
 The CP server name should not contain any special characters. CP server name can include alphanumeric characters, underscore, and hyphen.
 - Port number for the CP server
 Allocate a TCP/IP port for use by the CP server. Valid port range is between 49152 and 65535. The default port number is 14250.
 - Virtual IP address, network interface, netmask, and networkhosts for the CP server
 You can configure multiple virtual IP addresses for the CP server.

Installing the CP server using the installer

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

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

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

CP server setup uses a single system	<p>Install and configure VCS to create a single-node VCS cluster.</p> <p>During installation, make sure to select all RPMs for installation. The VRTScps RPM is installed only if you select to install all RPMs.</p> <p>See the <i>Veritas Cluster Server Installation Guide</i> for instructions on installing and configuring VCS.</p> <p>Proceed to configure the CP server.</p> <p>See “Configuring the CP server using the configuration utility” on page 87.</p> <p>See “Configuring the CP server manually” on page 96.</p>
CP server setup uses multiple systems	<p>Install and configure SFHA to create an SFHA cluster. This makes the CP server highly available.</p> <p>Meet the following requirements for CP server:</p> <ul style="list-style-type: none"> ■ During installation, make sure to select all RPMs for installation. The VRTScps RPM is installed only if you select to install all RPMs. ■ During configuration, configure disk-based fencing (scsi3 mode). <p>Proceed to set up shared storage for the CP server database.</p>

Configuring the CP server cluster in secure mode

You must configure security on the CP server only if you want to secure the communication between the CP server and the Storage Foundation High Availability (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 RPM.

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 87.

For CP servers on an SFHA cluster: See [“To configure the CP server on an SFHA cluster”](#) on page 91.

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

- 1 Verify that the VRTScps RPM 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: eth1
Enter a valid network interface for virtual IP 10.209.83.87
on mycps1.symantecexample.com: eth1
```

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
eth1 on system mycps1? [y/n] : y
Enter a valid IP address to configure NetworkHosts for
NIC eth1 on system mycps1 : 10.209.83.86
Do you want to add another Network Host ?[y/n] : n
```

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

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

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

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

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

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

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

For example:

```
# hagrps -state CPSSG

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

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

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

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

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

See [“Sample configuration files for CP server”](#) on page 347.

To configure the CP server on an SFHA cluster

- 1 Verify that the `VRTScps` RPM 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 : eth1
```

- 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 eth1 on system mycps1? [y/n] : y

Enter a valid IP address to configure NetworkHosts for NIC eth1 on system mycps1 : 10.209.83.86

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

17 Enter the netmask for each virtual IP address.

Enter the netmask for virtual IP 10.209.83.85 :
255.255.252.0

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

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

Enter the name of diskgroup for cps database :
cps_dg

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

Enter the name of volume created on diskgroup cps_dg :
cps_volume

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

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

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

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

For example:

```
# hagrp -state CPSSG

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

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

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

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

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

See [“Sample configuration files for CP server”](#) on page 347.

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 347.

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/VRTSvcs/db` (default location for CP server database)
- 2 Run the `cpsadm` command to check if the `vxcpserv` process is listening on the configured Virtual IP.

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

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

Configuring 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 10-1](#) lists the tasks that are involved in configuring SFHA using the script-based installer.

Table 10-1 Tasks to configure SFHA using the script-based installer

Task	Reference
Start the software configuration	See “Starting the software configuration” on page 101.
Specify the systems where you want to configure SFHA	See “Specifying systems for configuration” on page 102.
Configure the basic cluster	See “Configuring the cluster name” on page 103. See “Configuring private heartbeat links” on page 103.
Configure virtual IP address of the cluster (optional)	See “Configuring the virtual IP of the cluster” on page 106.
Configure the cluster in secure mode (optional)	See “Configuring the cluster in secure mode” on page 107.
Add VCS users (required if you did not configure the cluster in secure mode)	See “Adding VCS users” on page 110.
Configure SMTP email notification (optional)	See “Configuring SMTP email notification” on page 111.
Configure SNMP email notification (optional)	See “Configuring SNMP trap notification” on page 113.
Configure global clusters (optional) Note: You must have enabled Global Cluster Option when you installed SFHA.	See “Configuring global clusters” on page 115.
Complete the software configuration	See “Completing the SFHA configuration” on page 116.

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
- Checks whether SFHA is installed

- Exits if Veritas Storage Foundation 6.0 is not installed
- 3 Review the installer output about the I/O fencing configuration and confirm whether you want to configure fencing in enabled mode.

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

See [“About planning to configure I/O fencing”](#) on page 75.

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 384.

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.
 Make sure that you activated the NICs for the installer to be able to detect and automatically configure the heartbeat links.
 Skip to step 5.

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

The installer discovers and lists the network interface cards.

You must not enter the network interface card that is used for the public network (typically eth0.)

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

[b,q,?] **eth1**

eth1 has an IP address configured on it. It could be a public NIC on galaxy.

Are you sure you want to use eth1 for the first private heartbeat link? [y,n,q,b,?] (n) **y**

Would you like to configure a second private heartbeat link?

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

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

[b,q,?] **eth2**

eth2 has an IP address configured on it. It could be a public NIC on galaxy.

Are you sure you want to use eth2 for the second private heartbeat link? [y,n,q,b,?] (n) **y**

Would you like to configure a third private heartbeat link?

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

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

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

```

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

```

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

```

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

```

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

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

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

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

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

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

- 6 Enter a unique cluster ID:

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

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

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

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

Configuring the virtual IP of the cluster

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

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

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

To configure the virtual IP of the cluster

- 1 Review the required information to configure the virtual IP of the cluster.
- 2 When the system prompts whether you want to configure the virtual IP, enter `y`.
- 3 Confirm whether you want to use the discovered public NIC on the first system. Do one of the following:
 - If the discovered NIC is the one to use, press `Enter`.
 - If you want to use a different NIC, type the name of a NIC to use and press `Enter`.

```
Active NIC devices discovered on galaxy: eth1
Enter the NIC for Virtual IP of the Cluster to use on galaxy:
[b,q,?] (eth1)
```

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 eth1 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 10-2](#) lists the tasks that you must perform to configure a secure cluster.

Table 10-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See “ Configuring the first node ” on page 108.
Configure security on the remaining nodes	See “ Configuring the remaining nodes ” on page 109.

Table 10-2 Configuring a secure cluster node by node (*continued*)

Task	Reference
Complete the manual configuration steps	See “Completing the secure cluster configuration” on page 109.

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 113.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (eth0)
Is eth0 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: eth0
```

```
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 115.

3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: eth0
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (eth0)
Is eth0 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: eth0

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: eth0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

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

For IPv6: Global Cluster Option configuration verification:

```
NIC: eth0
IP: 2001:454e:205a:110:203:baff:feee:10
Prefix: 64
```

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

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 298.

Verifying the NIC configuration

The installer verifies on all the nodes if all NICs have `PERSISTENT_NAME` set correctly.

If the persistent interface names are not configured correctly for the network devices, the installer displays the following messages:

```
PERSISTENT_NAME is not set for all the NICs.
You need to set them manually before the next reboot.
```

Set the `PERSISTENT_NAME` for all the NICs.

Warning: If the installer finds the network interface name to be different from the name in the configuration file, then the installer exits.

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 118.

See [“Updating product licenses using vxlicinst”](#) on page 118.

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         = xxxxxx
License Type          = PERMANENT
OEM ID                = xxxxxx

Features :=
Platform              = Linux
Version               = 6.0
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 119.

To update product licenses

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

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

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 SFHA using the Web-based installer

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

By default, the communication between the systems is selected as SSH. If SSH is used for communication between systems, the SSH commands execute without prompting for passwords or confirmations.

You can click **Quit** to quit the Web-installer at any time during the configuration process.

To configure SFHA on a cluster

- 1 Start the Web-based installer.

See [“Starting the Veritas Web-based installer”](#) on page 61.

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

Task	Configure a Product
Product	Storage Foundation and High Availability

Click **Next**.

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

Example: **galaxy nebula**

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.

Click **Next** after the installer completes the system verification successfully.

- 4 In the Confirmation dialog box that appears, choose whether or not to configure I/O fencing.

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

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

See [“Configuring SFHA for data integrity using the Web-based installer”](#) on page 141.

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

- 5 On the Set Cluster Name/ID page, specify the following information for the cluster.

Cluster Name	Enter a unique cluster name.
Cluster ID	Enter a unique cluster ID. Note that you can have the installer check to see if the cluster ID is unique. Symantec recommends that you use the installer to check for duplicate cluster IDs in multi-cluster environments.
Check duplicate cluster ID	Select the check box if you want the installer to verify if the given cluster ID is unique in your private network. The verification is performed after you specify the heartbeat details in the following pages. The verification takes some time to complete.
LLT Type	Select an LLT type from the list. You can choose to configure LLT over UDP or over Ethernet. If you choose Auto detect over Ethernet , the installer auto-detects the LLT links over Ethernet. Verify the links and click Yes in the Confirmation dialog box. Skip to step 7. If you click No , you must manually enter the details to configure LLT over Ethernet.
Number of Heartbeats	Choose the number of heartbeat links you want to configure.
NICs	Choose the NICs that you want to configure.
Additional Low Priority Heartbeat NIC	Select the check box if you want to configure a low priority link. The installer configures one heartbeat link as low priority link.
Unique Heartbeat NICs per system	For LLT over Ethernet, select the check box if you do not want to use the same NIC details to configure private heartbeat links on other systems. For LLT over UDP, this check box is selected by default.

Click **Next**.

- 6 On the Set Cluster Heartbeat page, select the heartbeat link details for the LLT type you chose on the Set Cluster Name/ID page.

For **LLT over Ethernet**: Do the following:

- If you are using the same NICs on all the systems, select the NIC for each private heartbeat link.
- If you had selected **Unique Heartbeat NICs per system** on the Set Cluster Name/ID page, provide the NIC details for each system.

For **LLT over UDP**: Select the NIC, Port, and IP address for each private heartbeat link. You must provide these details for each system.

Click **Next**.

- 7 On the Optional Configuration page, decide the optional VCS features that you want to configure. Click the corresponding tab to specify the details for each option:

Virtual IP

- Select the **Configure Virtual IP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select the interface on which you want to configure the virtual IP.
- Enter a virtual IP address and value for the netmask. You can use an IPv4 or an IPv6 address.

VCS Users

- Reset the password for the Admin user, if necessary.
- Select the **Configure VCS users** option.
- Click **Add** to add a new user.
Specify the user name, password, and user privileges for this user.

SMTP

- Select the **Configure SMTP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
- In the **SMTP Server** box, enter the domain-based hostname of the SMTP server. Example: smtp.yourcompany.com
- In the **Recipient** box, enter the full email address of the SMTP recipient. Example: user@yourcompany.com.
- In the **Event** list box, select the minimum security level of messages to be sent to each recipient.
- Click **Add** to add more SMTP recipients, if necessary.

SNMP

- Select the **Configure SNMP** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- If all the systems use the same NIC, select the NIC for the VCS Notifier to be used on all systems. If not, select the NIC to be used by each system.
- In the **SNMP Port** box, enter the SNMP trap daemon port: (162).
- In the **Console System Name** box, enter the SNMP console system name.
- In the **Event** list box, select the minimum security level of messages to be sent to each console.
- Click **Add** to add more SNMP consoles, if necessary.

GCO

If you installed a valid HA/DR license, you can now enter the wide-area heartbeat link details for the global cluster that you would set up later.

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

- Select the **Configure GCO** check box.
- If each system uses a separate NIC, select the **Configure NICs for every system separately** check box.
- Select a NIC.
- Enter a virtual IP address and value for the netmask. You can use an IPv4 or an IPv6 address.

Security

To configure a secure SFHA cluster, select the **Configure secure cluster** check box.

If you want to perform this task later, do not select the **Configure secure cluster** check box. You can use the `-security` option of the `installsfha`.

Click **Next**.

- 8 On the Stop Processes page, click **Next** after the installer stops all the processes successfully.
- 9 On the Start Processes page, click **Next** after the installer performs the configuration based on the details you provided and starts all the processes successfully.

If you did not choose to configure I/O fencing in step 4, then skip to step 11. Go to step 10 to configure fencing.

- 10 On the Select Fencing Type page, choose the type of fencing configuration:

Configure Choose this option to configure server-based I/O fencing.

**Coordination Point
client based fencing**

**Configure disk based
fencing** Choose this option to configure disk-based I/O fencing.

Based on the fencing type you choose to configure, follow the installer prompts.

See [“Configuring SFHA for data integrity using the Web-based installer”](#) on page 141.

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

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

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

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

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

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

```
# fdisk -l
```

- 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 sdr
```

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

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

The `vxfcntlshdw` 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 127.

- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 128.
- Testing the shared disks for SCSI-3
See [“Testing the disks using `vxfsntsthdw` utility”](#) on page 129.

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
libvxhitachi.so	HITACHI	DF350, DF400, DF400F, DF500, DF500F
libvxxp1281024.so	HP	All
libvxxp12k.so	HP	All
libvxdds2a.so	DDN	S2A 9550, S2A 9900, S2A 9700
libvxpurple.so	SUN	T300
libvxxiotechE5k.so	XIOTECH	ISE1400
libvxcopan.so	COPANSYS	8814, 8818
libvxibm8k.so	IBM	2107

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

```
# vxdisk scandisks
```

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

Verifying that the nodes have access to the same disk

Before you test the disks that you plan to use as shared data storage or as coordinator disks using the vxfcntl utility, you must verify that the systems see the same disk.

To verify that the nodes have access to the same disk

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

From node A, enter:

```
# vxfenadm -i /dev/sdx
```

```
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
```

```
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/sdy` path.

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

```
# vxfenadm -i /dev/sdz
```

```
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
```

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

Testing the disks using `vxfcntlsthdw` utility

This procedure uses the `/dev/sdx` disk in the steps.

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

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

```
The disk /dev/sdx 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 `vxfcntlsthdw` utility

- 1 Make sure system-to-system communication functions properly.
- 2 From one node, start the utility.

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

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

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

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

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

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

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

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

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

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

Enter the disk name to be checked for SCSI-3 PGR on node galaxy in the format:

for dmp: /dev/vx/rdmp/sdx

for raw: /dev/sdx

Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/sdr

Enter the disk name to be checked for SCSI-3 PGR on node nebula in the format:

for dmp: /dev/vx/rdmp/sdx

for raw: /dev/sdx

Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/sdr

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/sdx have PASSED

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

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

Configuring disk-based I/O fencing using installsfha

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 installsfha

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

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

The installsfha 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 is configured properly.

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

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

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

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

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

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

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

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

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

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
 - Enter the disk group name.
- 6 Verify that the coordinator disks you chose meet the I/O fencing requirements. You must verify that the disks are SCSI-3 PR compatible using the `vxfcntlsthdw` utility and then return to this configuration program.
- See [“Checking shared disks for I/O fencing”](#) on page 126.
- 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 148.

Setting up disk-based I/O fencing manually

[Table 11-1](#) lists the tasks that are involved in setting up I/O fencing.

Table 11-1 Tasks to set up I/O fencing manually

Task	Reference
Initializing disks as VxVM disks	See “Initializing disks as VxVM disks” on page 126.
Identifying disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 135.
Checking shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 126.
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 135.
Creating I/O fencing configuration files	See “Creating I/O fencing configuration files” on page 136.
Modifying SFHA configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 137.
Configuring CoordPoint agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 148.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 138.

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 126.

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 126.

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 `sdx`, `sdz`, and `sdz`.

To create the `vxfencoorddg` disk group

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

```
# vxdg init vxfencoorddg sdx sdy sdz
```

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

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

- 1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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


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

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

- 4 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/sysconfig/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.

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

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

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

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

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

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

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

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

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

For example, on each remaining node, enter:

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

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

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

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

- Start VCS.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

- 1 On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is SCSI3 and the SCSI3 disk policy is dmp:

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

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

```
 * 0 (galaxy)
  1 (nebula)
```

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

- 2 Verify that the disk-based I/O fencing is using the specified disks.

```
# vxfenconfig -l
```

Setting up server-based I/O fencing using installsfha

See [“Setting up server-based I/O fencing using installsfha”](#) on page 139.

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 installsfha in a non-SCSI-3 PR-compliant setup

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

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

The installsfha 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 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 Storage Foundation High Availability nodes that host the applications for high availability.

- 8 Verify and confirm the CP server information that you provided.
- 9 Verify and confirm the Storage Foundation High Availability 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 Storage Foundation High Availability with the CP server.
 - Adds CP server user to the CP server.
 - Adds Storage Foundation High Availability to the CP server user.
 - Updates the following configuration files on each node of the Storage Foundation High Availability
 - `/etc/vxfenmode` file
 - `/etc/vxenviron` file
 - `/etc/sysconfig/vxfen` 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 Storage Foundation High Availability.
- 11 Confirm whether you want to send the installation information to Symantec.
- 12 After the installer configures I/O fencing successfully, note the location of summary, log, and response files that installer creates.
- The files provide useful information which can assist you with the configuration, and can also assist future configurations.

Configuring SFHA for data integrity using the Web-based installer

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

See [“Configuring SFHA using the Web-based installer”](#) on page 119.

See [“About planning to configure I/O fencing”](#) on page 75.

To configure SFHA for data integrity

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 61.
- 2 On the Select a task and a product page, select the task and the product as follows:

Task	I/O fencing configuration
Product	Storage Foundation and High Availability

Click **Next**.

- 3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.
- 4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

The installer performs the initial system verification. It checks for the system communication. It also checks for release compatibility, installed product version, platform version, and performs product prechecks.
- 5 On the Select Fencing Type page, choose whether to configure disk-based fencing or server-based fencing.

If you chose to configure disk-based fencing, go to step [7](#).
If you chose to configure server-based fencing, go to step [10](#).
- 6 In the Confirmation dialog box that appears, confirm whether your storage environment supports SCSI-3 PR.

You can configure non-SCSI-3 server-based fencing in a virtual environment that is not SCSI-3 PR compliant.
- 7 On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.
Click **Next**.
- 8 On the Configure Fencing page, specify the following information:

Select a Disk Group Select the **Create a new disk group** option or select one of the disk groups from the list.

- If you selected one of the disk groups that is listed, choose the fencing disk policy for the disk group.
Go to step [16](#).
- If you selected the **Create a new disk group** option, make sure you have SCSI-3 PR enabled disks, and click **Yes** in the confirmation dialog box.
Click **Next**. Go to step [9](#).

9 On the Create New DG page, specify the following information:

New Disk Group Name Enter a name for the new coordinator disk group you want to create.

Select Disks Select at least three disks to create the coordinator disk group.

If you want to select more than three disks, make sure to select an odd number of disks.

Fencing Disk Policy Choose the fencing disk policy for the disk group.

Go to step [16](#).

10 On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.

Click **Next**.

11 Provide the following details for each of the CP servers:

- Enter the virtual IP addresses or host names of the virtual IP address. The installer assumes these values to be identical as viewed from all the application cluster nodes.
- Enter the port that the CP server must listen on.
- Click **Next**.

12 If your server-based fencing configuration also uses disks as coordination points, perform the following steps:

- If you have not already checked the disks for SCSI-3 PR compliance, check the disks now, and click **OK** in the dialog box.
- If you do not want to use the default coordinator disk group name, enter a name for the new coordinator disk group you want to create.

- Select the disks to create the coordinator disk group.
 - Choose the fencing disk policy for the disk group.
- 13 In the Confirmation dialog box that appears, confirm whether the coordination points information you provided is correct, and click **Yes**.
 - 14 Verify and confirm the I/O fencing configuration information.
 The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.
 - 15 Configure the CP agent on the SFHA (application cluster), and click **Next**.
 - 16 Click **Next** to complete the process of configuring I/O fencing.
 On the Completion page, view the summary file, log file, or response file, if needed, to confirm the configuration.
 - 17 Select the checkbox to specify whether you want to send your installation information to Symantec.
 Click **Finish**. The installer prompts you for another task.

Setting up server-based I/O fencing manually

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

Table 11-2 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 Storage Foundation High Availability” on page 145.
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 137.
Configuring Coordination Point agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 148.
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 150.

Preparing the CP servers manually for use by the Storage Foundation High Availability

Use this procedure to manually prepare the CP server for use by the Storage Foundation High Availability or clusters.

[Table 11-3](#) displays the sample values used in this procedure.

Table 11-3 Sample values in procedure

CP server configuration component	Sample name
CP server	mycps1
Node #1 - Storage Foundation High Availability	galaxy
Node #2 - Storage Foundation High Availability	nebula
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the Storage Foundation High Availability

- 1 Determine the cluster name and uuid on the Storage Foundation High Availability.

For example, issue the following commands on one of the Storage Foundation High Availability 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 Storage Foundation High Availability 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 Storage Foundation High Availability and nodes to each CP server.

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

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

```
Cluster clus1 added successfully
```

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

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

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

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

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

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

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

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

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

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

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

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

5 Add the users to the CP server.

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

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

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

6 Authorize the CP server user to administer the Storage Foundation High Availability. You must perform this task for the CP server users corresponding to each node in the Storage Foundation High Availability.

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

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

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

Configuring 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 Storage Foundation High Availability has been properly installed and configured with fencing enabled.
- 2 Create a parallel service group vxfen and add a coordpoint resource to the vxfen service group using the following commands:

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

- 3 Verify the status of the agent on the Storage Foundation High Availability 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 144.

- 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/vxenvron` file as follows:

```
data_disk_fencing=off
```

- 5 On each node, edit the `/etc/sysconfig/vxfen` file as follows:

```
vxfen_vxfnd_tmt=25
```

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

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

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

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

- Run the following command:

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

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

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

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

- Make the VCS configuration file writable:

```
# haconf -makerw
```

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

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

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

- Run the following command to verify the value:

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

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

The command should not list any resources.

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

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

- Make the VCS configuration file read-only

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


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

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

```
# /etc/init.d/vcs stop
```

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

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

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

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

```
# /etc/init.d/vcs start
```

Sample /etc/vxfenmode file for non-SCSI-3 fencing

```
=====
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps        - use a coordination point server with optional script
#              controlled scsi3 disks
#
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is required
```

```
# only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
# scsi3_disk_policy=dmp

#
# Seconds for which the winning sub cluster waits to allow for the
# losing subcluster to panic & drain I/Os. Useful in the absence of
# SCSI3 based data disk fencing
loser_exit_delay=55

#
# Seconds for which vxfend process wait for a customized fencing
# script to complete. Only used with vxfen_mode=customized
vxfen_script_timeout=25

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

#
# Specify 3 or more odd number of coordination points in this file,
# one in its own line. They can be all-CP servers, all-SCSI-3
# compliant coordinator disks, or a combination of CP servers and
# SCSI-3 compliant coordinator disks. Please ensure that the CP
# server coordination points are numbered sequentially and in the
# same order on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as:
#
# cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames over
# different subnets, all of the IPs/names can be specified in a
```

```
# comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,...,
#  [<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
# is the serial number of the CPS as a coordination point; must
# start with 1.
# <vip>
# is the virtual IP address of the CPS, must be specified in
# square brackets ("[]").
# <vhn>
# is the virtual hostname of the CPS, must be specified in square
# brackets ("[]").
# <port>
# is the port number bound to a particular <vip/vhn> of the CPS.
# It is optional to specify a <port>. However, if specified, it
# must follow a colon (":") after <vip/vhn>. If not specified, the
# colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>, a
# default port can be specified as follows:
#
# port=<default_port>
#
# Where <default_port> is applicable to all the <vip/vhn>s for
# which a <port> is not specified. In other words, specifying <port>
# with a <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[mycps1.company.com]
# cps2=[192.168.0.25]
# cps3=[mycps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [mycps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
```

```
# [192.168.0.23]
# [mycps1.company.com]
# [192.168.0.25]
# - if default port 57777 were not specified, port 14250 would be used
# for all remaining <vip/vhn>s:
# [192.168.0.23]
# [mycps1.company.com]
# [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoorddg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
cps1=[mycps1.company.com]
cps2=[mycps2.company.com]
cps3=[mycps3.company.com]
port=14250
=====
```

Enabling or disabling the preferred fencing policy

You can enable or disable the preferred fencing feature for your I/O fencing configuration.

You can enable preferred fencing to use system-based race policy or group-based race policy. If you disable preferred fencing, the I/O fencing configuration uses the default count-based race policy.

See [“About preferred fencing”](#) on page 27.

To enable preferred fencing for the I/O fencing configuration

- 1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

- 2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

- 3 To enable system-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as System.

```
# haclus -modify PreferredFencingPolicy System
```

- Set the value of the system-level attribute FencingWeight for each node in the cluster.

For example, in a two-node cluster, where you want to assign galaxy five times more weight compared to nebula, run the following commands:

```
# hasys -modify galaxy FencingWeight 50  
# hasys -modify nebula FencingWeight 10
```

- Save the VCS configuration.

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

- 4 To enable group-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Group.

```
# haclus -modify PreferredFencingPolicy Group
```

- Set the value of the group-level attribute `Priority` for each service group. For example, run the following command:

```
# hagrps -modify service_group Priority 1
```

Make sure that you assign a parent service group an equal or lower priority than its child service group. In case the parent and the child service groups are hosted in different subclusters, then the subcluster that hosts the child service group gets higher preference.

- Save the VCS configuration.

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

- 5 To view the fencing node weights that are currently set in the fencing driver, run the following command:

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

- 1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

- 2 Make sure that the cluster-level attribute `UseFence` has the value set to `SCSI3`.

```
# haclus -value UseFence
```

- 3 To disable preferred fencing and use the default race policy, set the value of the cluster-level attribute `PreferredFencingPolicy` as `Disabled`.

```
# haconf -makerw
```

```
# haclus -modify PreferredFencingPolicy Disabled
```

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

Upgrading Veritas Storage Foundation and High Availability

- [Chapter 12. Preparing to upgrade](#)
- [Chapter 13. Upgrading Storage Foundation and High Availability](#)
- [Chapter 14. Performing a rolling upgrade](#)
- [Chapter 15. Performing a phased upgrade](#)
- [Chapter 16. Performing post-upgrade tasks](#)

Preparing to upgrade

This chapter includes the following topics:

- [About upgrading](#)
- [About the different ways that you can upgrade](#)
- [Supported upgrade paths](#)
- [About using the installer to upgrade when the root disk is encapsulated](#)
- [Preparing to upgrade](#)

About upgrading

There are many types of upgrades available. Before you start to upgrade, review the types of upgrades for the Veritas products.

See [“About the different ways that you can upgrade”](#) on page 161.

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

See [“Supported upgrade paths”](#) on page 161.

After you determine the type of upgrade that you want to perform and its upgrade paths, review the steps to prepare for the upgrade.

If you want to upgrade Coordination Point (CP) server systems that use Veritas Cluster Server (VCS) or Storage Foundation High Availability (SFHA) to 6.0, make sure you upgraded all application clusters to 6.0. Then, upgrade VCS or SFHA on the CP server systems.

About the different ways that you can upgrade

Symantec offers you several different ways to upgrade. You need to decide which upgrade method best suits your environment, your expertise, and the downtime required.

Table 12-1 Review this table to determine how you want to perform the upgrade

Upgrade types and considerations	Methods available for upgrade
Typical upgrades—use a Veritas provided tool or you can perform the upgrade manually. Requires some server downtime.	<p>Script-based—you can use this to upgrade for the supported upgrade paths</p> <p>Web-based—you can use this to upgrade for the supported upgrade paths</p> <p>Manual—you can use this to upgrade from the previous release</p> <p>Response file—you can use this to upgrade from the previous release</p>
Rolling upgrade—use a Veritas provided tool or you can perform the upgrade manually. Requires least amount of server downtime.	<p>Script-based—you can use this to upgrade from the previous release</p> <p>Web-based—you can use this to upgrade from the previous release</p>
Phased upgrades—use a Veritas provided tool and some manual steps. Requires less server downtime than a regular upgrade.	Script-based with some manual steps—you can use this to upgrade from the previous release
Native operating system upgrade—use the upgrade software that comes with the operating system. Note that not all operating systems support native upgrades.	<p>Operating system specific methods</p> <p>Operating system upgrades</p>

Supported upgrade paths

[Table 12-2](#) describes upgrading to 6.0 on RHEL 5 x64.

Table 12-2 RHEL 5 x64 upgrades using the script- or Web-based installer

Veritas software versions	RHAS 2.1 or RHEL 3	RHEL 4	RHEL 5
3.4.x 4.0	No upgrade path exists. Uninstall the product. Upgrade to RHEL 5 U5. Use the installer to install 6.0.	N/A	N/A
4.1 4.1 MP1 4.1 MP2 4.1 MP3 4.1 MP4	N/A	No upgrade path exists. Uninstall the product. Upgrade to RHEL 5 U5. Use the installer to install 6.0.	No upgrade path exists. Uninstall the product. Upgrade to RHEL 5 U5. Use the installer to install 6.0.
5.0 5.0 MP1 5.0 MP2	N/A	No upgrade path exists. Uninstall the product. Upgrade to RHEL 5 U5. Use the installer to install 6.0.	N/A
5.0 MP3 5.0 MP4	N/A	No upgrade path exists. Uninstall the product. Upgrade to RHEL 5 U5. Use the installer to install 6.0.	Upgrade to RHEL5 U5. Use the installer to upgrade to 6.0.
5.1 5.1 RPx 5.1 PR1 5.1 SP1 5.1 SP1 RPx 5.1 SP1 PR3	N/A	N/A	Upgrade to RHEL5 U5. Use the installer to upgrade to 6.0.

Table 12-3 describes upgrading to 6.0 on RHEL 6 x64.

Table 12-3 RHEL6 x64 upgrades using the script- or Web-based installer

Veritas software versions	RHEL 4	RHEL5	RHEL 6
4.1 4.1 MP1 4.1 MP2 4.1 MP3 4.1 MP4	No upgrade path exists. Uninstall the product. Upgrade to RHEL 6 U1. Use the installer to install to 6.0.	No upgrade path exists. Uninstall the product. Upgrade to RHEL 6 U1. Use the installer to install to 6.0.	N/A
5.0 5.0 MP1 5.0 MP2	No upgrade path exists. Uninstall the product. Upgrade to RHEL 6 U1. Use the installer to install to 6.0.	N/A	N/A
5.0 MP3 5.0 MP4	No upgrade path exists. Uninstall product. Upgrade to RHEL 6 U1. Use the installer to install to 6.0.	No upgrade path exists. Uninstall product. Upgrade to RHEL 6 U1. Use the installer to install to 6.0.	N/A
5.1 5.1 RPx 5.1 PR1 5.1 SP1 5.1 SP1 RPx	N/A	Upgrade to RHEL U5, 6 U1, or 7. Use the installer to install to 6.0.	Upgrade to RHEL 6 U1 and then use the installer to upgrade to 6.0.
5.1 SP1 PR2	N/A	N/A	Upgrade to RHEL 6 U1 and then use the installer to upgrade to 6.0.

Table 12-4 describes upgrading to 6.0 on SLES 10 x64.

Table 12-4 SLES 10 x64 upgrades using the script- or Web-based installer

Veritas software versions	SLES 8	SLES 9	SLES 10
3.4.x	No upgrade path exists. Uninstall the product. Upgrade to SLES10 SP4. Use the installer to install 6.0.	N/A	N/A
4.1 4.1 MP1 4.1 MP2	N/A	No upgrade path exists. Uninstall the product. Upgrade to SLES10 SP4. Use the installer to install 6.0.	N/A
4.1 MP3 4.1 MP4	N/A	No upgrade path exists. Uninstall the product. Upgrade to SLES10 SP4. Use the installer to install 6.0.	No upgrade path exists. Uninstall the product. Upgrade to SLES10 SP4. Use the installer to install 6.0.
5.0 5.0 MP1 5.0 MP2	N/A	No upgrade path exists. Uninstall the product. Upgrade to SLES10 SP4. Use the installer to install 6.0.	N/A
5.0 MP3 5.0 MP4	N/A	No upgrade path exists. Uninstall the product. Upgrade to SLES10 SP4. Use the installer to install 6.0.	Upgrade to SLES10 SP4. Use the installer to upgrade to 6.0.
5.0 RU4 5.1 5.1 RPx 5.1 PR1 5.1 SP1 5.1 SP1 RPx	N/A	N/A	Upgrade to SLES10 SP4. Use the installer to upgrade to 6.0.

Table 12-5 describes upgrading to 6.0 on SLES 11 x64.

Table 12-5 SLES 11 x64 upgrades using the script- or Web-based installer

Veritas software versions	SLES 9	SLES 10	SLES 11
4.1 4.1 MP1 4.1 MP2 5.0 5.0 MP1 5.0 MP2	No upgrade path exists. Uninstall the product. Upgrade to SLES11 SP1. Use the installer to install 6.0.	N/A	N/A
4.1 MP3 4.1 MP4 5.0 MP3	No upgrade path exists. Uninstall the product. Upgrade to SLES11 SP1. Use the installer to install 6.0.	No upgrade path exists. Uninstall the product. Upgrade to SLES11 SP1. Use the installer to install 6.0.	N/A
5.0 RU4	N/A	No upgrade path exists. Uninstall the product. Upgrade to SLES11 SP1. Use the installer to install 6.0.	N/A
5.0 MP4	No upgrade path exists. Uninstall the product. Upgrade to SLES11 SP1. Use the installer to install 6.0.	No upgrade path exists. Uninstall the product. Upgrade to SLES11 SP1. Use the installer to install 6.0. all 6.0.	Upgrade to SLES11 SP1. Use the installer to upgrade to 6.0.
5.0 RU1	N/A	N/A	Upgrade to SLES11 SP1. Use the installer to upgrade to 6.0.
5.1 5.1 RPx 5.1 PR1 5.1 SP1 5.1 SP1 RPx	N/A	N/A	Upgrade to SLES11 SP1. Use the installer to upgrade to 6.0.

About using the installer to upgrade when the root disk is encapsulated

When you use the installer to upgrade from a previous version of SFHA and the system where you plan to upgrade has an encapsulated root disk, you may have to unencapsulate it.

Table 12-6 Upgrading using installer when the root disk is encapsulated

Starting version	Ending version	Action required
5.0 MP3	6.0	You need to unencapsulate the root disk. The installer exits.
5.1 or 5.1 RPx	6.0	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 SP1 or 5.1 SP1 RPx	6.0	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Preparing to upgrade

Before you upgrade, you need to prepare the systems and storage. Review the following procedures and perform the appropriate tasks.

Getting ready for the upgrade

Complete the following tasks before you perform the upgrade:

- Review the Symantec Technical Support website for additional information: <http://www.symantec.com/techsupp/>
- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the RPMs, for example `/packages/Veritas` when the root file

system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.

Do not put the files under `/tmp`, which is erased during a system reboot.

Do not put the files on a file system that is inaccessible prior to running the upgrade script.

You can use a Veritas-supplied disc for the upgrade as long as modifications to the upgrade script are not required.

- For any startup scripts in `/etc/rcS.d`, comment out any application commands or processes that are known to hang if their file systems are not present.
- Make sure that the current operating system supports version 6.0 of the product. If the operating system does not support it, plan for a staged upgrade.
- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Veritas products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/fstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/fstab` and not mounted during the upgrade. Active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure the file systems are clean before upgrading.
- Upgrade arrays (if required).
- If replication using VVR is configured, Symantec recommends that the disk group version is at least 110 prior to upgrading.

```
# vxvg list diskgroup
```

- If replication using VVR is configured, make sure the size of the SRL volume is greater than 110 MB.
Refer to the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.
- If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date on all the hosts.

```
# vxrlink -g diskgroup status rlink_name
```

Note: Do not continue until the primary RLINKs are up-to-date.

- If VCS is used to manage VVR replication, follow the preparation steps to upgrade VVR and VCS agents.

- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. Removing the disk in this manner offers a failback point.
- Determine if the root disk is encapsulated.
See [“Determining if the root disk is encapsulated”](#) on page 169.

Creating backups

Save relevant system information before the upgrade.

To create backups

- 1 Log in as superuser.
- 2 Before the upgrade, ensure that you have made backups of all data that you want to preserve.
- 3 Back up information in files such as `/boot/grub/menu.lst`, `/etc/grub.conf` or `/etc/lilo.conf`, and `/etc/fstab`.
- 4 Copy the `fstab` file to `fstab.orig`:


```
# cp /etc/fstab /etc/fstab.orig
```
- 5 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 6 If you are installing the high availability version of the Veritas Storage Foundation 6.0 software, follow the guidelines given in the *Veritas Cluster Server Installation Guide* and *Veritas Cluster Server Release Notes* for information on preserving your VCS configuration across the installation procedure.

Tasks for upgrading the Storage Foundation for Databases (SFDB)

Tasks for upgrading SFDB tools to version 6.0:

- Preparing to migrate the repository database before upgrading from 5.0x or earlier to 6.0
See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 169.
- Migrating the repository database after upgrading from 5.0.x or earlier to 6.0
See [“Post upgrade tasks for migrating the SFDB repository database”](#) on page 213.

Caution: If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 6.0: upgrade the Oracle binaries and database to version 11.1.0.7 before moving to 6.0.

Determining if the root disk is encapsulated

Check if the system's root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system.

If the root disk is encapsulated, follow the appropriate upgrade procedures.

Pre-upgrade tasks for migrating the SFDB repository database

If you are upgrading from 5.1 or 5.1SP1 to 6.0, no upgrade steps are required for the SFDB tools. The migration process is automatic.

Note: The `Sfua_Base` repository resource group will be removed from the `main.cf` file. It is not required as a separate service group for SFHA 6.0.

Perform the following before upgrading SFHA.

To prepare to migrate the repository database

- ◆ Resynchronize all existing snapshots before upgrading. As Oracle user, enter:

```
$ /opt/VRTS/bin/dbed_vmsnap -S $ORACLE_SID \  
-f SNAPPLAN -o resync
```

Warning: The Database Flashsnap clone database will not be able to be carried over after upgrading. You must create a new Database Flashsnap clone database after upgrading to 6.0.

Pre-upgrade planning for Veritas Volume Replicator

Before installing or upgrading Veritas Volume Replicator (VVR):

- Confirm that your system has enough free disk space to install VVR.

- Make sure you have root permissions. You must have root permissions to perform the install and upgrade procedures.

See the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for more information.

See the *Getting Started Guide* for more information on the documentation.

Planning an upgrade from the previous VVR version

If you plan to upgrade VVR from the previous VVR version, you can upgrade VVR with reduced application downtime by upgrading the hosts at separate times. While the Primary is being upgraded, the application can be migrated to the Secondary, thus reducing downtime. The replication between the (upgraded) Primary and the Secondary, which have different versions of VVR, will still continue. This feature facilitates high availability even when the VVR upgrade is not complete on both the sites. Symantec recommends that the Secondary hosts be upgraded before the Primary host in the Replicated Data Set (RDS).

See the *Storage Foundation and High Availability Release Notes* for information regarding VVR support for replicating across Storage Foundation versions

Replicating between versions is intended to remove the restriction of upgrading the Primary and Secondary at the same time. VVR can continue to replicate an existing RDS with Replicated Volume Groups (RVGs) on the systems that you want to upgrade. When the Primary and Secondary are at different versions, VVR does not support changing the configuration with the `vradmin` command or creating a new RDS.

Also, if you specify TCP as the network protocol, the VVR versions on the Primary and Secondary determine whether the checksum is calculated. As shown in [Table 12-7](#), if either the Primary or Secondary are running a version of VVR prior to 6.0, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 6.0, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 12-7 VVR versions and checksum calculations

VVR prior to 6.0 (DG version <= 140)	VVR 6.0 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes
Primary and Secondary		Yes
	Primary and Secondary	No

Note: When replicating between versions of VVR, avoid using commands associated with new features. The earlier version may not support new features and problems could occur.

If you do not need to upgrade all the hosts in the RDS simultaneously, you can use replication between versions after you upgrade one host. You can then upgrade the other hosts in the RDS later at your convenience.

Note: If you have a cluster setup, you must upgrade all the nodes in the cluster at the same time.

Planning and upgrading VVR to use IPv6 as connection protocol

Storage Foundation and High Availability supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

- VVR continues to support replication between IPv4-only nodes with IPv4 as the internet protocol
- VVR supports replication between IPv4-only nodes and IPv4/IPv6 dual-stack nodes with IPv4 as the internet protocol
- VVR supports replication between IPv6-only nodes and IPv4/IPv6 dual-stack nodes with IPv6 as the internet protocol
- VVR supports replication between IPv6 only nodes
- VVR supports replication to one or more IPv6 only nodes and one or more IPv4 only nodes from a IPv4/IPv6 dual-stack node
- VVR supports replication of a shared disk group only when all the nodes in the cluster that share the disk group are at IPv4 or IPv6

Preparing to upgrade VVR when VCS agents are configured

To prepare to upgrade VVR when VCS agents for VVR are configured, perform the following tasks in the order presented:

- [Freezing the service groups and stopping all the applications](#)
- [Preparing for the upgrade when VCS agents are configured](#)

Freezing the service groups and stopping all the applications

This section describes how to freeze the service groups and stop all applications.

To freeze the service groups and stop applications

Perform the following steps for the Primary and Secondary clusters:

- 1 Log in as the superuser.
- 2 Make sure that `/opt/VRTS/bin` is in your PATH so that you can execute all the product commands.
- 3 Before the upgrade, cleanly shut down all applications.
 - OFFLINE all application service groups that do not contain RVG resources. Do not OFFLINE the service groups containing RVG resources.
 - If the application resources are part of the same service group as an RVG resource, then OFFLINE only the application resources. In other words, ensure that the RVG resource remains ONLINE so that the private disk groups containing these RVG objects do not get deported.

Note: You must also stop any remaining applications not managed by VCS.

- 4 On any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 5 On any node in the cluster, list the groups in your configuration:

```
# hagrps -list
```

- 6 On any node in the cluster, freeze all service groups except the ClusterService group by typing the following command for each group name displayed in the output from step 5.

```
# hagrps -freeze group_name -persistent
```

Note: Make a note of the list of frozen service groups for future use.

- 7 On any node in the cluster, save the configuration file (`main.cf`) with the groups frozen:

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

Note: Continue only after you have performed steps 3 to step 7 for each cluster.

- 8 Display the list of service groups that have RVG resources and the nodes on which each service group is online by typing the following command on any node in the cluster:

```
# hares -display -type RVG -attribute State
Resource      Attribute      System      Value
VVRGp         State          system02    ONLINE
ORAGrp        State          system02    ONLINE
```

Note: For the resources that are ONLINE, write down the nodes displayed in the System column of the output.

- 9 Repeat step 8 for each cluster.
- 10 For private disk groups, determine and note down the hosts on which the disk groups are imported.

See [“Determining the nodes on which disk groups are online”](#) on page 173.

Determining the nodes on which disk groups are online

For private disk groups, determine and note down the hosts on which the disk groups containing RVG resources are imported. This information is required for restoring the configuration after the upgrade.

To determine the online disk groups

- 1 On any node in the cluster, list the disk groups in your configuration, and note down the disk group names listed in the output for future use:

```
# hares -display -type RVG -attribute DiskGroup
```

Note: Write down the list of the disk groups that are under VCS control.

- 2 For each disk group listed in the output in step 1, list its corresponding disk group resource name:

```
# hares -list DiskGroup=diskgroup Type=DiskGroup
```

- 3 For each disk group resource name listed in the output in step 2, get and note down the node on which the disk group is imported by typing the following command:

```
# hares -display dg_resname -attribute State
```

The output displays the disk groups that are under VCS control and nodes on which the disk groups are imported.

Preparing for the upgrade when VCS agents are configured

If you have configured the VCS agents, it is recommended that you take backups of the configuration files, such as `main.cf` and `types.cf`, which are present in the `/etc/VRTSvcs/conf/config` directory.

To prepare a configuration with VCS agents for an upgrade

- 1 List the disk groups on each of the nodes by typing the following command on each node:

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

The output displays a list of the disk groups that are under VCS control and the disk groups that are not under VCS control.

Note: The disk groups that are not locally imported are displayed in parentheses.

- 2 If any of the disk groups have not been imported on any node, import them. For disk groups in your VCS configuration, you can import them on any node. For disk groups that are not under VCS control, choose an appropriate node on which to import the disk group. Enter the following command on the appropriate node:

```
# vxdg -t import diskgroup
```

- 3 If a disk group is already imported, then recover the disk group by typing the following command on the node on which it is imported:

```
# vxrecover -bs
```

- 4 Verify that all the Primary RLINKs are up to date.

```
# vxrlink -g diskgroup status rlink_name
```

Note: Do not continue until the Primary RLINKs are up-to-date.

Upgrading the array support

The Storage Foundation 6.0 release includes all array support in a single RPM, VRTSaslapm. The array support RPM includes the array support previously included in the VRTSvxvm RPM. The array support RPM also includes support previously packaged as external array support libraries (ASLs) and array policy modules (APMs).

See the 6.0 Hardware Compatibility List for information about supported arrays.

See [“Hardware compatibility list \(HCL\)”](#) on page 33.

When you upgrade Storage Foundation products with the product installer, the installer automatically upgrades the array support. If you upgrade Storage Foundation products with manual steps, you should remove any external ASLs or APMs that were installed previously on your system. Installing the VRTSvxvm RPM exits with an error if external ASLs or APMs are detected.

After you have installed Storage Foundation 6.0, Symantec provides support for new disk arrays through updates to the `VRTSaslapm` RPM.

For more information about array support, see the *Veritas Storage Foundation Administrator's Guide*.

Upgrading Storage Foundation and High Availability

This chapter includes the following topics:

- [Upgrading Storage Foundation and High Availability from previous versions to 6.0](#)
- [Upgrading Storage Foundation and High Availability using the Veritas Web-based installer](#)
- [Upgrading Veritas Volume Replicator](#)

Upgrading Storage Foundation and High Availability from previous versions to 6.0

If you are running an earlier release of Storage Foundation and High Availability, you can upgrade to the latest version using the procedures described in this chapter.

For a cluster, use the appropriate procedures to upgrade Veritas Storage Foundation High Availability.

See [“Upgrading Storage Foundation and High Availability using the script-based installer”](#) on page 178.

If you need to upgrade your kernel with Veritas Storage Foundation 6.0 already installed, use the kernel upgrade procedure.

See the *Veritas Storage Foundation Administrator's Guide* for information about upgrading the kernel.

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

Use this procedure to upgrade Storage Foundation and High Availability (SFHA).

To upgrade Veritas Storage Foundation and High Availability

- 1 Log in as superuser.
- 2 Take all service groups offline.

List all service groups:

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

For each service group listed, take it offline:

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

- 3 Enter the following commands on each node to freeze HA service group operations:

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

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

```
# df -F | grep vxfs
```

- 5 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name  
# umount /filesystem
```

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

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

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

Perform the following steps in the order listed:

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

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

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

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

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

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

- If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
 - Repeat this step to verify that the unclean file system is now clean.
- 7 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.
 - 8 Stop all the volumes by entering the following command for each disk group:


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

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

```
# vxprint -Aht -e v_open
```
 - 9 Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to recreate these entries in the `/etc/fstab` file on the freshly installed system.
 - 10 Perform any necessary preinstallation checks.

- 11 To invoke the common installer, run the `installer` command on the disc as shown in this example:

```
# cd /cdrom/cdrom0
# ./installer
```

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

```
Enter the 64 bit RHEL5 system names separated
by spaces : [q, ?] host1
           host2
```

Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.

During the system verification phase, the installer checks if the boot disk is encapsulated and the upgrade's path. If the upgrade is not supported, you need to un-encapsulate the boot disk.

- 14 The installer asks if you agree with the terms of the End User License Agreement. Press `y` to agree and continue.
- 15 The installer lists the RPMs to install or upgrade. You are prompted to confirm that you are ready to upgrade.
- 16 The installer discovers if any of the systems that you are upgrading have mirrored encapsulated boot disks. You now have the option to create a backup of the systems' root disks before the upgrade proceeds. If you want to split the mirrors on the encapsulated boot disks to create the backup, answer `y`.
- 17 The installer then prompts you to name the backup root disk. Enter the name for the backup and mirrored boot disk or press **Enter** to accept the default.

Note: The split operation can take some time to complete.

- 18 You are prompted to start the split operation. Press `y` to continue.
- 19 The installer lists the RPMs that it installs or upgrades.
- 20 The Storage Foundation and High Availability software is verified and configured.

The software processes start.

- 21 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the "Administering Disks" chapter of the *Veritas Storage Foundation Administrator's Guide*.

Note: Upgrades from version 5.0 MP3 of the Veritas software do not require the re-encapsulation and mirroring of the root disk on each node.

See ["About using the installer to upgrade when the root disk is encapsulated"](#) on page 166.

- 22 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node that you recorded in step 9.
- 23 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.
- 24 Make the VCS configuration writable again from any node in the upgraded group:

```
# haconf -makerw
```

- 25 Enter the following command on each node in the upgraded group to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

- 26 Make the configuration read-only:

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

- 27 Bring all of the VCS service groups, such as failover groups, online on the required node using the below command:

```
# hagrps -online groupname -sys nodename
```

- 28 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

- 29 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem
# mount /checkpoint_name
```

- 30 You can perform the following optional configuration steps:

- If you want to use features of Veritas Storage Foundation 6.0 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
See [“Upgrading VxVM disk group versions”](#) on page 220.

Upgrading Storage Foundation and High Availability using the Veritas Web-based installer

This section describes upgrading SFHA with the Veritas Web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems.

To upgrade SFHA

- 1 Perform the required steps to save any data that you wish to preserve. For example, make configuration file backups.
- 2 If you are upgrading a high availability (HA) product, take all service groups offline. List all service groups:

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

For each service group listed, take it offline:

```
# /opt/VRTSvcs/bin/hagrp -offline service_group -all
```

- 3 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 61.
- 4 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.
- 5 The installer discovers if any of the systems that you are upgrading have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the book disk group. To create the backup, check the **Split mirrors on all the systems** box. Check the appropriate box to use the same name for the backup disk group on all systems--you can use the default name or choose a new one. Check the systems where you want to create the backup. When you are ready, click the **Next** button.

- 6 Click **Next** to complete the upgrade.
After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.
- 7 After the upgrade, if the product is not configured, the Web-based installer asks: "Do you want to configure this product?" If the product is already configured, it will not ask any questions.
- 8 If you want to upgrade VCS or SFHA 5.1 on the CP server systems to version Veritas Storage Foundation 6.0, make sure that you upgraded all application clusters to version Veritas Storage Foundation 6.0. Then, upgrade VCS or SFHA on the CP server systems. For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.
- 9 Only perform this step if you have split the mirrored root disk to back it up. After a successful reboot, verify the upgrade and re-join the backup disk group into the upgraded boot disk group. If the upgrade fails, revert the upgrade boot disk group to the backup disk group.

See ["Re-joining the backup boot disk group into the current disk group"](#) on page 212.

See ["Reverting to the backup boot disk group after an unsuccessful upgrade"](#) on page 212.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native operating system accounts.

Upgrading Veritas Volume Replicator

If a previous version of Veritas Volume Replicator (VVR) is configured, the product installer upgrades VVR automatically when you upgrade the Storage Foundation products.

When upgrading from 4.1 or later, you have the option to upgrade without disrupting replication.

See ["Upgrading VVR without disrupting replication"](#) on page 183.

Upgrading VVR without disrupting replication

This section describes the upgrade procedure from an earlier version of VVR to the current version of VVR when replication is in progress, assuming that you do not need to upgrade all the hosts in the RDS simultaneously.

You may also need to set up replication between versions.

See ["Planning an upgrade from the previous VVR version"](#) on page 170.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade can be performed either on the Primary or on the Secondary. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS. This section includes separate sets of steps, for the Primary upgrade and for the Secondary upgrade.

Note: If you have a cluster setup, you must upgrade all the nodes in the cluster at the same time.

Upgrading VVR on the Secondary

Follow these instructions to upgrade the Secondary hosts.

To upgrade the Secondary

- 1 Stop replication to the Secondary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvgname
```

- 2 Upgrade from VVR 4.1 or later to VVR 6.0 on the Secondary.
- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname sec_hostname
```

Upgrading VVR on the Primary

After you upgrade the Secondary, use the Veritas product installer to upgrade the Primary.

To upgrade the Primary

- 1 Stop replication to the Primary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvgname
```

- 2 Upgrade from VVR 4.1 or later to VVR 6.0 on the Secondary.
- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname  
sec_hostname
```


See [“Planning an upgrade from the previous VVR version”](#) on page 170.

Performing a rolling upgrade

This chapter includes the following topics:

- [About rolling upgrades](#)
- [Performing a rolling upgrade using the script-based installer](#)
- [Performing a rolling upgrade of SFHA using the Web-based installer](#)

About rolling upgrades

The rolling upgrade minimizes downtime for highly available clusters to the amount of time that it takes to perform a service group failover. The rolling upgrade has two main phases where the installer upgrades kernel RPMs in phase 1 and VCS and VCS agent RPMs in phase 2.


Note: You need to perform a rolling upgrade on a completely configured cluster.

The following is an overview of the flow for a rolling upgrade:

1. The installer performs prechecks on the cluster.
2. The installer moves service groups to free nodes for the first phase of the upgrade as is needed. Application downtime occurs during the first phase as the installer moves service groups to free nodes for the upgrade. The only downtime that is incurred is the normal time required for the service group to fail over.
3. The installer performs the second phase of the upgrade on all of the nodes in the cluster. The second phase of the upgrade includes downtime of the Veritas Cluster Server (VCS) engine HAD, but does not include application downtime.

Figure 14-1 illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 14-1 Example of the installer performing a rolling upgrade



The following limitations apply to rolling upgrades:

- Rolling upgrades are not compatible with phased upgrades.
- Do not mix rolling upgrades and phased upgrades.
- You can perform a rolling upgrade from 5.1 and later versions.

Performing a rolling upgrade using the script-based installer

Before you start the rolling upgrade, make sure that Veritas Cluster Server (VCS) is running.

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.0, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in [Patch Lookup](#) on the [SORT](#) website.

To perform a rolling upgrade

- 1 Complete the preparatory steps on the first sub-cluster.
- 2 Log in as superuser and mount the SFHA Veritas Storage Foundation 6.0 installation media.
- 3 From root, start the installer.

```
# ./installer
```
- 4 From the menu select **G) Upgrade a product**; then select **2) Rolling Upgrade**;
- 5 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.
- 6 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.
- 7 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.
- 8 Review the EULA, and type **y** if you agree to its terms.

- 9 If the upgrade is from the 5.1 SP1 release or later and the boot disk is encapsulated and mirrored, you can create a backup boot disk.

If you choose to create a backup boot disk, type **y**. Provide a backup name for the boot disk group or accept the default name. The installer then creates a backup copy of the boot disk group.

See [“Re-joining the backup boot disk group into the current disk group”](#) on page 212.

See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 212.

The installer lists the RPMs to upgrade on the selected node or nodes.
- 10 The installer prompts you to stop the applicable processes. Type **y** to continue.

The installer fails over failover service groups to the node or nodes that are not upgraded at this time. The downtime is the time that it normally takes for the service group's failover. The installer stops parallel service groups on the nodes that are to be upgraded.
- 11 The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs.

It performs the configuration for the upgrade and re-starts processes.

In case of failure in the startup of some of the processes, you may need to reboot the nodes and manually check the cluster's status.
- 12 Complete the preparatory steps on the nodes that you have not yet upgraded.
- 13 The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade.

If the installer reboots nodes, restart the installer.

The installer repeats step 6 through step 11.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.
- 14 When phase 1 of the rolling upgrade completes, begin phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.

If the boot disk is encapsulated, you need to reboot nodes when phase 1 completes.
- 15 The installer determines the remaining RPMs to upgrade. Press **Enter** to continue.

- 16 The installer stops Veritas Cluster Server (VCS) processes. Type **y** to continue.
The installer performs prechecks, uninstalls old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.
- 17 Type **y** or **n** to help Symantec improve the automated installation.
- 18 If you have network connection to the Internet, the installer checks for updates.
If updates are discovered, you can apply them now.
- 19 To upgrade VCS or Storage Foundation High Availability (SFHA) on the Coordination Point (CP) server systems to version 6.0, upgrade all the application clusters to 6.0. You then upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, refer to the appropriate installation guide.

Performing a rolling upgrade of SFHA using the Web-based installer

This section describes using the Veritas Web-based installer to perform a rolling upgrade. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

See [“About rolling upgrades”](#) on page 186.

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.0, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in [Patch Lookup](#) on the [SORT](#) website.

To start the rolling upgrade—phase 1

- 1 Complete the preparatory steps on the first sub-cluster.
- 2 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 3 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 61.
- 4 In the Task pull-down menu, select **G) Upgrade a product;** then select **2)Rolling Upgrade;**

Click the **Next** button to proceed.

- 5 Review the systems that the installer has chosen to start the rolling upgrade. These systems are chosen to minimize downtime during the upgrade.

Click **Yes** to proceed.

The installer validates systems. If it throws an error, address the error and return to the installer.

- 6 Review the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.

- 7 The installer stops all processes. Click **Next** to proceed.

- 8 The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted. The installer starts all the relevant processes and brings all the service groups online.

If the installer reboots nodes, restart the installer.

- 9 Complete the preparatory steps on the nodes that you have not yet upgraded.

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

- 11 When prompted, start phase 2. Click **Yes** to continue with the rolling upgrade.

You may need to restart the Web-based installer to perform phase 2.

See [“Starting the Veritas Web-based installer”](#) on page 61.

To upgrade the non-kernel components—phase 2

- 1 In the Task pull-down menu, make sure that **Rolling Upgrade** is selected.

Click the **Next** button to proceed.

- 2 The installer detects the information of cluster and the state of rolling upgrade.

The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.

- 3 Review the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.

- 4 The installer validates systems. If it throws an error, address the error and return to the installer. Click **Next** to proceed.

- 5 The installer stops all processes. Click **Next** to proceed.

- 6 The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted. The installer starts all the relevant processes and brings all the service groups online.
- 7 If you have network connection to the Internet, the installer checks for updates. If updates are discovered, you can apply them now.

The upgrade is complete.

Performing a phased upgrade

This chapter includes the following topics:

- [About phased upgrade](#)
- [Performing a phased upgrade](#)

About phased upgrade

Perform a phased upgrade to minimize the downtime for the cluster. Depending on the situation, you can calculate the approximate downtime as follows:

You can fail over all your service groups to the nodes that are up. Downtime equals the time that is taken to offline and online the service groups.

You have a service group that you cannot fail over to a node that runs during upgrade. Downtime for that service group equals the time that is taken to perform an upgrade and restart the node.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade.

Planning for a phased upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group.

Some rough guidelines follow:

- Split the cluster into two sub-clusters of equal or near equal size.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.
- Before you start the upgrade, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.


Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules.
- When you start the installer, only select SFHA.
- While you perform the upgrades, do not add or remove service groups to any of the nodes.
- After you upgrade the first half of your cluster (the first subcluster), you need to set up password-less SSH. Create the connection between an upgraded node in the first subcluster and a node from the other subcluster. The node from the other subcluster is where you plan to run the installer and also plan to upgrade.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.

Phased upgrade example

In this example, you have a secure cluster that you have configured to run on four nodes: `node01`, `node02`, `node03`, and `node04`. You also have four service groups: `sg1`, `sg2`, `sg3`, and `sg4`. For the purposes of this example, the cluster is split into two subclusters. The nodes `node01` and `node02` are in the first subcluster, which you first upgrade. The nodes `node03` and `node04` are in the second subcluster, which you upgrade last.

Figure 15-1 Example of phased upgrade set up

Each service group is running on the nodes as follows:

- sg1 and sg2 are parallel service groups and run on all the nodes.
- sg3 and sg4 are failover service groups. sg3 runs on node01 and sg4 runs on node02.

In your system list, you have each service group that fails over to other nodes as follows:

- sg1 and sg2 are running on all the nodes.
- sg3 and sg4 can fail over to any of the nodes in the cluster.

Phased upgrade example overview

This example's upgrade path follows:

- Move all the failover service groups from the first subcluster to the second subcluster.
- Take all the parallel service groups offline on the first subcluster.
- Upgrade the operating system on the first subcluster's nodes, if required.
- On the first subcluster, start the upgrade using the installation program.
- Get the second subcluster ready.
- Activate the first subcluster.
- Upgrade the operating system on the second subcluster's nodes, if required.
- On the second subcluster, start the upgrade using the installation program.
- Activate the second subcluster.

See [“Performing a phased upgrade”](#) on page 196.

Performing a phased upgrade

This section explains how to perform a phased upgrade of SFHA on four nodes with four service groups. Note that in this scenario, VCS and the service groups cannot stay online on the second subcluster during the upgrade of the second subcluster. Do not add, remove, or change resources or service groups on any nodes during the upgrade. These changes are likely to get lost after the upgrade.

An example of a phased upgrade follows. It illustrates the steps to perform a phased upgrade. The example makes use of a secure SFHA cluster.

You can perform a phased upgrade from SFHA 5.1 or other supported previous versions to SFHA 6.0.

See [“About phased upgrade”](#) on page 193.

See [“Phased upgrade example”](#) on page 194.

Moving the service groups to the second subcluster

Perform the following steps to establish the service group's status and to switch the service groups.

To move service groups to the second subcluster

- 1 On the first subcluster, determine where the service groups are online.

```
# hagr -state
```

The output resembles:

```
#Group   Attribute System Value
sg1      State    node01 | ONLINE |
sg1      State    node02 | ONLINE |
sg1      State    node03 | ONLINE |
sg1      State    node04 | ONLINE |
sg2      State    node01 | ONLINE |
sg2      State    node02 | ONLINE |
sg2      State    node03 | ONLINE |
sg2      State    node04 | ONLINE |
sg3      State    node01 | ONLINE |
sg3      State    node02 | OFFLINE |
sg3      State    node03 | OFFLINE |
sg3      State    node04 | OFFLINE |
sg4      State    node01 | OFFLINE |
sg4      State    node02 | ONLINE |
sg4      State    node03 | OFFLINE |
sg4      State    node04 | OFFLINE |
```

- 2 Offline the parallel service groups (sg1 and sg2) from the first subcluster. Switch the failover service groups (sg3 and sg4) from the first subcluster (node01 and node02) to the nodes on the second subcluster (node03 and node04).

```
# hagr -offline sg1 -sys node01
# hagr -offline sg2 -sys node01
# hagr -offline sg1 -sys node02
# hagr -offline sg2 -sys node02
# hagr -switch sg3 -to node03
# hagr -switch sg4 -to node04
```

- 3 On the nodes in the first subcluster, unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -h
Filesystem              Size  Used Avail Use% Mounted on
/dev/sda1                26G   3.3G   22G   14% /
udev                   1007M   352K 1006M    1% /dev
tmpfs                   4.0K     0   4.0K    0% /dev/vx
/dev/vx/dsk/dg2/dg2vol1
                        3.0G   18M   2.8G    1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2
                        1.0G   18M   944M    2% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3
                        10G   20M   9.4G    1% /mnt/dg2/dg2vol3
# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

- 4 On the nodes in the first subcluster, stop all VxVM volumes (for each disk group) that VCS does not manage.
- 5 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

- 6 Freeze the nodes in the first subcluster.

```
# hasys -freeze -persistent node01
# hasys -freeze -persistent node02
```

- 7 Dump the configuration and make it read-only.

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

- 8 Verify that the service groups are offline on the first subcluster that you want to upgrade.

```
# hagrps -state
```

Output resembles:

```
#Group Attribute System Value
sg1 State node01 |OFFLINE|
sg1 State node02 |OFFLINE|
sg1 State node03 |ONLINE|
sg1 State node04 |ONLINE|
sg2 State node01 |OFFLINE|
sg2 State node02 |OFFLINE|
sg2 State node03 |ONLINE|
sg2 State node04 |ONLINE|
sg3 State node01 |OFFLINE|
sg3 State node02 |OFFLINE|
sg3 State node03 |ONLINE|
sg3 State node04 |OFFLINE|
sg4 State node01 |OFFLINE|
sg4 State node02 |OFFLINE|
sg4 State node03 |OFFLINE|
sg4 State node04 |ONLINE|
```

- 9 Perform this step on the nodes (node01 and node02) in the first subcluster if the cluster uses I/O Fencing. Use an editor of your choice and change the following:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `scsi3` to `disabled`. Ensure that the line in the `vxfenmode` file resembles:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. Ensure that the line in the `main.cf` file resembles:

```
UseFence = NONE
```

- 10 Back up the llttab, llthosts, gabtab, types.cf, main.cf and AT configuration files on the first subcluster.

```
# cp /etc/llttab /etc/llttab.bkp
# cp /etc/llthosts /etc/llthosts.bkp
# cp /etc/gabtab /etc/gabtab.bkp
# cp /etc/VRTSvcs/conf/config/main.cf \
    /etc/VRTSvcs/conf/config/main.cf.bkp
# cp /etc/VRTSvcs/conf/config/types.cf \
    /etc/VRTSvcs/conf/config/types.cf.bkp
# /opt/VRTSat/bin/vssat showbackuplist
B|/var/VRTSat/.VRTSat/profile/VRTSatlocal.conf
B|/var/VRTSat/.VRTSat/profile/certstore
B|/var/VRTSat/ABAuthSource
B|/etc/vx/vss/VRTSat.conf
Quiescing ...
Snapshot Directory :/var/VRTSatSnapShot
```

Upgrading the operating system on the first subcluster

You can perform the operating system upgrade on the first subcluster, if required. Refer to the operating system's documentation for more information.

Upgrading the first subcluster

You now navigate to the installer program and start it.

To start the installer for the phased upgrade

- 1 Confirm that you are logged on as the superuser and you mounted the product disc.
- 2 Navigate to the folder that contains installsfha.

```
# cd storage_foundation_high_availability
```

- 3 Make sure that you can ssh or rsh from the node where you launched the installer to the nodes in the second subcluster without requests for a password.

- 4 Start the `installsfha` program, specify the nodes in the first subcluster (`node1` and `node2`).

```
# ./installsfha -upgrade node1 node2
```

The program starts with a copyright message and specifies the directory where it creates the logs. It performs a system verification and outputs upgrade information.

- 5 Enter **y** to agree to the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement
as specified in the storage_foundation_high_availability/
EULA/lang/EULA_SFHA_Ux_6.0.pdf
file present on media? [y,n,q,?] y
```

- 6 The installer displays the list of RPMs that get removed, installed, and upgraded on the selected systems.
- 7 When you are prompted, reply **y** to stop appropriate processes.

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

The installer stops processes, uninstalls RPMs, and installs RPMs.

The upgrade is finished on the first subcluster. Do not reboot the nodes in the first subcluster until you complete the [Preparing the second subcluster](#) procedure.

Preparing the second subcluster

Perform the following steps on the second subcluster before rebooting nodes in the first subcluster.

To prepare to upgrade the second subcluster

1 Get the summary of the status of your resources.

```
# hastatus -summ
-- SYSTEM STATE
-- System                State                Frozen

A  node01                EXITED                1
A  node02                EXITED                1
A  node03                RUNNING              0
A  node04                RUNNING              0

-- GROUP STATE
-- Group                System  Probed    AutoDisabled  State

B  SG1                  node01  Y         N              OFFLINE
B  SG1                  node02  Y         N              OFFLINE
B  SG1                  node03  Y         N              ONLINE
B  SG1                  node04  Y         N              ONLINE
B  SG2                  node01  Y         N              OFFLINE
B  SG2                  node02  Y         N              OFFLINE
B  SG2                  node03  Y         N              ONLINE
B  SG2                  node04  Y         N              ONLINE
B  SG3                  node01  Y         N              OFFLINE
B  SG3                  node02  Y         N              OFFLINE
B  SG3                  node03  Y         N              ONLINE
B  SG3                  node04  Y         N              OFFLINE
B  SG4                  node01  Y         N              OFFLINE
B  SG4                  node02  Y         N              OFFLINE
B  SG4                  node03  Y         N              OFFLINE
B  SG4                  node04  Y         N              ONLINE
```

2 Unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -h
Filesystem                Size  Used Avail Use% Mounted on
/dev/sda1                  26G   3.3G   22G   14% /
udev                      1007M 352K 1006M   1% /dev
tmpfs                      4.0K    0   4.0K   0% /dev/vx
/dev/vx/dsk/dg2/dg2vol1   3.0G   18M   2.8G   1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2   1.0G   18M   944M   2% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3   10G   20M   9.4G   1% /mnt/dg2/dg2vol3

# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

3 Stop all VxVM volumes (for each disk group) that VCS does not manage.**4** Make the configuration writable on the second subcluster.

```
# haconf -makerw
```

5 Unfreeze the service groups.

```
# hagrps -unfreeze sg1 -persistent
# hagrps -unfreeze sg2 -persistent
# hagrps -unfreeze sg3 -persistent
# hagrps -unfreeze sg4 -persistent
```

6 Dump the configuration and make it read-only.

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

7 Take the service groups offline on node03 and node04.

```
# hagrps -offline sg1 -sys node03
# hagrps -offline sg1 -sys node04
# hagrps -offline sg2 -sys node03
# hagrps -offline sg2 -sys node04
# hagrps -offline sg3 -sys node03
# hagrps -offline sg4 -sys node04
```

8 Verify the state of the service groups.

```
# hagrps -state
#Group      Attribute  System  Value
SG1         State     node01  |OFFLINE|
SG1         State     node02  |OFFLINE|
SG1         State     node03  |OFFLINE|
SG1         State     node04  |OFFLINE|
SG2         State     node01  |OFFLINE|
SG2         State     node02  |OFFLINE|
SG2         State     node03  |OFFLINE|
SG2         State     node04  |OFFLINE|
SG3         State     node01  |OFFLINE|
SG3         State     node02  |OFFLINE|
SG3         State     node03  |OFFLINE|
SG3         State     node04  |OFFLINE|
```

9 Perform this step on node03 and node04 if the cluster uses I/O Fencing. Use an editor of your choice and change the following:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `scsi3` to `disabled`. Ensure that the line in the `vxfenmode` file resembles:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. Ensure that the line in the `main.cf` file resembles:

```
UseFence = NONE
```

- 10 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

```
# hastop -local
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

- 11 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 are not added.

```
# /etc/init.d/vxfen status
VXFEN module is not loaded

# /etc/init.d/gab status
GAB module is not loaded

# /etc/init.d/llt status
LLT module is not loaded
```

Activating the first subcluster

Get the first subcluster ready for the service groups.

To activate the first subcluster

- 1 Perform this step on node01 and node02 if the cluster uses I/O Fencing. Use an editor of your choice and revert the following to an enabled state before you reboot the first subcluster's nodes:
 - In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `NONE` to `SCSI3`. Ensure that the line in the `main.cf` file resembles:

```
UseFence = SCSI3
```
 - In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `disabled` to `scsi3`. Ensure that the line in the `vxfenmode` file resembles:

```
vxfen_mode=scsi3
```
- 2 Reboot the node01 and node02 in the first subcluster.

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

- 3 Seed node01 and node02 in the first subcluster.

```
# gabconfig -xc
```

- 4 Start VCS on node01 and node02. On each node run:

```
# hstart
```

- 5 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

- 6 Unfreeze the nodes in the first subcluster.

```
# hasys -unfreeze -persistent node01
```

```
# hasys -unfreeze -persistent node02
```

- 7 Dump the configuration and make it read-only.

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

- 8 Bring the service groups online on node01 and node02.

```
# hagrps -online sg1 -sys node01
```

```
# hagrps -online sg1 -sys node02
```

```
# hagrps -online sg2 -sys node01
```

```
# hagrps -online sg2 -sys node02
```

```
# hagrps -online sg3 -sys node01
```

```
# hagrps -online sg4 -sys node02
```

Upgrading the operating system on the second subcluster

You can perform the operating system upgrade on the second subcluster, if required. Refer to the operating system's documentation for more information.

Before you perform the operating system upgrade, make sure to disable VCS, VXFEN, GAB, and LLT.

To disable VCS, VXFEN, GAB, and LLT

- 1 On the second subcluster, perform the following commands:

```
# chkconfig vcs off
# chkconfig vxfen off
# chkconfig gab off
# chkconfig llt off
```

- 2 For a cluster that uses secure mode, create a password-less SSH connection. The connection is from the node where you plan to run the installer to one of the nodes that you have already upgraded.

To enable VCS, VXFEN, GAB and LLT

- ◆ On second subcluster, perform following commands:

```
# chkconfig llt on
# chkconfig gab on
# chkconfig vxfen on
# chkconfig vcs on
```

Upgrading the second subcluster

Perform the following procedure to upgrade the second subcluster (node03 and node04).

To start the installer to upgrade the second subcluster

- 1 Confirm that you are logged on as the superuser and you mounted the product disc.
- 2 Navigate to the folder that contains `installsfha`.

```
# cd storage_foundation_high_availability
```

- 3 Confirm that SFHA is stopped on node03 and node04. Start the `installsfha` program, specify the nodes in the second subcluster (node3 and node4).

```
# ./installsfha -upgrade node3 node4
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 4 The installer displays the list of RPMs that get removed, installed, and upgraded on the selected systems.

- 5 When you are prompted, reply **y** to stop appropriate processes.

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

The installer stops processes, uninstalls RPMs, and installs RPMs.

- 6 Enter **y** to agree to the End User License Agreement (EULA).

```
Do you agree with the terms of the End User License Agreement
as specified in the storage_foundation_high_availability/
EULA/lang/EULA_SFHA_Ux_6.0.pdf
file present on media? [y,n,q,?] y
```

- 7 Monitor the installer program answering questions as appropriate until the upgrade completes.

Finishing the phased upgrade

Complete the following procedure to complete the upgrade.

To finish the upgrade

- 1 Verify that the cluster UUID is the same on the nodes in the second subcluster and the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl
-clus -display node1 [node2 ...]
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first subcluster to the nodes in the second subcluster. For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus
-copy -from_sys node01 -to_sys node03 node04
```

- 2 Perform this step on node03 and node04 if the cluster uses I/O Fencing. Use an editor of your choice and revert the following to an enabled state before you reboot the second subcluster's nodes:

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from disabled to `scsi3`. Ensure that the line in the `vxfenmode` file resembles:

```
vxfen_mode=scsi3
```

- 3 Reboot the node03 and node04 in the second subcluster.

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

The nodes in the second subcluster join the nodes in the first subcluster.

4 Check to see if SFHA and High Availability and its components are up.

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen  nxxxxnn membership 0123
Port b gen  nxxxxnn membership 0123
Port h gen  nxxxxnn membership 0123
```

5 Run an `hastatus -sum` command to determine the status of the nodes, service groups, and cluster.

```
# hastatus -sum

-- SYSTEM STATE
-- System          State          Frozen

A  node01          RUNNING        0
A  node02          RUNNING        0
A  node03          RUNNING        0
A  node04          RUNNING        0

-- GROUP STATE
-- Group   System   Probed   AutoDisabled   State
B  sg1     node01   Y          N              ONLINE
B  sg1     node02   Y          N              ONLINE
B  sg1     node03   Y          N              ONLINE
B  sg1     node04   Y          N              ONLINE
B  sg2     node01   Y          N              ONLINE
B  sg2     node02   Y          N              ONLINE
B  sg2     node03   Y          N              ONLINE
B  sg2     node04   Y          N              ONLINE
B  sg3     node01   Y          N              ONLINE
B  sg3     node02   Y          N              OFFLINE
B  sg3     node03   Y          N              OFFLINE
B  sg3     node04   Y          N              OFFLINE
B  sg4     node01   Y          N              OFFLINE
B  sg4     node02   Y          N              ONLINE
B  sg4     node03   Y          N              OFFLINE
B  sg4     node04   Y          N              OFFLINE
```

6 After the upgrade is complete, start the VxVM volumes (for each disk group) and mount the VxFS file systems.

In this example, you have performed a phased upgrade of SFHA. The service groups were down when you took them offline on node03 and node04, to the time SFHA brought them online on node01 or node02.

Note: If you want to upgrade Coordination Point (CP) server systems that use Veritas Cluster Server (VCS) or Storage Foundation High Availability (SFHA) to 6.0, make sure that you upgraded all application clusters to version 6.0. Then, upgrade VCS or SFHA on the CP server systems. For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.

Performing post-upgrade tasks

This chapter includes the following topics:

- [Optional configuration steps](#)
- [Re-joining the backup boot disk group into the current disk group](#)
- [Reverting to the backup boot disk group after an unsuccessful upgrade](#)
- [Post upgrade tasks for migrating the SFDB repository database](#)
- [Recovering VVR if automatic upgrade fails](#)
- [Post-upgrade tasks when VCS agents for VVR are configured](#)
- [Upgrading disk layout versions](#)
- [Upgrading VxVM disk group versions](#)
- [Updating variables](#)
- [Setting the default disk group](#)
- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [Verifying the Storage Foundation and High Availability upgrade](#)

Optional configuration steps

After the upgrade is complete, additional tasks may need to be performed.

You can perform the following optional configuration steps:

- If Veritas Volume Replicator (VVR) is configured, do the following steps in the order shown:
 - Reattach the RLINKs.
 - Associate the SRL.
- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Veritas Storage Foundation Administrator's Guide*.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
 See ["Upgrading VxVM disk group versions"](#) on page 220.

Re-joining the backup boot disk group into the current disk group

Perform this procedure to rejoin the backup boot disk if you split the mirrored boot disk during upgrade. After a successful upgrade and reboot, you no longer need to keep the boot disk group backup.

To re-join the backup boot disk group

- ◆ Re-join the *backup_bootdg* disk group to the boot disk group.

```
# /etc/vx/bin/vxrootadm -Y join backup_bootdg
```

where the `-Y` option indicates a silent operation, and *backup_bootdg* is the name of the backup boot disk group that you created during the upgrade.

Reverting to the backup boot disk group after an unsuccessful upgrade

Perform this procedure if your upgrade was unsuccessful and you split the mirrored boot disk to back it up during upgrade. You can revert to the backup that you created when you upgraded.

To revert the backup boot disk group after an unsuccessful upgrade

- 1 To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

```
# vxprint
```

- 2 Use the `vxdg` command to find the boot disk group where you are currently booted.

```
# vxdg bootdg
```

- 3 Boot the operating system from the backup boot disk group.
- 4 Join the original boot disk group to the backup disk group.

```
# /etc/vx/bin/vxrootadm -Y join original_bootdg
```

where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

Post upgrade tasks for migrating the SFDB repository database

Database Storage Checkpoints that have been created by using the SFDB tools before upgrade are visible using the `vxsfadm` CLI, and you can mount these Database Storage Checkpoints and roll back to them, if required. However, creating clones by using migrated Database Storage Checkpoints is not supported.

If you want to continue using previously created FlashSnap snapplans to take snapshots, you must validate them by using the `-o validate` option of the `vxsfadm` command.

- Rename startup script after upgrading from 5.0x and before migrating the SFDB repository
 See [“After upgrading from 5.0.x and before migrating SFDB”](#) on page 216.
- Migrate from a 5.0x SFDB repository database to 6.0
 See [“Migrating from a 5.0 repository database to 6.0”](#) on page 213.

Migrating from a 5.0 repository database to 6.0

Perform the following on one node only.

To migrate from a 5.0 repository database to 6.0

- 1 Rename the startup script NO_S*vxdbms3 to S*vxdbms3.

See “[After upgrading from 5.0.x and before migrating SFDB](#)” on page 216.

- 2 As root, dump out the old Sybase ASA repository. If you are using SFHA or SF Oracle RAC, you only need to do this on one node.

```
# /opt/VRTSdbed/migrate/sfua_rept_migrate
```

- 3 On the same node that you ran `sfua_rept_migrate` run the following command as Oracle user. For each Oracle instance, migrate the old repository data to the SQLite repository.

```
$ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME -G \
Oracle_service_group
```

- 4 By default, the repository is created on the file system which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:

- Repository path has to be a directory writable by Oracle user.
- If you are using SFHA, the repository must be accessible by all nodes. You can put it in a resource group under VCS control so it can be failed over together with the Oracle database.
- The update commands will not be able to verify accessibility of the repository path and will fail if you have not set up the path correctly.

Create an alternate repository path.

```
$ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME \
-G Oracle_service_group -R Alternate_path
```

- 5 If you are using Database Flashsnap for off-host processing, and if you have a repository on the secondary host that you need to migrate: perform the previous steps on the secondary host.

- 6 On the primary host, edit your snapplans to remove the "SNAPSHOT_DG=SNAP_*" parameter and add "SNAPSHOT_DG_PREFIX=SNAP_*". The parameter can be any PREFIX value and not necessarily "SNAP_*".

For example:

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=system1_data
SNAPSHOT_DG=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=judge_data
SNAPSHOT_DG_PREFIX=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

7 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/vxsfadm -s flashsnap \  
-a oracle -c SNAPPLAN -o validate
```

This completes the migration of the repository for Database Storage Checkpoints and Database Tiered Storage parameters.

To begin using the Storage Foundation for Databases (SFDB) tools:

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

After upgrading from 5.0.x and before migrating SFDB

When upgrading from SFHA version 5.0 to SFHA 6.0 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by sfua_rept_migrate. Thus when sfua_rept_migrate is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found  
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.  
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

To prevent S*vxdbms3 startup script error

- ◆ Rename the startup script NO_S*vxdbms3 to S*vxdbms3.

Recovering VVR if automatic upgrade fails

If the upgrade fails during the configuration phase, after displaying the VVR upgrade directory, the configuration needs to be restored before the next attempt. Run the scripts in the upgrade directory in the following order to restore the configuration:

```
# restoresrl  
# adddcn  
# srlprot  
# attrlink  
# start.rvg
```

After the configuration is restored, the current step can be retried.

Post-upgrade tasks when VCS agents for VVR are configured

The following lists post-upgrade tasks with VCS agents for VVR:

- [Unfreezing the service groups](#)
- [Restoring the original configuration when VCS agents are configured](#)

Unfreezing the service groups

This section describes how to unfreeze services groups and bring them online.

To unfreeze the service groups

- 1 On any node in the cluster, make the VCS configuration writable:

```
# haconf -makerw
```

- 2 Edit the `/etc/VRTSvcs/conf/config/main.cf` file to remove the deprecated attributes, SRL and RLinks, in the RVG and RVGShared resources.

- 3 Verify the syntax of the main.cf file, using the following command:

```
# hacf -verify
```

- 4 Unfreeze all service groups that you froze previously. Enter the following command on any node in the cluster:

```
# hagrps -unfreeze service_group -persistent
```

- 5 Save the configuration on any node in the cluster.

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

- 6 If you are upgrading in a shared disk group environment, bring online the RVGShared groups with the following commands:

```
# hagrps -online RVGShared -sys masterhost
```

- 7 Bring the respective IP resources online on each node.

See [“Preparing for the upgrade when VCS agents are configured”](#) on page 174.

Type the following command on any node in the cluster.

```
# hares -online ip_name -sys system
```

This IP is the virtual IP that is used for replication within the cluster.

- 8 In shared disk group environment, online the virtual IP resource on the master node.

Restoring the original configuration when VCS agents are configured

This section describes how to restore a configuration with VCS configured agents.

Note: Restore the original configuration only after you have upgraded VVR on all nodes for the Primary and Secondary cluster.

To restore the original configuration

- 1 Import all the disk groups in your VVR configuration.

```
# vxdg -t import diskgroup
```

Each disk group should be imported onto the same node on which it was online when the upgrade was performed. The reboot after the upgrade could result in another node being online; for example, because of the order of the nodes in the AutoStartList. In this case, switch the VCS group containing the disk groups to the node on which the disk group was online while preparing for the upgrade.

```
# hagr -switch grpname -to system
```

- 2 Recover all the disk groups by typing the following command on the node on which the disk group was imported in step 1.

```
# vxrecover -bs
```

- 3 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

- 4 On all nodes that are Secondary hosts of VVR, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup shrinkto volume_name volume_length
```

where *volume_length* is the length of the volume on the Primary.

Note: Do not continue until you complete this step on all the nodes in the Primary and Secondary clusters on which VVR is upgraded.

- 5 Restore the configuration according to the method you used for upgrade:

If you upgraded with the VVR upgrade scripts

Complete the upgrade by running the `vvr_upgrade_finish` script on all the nodes on which VVR was upgraded. We recommend that you first run the `vvr_upgrade_finish` script on each node that is a Secondary host of VVR.

Perform the following tasks in the order indicated:

- To run the `vvr_upgrade_finish` script, type the following command:

```
# /disc_path/scripts/vvr_upgrade_finish
```

where `disc_path` is the location where the Veritas software disc is mounted.

- Attach the RLINKs on the nodes on which the messages were displayed:

```
# vxrlink -g diskgroup -f att rlink_name
```

If you upgraded with the product installer

Use the Veritas product installer and select Start an Installed Product. Or use the installation script with the `-start` option.

- 6 Bring online the RVGLogowner group on the master:

```
# hagr -online RVGLogownerGrp -sys masterhost
```

- 7 If you plan on using IPv6, you must bring up IPv6 addresses for virtual replication IP on primary/secondary nodes and switch from using IPv4 to IPv6 host names or addresses, enter:

```
# vradm changeip newpri=v6 newsec=v6
```

where `v6` is the IPv6 address.

- 8 Restart the applications that were stopped.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 7, 8, and 9. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Disk layout Version 6 has been deprecated and you cannot cluster mount an existing file system that has disk layout Version 6. To upgrade a cluster file system with disk layout Version 6, you must local mount the file system and then upgrade the file system using the `vxupgrade` utility to a later version.

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

Support for disk layout Version 4 has been removed. You must upgrade any existing file systems with disk layout Version 4 to disk layout Version 7 or later using the `vxfsconvert` command.

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

Note: Symantec recommends that you upgrade existing file systems to the highest supported disk layout version prior to upgrading to this release.

You can check which disk layout version your file system has by using the following command:

```
# fstyp -v /dev/vx/dsk/dg1/vol1 | grep -i version
```

For more information about disk layout versions, see the *Veritas Storage Foundation Administrator's Guide*.

Upgrading VxVM disk group versions

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions. VxVM can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks or use the features, upgrade the existing disk groups.

For 6.0, the Veritas Volume Manager disk group version is different than in previous VxVM releases. Symantec recommends that you upgrade the disk group version if you upgraded from a previous VxVM release.

After upgrading to SFHA 6.0, you must upgrade any existing disk groups that are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

Use the following command to find the version of a disk group:

```
# vxvg list diskgroup
```

To upgrade a disk group to the current disk group version, use the following command:

```
# vxvg upgrade diskgroup
```

For more information about disk group versions, see the *Veritas Storage Foundation Administrator's Guide*.

Updating variables

In `/etc/profile`, update the `PATH` and `MANPATH` variables as needed.

`MANPATH` could include `/opt/VRTS/man` and `PATH /opt/VRTS/bin`.

Setting the default disk group

In releases prior to Volume Manager 4.0, the default disk group was `rootdg` (the root disk group). For Volume Manager to function, the `rootdg` disk group had to exist and it had to contain at least one disk.

This requirement no longer exists; however, you may find it convenient to create a system-wide default disk group. The main benefit of creating a default disk group is that VxVM commands default to the default disk group. You do not need to use the `-g` option.

You can set the name of the default disk group after installation by running the following command on a system:

```
# vxctl defaultdg diskgroup
```

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

About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as Oracle Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

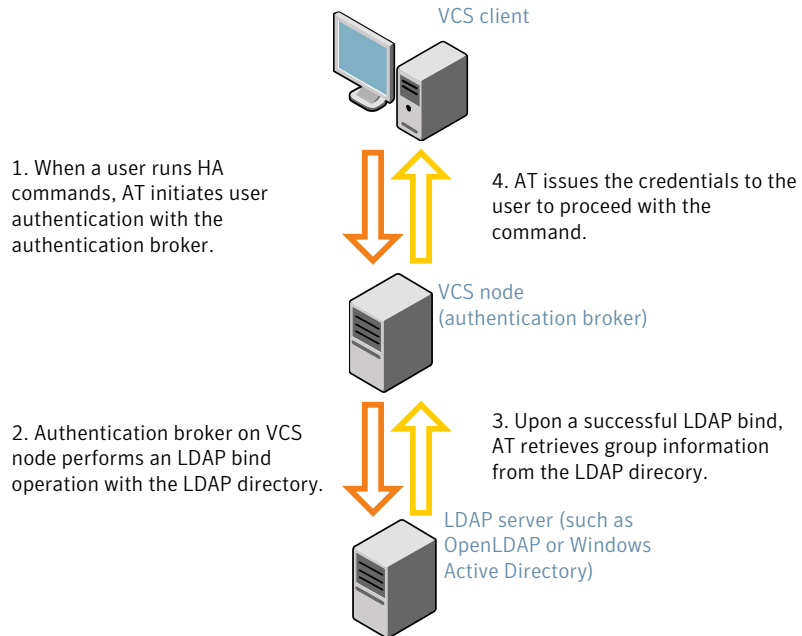
See [“Enabling LDAP authentication for clusters that run in secure mode”](#) on page 223.

If you have not already added VCS users during installation, you can add the users later.

See the *Veritas Cluster Server Administrator's Guide* for instructions to add VCS users.

[Figure 16-1](#) depicts the SFHA cluster communication with the LDAP servers when clusters run in secure mode.

Figure 16-1 Client communication with LDAP servers



The LDAP schema and syntax for LDAP commands (such as `ldapadd`, `ldapmodify`, and `ldapsearch`) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
 - UserObjectClass (the default is `posixAccount`)
 - UserObject Attribute (the default is `uid`)
 - User Group Attribute (the default is `gidNumber`)
 - Group Object Class (the default is `posixGroup`)
 - GroupObject Attribute (the default is `cn`)
 - Group GID Attribute (the default is `gidNumber`)
 - Group Membership Attribute (the default is `memberUid`)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, `UserBaseDN=ou=people,dc=comp,dc=com`)

- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 6.1.6.0 or later.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion
vssat version: 6.1.6.0
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

To enable OpenLDAP authentication for clusters that run in secure mode

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat addldapdomain \
--domainname "MYENTERPRISE.symantecdomain.com"\
--server_url "ldap://my_openldap_host.symantecexample.com"\
--user_base_dn "ou=people,dc=symantecdomain,dc=myenterprise,dc=com"\
--user_attribute "cn" --user_object_class "account"\
--user_gid_attribute "gidNumber"\
--group_base_dn "ou=group,dc=symantecdomain,dc=myenterprise,dc=com"\
--group_attribute "cn" --group_object_class "posixGroup"\
--group_gid_attribute "member"\
--admin_user "cn=manager,dc=symantecdomain,dc=myenterprise,dc=com"\
--admin_user_password "password"--auth_type "FLAT"
```

- 2 Verify that you can successfully authenticate an LDAP user on the SFHA nodes.

You must have a valid LDAP user ID and password to run the command. In the following example, authentication is verified for the MYENTERPRISE domain for the LDAP user, `vcsadmin1`.

```
galaxy# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat authenticate
--domain ldap:MYENTERPRISE.symantecdomain.com
--prplname vcsadmin1 --broker galaxy:14149
```

Enter password for vcsadmin1: #####

```
authenticate
-----
-----
```

```
Authenticated User vcsadmin1
-----
```


3 Add the LDAP user to the main.cf file.

```
# haconf makerw
# hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
# haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
# hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

4 Verify that the main.cf file has the following lines:

```
# cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster clus1 (
  SecureClus = 1
  Administrators = {
    "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" }
  AdministratorGroups = {
    "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com " }
  )
...
...
```

5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=myenterprise.symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh) or the Korn shell (ksh), run the following commands:

```
# export VCS_DOMAIN=myenterprise.symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute    Value
galaxy       Attribute    RUNNING
nebula       Attribute    RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFHA node using the VCS Cluster Manager (Java Console).

- 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

To enable Windows Active Directory authentication for clusters that run in secure mode

- 1 Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -d \  
-s domain_controller_name_or_ipaddress \  
-u domain_user -g domain_group
```

For example:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-d -s 192.168.20.32 -u Administrator -g "Domain Admins"  
Search User provided is invalid or Authentication is required to  
proceed further.  
Please provide authentication information for LDAP server.
```

```
Username/Common Name: symantecdomain\administrator  
Password:
```

Attribute file created.

- 2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-c -d windows_domain_name
```

For example:

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \  
-c -d symantecdomain.com  
Attribute list file not provided, using default AttributeList.txt.  
CLI file name not provided, using default CLI.txt.
```

CLI for `addldapdomain` generated.

- 3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -x
```

- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat listldapdomains
```

```
Domain Name :          symantecdomain.com
Server URL :          ldap://192.168.20.32:389
SSL Enabled :         No
User Base DN :        CN=people,DC=symantecdomain,DC=com
User Object Class :   account
User Attribute :      cn
User GID Attribute :  gidNumber
Group Base DN :       CN=group,DC=symantecdomain,DC=com
Group Object Class :  group
Group Attribute :     cn
Group GID Attribute : cn
Auth Type :           FLAT
Admin User :
Admin User Password :
Search Scope :        SUB
```

- 5 Set the VCS_DOMAIN and VCS_DOMAINTYPE environment variables as follows:

- VCS_DOMAIN=symantecdomain.com
- VCS_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh) or the Korn shell (ksh), run the following commands:

```
# export VCS_DOMAIN=symantecdomain.com
# export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
# halogin vcsadmin1 password
# hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System      Attribute  Value
galaxy       Attribute  RUNNING
nebula       Attribute  RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the SFHA node using the VCS Cluster Manager (Java Console).

- 7 To enable LDAP authentication on other nodes in the cluster, perform the procedure on each of the nodes in the cluster.

Verifying the Storage Foundation and High Availability upgrade

Refer to the section about verifying the installation to verify the upgrade.

See [“Verifying that the products were installed”](#) on page 231.

Verifying the Veritas Storage Foundation and High Availability installation

- [Chapter 17. 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](#)
- [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 295.

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 command to check which RPMs have been installed.

```
# rpm -qa | grep VRTS
```

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 RPMs, and the status (success or failure) of each RPM. 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 `vxiod`, `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.

Verifying the LLT, GAB, and VCS configuration files

Make sure that the LLT, GAB, and VCS configuration files contain the information you provided during VCS installation and configuration.

To verify the LLT, GAB, and VCS configuration files

- 1 Navigate to the location of the configuration files:
 - LLT
/etc/llthosts
/etc/llttab
 - GAB
/etc/gabtab
 - VCS

```
/etc/VRTSvcs/conf/config/main.cf
```

- 2 Verify the content of the configuration files.
 See [“About the LLT and GAB configuration files”](#) on page 335.
 See [“About the VCS configuration files”](#) on page 339.

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 234.
- 4 Verify GAB operation.
- 5 Verify the cluster operation.
 See [“Verifying the cluster”](#) on page 236.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

- 1 Log in as superuser on the node galaxy.
- 2 Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node           State      Links
*0 galaxy      OPEN      2
  1 nebula     OPEN      2
```

Each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
Node           State   Links
* 0 galaxy     OPEN   2
  1 nebula     OPEN   2
  2 saturn     OPEN   1
```

- 3 Log in as superuser on the node nebula.
- 4 Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
Node           State   Links
  0 galaxy     OPEN   2
*1 nebula     OPEN   2
```

- 5 To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv active
```

The output on galaxy resembles:

```
Node           State   Link   Status   Address
*0 galaxy     OPEN
                eth1 UP    08:00:20:93:0E:34
                eth2 UP    08:00:20:93:0E:38
  1 nebula     OPEN
                eth1 UP    08:00:20:8F:D1:F2
                eth2 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.

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    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

galaxy    ConfigChecksum

galaxy    ConfigDiskState      CURRENT

galaxy    ConfigFile            /etc/VRTSvcs/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         eth1 UP eth2 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 Linux:galaxy,#1 SMP Fri Jul 8
17:36:59 EDT
2011,2.6.18-274.e15,x86_64
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 18. Adding a node to a cluster](#)
- [Chapter 19. 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 using the product installer](#)
- [Adding a node using the Web-based installer](#)
- [Adding the node to a cluster manually](#)
- [Configuring server-based fencing on the new node](#)
- [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.

The following procedure provides a summary of the tasks required to add a node to an existing SFHA cluster.

Table 18-1 Tasks for adding a node to a cluster

Step	Description
Complete the prerequisites and preparatory tasks before adding a node to the cluster.	See “Before adding a node to a cluster” on page 242.

Table 18-1 Tasks for adding a node to a cluster (*continued*)

Step	Description
Add a new node to the cluster: <ul style="list-style-type: none"> ■ Using the product installer ■ Using the Web installer ■ Manually 	See “Adding a node to a cluster using the product installer” on page 244. See “Adding a node using the Web-based installer” on page 246. See “Adding the node to a cluster manually” on page 247.
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 263.

The example procedures describe how to add a node to an existing cluster with two nodes.

Before adding a node to a cluster

Before preparing to add the node to an existing cluster, perform the required preparations.


- Verify hardware and software requirements are met.
- Set up the hardware.
- Prepare the new node.

To verify hardware and software requirements are met

- 1 Review the hardware and software requirements.
 See [“Hardware compatibility list \(HCL\)”](#) on page 33.
- 2 Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
- 3 Verify the existing cluster is running SFHA 6.0 on all nodes.

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in [Figure 18-1](#).

Figure 18-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the SFHA private Ethernet controllers.
 - When you add nodes to a cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a two-node cluster, so you might have to swap out the cable for a switch or hub.
 - If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.
- Figure 18-1 illustrates a new node being added to an existing two-node cluster using two independent hubs.
- 2 Connect the system to the shared storage as required. Make sure that you meet the following requirements:
 - The new node must be connected to the same shared storage devices as the existing nodes.
 - The node must have private network connections to two independent switches for the cluster.

- The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

Prepare the new node before you add it to an existing cluster.

To prepare the new node

- 1 Verify that the new node meets installation requirements.

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

You can also use the Web-based installer for the 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.

See [“About installation and configuration methods”](#) on page 29.

Adding a node to a cluster using the product installer

You can add a node using the `-addnode` option with the product installer.

The product installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:
 - `/etc/llttab`
 - `/etc/VRTSvcs/conf/sysname`
- Copies the following files on the new node:
 - `/etc/llthosts`
 - `/etc/gabtab`
 - `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files on the new node:
 - `/etc/llthosts`
 - `/etc/gabtab`
 - `/etc/VRTSvcs/conf/config/main.cf`
 - `/etc/vxfenmode`
 - `/etc/vxfendg`

```

/etc/vx/.uuids/clusuuid
/etc/sysconfig/llt
/etc/sysconfig/gab
/etc/sysconfig/vxfen

```

- Generates security credentials on the new node if the CPS server of existing cluster is secure
- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.
Configures disk-based fencing.

At the end of the process, the new node joins the existing cluster.

Note: If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the product installer with the `-addnode` option.

```

# cd /opt/VRTS/install
# ./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 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 names of the systems that you want to add as new nodes to the cluster.

```

Enter the system names separated by spaces
to add to the cluster: saturn

```

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

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

Note: The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

```
Enter the NIC for the first private heartbeat
link on saturn: [b,q,?] eth1
```

```
Enter the NIC for the second private heartbeat
link on saturn: [b,q,?] eth2
```

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

- 7 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The product installer verifies the network interface settings and displays the information.
- 8 Review and confirm the information.
- 9 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: eth3
```

- 10 If the existing cluster uses server-based fencing in secure mode, the installer will configure server-based fencing in secure mode on the new nodes.

The installer then starts all the required Veritas processes and joins the new node to cluster.
- 11 Confirm that the new node has joined the existing cluster using `lltstat -n` and `gabconfig -a` commands.

Adding a node using the Web-based installer

You can use the Web-based installer to add a node to a cluster.

To add a node to a cluster using the Web-based installer

- 1 From the Task pull-down menu, select **Add a Cluster node**.
 From the product pull-down menu, select the product.
 Click the **Next** button.
- 2 Click **OK** to confirm the prerequisites to add a node.
- 3 In the System Names field enter a name of a node in the cluster where you plan to add the node and click **OK**.
 The installer program checks inter-system communications and compatibility. If the node fails any of the checks, review the error and fix the issue.
 If prompted, review the cluster's name, ID, and its systems. Click the **Yes** button to proceed.
- 4 In the System Names field, enter the names of the systems that you want to add to the cluster as nodes. Separate system names with spaces. Click the **Next** button.
 The installer program checks inter-system communications and compatibility. If the system fails any of the checks, review the error and fix the issue.
 Click the **Next** button. If prompted, click the **Yes** button to add the system and to proceed.
- 5 From the heartbeat NIC pull-down menus, select the heartbeat NICs for the cluster. Click the **Next** button.
- 6 Once the addition is complete, review the log files. Optionally send installation information to Symantec. Click the **Finish** button to complete the node's addition to the cluster.

Adding the node to a cluster manually

Perform the following tasks in order after you install to add the node to the cluster manually.

Table 18-2 Tasks for adding a cluster manually

Task	Description
Start the Volume Manager.	See "Starting Volume Manager on the new node" on page 255.

Table 18-2 Tasks for adding a cluster manually (*continued*)

Task	Description
Add the cluster files.	<p>Configure LLT, GAB, and VCSMM.</p> <p>Configure LLT and GAB.</p> <p>See “Configuring cluster files on the new node” on page 255.</p>
Generate security credentials.	<p>If the CP server of the existing cluster is secure, generate security credentials on the new node.</p> <p>See “Setting up the node to run in secure mode” on page 257.</p>
Set up fencing.	<p>Configure fencing for the new node to match the fencing configuration on the existing cluster.</p> <p>If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node.</p> <p>See “Starting fencing on the new node” on page 260.</p>
Start VCS.	<p>See “Starting VCS after adding the new node” on page 263.</p>
Configure Cluster Volume Manager (CVM) and Cluster File System (CFS).	<p>If CVM and CFS are configured on the existing cluster, configure them on the new node.</p>
Configure the ClusterService group.	<p>If the ClusterService group is configured on the existing cluster, add the node to the group.</p> <p>See “Configuring the ClusterService group for the new node” on page 260.</p>

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfha` program.

Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `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 cluster files on the new node

To configure LLT and GAB on the new node

- 1 Do not apply for SUSE Linux.
- 2 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
```

- 3 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.
- 4 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 5 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

```
/sbin/gabconfig -c -nN
```

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

- 6 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 7 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

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

```
# /etc/init.d/llt start  
# /etc/init.d/gab start  
# /etc/init.d/odm restart
```

- 10 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a  
GAB Port Memberships  
=====
```

Port a gen df204 membership 012
Port d gen df207 membership 012

Configuring cluster files on the new node

Perform the steps in the following procedure to configure cluster files on the new node.

To configure cluster files on the new node

- 1 If you are using Red Hat Linux, modify the file `/etc/sysctl.conf` on the new system to set the shared memory and other parameter required by your application; refer to the your application documentation for details. The value of the shared memory parameter is put to effect when the system restarts.

Do not apply this step for SUSE Linux.

- 2 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
```

- 3 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.
- 4 Create an `/etc/llttab` file on the new system.

For example:

```
set-node saturn
set-cluster 101

link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - ether - -
```

- 5 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.
- 6 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.

- 7 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.

- 8 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 9 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

Start the cluster components

- 1 Start LLT.

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

- 2 Start GAB.

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

- 3 Start disk-based fencing.

- 4 Start VCSMM.

```
# /etc/init.d/vcsmm start
```

```
# /etc/init.d/vxgms start
```

```
# /etc/init.d/vxglm start
```

- 5 Start LMX.

- 6 Start ODM.

```
# /etc/init.d/vxodm restart
```

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

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

# /etc/init.d/odm restart
```

- 8 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen df204 membership 012
Port d gen df207 membership 012
```

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/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:

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/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:

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:

```
# hastart
```

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.

```
# hagrp -modify ClusterService SystemList -add saturn 2
```

```
# hagrp -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 eth0 -sys saturn
# hares -modify gconic Device eth0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `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 cluster files on the new node

To configure LLT and GAB on the new node

- 1 Do not apply for SUSE Linux.
- 2 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
```

- 3 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.

- 4 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 5 Use `vi` or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

```
/sbin/gabconfig -c -nN
```

Where `N` represents the number of systems in the cluster including the new node. For a three-system cluster, `N` would equal 3.

- 6 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 7 Use `vi` or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# uuidconfig.pl -rsh -clus -copy \  
-from_sys galaxy -to_sys saturn
```

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

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

# /etc/init.d/gab start

# /etc/init.d/odm restart
```


10 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen df204 membership 012
Port d gen df207 membership 012
```

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

[Table 18-3](#) uses the following information for the following command examples.

Table 18-3 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
saturn	saturn.nodes.example.com	The new node that you are adding to the cluster.

Configuring the authentication broker on node saturn

To configure the authentication broker on node saturn

- 1 Extract the embedded authentication files and copy them to temporary directory:

```
# mkdir -p /var/VRTSvcs/vcsauth/bkup  
# cd /tmp; gunzip -c /opt/VRTSvcs/bin/VxAT.tar.gz | tar xvf -
```

- 2 Edit the setup file manually:

```
# /cat /etc/vx/.uuids/clusuuid 2>&1
```

The output is a string denoting the UUID. This UUID (without { and }) is used as the ClusterName for the setup file.

```
{UUID}
```

```
# cat /tmp/eat_setup 2>&1
```

The file content must resemble the following example:

```
AcceptorMode=IP_ONLY  
  
BrokerExeName=vcsauthserver  
  
ClusterName=UUID  
  
DataDir=/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER  
  
DestDir=/opt/VRTSvcs/bin/vcsauth/vcsauthserver  
  
FipsMode=0  
  
IPPort=14149  
  
RootBrokerName=vcsroot_uuid  
  
SetToRBPlusABorNot=0  
  
SetupPDRs=1  
  
SourceDir=/tmp/VxAT/version
```

- 3 Set up the embedded authentication file:

```
# cd /tmp/VxAT/version/bin/edition_number; \  
./broker_setup.sh/tmp/eat_setup  
  
/opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssregctl -s -f  
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER/root/.VRTSat/profile \  
/VRTSatlocal.conf -b 'Security\Authentication \  
\Authentication Broker' -k UpdatedDebugLogFileName \  
-v /var/VRTSvcs/log/vcsauthserver.log -t string
```

- 4 Copy the broker credentials from one node in the cluster to saturn by copying the entire `bkup` directory.

The `bkup` directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/  
  
# ls  
  
CMDSERVER  CPSADM  CPSEVER  HAD  VCS_SERVICES  WAC
```

- 5 Import the `VCS_SERVICES` domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \  
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER -f /var/VRTSvcs/vcsauth/bkup \  
/VCS_SERVICES -p password
```

- 6 Import the credentials for `HAD`, `CMDSERVER`, `CPSADM`, `CPSEVER`, and `WAC`.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \  
/var/VRTSvcs/vcsauth/data/VCS_SERVICES -f /var/VRTSvcs/vcsauth/bkup \  
/HAD -p password
```

- 7 Start the `vcsauthserver` process on saturn.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh
```

8 Perform the following tasks:

```
# mkdir /var/VRTSvcs/vcsauth/data/CLIENT
# mkdir /var/VRTSvcs/vcsauth/data/TRUST
# export EAT_DATA_DIR='/var/VRTSvcs/vcsauth/data/TRUST'
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat setuptrust -b \
localhost:14149 -s high
```

9 Create the `/etc/VRTSvcs/conf/config/.secure` file:

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

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1** For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

For SF Sybase CE:

For disk-based fencing, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/sysconfig/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:

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 eth0 -sys saturn
```

```
# hares -modify gconic Device eth0 -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 Storage Foundation High Availability, 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 Storage Foundation High Availability:

```
# 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:

```
# hastart
```

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 19-1](#) 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 19-1 Tasks that are involved in removing a node

Task	Reference
<ul style="list-style-type: none"> ■ 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 266.
<ul style="list-style-type: none"> ■ 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 267.
Modify the llthosts(4) and gabtab(4) files to reflect the change.	See “Modifying configuration files on each remaining node” on page 270.
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 270.

Table 19-1 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 271.
On the node departing the cluster: <ul style="list-style-type: none"> ■ 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. 	See “Unloading LLT and GAB and removing VCS on the departing node” on page 271.

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
```

- 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/llhosts` 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 SFHA as part of the Storage Foundation and High Availability products, you may have to delete other dependent RPMs before you can delete all of the following ones.

To stop LLT and GAB and remove SFHA

- 1 If you had configured I/O fencing in enabled mode, then stop I/O fencing.

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

- 2 Stop GAB and LLT:

```
# /etc/init.d/gab stop
```

```
# /etc/init.d/llt stop
```

- 3 To determine the RPMs to remove, enter:

```
# rpm -qa | grep VRTS
```

- 4 To permanently remove the VCS RPMs from the system, use the `rpm -e` command. Start by removing the following RPMs, which may have been optionally installed, in the order shown:

```
# rpm -e VRTSvcsea
# rpm -e VRTSatServer
# rpm -e VRTSatClient
# rpm -e VRTSvcshr
# rpm -e VRTSvcsg
# rpm -e VRTScps
# rpm -e VRTSvcS
# rpm -e VRTSamf
# rpm -e VRTSvxfen
# rpm -e VRTSgab
# rpm -e VRTSllt
# rpm -e VRTSspt
# rpm -e VRTSsfcp160
# rpm -e VRTSperl
# rpm -e VRTSvlic
```

- 5 Remove the LLT and GAB configuration files.


```
# rm /etc/llttab  
# rm /etc/gabtab  
# rm /etc/llthosts
```

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 263.

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 286.

Uninstalling Veritas Storage Foundation and High Availability

- [Chapter 20. Uninstalling Storage Foundation and High Availability](#)

Uninstalling Storage Foundation and High Availability

This chapter includes the following topics:

- [Removing VxFS file systems](#)
- [Removing rootability](#)
- [Moving volumes to disk partitions](#)
- [Disabling VCS agents for VVR the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling SFHA RPMs using the script-based installer](#)
- [Uninstalling SFHA with the Veritas Web-based installer](#)
- [Removing license files \(Optional\)](#)
- [Removing the CP server configuration using the removal script](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product](#)

Removing VxFS file systems

The VxFS RPM cannot be removed if there are any mounted VxFS file systems. Unmount all VxFS file systems before removing the RPM. After you remove the VxFS RPM, VxFS file systems are not mountable or accessible until another VxFS

RPM is installed. It is advisable to back up VxFS file systems before installing a new VxFS RPM. If VxFS will not be installed again, all VxFS file systems must be converted to a new file system type.

To remove VxFS file systems

- 1 Check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

- 2 Make backups of all data on the file systems that you wish to preserve, or recreate them as non-VxFS file systems on non-VxVM volumes or partitions.
- 3 Unmount all Storage Checkpoints and file systems:

```
# umount /checkpoint_name
```

```
# umount /filesystem
```

- 4 Comment out or remove any VxFS file system entries from the `/etc/fstab` file.

Removing rootability

Perform this procedure if you configured rootability by encapsulating the root disk.

To remove rootability

- 1 Check if the system's root disk is under VxVM control by running this command:

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- 2 Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Warning: Do not remove the plexes that correspond to the original root disk partitions.

- 3 Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices:

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

Moving volumes to disk partitions

All volumes must be moved to disk partitions.

This can be done using one of the following procedures:

- Back up the system fully onto tape and then recover from it.
- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Use VxVM to move volumes incrementally onto disk partitions as described in the following section.

Moving volumes onto disk partitions using VxVM

Use the following procedure to move volumes onto disk partitions.

To move volumes onto disk partitions

- 1 Evacuate disks using the `vxdiskadm` program, VEA, or the `vxevac` script. You should consider the amount of target disk space required for this before you begin.

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 using the following commands:

```
# vxdg -g diskgroup rmdisk disk _media_name
# vxdisk rm disk_access_name
```

- 3 Decide which volume to move first. If the volume to be moved 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 data on the volume is synced.
- 5 Create a partition on free disk space of the same size as the volume. If there is not enough free space for the partition, a new disk must be added to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.
- 6 Copy the data on the volume onto the newly created disk partition using a command similar to the following:

```
# dd if=/dev/vx/dsk/diskgroup/volume-name of=/dev/sdb2
```

where `sdb` is the disk outside of VxVM and `2` is the newly created partition on that disk.

- 7 Replace the entry for that volume (if present) in `/etc/fstab` with an entry for the newly created partition.
- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Stop the volume and remove it from VxVM using the following commands:

```
# vxvol -g diskgroup -f stop volume_name
# vxedit -g diskgroup -rf rm volume_name
```

- 10 Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

```
# vxprint -F "%sdnum" disk_media_name
```

- 11 If the output is not 0, there are still some subdisks on this disk that must be subsequently removed. If the output is 0, remove the disk from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name  
# vxdisk rm disk_access_name
```

- 12 The free space now created can be used for adding the data in the next volume to be removed.
- 13 After all volumes have been converted into disk partitions successfully, reboot the system. After the reboot, none of the volumes should be open. To verify that none of the volumes are open, use the following command:

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

- 14 If any volumes remain open, repeat the steps listed above.

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 RPMs using the script-based installer

Use the following procedure to remove SFHA products.

Not all RPMs 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 with a previous version of SFHA.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 354.

To shut down and remove the installed SFHA RPMs

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/fstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems.

```
# umount /mount_point
```

- 3 If the VxVM RPM (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

- 4 Make sure you have performed all of the prerequisite steps.

- 5 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 6 Move to the `/opt/VRTS/install` directory and run the uninstall script.

```
# cd /opt/VRTS/install
```

For Veritas Storage Foundation 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 RPMs are uninstalled.

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

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

Uninstalling SFHA with the Veritas Web-based installer

This section describes how to uninstall using the Veritas Web-based installer.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout Version in SFHA 6.0 with with a previous version of SFHA.

To uninstall SFHA

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 61.
- 3 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 4 Select **Storage Foundation High Availability** from the Product drop-down list, and click **Next**.
- 5 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Next**.
- 6 After the validation completes successfully, click **Next** to uninstall SFHA on the selected system.
- 7 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 8 After the installer stops the processes, the installer removes the products from the specified system.
Click **Next**.
- 9 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.
- 10 Click **Finish**.

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

Removing license files (Optional)

Optionally, you can remove the license files.

To remove the VERITAS license files

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

```
# /sbin/vxlicrep
```

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

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

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

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

Removing the CP server configuration using the removal script

This section describes how to remove the CP server configuration from a node or a cluster that hosts the CP server.

Warning: Ensure that no 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) :
```

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

- 6 Answer **y** at the prompt to confirm the deletion of the CP server database.

```
Warning: This database won't be available if CP server
is reconfigured on the cluster. Are you sure you want to
proceed with the deletion of database? (y/n) (Default:n) :
```

- 7 Answer **y** to delete the CP server configuration file and log files.

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 8 Run the `hagrp -state` command to ensure that the CPSSG service group has been removed from the node. For example:

```
root@mycps1.symantecexample.com # hagrp -state CPSSG
```

```
VCS WARNING V-16-1-40131 Group CPSSG does not exist
in the local cluster
```

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file disables the SFDB tools.

To remove the SFDB repository

- 1 Identify the SFDB repositories created on the host.

```
# cat /var/vx/vxdba/rep_loc

{
  "sfae_rept_version" : 1,
  "oracle" : {
    "SFAEDB" : {
      "location" : "/data/sfaedb/.sfae",
      "old_location" : "",
      "alias" : [
        "sfaedb"
      ]
    }
  }
}
```

- 2 Remove the directory identified by the `location` key.

```
# rm -rf /data/sfaedb/.sfae
```

- 3 Remove the repository location file.

```
# rm -rf /var/vx/vxdba/rep_loc
```

This completes the removal of the SFDB repository.

Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Response files](#)
- [Appendix C. Tunable files for installation](#)
- [Appendix D. Configuring I/O fencing using a response file](#)
- [Appendix E. 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 Storage Foundation High Availability setup diagrams for CP server-based I/O fencing](#)
- [Appendix K. Configuring LLT over UDP using IPv4](#)
- [Appendix L. Configuring LLT over UDP using IPv6](#)
- [Appendix M. Compatibility issues when installing Storage Foundation High Availability with other products](#)

Installation scripts

This appendix includes the following topics:

- [About Veritas Storage Foundation and High Availability Solutions installation scripts](#)
- [Installation script options](#)
- [About using the postcheck option](#)

About Veritas Storage Foundation and High Availability Solutions installation scripts

Veritas Storage Foundation and High Availability Solutions products 6.0 provides several installation scripts. You can find these scripts at the root of the product media in the scripts directory.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from the Symantec download site, which does not include the installer, use the appropriate product installation script.

See [“Downloading the Storage Foundation and High Availability software”](#) on page 30.

The following product installation scripts are available:

Table A-1 Product installation scripts

Product installation script	Veritas product name
<code>installvcs</code>	Veritas Cluster Server (VCS)
<code>installsf</code>	Veritas Storage Foundation (SF)

Table A-1 Product installation scripts (*continued*)

Product installation script	Veritas product name
<code>installsfha</code>	Veritas Storage Foundation and High Availability (SFHA)
<code>installsfcfsha</code>	Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)
<code>installsfrac</code>	Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)
<code>installsfsybasece</code>	Veritas Storage Foundation for Sybase ASE CE (SF Sybase CE)
<code>installvm</code>	Veritas Volume Manager
<code>installfs</code>	Veritas File System
<code>installdmp</code>	Veritas Dynamic Multi-Pathing
<code>installsvs</code>	Symantec VirtualStore

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

Installation script options

[Table A-2](#) shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts, except where otherwise noted.

See [“About Veritas Storage Foundation and High Availability Solutions installation scripts”](#) on page 289.

Table A-2 Available command line options

Command Line Option	Function
<code>system1 system2...</code>	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
<code>--addnode</code>	Adds a node to a high availability cluster.

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

Command Line Option	Function
<code>-allpkgs</code>	Displays all RPMs and patches required for the specified product. The RPMs and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
<code>-comcleanup</code>	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
<code>-configure</code>	Configures the product after installation.
<code>-fencing</code>	Configures I/O fencing in a running cluster.
<code>-hostfile <i>full_path_to_file</i></code>	Specifies the location of a file that contains a list of hostnames on which to install.
<code>-install</code>	The <code>-install</code> option is used to install products on systems.
<code>-installallpkgs</code>	Specifies that all RPMs are installed.
<code>-installminpkgs</code>	Specifies that the minimum RPM set is installed.
<code>-installrecpkgs</code>	Specifies that the required RPM set is installed.
<code>-keyfile <i>ssh_key_file</i></code>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i <i>ssh_key_file</i></code> to every SSH invocation.
<code>-kickstart <i>dir_path</i></code>	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec rpms in the correct order for installing, in a format that can be used for Kickstart installations. The <code><i>dir_path</i></code> indicates the path to the directory in which to create the file.
<code>-license</code>	Registers or updates product licenses on the specified systems.

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

Command Line Option	Function
<code>-logpath <i>log_path</i></code>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
<code>-makeresponsefile</code>	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.
<code>-minpkgs</code>	Displays the minimal RPMs and patches required for the specified product. The RPMs and patches are listed in correct installation order. Optional RPMs are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.
<code>-nolic</code>	Allows installation of product RPMs without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
<code>-pkginfo</code>	Displays a list of RPMs and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the <code>-pkginfo</code> option with the <code>installvcs</code> script to display VCS RPMs.
<code>-pkgpath <i>package_path</i></code>	Designates the path of a directory that contains all RPMs to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
<code>-pkgset</code>	Discovers and displays the RPM group (minimum, recommended, all) and RPMs that are installed on the specified systems.
<code>-pkgtable</code>	Displays product's RPMs in correct installation order by group.
<code>-postcheck</code>	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.

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

Command Line Option	Function
<code>-precheck</code>	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
<code>-recpkgs</code>	Displays the recommended RPMs and patches required for the specified product. The RPMs and patches are listed in correct installation order. Optional RPMs are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.
<code>-redirect</code>	Displays progress details without showing the progress bar.
<code>-requirements</code>	The <code>-requirements</code> option displays required OS version, required RPMs and patches, file system space, and other system requirements in order to install the product.
<code>-responsefile <i>response_file</i></code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<code>-rolling_upgrade</code>	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
<code>-rollingupgrade_phase1</code>	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs 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 RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."

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

Command Line Option	Function
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See “About configuring secure shell or remote shell communication modes before installing products” on page 354.
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-setunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the <code>-tunablesfile</code> option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-tmppath <i>tmp_path</i>	Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.
-uninstall	The <code>-uninstall</code> 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.

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

Command Line Option	Function
<code>-version</code>	Checks and reports the installed products and their versions. Identifies the installed and missing RPMs and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPMs and patches where applicable. Lists the installed patches, hotfixes, and available updates for the installed product if an Internet connection is available.
<code>-yumgroupxml</code>	The <code>-yumgroupxml</code> option is used to generate a yum group definition XML file. The <code>createrepo</code> command can use the file on Redhat Linux to create a yum group for automated installation of all rpms for a product. An available location to store the XML file should be specified as a complete path. The <code>-yumgroupxml</code> option is supported on Redhat Linux only.

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 231.

Response files

This appendix includes the following topics:

- [About response files](#)
- [Installing SFHA using response files](#)
- [Configuring SFHA using response files](#)
- [Upgrading 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](#)
- [Sample response file for SF upgrade](#)
- [Sample response file for SFHA upgrade](#)

About response files

The installer or product installation script generates a response file during any installation, configuration, upgrade, or uninstall procedure. The response file contains the configuration information that you entered during the procedure. When the procedure completes, the installation script displays the location of the response files.

You can use the response file for future installation procedures by invoking an installation script with the `-responsefile` option. The response file passes arguments to the script to automate the installation of that product. You can edit the file to automate installation and configuration of additional systems.

You can generate a response file using the `-makeresponsefile` option.

See “[Installation script options](#)” on page 290.

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 RPMs 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 303.

- 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.

Upgrading SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA upgrade on one system to upgrade SFHA on other systems. You can also create a response file using the `makeresponsefile` option of the installer.

To perform automated SFHA upgrade

- 1 Make sure the systems where you want to upgrade SFHA meet the upgrade requirements.
- 2 Make sure the pre-upgrade tasks are completed.
- 3 Copy the response file to one of the systems where you want to upgrade SFHA.
- 4 Edit the values of the response file variables as necessary.

- 5 Mount the product disc and navigate to the folder that contains the installation program.
- 6 Start the upgrade from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file  
# ./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 RPMs. 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}	<p>Defines a location, typically an NFS mount, from which all remote systems can install product RPMs. The location must be accessible from all target systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{tmppath}	<p>Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /var/tmp.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{rsh}	<p>Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{donotinstall} {RPM}	<p>Instructs the installation to not install the optional RPMs in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{donotremove} {RPM}	<p>Instructs the uninstallation to not remove the optional RPMs in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{opt}{logpath}	<p>Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
\$CFG{opt}{prodmode}	<p>List of modes for product</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>

Table B-1 Response file variables specific to installing, upgrading, or uninstalling SFHA (*continued*)

Variable	Description
CFG{opt}{upgrade}	Upgrades all RPMs installed, without configuration. List or scalar: list Optional or required: optional
CFG{opt}{uninstall}	Uninstalls SFHA RPMs. 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 RPMs 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 <i>ssh</i> as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is <code>/opt/VRTS/install/logs</code> . Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group

(csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtprecp, and smtpsrsev), the SNMP trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Table B-3 lists the response file variables that specify the required information to configure a basic 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	<p>Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links.</p> <p>You must enclose the system name within double quotes.</p> <p>(Required)</p>
CFG{vcs_lltlinklowpri#} {"system"}	Scalar	<p>Defines a low priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication.</p> <p>If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on.</p> <p>You must enclose the system name within double quotes.</p> <p>(Optional)</p>

[Table B-5](#) lists the response file variables that specify the required information to configure LLT over UDP.

Table B-5 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltverudp}=1	Scalar	<p>Indicates whether to configure heartbeat link using LLT over UDP.</p> <p>(Required)</p>

Table B-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG {vcs_udplinklowpri<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table B-5 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table B-6 lists the response file variables that specify the required information to configure virtual IP for 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. <ul style="list-style-type: none"> ■ 1—Configure the first node ■ 2—Configure the other node
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.

Table B-8 lists the response file variables that specify the required information to configure VCS users.

Table B-8 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	List of encoded passwords for VCS users The value in the list can be "Administrators Operators Guests" Note: The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list. (Optional)
CFG{vcs_username}	List	List of names of VCS users (Optional)

Table B-8 Response file variables specific to configuring VCS users (*continued*)

Variable	List or Scalar	Description
CFG{vcs_userpriv}	List	List of privileges for VCS users Note: The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list. (Optional)

Table B-9 lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table B-9 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)
CFG{vcs_smtpsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

Table B-10 lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table B-10 Response file variables specific to configuring VCS notifications using SNMP

Variable	List or Scalar	Description
CFG{vcs_snmpport}	Scalar	Defines the SNMP trap daemon port (default=162). (Optional)
CFG{vcs_snmpcons}	List	List of SNMP console system names (Optional)
CFG{vcs_snmpcsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

[Table B-11](#) lists the response file variables that specify the required information to configure 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{upi}="SF";  
$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(JlmElgLimHmKumGlj bQOsOUUnVQoOUUnTQoSOSnUQuOUUnPQtOS) ];  
$CFG{vcs_userpriv}=[ qw(Administrators Operators) ];  
$CFG{vcs_lltlink1}{system01}="eth1";  
$CFG{vcs_lltlink2}{system01}="eth2";  
$CFG{vcs_lltlink1}{system02}="eth1";  
$CFG{vcs_lltlink2}{system02}="eth2";  
$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 #
#####

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;
```

Sample response file for SF upgrade

The following example shows a response file for upgrading Storage Foundation.

```
our %CFG;

$CFG{accepteula}=1;
$CFG{opt}{upgrade}=1;
$CFG{systems}=[ qw(system01) ];

1;
```

Sample response file for SFHA upgrade

The following example shows a response file for upgrading Storage Foundation High Availability.

```
our %CFG;
$CFG{accepteula}=1;
$CFG{mirrordgname}{backupdg_01}="backupdg";
$CFG{mirrordgname}{backupdg_02}="mirrorrootdg";
$CFG{opt}{upgrade}=1;
$CFG{splitmirror}{system01}=1;
```

```
$CFG{splitmirror}{system02}=1;  
$CFG{systems}=[ qw(system01 system02) ];  
$CFG{vcs_allowcomms}=1;  
1;
```

The `vcs_allowcomms` variable is set to 0 if it is a single-node cluster, and the `llt` and `gab` processes are not started before upgrade.

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 316.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -setttunables [  
system1 system2 ...]
```

See [“Setting tunables with no other installer-related operations”](#) on page 317.

- 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 318.

See [“About response files”](#) on page 297.

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

See [“Tunables value parameter definitions”](#) on page 320.

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 320.

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 319.
- 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 320.

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 319.
- 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 320.

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 319.
- 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 320.

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_support	(Veritas Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_path_age	(Veritas Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_pathswitch_blks_shift	(Veritas Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_idle_lun	(Veritas Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

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

Tunable	Description
dmp_probe_threshold	(Veritas Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_cycles	(Veritas Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_interval	(Veritas Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_policy	(Veritas Dynamic Multi-Pathing) The policy used by DMP path restoration thread. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_state	(Veritas Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_retry_count	(Veritas Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_scsi_timeout	(Veritas Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_sfg_threshold	(Veritas Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_stat_interval	(Veritas Dynamic Multi-Pathing) The time interval between gathering DMP statistics. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer sets only the system default value of max_diskq. Refer to the tune(4) manual page for setting this tunable for a specified block device.

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

Tunable	Description
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer sets only the system default value of read_ahead. Refer to the tunefstab(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 tunefstab(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 tunefstab(4) manual page for setting this tunable for a specified block device.
vol_checkpt_default	(Veritas File System) Size of VxVM checkpoints (sectors). This tunable requires system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for VERITAS Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for VERITAS Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires system reboot to take effect.
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).

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

Tunable	Description
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of VERITAS Volume Replicator (bytes).
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (sectors). This tunable requires system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (sectors). This tunable requires system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Veritas Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by VERITAS Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions that can exist on a non-sequential DRL. This tunable requires system reboot to take effect.
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (sectors). This tunable requires system reboot to take effect.

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

Tunable	Description
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_jobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_jobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires system reboot to take effect.
voliot_jobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires system reboot to take effect.
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires system reboot to take effect.

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

Tunable	Description
vx_era_nthreads	(Veritas File System) Maximum number of threads VxFS will detect read_ahead patterns on. This tunable requires system reboot to take effect.
vx_bc_bufhwm	(Veritas File System) VxFS metadata buffer cache high water mark. This tunable requires system reboot to take effect.
vxfs_mbuf	(Veritas File System) Maximum memory used for VxFS buffer cache. This tunable requires system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer sets only the system default value of write_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer sets only the system default value of write_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.

Configuring I/O fencing using a response file

This appendix includes the following topics:

- [Configuring I/O fencing using response files](#)
- [Response file variables to configure disk-based I/O fencing](#)
- [Sample response file for configuring disk-based I/O fencing](#)
- [Response file variables to configure server-based I/O fencing](#)
- [Sample response file for configuring 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 75.

- 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 329.
 See “[Sample response file for configuring server-based I/O fencing](#)” on page 332.
- 4 Edit the values of the response file variables as necessary.
 See “[Response file variables to configure disk-based I/O fencing](#)” on page 328.
 See “[Response file variables to configure server-based I/O fencing](#)” on page 330.
- 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. <ul style="list-style-type: none"> ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode ■ 4—Fencing migration when the cluster is online (Required)

Table D-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG {fencing_scsi3_disk_policy}	Scalar	Specifies the I/O fencing mechanism. This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the <code>fencing_scsi3_disk_policy</code> variable and either the <code>fencing_dgname</code> variable or the <code>fencing_newdg_disks</code> variable. (Optional)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.
CFG{fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.

Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions.

See [“Response file variables to configure disk-based I/O fencing”](#) on page 328.

```
#
# 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(sdx sdy sdz) ];
$CFG{fencing_option}=2;
```

Response file variables to configure server-based I/O fencing

You can use a coordination point server-based fencing response file to configure server-based customized I/O fencing.

[Table D-2](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table D-2 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_config_cpagent}	<p>Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not.</p> <p>Enter "0" if you do not want to configure the Coordination Point agent using the installer.</p> <p>Enter "1" if you want to use the installer to configure the Coordination Point agent.</p>

Table D-2 Coordination point server (CP server) based fencing response file definitions *(continued)*

Response file field	Definition
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. Note: This field is obsolete if the <code>fencing_config_cpagent</code> field is given a value of '0'.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_reusedg}	This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks). Enter either a "1" or "0". Entering a "1" indicates reuse, and entering a "0" indicates do not reuse. When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text, such as " <code>CFG{fencing_reusedg}=0</code> " or " <code>CFG{fencing_reusedg}=1</code> " before proceeding with a silent installation.
CFG {fencing_dgname}	The name of the disk group to be used in the customized fencing, where at least one disk is being used.
CFG {fencing_disks}	The disks being used as coordination points if any.
CFG {fencing_ncp}	Total number of coordination points being used, including both CP servers and disks.
CFG {fencing_ndisks}	The number of disks being used.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_ports}	The port that the virtual IP address or the fully qualified host name of the CP server listens on.
CFG {fencing_scsi3_disk_policy}	The disk policy that the customized fencing uses. The value for this field is either "raw" or "dmp"

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing:

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cps_vips}{"10.200.117.145"}=[ qw(10.200.117.145) ];
$CFG{fencing_dgname}="vxfencoorddg";
$CFG{fencing_disks}=[ qw(emc_clariion0_37 emc_clariion0_13) ];
$CFG{fencing_scsi3_disk_policy}="raw";
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=2;
$CFG{fencing_ports}{"10.200.117.145"}=14250;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="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 25.

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/sysconfig/llt	<p>This file stores the start and stop environment variables for LLT:</p> <ul style="list-style-type: none"> ■ LLT_START—Defines the startup behavior for the LLT module after a system reboot. Valid values include: <ul style="list-style-type: none"> 1—Indicates that LLT is enabled to start up. 0—Indicates that LLT is disabled to start up. ■ LLT_STOP—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include: <ul style="list-style-type: none"> 1—Indicates that LLT is enabled to shut down. 0—Indicates that LLT is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of SFHA configuration.</p>
/etc/llthosts	<p>The file <code>llthosts</code> is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.</p> <p>For example, the file <code>/etc/llthosts</code> contains the entries that resemble:</p> <pre style="margin-left: 40px;"> 0 galaxy 1 nebula </pre>

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

File	Description
<code>/etc/llttab</code>	<p>The file <code>llttab</code> contains the information that is derived during installation and used by the utility <code>lltconfig (1M)</code>. After installation, this file lists the LLT network links that correspond to the specific system.</p> <pre style="margin-left: 40px;">set-node galaxy set-cluster 2 link eth1 eth1 - ether - - link eth2 eth2 - ether - -</pre> <p>For example, the file <code>/etc/llttab</code> contains the entries that resemble:</p> <pre style="margin-left: 40px;">set-node galaxy set-cluster 2 link eth1 eth-00:04:23:AC:12:C4 - ether - - link eth2 eth-00:04:23:AC:12:C5 - ether - -</pre> <p>If you use aggregated interfaces, then the file contains the aggregated interface name instead of the <code>eth-MAC_address</code>.</p> <p>The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the <code>link</code> command. These lines identify the two network cards that the LLT protocol uses.</p> <p>If you configured a low priority link under LLT, the file also includes a "link-lowpri" line.</p> <p>Refer to the <code>llttab(4)</code> manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the <code>llttab</code> file.</p>

[Table E-2](#) lists the GAB configuration files and the information that these files contain.

Table E-2 GAB configuration files

File	Description
<code>/etc/sysconfig/gab</code>	<p>This file stores the start and stop environment variables for GAB:</p> <ul style="list-style-type: none"> ■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: <ul style="list-style-type: none"> 1—Indicates that GAB is enabled to start up. 0—Indicates that GAB is disabled to start up. ■ GAB_STOP—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include: <ul style="list-style-type: none"> 1—Indicates that GAB is enabled to shut down. 0—Indicates that GAB is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of SFHA configuration.</p>
<code>/etc/gabtab</code>	<p>After you install SFHA, the file <code>/etc/gabtab</code> contains a <code>gabconfig(1)</code> command that configures the GAB driver for use.</p> <p>The file <code>/etc/gabtab</code> contains a line that resembles:</p> <pre style="margin-left: 40px;"><code>/sbin/gabconfig -c -nN</code></pre> <p>The <code>-c</code> option configures the driver for use. The <code>-nN</code> specifies that the cluster is not formed until at least <i>N</i> nodes are ready to form the cluster. Symantec recommends that you set <i>N</i> to be the total number of nodes in the cluster.</p> <p>Note: Symantec does not recommend the use of the <code>-c -x</code> option for <code>/sbin/gabconfig</code>. Using <code>-c -x</code> can lead to a split-brain condition.</p>

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
<code>/etc/sysconfig/amf</code>	<p>This file stores the start and stop environment variables for AMF:</p> <ul style="list-style-type: none"> ■ 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: <ul style="list-style-type: none"> 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: <ul style="list-style-type: none"> 1—Indicates that AMF is enabled to shut down. (default) 0—Indicates that AMF is disabled to shut down.
<code>/etc/amftab</code>	<p>After you install VCS, the file <code>/etc/amftab</code> contains a <code>amfconfig(1)</code> command that configures the AMF driver for use.</p> <p>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:</p> <pre><code>/opt/VRTSamf/bin/amfconfig -c</code></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 VCS clusters”](#) on page 341.
See [“Sample main.cf file for global clusters”](#) on page 342.
- `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/sysconfig/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 Linux systems.

Sample main.cf file for VCS clusters

The following sample main.cf file is for a cluster in secure mode.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTTypes.cf"
include "Db2udbTypes.cf"
include "SybaseTypes.cf"

cluster vcs_cluster2 (
    UserNames = { admin = cDRpdxPmHpzS, smith = dKLhKJkHLh }
    ClusterAddress = "192.168.1.16"
    Administrators = { admin, smith }
    CounterInterval = 5
    SecureClus = 1
)

system galaxy (
)

system nebula (
)

group ClusterService (
    SystemList = { galaxy = 0, nebula = 1 }
    UserStrGlobal = "LocalCluster@https://10.182.2.76:8443;"
    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

IP webip (
    Device = eth0
    Address = "192.168.1.16"
    NetMask = "255.255.240.0"
)

NIC csgnic (
    Device = eth0
    NetworkHosts = { "192.168.1.17", "192.168.1.18" }
)
```

```
NotifierMngr ntfr (
    SnmpConsoles = { "saturn" = Error, "jupiter" = SevereError }
    SntpServer = "smtp.example.com"
    SntpRecipients = { "ozzie@example.com" = Warning,
                     "harriet@example.com" = Error }
)

webip requires csgnic
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.

```
.
.
group ClusterService (
    SystemList = { galaxy = 0, nebula = 1 }

    UserStrGlobal = "LocalCluster@https://10.182.2.78:8443;"

    AutoStartList = { galaxy, nebula }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
)

```

```
.  
.
```

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 = eth0  
    Address = "10.182.13.50"  
    NetMask = "255.255.240.0"  
)  
  
NIC csgnic (  

```

```
Device = eth0
NetworkHosts = { "10.182.13.1" }
)

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

Table E-4 I/O fencing configuration files (*continued*)

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

Table E-4 I/O fencing configuration files (*continued*)

File	Description
<code>/etc/vxfentab</code>	<p>When I/O fencing starts, the vxfen startup script creates this <code>/etc/vxfentab</code> file on each node. The startup script uses the contents of the <code>/etc/vxfendg</code> and <code>/etc/vxfenmode</code> files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.</p> <p>Note: The <code>/etc/vxfentab</code> file is a generated file; do not modify this file.</p> <p>For disk-based I/O fencing, the <code>/etc/vxfentab</code> file on each node contains a list of all paths to each coordinator disk. An example of the <code>/etc/vxfentab</code> file in a disk-based fencing configuration on one node resembles as follows:</p> <ul style="list-style-type: none"> ■ Raw disk: <ul style="list-style-type: none"> <code>/dev/sdx</code> <code>/dev/sdy</code> <code>/dev/sdz</code> ■ DMP disk: <ul style="list-style-type: none"> <code>/dev/vx/rdmp/sdx3</code> <code>/dev/vx/rdmp/sdy3</code> <code>/dev/vx/rdmp/sdz3</code> <p>For server-based fencing, the <code>/etc/vxfentab</code> file also includes the security settings information.</p> <p>For server-based fencing with single CP server, the <code>/etc/vxfentab</code> file also includes the <code>single_cp</code> settings information.</p>

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 353.

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 348.
- 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 350.

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 = eth0
    Address = "10.209.3.1"
```

```

        NetMask = "255.255.252.0"
    )

IP cpsvip2 (
    Critical = 0
    Device @mycps1 = eth1
    Address = "10.209.3.2"
    NetMask = "255.255.252.0"
)

NIC cpsnic1 (
    Critical = 0
    Device @mycps1 = eth0
    PingOptimize = 0
    NetworkHosts @mycps1 = { "10.209.3.10" }
)

NIC cpsnic2 (
    Critical = 0
    Device @mycps1 = eth1
    PingOptimize = 0
)

Process vxcperv (
    PathName = "/opt/VRTScps/bin/vxcperv"
    ConfInterval = 30
    RestartLimit = 3
)

Quorum quorum (
    QuorumResources = { cpsvip1, cpsvip2 }
)

cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
vxcperv requires quorum

// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1

```

```
// {  
//   NIC cpsnic1  
// }  
// IP cpsvip2  
// {  
//   NIC cpsnic2  
// }  
// Process vxcpserv  
// {  
//   Quorum quorum  
// }  
// }
```

Sample main.cf file for CP server hosted on a two-node SFHA cluster

The following is an example of a main.cf, where the CP server is hosted on an SFHA cluster.

For this CP server hosted on an SFHA cluster main.cf, note the following values:

- Cluster name: cps1
- Nodes in the cluster: mycps1, mycps2

```
include "types.cf"  
include "CFSTypes.cf"  
include "CVMTypes.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 = eth0
        Device @mycps2 = eth0
        Address = "10.209.81.88"
        NetMask = "255.255.252.0"
    )

    IP cpsvip2 (
        Critical = 0
        Device @mycps1 = eth1
        Device @mycps2 = eth1
        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 = eth0
        Device @mycps2 = eth0
        PingOptimize = 0
        NetworkHosts @mycps1 = { "10.209.81.10" }
    )

```

```

NIC cpsnic2 (
    Critical = 0
    Device @mycps1 = eth1
    Device @mycps2 = eth1
    PingOptimize = 0
)

Process vxcpserv (
    PathName = "/opt/VRTScps/bin/vxcpserv"
)

Quorum quorum (
    QuorumResources = { cpsvip1, cpsvip2 }
)

Volume cpsvol (
    Volume = cps_volume
    DiskGroup = cps_dg
)

cpsmount requires cpsvol
cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
cpsvol requires cpsdg
vxcpserv requires cpsmount
vxcpserv requires quorum

// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1
//   {
//     NIC cpsnic1
//   }
// IP cpsvip2
//   {
//     NIC cpsnic2
//   }
// Process vxcpserv
//   {

```



```
//      Quorum quorum
//      Mount cpsmount
//      {
//          Volume cpsvol
//          {
//              DiskGroup cpsdg
//          }
//      }
//  }
```

Sample CP server configuration (/etc/vxcps.conf) file output

The following is an example of a coordination point server (CP server) configuration file `/etc/vxcps.conf` output.

```
## The vxcps.conf file determines the
## configuration for Veritas CP Server.
cps_name=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 rsh for Linux](#)

About configuring secure shell or remote shell communication modes before installing products

Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installer is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install Veritas software.

You can install products to remote systems using either secure shell (`ssh`) or remote shell (`rsh`). Symantec recommends that you use `ssh` as it is more secure than `rsh`.

This section contains an example of how to set up `ssh` password free communication. The example sets up `ssh` between a source system (`system1`) that contains the installation directories, and a target system (`system2`). This procedure also applies to multiple target systems.

Note: The script- and Web-based installers support establishing password less communication for you.


Manually configuring and passwordless ssh

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

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

Figure 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 /root
```

- 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 (/root/.ssh/id_dsa):
```

- 3 Press Enter to accept the default location of `/root/.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 Output similar to the following lines appears.

```
Your identification has been saved in /root/.ssh/id_dsa.  
Your public key has been saved in /root/.ssh/id_dsa.pub.  
The key fingerprint is:  
1f:00:e0:c2:9b:4e:29:b4:0b:6e:08:f8:50:de:48:d2 root@system1
```

To append the public key from the source system to the `authorized_keys` file on the target system, using secure file transfer

- 1 From the source system (system1), move the public key to a temporary file on the target system (system2).

Use the secure file transfer program.

In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

Use the following command for secure file transfer:

```
system1 # sftp system2
```

If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

```
Connecting to system2 ...
The authenticity of host 'system2 (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
Are you sure you want to continue connecting (yes/no)?
```

- 2 Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'system2,10.182.00.00'
(DSA) to the list of known hosts.
root@system2 password:
```

- 3 Enter the root password of system2.
- 4 At the `sftp` prompt, type the following command:

```
sftp> put /root/.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /root/.ssh/id_dsa.pub to /root/id_dsa.pub
```

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

```
sftp> quit
```

- 6 Add the `id_dsa.pub` keys to the `authorized_keys` file on the target system. 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:
```

Type the following commands on system2:

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

- 7 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 /root/.ssh/id_dsa.pub >> /root/.ssh/authorized_keys
```

- 8 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: /root/.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 rsh for Linux

The following section describes how to enable remote shell.

Veritas recommends configuring a secure shell environment for Veritas product installations.

See [“Manually configuring and passwordless ssh”](#) on page 355.

See the operating system documentation for more information on configuring remote shell.

To enable rsh

- 1 To ensure that the `rsh` and `rsh-server` packages are installed, type the following command:

```
# rpm -qa | grep -i rsh
```

If it is not already in the file, type the following command to append the line "rsh" to the `/etc/securetty` file:

```
# echo "rsh" >> /etc/securetty
```

- 2 Modify `/etc/init.d/rsh` `disable=no`.

- 3 In the `/etc/pam.d/rsh` file, change the "auth" type from "required" to "sufficient":

```
auth    sufficient
```

- 4 Add the "promiscuous" flag into `/etc/pam.d/rsh` and `/etc/pam.d/rlogin` after item "pam_rhosts_auth.so".
- 5 To enable the rsh server, type the following command:

```
# chkconfig rsh on
```

- 6 Modify the `.rhosts` file. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system. This file also contains the name of a user having access to the local system. For example, if the root user must remotely access `system1` from `system2`, add an entry for `system2.companyname.com` to the `.rhosts` file on `system1` by typing the following command:

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

- 7 Install the Veritas product.

To disable rsh

- 1 Remove the "rsh" entry in the `/etc/securetty` file.
- 2 Disable the rsh server by typing the following command:

```
# chkconfig rsh off
```

- 3 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 RPMs](#)
- [Veritas Cluster Server installation RPMs](#)
- [Veritas Storage Foundation obsolete and reorganized installation RPMs](#)

Storage Foundation and High Availability installation RPMs

[Table G-1](#) shows the RPM name and contents for each English language RPM for Storage Foundation and High Availability. The table also gives you guidelines for which RPMs 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) RPMs, the combined functionality is called Storage Foundation and High Availability and High Availability.

See [“Veritas Cluster Server installation RPMs”](#) on page 364.

Table G-1 Storage Foundation and High Availability RPMs

RPMs	Contents	Configuration
VRTSaslapm	Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries Required for the support and compatibility of various storage arrays.	Minimum
VRTSperl	Perl 5.12.2 for Veritas	Minimum
VRTSvlic	Veritas License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Minimum
VRTSvxfv	Veritas File System binaries Required for VxFS file system support.	Minimum
VRTSvxvm	Veritas Volume Manager binaries, scripts, and utilities. Required for VxVM volume manager support.	Minimum
VRTSdbed	Veritas Storage Foundation for Databases	Recommended
VRTSob	Veritas Enterprise Administrator	Recommended
VRTSodm	Veritas ODM Driver for VxFS Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager allows Oracle 9i and 10g to improve performance and manage system bandwidth.	Recommended

Table G-1 Storage Foundation and High Availability RPMs (*continued*)

RPMs	Contents	Configuration
VRTSsfcp160	<p>Veritas Storage Foundation Common Product Installer</p> <p>The Storage Foundation Common Product installer RPM contains the installer libraries and product scripts that perform the following:</p> <ul style="list-style-type: none"> ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ removing nodes ■ etc. <p>You can use these script to simplify the native operating system installations, configurations, and upgrades.</p>	Minimum
VRTSsfmh	<p>Veritas Storage Foundation Managed Host</p> <p>Discovers configuration information on a Storage Foundation managed host. This information is stored on a central database, which is not part of this release. You must download the database separately at:</p>	Recommended
VRTSspt	Veritas Software Support Tools	Recommended
VRTSsvcsdr	Contains the binaries for Veritas Cluster Server disk reservation.	Recommended
VRTSfsadv	Minimum Veritas File System Advanced Solutions by Symantec (Solaris SPARC only).	Minimum
VRTSfssdk	<p>Veritas File System Software Developer Kit</p> <p>For VxFS APIs, the RPM contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.</p>	All

Table G-1 Storage Foundation and High Availability RPMs (*continued*)

RPMs	Contents	Configuration
VRTSivmconv	Symantec Logical Volume Manager (LVM) volume converter Converts offline Linux LVM managed volumes to VxVM volumes by rearranging media contents.	All

Veritas Cluster Server installation RPMs

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

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

See [“Storage Foundation and High Availability installation RPMs”](#) on page 361.

Table G-2 VCS installation RPMs

RPM	Contents	Configuration
VRTSgab	Veritas Cluster Server group membership and atomic broadcast services	Minimum
VRTSilt	Veritas Cluster Server low-latency transport	Minimum
VRTSamf	Veritas Cluster Server Asynchronous Monitoring Framework	Minimum
VRTSvc	Veritas Cluster Server	Minimum
VRTSvcags	Veritas Cluster Server Bundled Agents	Minimum
VRTSvcxfen	Veritas I/O Fencing	Minimum
VRTSvcsea	Consolidated database and enterprise agent RPMs	Recommended

Table G-2 VCS installation RPMs (*continued*)

RPM	Contents	Configuration
VRTScps	Veritas Coordination Point Server The Coordination Point Server is an alternate mechanism for I/O fencing. It implements I/O fencing through a client/server architecture and can provide I/O fencing for multiple VCS clusters.	All

Veritas Storage Foundation obsolete and reorganized installation RPMs

[Table G-3](#) lists the RPMs that are obsolete or reorganized for Storage Foundation and High Availability.

Table G-3 Veritas Storage Foundation obsolete and reorganized RPMs

RPM	Description
Obsolete and reorganized for 6.0	
VRTSat	Obsolete
VRTSatClient	Obsolete
VRTSatServer	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

Table G-3 Veritas Storage Foundation obsolete and reorganized RPMs
(continued)

RPM	Description
VRTSmh	Included in VRTSsfmh
VRTSobc33	Obsolete
VRTSobweb	Obsolete
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product RPMs	
VRTSacclib	<p>Obsolete</p> <p>The following information is for installations, upgrades, and uninstalls using the script- or Web-based installer.</p> <ul style="list-style-type: none"> ■ For fresh installations VRTSacclib is not installed. ■ For upgrades, VRTSacclib is not uninstalled. ■ For uninstalls, VRTSacclib is not uninstalled.
VRTSalloc	Obsolete
VRTScmccc	Obsolete
VRTScmcm	Obsolete
VRTScmcs	Obsolete
VRTScscm	Obsolete
VRTScscw	Obsolete
VRTScsocw	Obsolete
VRTScssim	Obsolete
VRTScutil	Obsolete
VRTSd2gui-common	Included in VRTSdbed

Table G-3 Veritas Storage Foundation obsolete and reorganized RPMs
(continued)

RPM	Description
VRTSdb2ed-common	Included in VRTSdbed
VRTSdbcom-common	Included in VRTSdbed
VRTSdbed-common	Included in VRTSdbed
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSdsa	Obsolete
VRTSfsman	Included in the product's main RPM.
VRTSfsmnd	Included in the product's main RPM.
VRTSfspro	Included in VRTSsfmh
VRTSmapro-common	Included in VRTSsfmh
VRTSodm-common	Included in VRTSodm
VRTSodm-platform	Included in VRTSodm
VRTSorgui-common	Obsolete
VRTSvcldb	Included in VRTSvcsea
VRTSvcsmn	Included in VRTSvc
VRTSvcsor	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvc
VRTSvdid	Obsolete
VRTSvmman	Included in the product's main RPM.
VRTSvmpro	Included in VRTSsfmh
VRTSvrpro	Included in VRTSob
VRTSvrw	Obsolete
VRTSvxfs-common	Included in VRTSvxfs
VRTSvxfs-platform	Included in VRTSvxfs

Table G-3 Veritas Storage Foundation obsolete and reorganized RPMs
(continued)

RPM	Description
VRTSvxmsa	Obsolete
VRTSvxvm-common	Included in VRTSvxvm
VRTSvxvm-platform	Included in VRTSvxvm
Documentation	All Documentation RPMs 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](#)
- [Upgrading Veritas Storage Foundation for Databases \(SFDB\) tools from 5.0x to 6.0 \(2184482\)](#)

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 RPM 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.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 354.

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.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 (2184482)

```

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 Linux 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.

Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0x to 6.0 (2184482)

The `sfua_rept_migrate` command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.0.

When upgrading from SFHA version 5.0 to SFHA 6.0 the `S*vxdbs3` startup script is renamed to `NO_S*vxdbs3`. The `S*vxdbs3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the `S*vxdbs3` startup script and gives the error message:

```

/sbin/rc3.d/S*vxdbs3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File: is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.

```

Workaround

Before running `sfua_rept_migrate`, rename the startup script `NO_S*vxdbs3` to `S*vxdbs3`.

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 51.

- 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 `vxfcntlsthwdw` utility fails when SCSI TEST UNIT READY command fails

While running the `vxfcntlsthwdw` 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 Storage Foundation High Availability (client cluster) node.

See [“Troubleshooting issues related to the CP server service group”](#) on page 376.

See [“Checking the connectivity of CP server”](#) on page 376.

See [“Issues during fencing startup on Storage Foundation High Availability nodes set up for server-based fencing”](#) on page 377.

See [“Issues during online migration of coordination points”](#) on page 377.

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 Storage Foundation High Availability (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 Storage Foundation High Availability nodes set up for server-based fencing

Table I-1 Fencing startup issues on Storage Foundation High Availability (client cluster) nodes

Issue	Description and resolution
<p><code>cpsadm</code> command on the Storage Foundation High Availability gives connection error</p>	<p>If you receive a connection error message after issuing the <code>cpsadm</code> command on the Storage Foundation High Availability, perform the following actions:</p> <ul style="list-style-type: none"> ■ Ensure that the CP server is reachable from all the Storage Foundation High Availability nodes. ■ Check that the Storage Foundation High Availability nodes use the correct CP server virtual IP or virtual hostname and the correct port number. Check the <code>/etc/vxfenmode</code> file. ■ Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.
<p>Authorization failure</p>	<p>Authorization failure occurs when the CP server's nodes or users are not added in the CP server configuration. Therefore, fencing on the Storage Foundation High Availability (client cluster) node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points.</p> <p>To resolve this issue, add the CP server node and user in the CP server configuration and restart fencing.</p> <p>See “Preparing the CP servers manually for use by the Storage Foundation High Availability” on page 145.</p>
<p>Authentication failure</p>	<p>If you had configured secure communication between the CP server and the Storage Foundation High Availability (client cluster) nodes, authentication failure can occur due to the following causes:</p> <ul style="list-style-type: none"> ■ Symantec Product Authentication Services (AT) is not properly configured on the CP server and/or the Storage Foundation High Availability. ■ The CP server and the Storage Foundation High Availability 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 Storage Foundation High Availability 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 Storage Foundation High Availability 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 Storage Foundation High Availability nodes to the CP server(s).
- Cluster, nodes, or users for the Storage Foundation High Availability 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 Storage Foundation High Availability 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 `vxfencoordg`. 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 vxfencoordg.
 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 `vxfencoordg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

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/lfttab` explicitly depending on the subnet for each link. See [“Broadcast address in the /etc/lfttab file”](#) on page 385.

- Make sure that each NIC has an IP address that is configured before configuring LLT.
- Make sure the IP addresses in the `/etc/l1ttab` 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 387.
- Set the broadcast address correctly for direct-attached (non-routed) links. See [“Sample configuration: direct-attached links”](#) on page 388.
- For the links that cross an IP router, disable broadcast features and specify the IP address of each link manually in the `/etc/l1ttab` file. See [“Sample configuration: links crossing IP routers”](#) on page 390.

Broadcast address in the `/etc/l1ttab` file

The broadcast address is set explicitly for each link in the following example.

- Display the content of the `/etc/l1ttab` file on the first node galaxy:

```
galaxy # cat /etc/l1ttab

set-node galaxy
set-cluster 1
link link1 udp - udp 50000 - 192.168.9.1 192.168.9.255
link link2 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/l1ttab` file on the second node nebula:

```
nebula # cat /etc/l1ttab

set-node nebula
set-cluster 1
link link1 udp - udp 50000 - 192.168.9.2 192.168.9.255
link link2 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/l1ttab` file

Review the link command information in this section for the `/etc/l1ttab` file. See the following information for sample configurations:

- See [“Sample configuration: direct-attached links”](#) on page 388.
- See [“Sample configuration: links crossing IP routers”](#) on page 390.

[Table K-1](#) describes the fields of the link command that are shown in the `/etc/llttab` file examples. Note that some of the fields differ from the command for standard LLT links.

Table K-1 Field description for link command in `/etc/llttab`

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example udp. A place holder string. On other unix platforms like Solaris or HP, this entry points to a device file (for example, <code>/dev/udp</code>). Linux does not have devices for protocols. So this field is ignored.
<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 387.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IP address</i>	IP address of the link on the local node.
<i>bcast-address</i>	<ul style="list-style-type: none"> ■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the `/etc/llttab` file

The `set-addr` command in the `/etc/llttab` file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See [“Sample configuration: links crossing IP routers”](#) on page 390.

[Table K-2](#) describes the fields of the `set-addr` command.

Table K-2 Field description for set-addr command in /etc/llttab

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.
<i>link tag-name</i>	The string that LLT uses to identify the link; for example link1, link2,....
<i>address</i>	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address      State
udp        0      0 *:32768                *:*
udp        0      0 *:956                  *:*
udp        0      0 *:tftp                 *:*
udp        0      0 *:sunrpc               *:*
udp        0      0 *:ipp                  *:*
```

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 udp - udp 50000 - 192.168.30.1 192.168.30.255  
link link2 udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: direct-attached links

[Figure K-1](#) depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure K-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 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 udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

Figure K-2 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure K-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1

link link1 udp - udp 50000 - 192.1.3.1 -
link link2 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/lltab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 udp - udp 50000 - 192.1.1.1 -
link link2 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

Veritas Storage Foundation 6.0 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

Manually configuring LLT over UDP using IPv6

The following checklist is to configure LLT over UDP:


- For UDP6, the multicast address is set to "-".
- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.

- Make sure the IPv6 addresses in the `/etc/llttab` files are consistent with the IPv6 addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
- 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 394.

Sample configuration: direct-attached links

Figure L-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure L-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
```

```
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 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 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

Sample configuration: links crossing IP routers

Figure L-2 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure L-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```

set-node Node1
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 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/lltab` file on Node 0 resembles:

```

set-node Node0
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 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 ApplicationHA 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 RPMs 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 ApplicationHA is already present

If you plan to install or upgrade Storage Foundation on systems where ApplicationHA has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not consider VCS as an installed product even though it uses the bundled VRTSvcS RPM.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not allow the installation or upgrade for products that use VCS. The following products cannot be installed or upgraded: VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SFCFSRAC, or SFSYBASECE.

- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer allows the installation or upgrade of VM, FS, SF, or DMP.
- When you uninstall Storage Foundation products where ApplicationHA is present, the installer does not uninstall VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspbx and VRTSsisco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSspbx, VRTSsisco, and VRTSat.

Index

A

- about
 - global clusters 28
- adding
 - users 110
- agents
 - disabling 279
- applications, stopping 174
- attributes
 - UseFence 137

B

- backup boot disk group 212
 - rejoining 212

C

- cables
 - cross-over Ethernet 243
- cluster
 - removing a node from 265
 - verifying operation 236
- command failures 374
- commands
 - hastatus 236
 - hasys 237
 - lltconfig 335
 - lltstat 234
 - vxdisksetup (initializing disks) 126
 - vxlicinst 118–119
 - vxlicrep 118
- configuration
 - restoring the original 218
- configuring
 - rsh 47
 - ssh 47
- configuring SFHA
 - script-based installer 99
- configuring VCS
 - adding users 110
 - event notification 111, 113

- configuring VCS (*continued*)
 - global clusters 115
 - secure mode 107
 - starting 101
- coordinator disks
 - DMP devices 26
 - for I/O fencing 26
 - setting up 135

D

- data disks
 - for I/O fencing 26
- disabling the agents 279
- disks
 - adding and initializing 126
 - coordinator 135
 - testing with vxfcntlsthdw 126
 - verifying node access 128

E

- Ethernet controllers 243

F

- freezing service groups 174

G

- GAB
 - description 23
- gabtab file
 - verifying after installation 335
- global clusters 28
 - configuration 115

H

- hastatus -summary command 236
- hasys -display command 237
- hubs
 - independent 243

I

- I/O fencing
 - checking disks 126
 - setting up 134
 - shared storage 126
- I/O fencing requirements
 - non-SCSI-3 39
- Installing
 - SFHA with the Web-based installer 63
- installing
 - post 116
 - yum 69

L

- license keys
 - adding with vxlicinst 118
 - replacing demo key 119
- licenses
 - information about 118
- licensing 35
 - add-on 34
 - CDS 35
 - full 34
- links
 - private network 335
- LLT
 - description 23
 - interconnects 24
 - verifying 234
- lltconfig command 335
- llthosts file
 - verifying after installation 335
- lltstat command 234
- llttab file
 - verifying after installation 335
- log files 375

M

- main.cf file
 - contents after installation 341
- main.cf files 347
- manual pages
 - potential problems 373
 - troubleshooting 373
- media speed 24
 - optimizing 23
- mounting
 - software disc 52

N

- nodes
 - adding application nodes
 - configuring GAB 249, 255
 - configuring LLT 249, 255
 - configuring VXFEN 249, 255
 - starting Volume Manager 248, 255
 - adding new nodes
 - configuring GAB 250
 - configuring LLT 250
 - configuring VXFEN 250
- non-SCSI-3 fencing
 - manual configuration 151
 - setting up 151
- non-SCSI-3 I/O fencing requirements 39
- non-SCSI3 fencing
 - setting up 139
 - using installsfha 139

O

- optimizing
 - media speed 23
- original configuration
 - restoring the 218

P

- PATH variable
 - VCS commands 234
- persistent reservations
 - SCSI-3 48
- phased 193
- phased upgrade 193, 195
 - example 194
- planning to upgrade VVR 169
- preinstallation 169
- preparing to upgrade VVR 174
- problems
 - accessing manual pages 373
 - executing file system commands 374

R

- rejoining
 - backup boot disk group 212
- removing
 - the Replicated Data Set 280
- removing a system from a cluster 265

- Replicated Data Set
 - removing the 280
- restoring the original configuration 218
- rolling upgrade 188
 - versions 186
- rsh 102
 - configuration 47

S

- script-based installer
 - SFHA configuration overview 99
- SCSI-3
 - persistent reservations 48
- SCSI-3 persistent reservations
 - verifying 134
- service groups
 - freezing 174
 - unfreezing 217
- SFHA
 - about 22
 - configuring 99
 - coordinator disks 135
- SFHA installation
 - verifying
 - cluster operations 234
 - GAB operations 234
 - LLT operations 234
- SMTP email notification 111
- SNMP trap notification 113
- ssh 102
 - configuration 47
- starting configuration
 - installvcs program 102
 - Veritas product installer 101
- stopping
 - applications 174
- Symantec Product Authentication Service 107
- system state attribute value 236

T

- troubleshooting
 - accessing manual pages 373
 - executing file system commands 374
- tunables file
 - about setting parameters 315
 - parameter definitions 320
 - preparing 319
 - setting for configuration 316

- tunables file (*continued*)
 - setting for installation 316
 - setting for upgrade 316
 - setting parameters 319
 - setting with no other operations 317
 - setting with un-integrated response file 318

U

- unfreezing service groups 217
- unsuccessful upgrade 212
- upgrade
 - phased 193, 195
- upgrading
 - phased 193
- upgrading VVR
 - from 4.1 170
 - planning 169
 - preparing 174

V

- VCS
 - command directory path variable 234
 - configuration files
 - main.cf 339
 - verifying
 - NIC configuration 117
- Veritas Operations Manager 22
- vradmin
 - delpri 281
 - stoprep 280
- VVR 4.1
 - planning an upgrade from 170
- vvr_upgrade_finish script 219
- vxdisksetup command 126
- vxlicinst command 118
- vxlicrep command 118
- vxplex
 - used to remove mirrors of root disk volumes 277

W

- Web-based installer 63

Y

- yum
 - installing 69