

Veritas Storage Foundation™ and High Availability Installation Guide

Solaris

5.1 Service Pack 1

Veritas Storage Foundation™ Installation Guide

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

Product version: 5.1 SP1

Document version: 5.1SP1.2

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

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

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

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

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

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

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

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

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

docs@symantec.com

About Symantec Connect

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

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

Contents

Technical Support	4
Section 1 Installation overview and planning	21
Chapter 1 Introducing Storage Foundation and High Availability Solutions	23
About Veritas products	23
About Veritas Storage Foundation and Storage Foundation High Availability	23
About Veritas Cluster Server	24
About Veritas high availability agents	24
About Veritas Volume Replicator	24
About Veritas Storage Foundation Cluster File System	25
About Veritas Storage Foundation for Oracle® RAC by Symantec	25
About Veritas graphical user interfaces	25
Veritas Operations Manager	25
About Storage Foundation and High Availability features	26
About Symantec Product Authentication Service	26
About LLT and GAB	26
About configuring SFHA clusters for data integrity	27
About I/O fencing components	28
About global clusters	30
Chapter 2 Planning to install the Storage Foundation and High Availability products	31
About planning for SFHA installation	31
About installation and configuration methods	32
Downloading the Storage Foundation and High Availability software	33
Chapter 3 System requirements	35
Release notes	35
Hardware compatibility list (HCL)	36

	Veritas File System requirements	36
	Cluster environment requirements for Sun Clusters	37
	Supported Solaris operating systems	37
	Disk space requirements	38
	Database requirements	38
	I/O fencing requirements	38
	Coordinator disk requirements for I/O fencing	39
	CP server requirements	39
	Number of nodes supported	42
Chapter 4	Licensing Veritas products	43
	About Veritas product licensing	43
	Setting or changing the product level for keyless licensing	44
	Installing Veritas product license keys	46
Section 2	Installation of Storage Foundation and High Availability products	47
Chapter 5	Preparing to install	49
	Installation preparation overview	49
	About configuring ssh or rsh using the Veritas installer	50
	Setting up shared storage	51
	Setting up shared storage: SCSI disks	51
	Setting up shared storage: Fibre Channel	54
	Creating the /opt directory	55
	Setting environment variables	55
	Mounting the product disc	55
	Assessing system preparedness	56
	Symantec Operations Readiness Tools	56
	Prechecking your systems using the Veritas installer	57
Chapter 6	Installing Storage Foundation and High Availability Solutions using the script-based installer	59
	About the Veritas installer	59
	Installing Storage Foundation using the installer	60
	Installing Storage Foundation and High Availability Solutions using the installer	62
	Installing language packages	65

Chapter 7	Installing Storage Foundation and High Availability Solutions using the Web-based installer	67
	About the Web-based installer	67
	Features not supported with Web-based installer	68
	Before using the Veritas Web-based installer	68
	Starting the Veritas Web-based installer	69
	Obtaining a security exception on Mozilla Firefox	69
	Performing a pre-installation check with the Veritas Web-based installer	70
	Installing SFHA with the Web-based installer	70
Chapter 8	Installing Storage Foundation and High Availability products using other methods	73
	Installing with JumpStart	73
	Overview of JumpStart installation tasks	73
	Generating the finish scripts	74
	Preparing installation resources	75
	Adding language pack information to the finish file	76
	Installing SFHA using the pkgadd command	77
Section 3	Configuration of Storage Foundation and High Availability products	81
Chapter 9	Preparing to configure Storage Foundation and High Availability	83
	Preparing to configure the clusters in secure mode	83
	Installing the root broker for the security infrastructure	87
	Creating authentication broker accounts on root broker system	88
	Creating encrypted files for the security infrastructure	89
	Preparing the installation system for the security infrastructure	91
	About planning to configure I/O fencing	92
	Typical SF HA cluster configuration with server-based I/O fencing	94
	Recommended CP server configurations	95
	Setting up the CP server	98
	Planning your CP server setup	98
	Installing the CP server using the installer	99
	Configuring the CP server cluster in secure mode	100

	Setting up shared storage for the CP server database	101
	Configuring the CP server using the configuration utility	102
	Configuring the CP server manually	110
	Verifying the CP server configuration	112
Chapter 10	Configuring Storage Foundation	113
	Configuring Storage Foundation using the installer	113
	Configuring Storage Foundation manually	113
	Configuring Veritas Volume Manager	113
	Configuring Veritas File System	122
	Configuring the SFDB repository database after installation	123
	Veritas Volume Replicator and Volume Manager setup after installation	124
Chapter 11	Configuring Storage Foundation and High Availability	125
	Configuring Storage Foundation and High Availability Solutions	125
	Required information for configuring Storage Foundation and High Availability Solutions	125
	Configuring Storage Foundation High Availability using the installer	126
	Configuring SFHA using the Web-based installer	145
	Configuring and starting Veritas Enterprise Administrator	151
Chapter 12	Configuring Storage Foundation High Availability for data integrity	153
	Setting up disk-based I/O fencing using installsfha	153
	Initializing disks as VxVM disks	153
	Checking shared disks for I/O fencing	154
	Configuring disk-based I/O fencing using installsfha	158
	Setting up disk-based I/O fencing manually	161
	Removing permissions for communication	161
	Identifying disks to use as coordinator disks	161
	Setting up coordinator disk groups	162
	Creating I/O fencing configuration files	163
	Modifying VCS configuration to use I/O fencing	164
	Verifying I/O fencing configuration	166
	Setting up server-based I/O fencing using installsfha	166
	Verifying the security configuration on the SF HA cluster to use CP server coordination point	167
	Configuring server-based I/O fencing using the installsfha	169

	Setting up server-based I/O fencing manually	178
	Preparing the CP servers manually for use by the SF HA cluster	178
	Configuring Coordination Point agent to monitor coordination points	182
	Verifying server-based I/O fencing configuration	184
	Enabling or disabling the preferred fencing policy	185
Section 4	Upgrading Storage Foundation and High Availability products	189
Chapter 13	Preparing to upgrade	191
	About upgrading	191
	About the different ways that you can upgrade	192
	Supported upgrade paths	193
	About using the installer to upgrade when the root disk is encapsulated	194
	Tasks for upgrading the Storage Foundation for Databases (SFDB) tools	195
	Preparing to upgrade	195
	Getting ready for the upgrade	195
	Creating backups	197
	Pre-upgrade tasks for migrating the SFDB repository database	198
	Determining if the root disk is encapsulated	199
	Preupgrade planning for Veritas Volume Replicator	199
	Preparing to upgrade VVR when VCS agents are configured	202
	Verifying that the file systems are clean	205
	Upgrading the array support	206
Chapter 14	Upgrading Storage Foundation or Storage Foundation and High Availability	209
	Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required	209
	Upgrading Veritas Storage Foundation to 5.1 SP1 using the product installer or manual steps	211
	Upgrading Veritas Storage Foundation with the product installer	212
	Upgrading Veritas Storage Foundation using manual steps	214
	Upgrading Veritas Storage Foundation to 5.1 SP1 using upgrade scripts (OS upgrade)	216

	Upgrading SFHA with the Veritas Web-based installer	220
	Upgrading the Solaris operating system	221
	Upgrading Veritas Volume Replicator	224
	Upgrading VVR without disrupting replication	224
	Upgrading language packages	225
Chapter 15	Performing a rolling upgrade	227
	Performing a rolling upgrade using the installer	227
	About rolling upgrades	227
	Prerequisites for a rolling upgrade	227
	Performing a rolling upgrade on kernel packages: phase 1	228
	Performing a rolling upgrade on non-kernel packages: phase 2	228
	Performing a rolling upgrade of SFHA using the Web-based installer	229
Chapter 16	Performing a phased upgrade	231
	About phased upgrade	231
	Prerequisites for a phased upgrade	231
	Planning for a phased upgrade	231
	Phased upgrade limitations	232
	Phased upgrade example	232
	Phased upgrade example overview	233
	Performing a phased upgrade	233
	Moving the service groups to the second subcluster	234
	Upgrading the operating system on the first subcluster	238
	Upgrading the first subcluster	238
	Preparing the second subcluster	240
	Activating the first subcluster	245
	Upgrading the operating system on the second subcluster	246
	Upgrading the second subcluster	247
	Finishing the phased upgrade	248
Chapter 17	Upgrading with Live Upgrade	253
	About Live Upgrade	253
	About Live Upgrade in a Veritas Volume Replicator (VVR) environment	254
	Supported upgrade paths for Live Upgrade	255
	Performing Live Upgrade in a Solaris zone environment	256
	Before you upgrade SFHA using Solaris Live Upgrade	257
	Upgrading SFHA and Solaris using Live Upgrade	261

Creating a new boot environment on the alternate boot disk	262
Upgrading SFHA using the installer for a Live Upgrade	263
Upgrading SF manually	265
Completing the Live Upgrade	268
Verifying Live Upgrade of SFHA	269
Upgrading Solaris using Live Upgrade	270
Removing and reinstalling SFHA using the installer	271
Upgrading SFHA using Live Upgrade	272
Administering boot environments	273
Reverting to the primary boot environment	273
Switching the boot environment for Solaris SPARC	273
Switching the boot environment for Solaris x86-64	275
 Chapter 18	
Performing post-upgrade tasks	279
Optional configuration steps	279
Post upgrade tasks for migrating the SFDB repository database	280
Migrating from a 5.0 repository database to 5.1 SP1	280
Migrating from a 4.x repository database to 5.1 SP1	284
Recovering VVR if automatic upgrade fails	287
Post-upgrade tasks when VCS agents for VVR are configured	287
Unfreezing the service groups	288
Restoring the original configuration when VCS agents are configured	288
Upgrading disk layout versions	290
Upgrading VxVM disk group versions	291
Updating variables	291
Setting the default disk group	291
Upgrading the Array Support Library	292
Adding JBOD support for storage arrays for which there is not an ASL available	292
Unsuppressing DMP for EMC PowerPath disks	293
Converting from QuickLog to Multi-Volume support	301
About enabling LDAP authentication for clusters that run in secure mode	303
Enabling LDAP authentication for clusters that run in secure mode	305
Verifying the Veritas Storage Foundation upgrade	311

Section 5	Verification of the installation or the upgrade	313
Chapter 19	Verifying the installation	315
	About using the postcheck option	315
	Performing a postcheck on a node	316
	Verifying that the products were installed	316
	Installation log files	316
	Using the installation log file	317
	Using the summary file	317
	Starting and stopping processes for the Veritas products	317
	Checking Veritas Volume Manager processes	318
	Checking Veritas File System installation	318
	Verifying Veritas File System kernel installation	318
	Verifying command installation	318
	Verifying the LLT, GAB, and VCS configuration files	319
	Verifying LLT, GAB, and cluster operation	319
	Verifying LLT	320
	Verifying the cluster	322
	Verifying the cluster nodes	323
Section 6	Adding and removing nodes	327
Chapter 20	Adding a node to a cluster	329
	About adding a node to a cluster	329
	Before adding a node to a cluster	330
	Meeting hardware and software requirements	330
	Setting up the hardware	330
	Preparing to add a node to a cluster	332
	Adding a node to a cluster	332
	Adding a node to a cluster using the SFHA installer	332
	Adding a node using the Web-based installer	336
	Adding the node to a cluster manually	337
	Configuring server-based fencing on the new node	345
	After adding the new node	347
	Updating the Storage Foundation for Databases (SFDB) repository after adding a node	347

Chapter 21	Removing a node from a cluster	349
	Removing a node from a cluster	349
	Verifying the status of nodes and service groups	350
	Deleting the departing node from SFHA configuration	351
	Modifying configuration files on each remaining node	354
	Removing the node configuration from the CP server	354
	Removing security credentials from the leaving node	355
	Unloading LLT and GAB and removing VCS on the departing node	355
Section 7	Uninstallation of Storage Foundation and High Availability products	359
Chapter 22	Uninstalling Storage Foundation and High Availability products	361
	About removing Veritas Storage Foundation	361
	Preparing to uninstall	362
	Preparing to remove Veritas Volume Manager	362
	Preparing to remove Veritas File System	370
	Disabling VCS agents for VVR the agents on a system	371
	Removing the Replicated Data Set	372
	Uninstalling SFHA packages using the script-based installer	374
	Uninstalling SFHA with the Veritas Web-based installer	375
	Uninstalling Storage Foundation using the pkgrm command	376
	Uninstalling the language packages using the pkgrm command	377
	Removing the CP server configuration using the removal script	378
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product	381
Section 8	Installation reference	383
Appendix A	Installation scripts	385
	About installation scripts	385
	Installation script options	386

Appendix B	Response files	393
	About response files	393
	Installing Storage Foundation or Storage Foundation and High Availability using response files	394
	Configuring SFHA using response files	395
	Upgrading Storage Foundation or Storage Foundation and High Availability using response files	395
	Uninstalling Storage Foundation or Storage Foundation and High Availability using response files	396
	Syntax in the response file	396
	Response file variables to install, upgrade, or uninstall Storage Foundation or Storage Foundation and High Availability	397
	Response file variables to configure SFHA	400
	Sample response file for SFHA configuration	409
	Sample response file for SFHA install	409
	Sample response file for SF upgrade	410
	Sample response file for SFHA upgrade	410
Appendix C	Configuring I/O fencing using a response file	411
	Configuring I/O fencing using response files	411
	Response file variables to configure disk-based I/O fencing	412
	Sample response file for configuring disk-based I/O fencing	413
	Response file variables to configure server-based I/O fencing	414
	Sample response file for configuring server-based I/O fencing	416
Appendix D	Configuration files	419
	About the LLT and GAB configuration files	419
	About the AMF configuration files	422
	About the VCS configuration files	423
	Sample main.cf file for VCS clusters	425
	Sample main.cf file for global clusters	426
	About I/O fencing configuration files	429
	Sample configuration files for CP server	431
	Sample main.cf file for CP server hosted on a single node that runs VCS	432
	Sample main.cf file for CP server hosted on a two-node SFHA cluster	434

Appendix E	Configuring the secure shell or the remote shell for communications	439
	About configuring secure shell or remote shell communication modes	
	before installing products	439
	Configuring and enabling ssh	440
	Restarting the ssh session	444
	Enabling and disabling rsh for Solaris	445
Appendix F	Storage Foundation and High Availability components	447
	Storage Foundation and High Availability installation packages	447
	Veritas Cluster Server installation packages	449
	Chinese language packages	450
	Japanese language packages	451
	Veritas Storage Foundation obsolete and reorganized installation packages	451
Appendix G	Troubleshooting installation issues	455
	Restarting the installer after a failed connection	455
	What to do if you see a licensing reminder	455
	Troubleshooting information	456
	Incorrect permissions for root on remote system	456
	Inaccessible system	458
	Upgrading Veritas Storage Foundation for Databases (SFDB) tools	
	from 5.0.x to 5.1SP1 (2184482)	458
	Workaround	458
Appendix H	Troubleshooting cluster installation	459
	Unmount failures	459
	Command failures	459
	Installer cannot create UUID for the cluster	460
	The vxfsntsthdw utility fails when SCSI TEST UNIT READY command fails	460
	Troubleshooting server-based I/O fencing	461
	Troubleshooting issues related to the CP server service group	461
	Checking the connectivity of CP server	462
	Troubleshooting server-based fencing on the SF HA cluster nodes	462

	Issues during fencing startup on SF HA cluster nodes set up for server-based fencing	463
	Issues during online migration of coordination points	465
	Troubleshooting server-based I/O fencing in mixed mode	466
	Checking keys on coordination points when vxfen_mechanism value is set to cps	470
	After upgrading from 5.0.x and before migrating SFDB	471
Appendix I	Sample SF HA cluster setup diagrams for CP server-based I/O fencing	473
	Configuration diagrams for setting up server-based I/O fencing	473
	Two unique client clusters served by 3 CP servers	473
	Client cluster served by highly available CPS and 2 SCSI-3 disks	474
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	476
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	478
Appendix J	Reconciling major/minor numbers for NFS shared disks	481
	Reconciling major/minor numbers for NFS shared disks	481
	Checking major and minor numbers for disk partitions	482
	Checking the major and minor number for VxVM volumes	485
Appendix K	Configuring LLT over UDP using IPv4	489
	Using the UDP layer for LLT	489
	When to use LLT over UDP	489
	Manually configuring LLT over UDP using IPv4	489
	Broadcast address in the /etc/llttab file	490
	The link command in the /etc/llttab file	491
	The set-addr command in the /etc/llttab file	491
	Selecting UDP ports	492
	Configuring the netmask for LLT	493
	Configuring the broadcast address for LLT	493
	Sample configuration: direct-attached links	494
	Sample configuration: links crossing IP routers	496

Appendix L	Configuring LLT over UDP using IPv6	499
	Using the UDP layer of IPv6 for LLT	499
	When to use LLT over UDP	499
	Manually configuring LLT over UDP using IPv6	499
	Sample configuration: direct-attached links	500
	Sample configuration: links crossing IP routers	502
Index		505

Installation overview and planning

- [Chapter 1. Introducing Storage Foundation and High Availability Solutions](#)
- [Chapter 2. Planning to install the Storage Foundation and High Availability products](#)
- [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 Veritas Storage Foundation and Storage Foundation 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, but with 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, 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 Oracle® RAC by Symantec

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.

Veritas Operations Manager

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

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

If you wish to manage a single cluster using Cluster Manager (Java Console), a version is available for download from http://go.symantec.com/vcsm_download. Veritas Cluster Server Management Console is no longer supported.

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. Veritas Operations Manager is not available on the Storage Foundation and High Availability Solutions release. 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.

About Storage Foundation and High Availability features

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

About Symantec Product Authentication Service

The Symantec Product Authentication Service protects communication channels among Symantec application clients and services through message integrity and confidentiality services.

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 the global message order that is required to maintain a synchronized state among the nodes. It monitors disk communications such as the VCS heartbeat utility.

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 a corrective action would cause a split-brain situation.

Some 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 92.

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 28.
- Coordination points—Act as a global lock during membership changes
See [“About coordination points”](#) on page 28.

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

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 be disks, 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 Volume Manager Administrator's Guide*.

■ **Coordination point servers**

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

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

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

About preferred fencing

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

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

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

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

See [“Enabling or disabling the preferred fencing policy”](#) on page 185.

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 the Storage Foundation and High Availability products

This chapter includes the following topics:

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

About planning for SFHA installation

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

<http://www.symantec.com/business/support/overview.jsp?pid=15107>

Document version: 5.1SP1.2.

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

Follow the preinstallation instructions if you are installing one of the Storage Foundation and High Availability products by Symantec.

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

- Veritas Storage Foundation Basic
- Veritas Storage Foundation (Standard and Enterprise Editions)
- Veritas Storage Foundation High Availability (HA) (Standard and Enterprise Editions)

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

About installation and configuration methods

You can install and configure SFHA with Veritas installation programs or with 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. Otherwise, installing with the installation script is identical to specifying SFHA from the installer menu.
- The Web-based Veritas installer
The installer provides an interface to manage the installation from a remote site using a standard Web browser.
In this release, there are some limitations in the Web-based installer.
See [“About the Web-based installer”](#) on page 67.
- Silent installation with response files
You can use any of the above options to generate a response file. You can then customize the response file for another system. Run the product installation script with the response file to install silently on one or more other systems.
See [“About response files”](#) on page 393.
- JumpStart
You can use the Veritas product installer or the product-specific installation script to generate a Jumpstart script file. Use the generated script to install Veritas packages from your JumpStart 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, you can use the following link. For other downloads, contact your Veritas representative for more information.

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

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.

See “[About installation scripts](#)” on page 385.

To download the software

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

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

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

See “[Disk space requirements](#)” on page 38.

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

```
# df -b filesystem
```

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

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

System requirements

This chapter includes the following topics:

- [Release notes](#)
- [Hardware compatibility list \(HCL\)](#)
- [Veritas File System requirements](#)
- [Cluster environment requirements for Sun Clusters](#)
- [Supported Solaris operating systems](#)
- [Disk space requirements](#)
- [Database requirements](#)
- [I/O fencing requirements](#)
- [Number of nodes supported](#)

Release notes

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

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

<http://www.symantec.com/business/support/overview.jsp?pid=15107>

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/TECH74012>

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

Veritas File System requirements

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

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

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

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

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

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

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

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

Cluster environment requirements for Sun Clusters

Use these steps if the configuration is with Sun cluster, which is a set of hosts that share a set of disks.

To configure a cluster

- 1 Obtain a license for the optional VxVM cluster feature for a Sun Cluster from your Oracle Customer Support channel.
- 2 If you plan to encapsulate the root disk group, decide where you want to place it for each node in the cluster. The root disk group, usually aliased as `bootdg`, contains the volumes that are used to boot the system. VxVM sets `bootdg` to the appropriate disk group if it takes control of the root disk. Otherwise `bootdg` is set to `nodg`. To check the name of the disk group, enter the command:

```
# vxvg bootdg
```

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

Supported Solaris operating systems

This release of the Veritas products is supported on the following Solaris operating systems:

- Solaris 9 (32-bit and 64-bit, SPARC) with Update 7, 8, and 9
Symantec VirtualStore is only supported on Solaris 9 (SPARC Platform 64-bit).

Note: In the next major release, Veritas products will not support Solaris 9.

- Solaris 10 (64-bit, SPARC or x86_64) with Update 6, 7, 8, and 9
Solaris 10 (SPARC and x86_64) with Update 9 requires VRTSVxvm patch 142629-08 (SPARC) or 142630-08 (x86_64)
Symantec VirtualStore is only supported on Solaris 10 (SPARC or X86 Platform 64-bit).

For the most up-to-date list of operating system patches, refer to the Release Notes for your product.

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

<http://entsupport.symantec.com/docs/334829>

For information about the use of this product in a VMware Environment on Solaris x64, refer to <http://entsupport.symantec.com/docs/289033>

Disk space requirements

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

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

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

Database requirements

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

<http://entsupport.symantec.com/docs/331625>

Note: SFHA does not support running SFDB tools with DB2 and Sybase.

I/O fencing requirements

Depending on whether you plan to configure disk-based fencing or server-based fencing, make sure that you meet the requirements for coordination points:

- Coordinator disks
See ["Coordinator disk requirements for I/O fencing"](#) on page 39.
- CP servers

See “[CP server 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 5.1SP1 clusters (application clusters) support CP servers which are hosted on the following VCS and SFHA versions:

- VCS 5.1 or 5.1SP1 single-node cluster
CP server requires LLT and GAB to be configured on the single-node VCS cluster that hosts CP server. This requirement also applies to any single-node application cluster that uses server-based fencing.
- SFHA 5.1 or 5.1SP1 cluster

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

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

See the *Veritas Cluster Server Installation Guide*.

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

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

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

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

Table 3-1 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: <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
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the CP servers.
RAM	Each CP server requires at least 512 MB.
CP server to client node physical link	A secure TCP/IP connection is required to connect the CP servers to the SFHA clusters (application clusters).

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

Table 3-2 CP server supported operating systems and versions

CP server	Operating system and version
CP server hosted on a VCS single-node cluster or on an SFHA cluster	<p>CP server supports any of the following operating systems:</p> <ul style="list-style-type: none"> ■ AIX 5.3 and 6.1 ■ HP-UX 11i v3 ■ Linux: <ul style="list-style-type: none"> ■ RHEL 5 ■ SLES 10 ■ SLES 11 ■ Solaris 9 and 10 <p>Review other details such as supported operating system levels and architecture for the supported operating systems.</p> <p>See “Supported Solaris operating systems” on page 37.</p> <p>For other supported operating systems, see the <i>Veritas Cluster Server Installation Guide</i> or the <i>Veritas Storage Foundation High Availability Installation Guide</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.
- 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 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 communications between the SF HA cluster and CP server, consider the following requirements and suggestions:

- In a secure communication environment, all CP servers that are used by the application cluster must be configured with security enabled. A configuration where the application cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.
- The CP server and application clusters should also use the same root broker. If the same root broker is not being used, then trust can be established between the cluster nodes and CP server for the secure communication. Trust can be established by the installer when configuring fencing.
- For non-secure communication between CP server and 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*.

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

For more updates on this support, see the Late-Breaking News TechNote on the Symantec Technical Support website:

<http://entsupport.symantec.com/docs/334829>

Licensing Veritas products

This chapter includes the following topics:

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

About Veritas product licensing

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

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

www.symantec.com/techsupp/

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

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

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

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

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

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.
See “[Setting or changing the product level for keyless licensing](#)” on page 44.
See the `vxkeyless (1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.
See “[Installing Veritas product license keys](#)” on page 46.
See the `vxlicinst (1m)` manual page.

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

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

Setting or changing the product level for keyless licensing

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

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

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

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

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

To set or change the product level

- 1 View the current setting for the product level.

```
# vxkeyless -v display
```

- 2 View the possible settings for the product level.

```
# vxkeyless displayall
```

- 3 Set the desired product level.

```
# vxkeyless -q set prod_levels
```

where *prod_levels* is a comma-separated list of keywords.

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas product license keys

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

<code>vxlicinst</code>	Installs a license key for a Symantec product
<code>vxlicrep</code>	Displays currently installed licenses
<code>vxlictest</code>	Identifies whether a Symantec product feature is licensed on the system

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

To install a new license

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

```
# cd /opt/VRTS/bin  
  
# ./vxlicinst -k xxxx-xxxx-xxxx-xxxx-xxxx-xxx
```

Installation of Storage Foundation and High Availability products

- [Chapter 5. Preparing to install](#)
- [Chapter 6. Installing Storage Foundation and High Availability Solutions using the script-based installer](#)
- [Chapter 7. Installing Storage Foundation and High Availability Solutions using the Web-based installer](#)
- [Chapter 8. Installing Storage Foundation and High Availability products using other methods](#)

Preparing to install

This chapter includes the following topics:

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

Installation preparation overview

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

Table 5-1 Installation overview

Installation task	Section
Obtain product licenses.	See “About Veritas product licensing” on page 43.
Download the software, or insert the product DVD.	See “Downloading the Storage Foundation and High Availability software” on page 33. See “Mounting the product disc” on page 55.
Set environment variables.	See “Setting environment variables” on page 55.
Create the /opt directory, if it does not exist.	See “Creating the /opt directory” on page 55.

Table 5-1 Installation overview (*continued*)

Installation task	Section
Configure the secure shell (ssh) on all nodes.	See “About configuring secure shell or remote shell communication modes before installing products” on page 439.
Verify that hardware, software, and operating system requirements are met.	See “Supported Solaris operating systems” on page 37. See “Release notes” on page 35.
Check that sufficient disk space is available.	See “Disk space requirements” on page 38.
Use the installer to install the products.	See “About the Veritas installer” on page 59.

About configuring ssh or rsh using the Veritas installer

The installer can configure 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 386.

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 the root broker is outside of the cluster that you plan to configure.
- When you add new nodes to an existing cluster.
- When the nodes are in a sub-cluster 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 439.

Setting up shared storage

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

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

See [“About planning to configure I/O fencing”](#) on page 92.

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

Setting up shared storage: SCSI disks

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

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

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

To set up shared storage

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

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

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

4 Find the paths to the host adapters:

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

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

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

```
{0} ok nvedit
```

As you edit the script, note the following points:

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

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

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

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

```
{0} ok nvstore
```

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

```
{0} ok printenv nvrामrc
```

You can re-edit the file to make corrections:

```
{0} ok nvedit
```

Or, discard the changes if necessary by entering:

```
{0} ok nvquit
```

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

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

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

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

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

Permit the system to continue booting.

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

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

Permit the system to continue booting.

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up shared storage

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

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

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

```
ok boot -r
```

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

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

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

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

Creating the /opt directory

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

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

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.

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

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

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

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

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

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

Mounting the product disc

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

To mount the product disc

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

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

Where c0t6d0s2 is the default address for the disc drive.

Assessing system preparedness

Symantec provides the following tools for assessing your system, to ensure that the system meets the requirements for installing Storage Foundation 5.1 SP1.

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 56.

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 Storage Foundation 5.1 SP1.

See [“Prechecking your systems using the Veritas installer”](#) on page 57.

Symantec Operations Readiness Tools

Symantec™ Operations Readiness Tools (SORT) is a set of Web-based tools that supports Symantec enterprise products. SORT increases operational efficiency and helps improve application availability.

Among its broad set of features, SORT provides patches, patch notifications, and documentation for Symantec enterprise products.

To access SORT, go to:

<http://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 Solutions using the script-based installer

This chapter includes the following topics:

- [About the Veritas installer](#)
- [Installing Storage Foundation using the installer](#)
- [Installing Storage Foundation and High Availability Solutions using the installer](#)
- [Installing language packages](#)

About the Veritas installer

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

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

See [“About installation scripts”](#) on page 385.

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.

Additional options are available for the installer.

See [“Installation script options”](#) on page 386.

Installing Storage Foundation using the installer

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

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

To install Storage Foundation

- 1 Set up the systems so that commands between systems execute without prompting for passwords or confirmations.
[See “About configuring secure shell or remote shell communication modes before installing products”](#) on page 439.
- 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 55.

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

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

- 4 From this directory, type the following command to install on the local system:

```
# ./installer
```

Use this command to install on remote systems if secure shell or remote shell communication modes are configured.

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

- 6 When the list of available products is displayed, select Storage Foundation, enter the corresponding number, and press Return.
- 7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).

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

- 8 You are prompted to enter the system names (in the following example, "host1") where you want to install the software. Enter the system name or names and then press Return.

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

Where *platform* indicates the operating system.

- 9 After the system checks complete, the installer displays a list of the packages to be installed. Press Return to continue with the installation.
- 10 The installer can configure remote shell or secure shell communications for you among systems, however each system needs to have remote shell or secure shell 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.
- 11 The installer may prompt for previous Veritas Volume Manager configurations.
- 12 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 must still have a valid license to install and use Veritas products. Keyless licensing requires that you manage the systems with a Management Server.

Note: If you are install Storage Foundation Basic, choose the first option to install the required license keys.

See [“About Veritas product licensing”](#) on page 43.

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

- 1) SF Standard
- 2) SF Enterprise
- b) Back to previous menu

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

14 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
```

15 The installation and configuration complete automatically. The product processes are started.

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

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

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

Installing Storage Foundation and High Availability Solutions using the installer

The following sample procedure is based on the installation of a Storage Foundation Enterprise High Availability (SF/HA) cluster with two nodes: "host1" and "host2."

To install Storage Foundation and High Availability products

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

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 439.

- 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 55.

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

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

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

```
# ./installer
```

- 5 Enter `Y` to install and press Return.
- 6 When the list of available products is displayed, select Storage Foundation 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).

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

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

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

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

```
1) Install minimal Storage Foundation HA packages -
   554 MB required
2) Install recommended Storage Foundation HA packages -
   798 MB required
3) Install all Storage Foundation HA packages -
   845 MB required
4) Display packages to be installed for each option
```

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

- 9 You are prompted to enter the system names (in the following example, "host1" and "host2") where you want to install the software. Enter the system name or names and then press Return.

```
Enter the platform system names separated by spaces: [q,?] host1 host2
```

Where *platform* indicates the operating system.

- 10 The installer can configure remote shell or secure shell communications for you among systems, however each system needs to have secure shell or remote shell 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.
- 11 After the system checks complete, the installer displays a list of the packages that will be installed. Press Enter to continue with the installation.
- 12 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 must still have a valid license to install and use Veritas products. Keyless licensing requires that you manage the systems with a Management Server.

See [“About Veritas product licensing”](#) on page 43.

- 13 You are prompted to enter the Standard or Enterprise product mode.
- 14 If you are going to use the Veritas Volume Replicator, enter **y** at the following prompt:

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

- 15 If you are going to use the Global Cluster Option, enter **y** at the following prompt:

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

16 The product installation completes.

Configure Storage Foundation and High Availability (SF and VCS) when prompted.

```
Would you like to configure SFHA on host1 host2? [y,n,q] (n) y
```

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

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

Note: You must configure Storage Foundation High Availability before you can use the product.

See [“Configuring Storage Foundation and High Availability Solutions”](#) on page 125.

17 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
```

18 View the log file, if needed, to confirm the installation.

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

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

Installing language packages

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

To install the language packages on the server**1** Make sure the VEA Service is not running.

```
# /opt/VRTS/bin/vxsvcctl status  
Current state of server : RUNNING
```

2 If the VEA Service is running, stop it by using the `vxsvcctl stop` command.

```
# /opt/VRTS/bin/vxsvcctl stop
```

3 Insert the "Language" disc into the DVD-ROM or CD-ROM drive. With Solaris volume management software, the disc is automatically mounted as `/cdrom/cdrom0`.

4 Install the language packages using the `install_lp` command.

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

5 Restart the VEA Service.

```
# /opt/VRTS/bin/vxsvcctrl start
```

Installing Storage Foundation and High Availability Solutions using the Web-based installer

This chapter includes the following topics:

- [About the Web-based installer](#)
- [Features not supported with Web-based installer](#)
- [Before using the Veritas Web-based installer](#)
- [Starting the Veritas Web-based installer](#)
- [Obtaining a security exception on Mozilla Firefox](#)
- [Performing a pre-installation check with the Veritas Web-based installer](#)
- [Installing SFHA with the Web-based installer](#)

About the Web-based installer

Use the Web-based installer's 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 `xprtlwid` 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 Internet Explorer or FireFox.

The Web installer creates log files whenever the Web installer is operating. While the installation processes are operating, the log files are located in a session-based directory under the `/var/tmp` directory. After the install process completes, the log files are located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and for 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 68.

See [“Starting the Veritas Web-based installer”](#) on page 69.

Features not supported with Web-based installer

In this release, the following features that can be performed using the script installer are not available in the Web-based installer:

- Configuring server-based I/O fencing
- Configuring non-SCSI3 I/O fencing in virtual environments where SCSI3 is not supported
- Installing language packages

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 Storage Foundation 5.1 SP1.
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.

Table 7-1 Web-based installer requirements (*continued*)

System	Function	Requirements
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, and later. ■ 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 `xprtlwid`, 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.

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


```
Secure Connection Failed
```

Obtain a security exception for your browser.
- 5 When prompted, enter `root` and root's password of the installation server.

Obtaining a security exception on Mozilla Firefox

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

To obtain a security exception

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

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

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

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

To perform a pre-installation check

- 1 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 69.
- 2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.
- 3 Select the product from the **Product** drop-down list, and click **Next**.
- 4 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Validate**.
- 5 The installer performs the precheck and displays the results.
- 6 If the validation completes successfully, the installer prompts you to begin the installation. Click **Yes** to install on the selected system. Click **No** to install later.
- 7 Click **Finish**. The installer prompts you for another task.

Installing 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 70.
- 2 Select **Install a Product** from the **Task** drop-down list.
- 3 Select **SFHA** from the Product drop-down list, and click **Next**.

- 4 On the License agreement page, read the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 5 Choose minimal, recommended, or all packages. Click **Next**.
- 6 Indicate the systems where you want to install. Separate multiple system names with spaces. Click **Validate**.
- 7 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.
- 8 After the validation completes successfully, click **Next** to install SFHA or SFCFSHA on the selected system.
- 9 After the installation completes, you must choose your licensing method. On the license page, select one of the following tabs:
 - Keyless licensing

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

For more information, go to the following website:

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

Complete the following information:

- Choose whether you want to install Standard or Enterprise mode.
- Choose whether you want to enable Veritas Volume Replicator. Click **Register**.
- Enter license key
If you have a valid license key, select this tab. Enter the license key for each system. Click **Register**.

- 10** For Storage Foundation, 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.

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

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

Click **Finish**.

Installing Storage Foundation and High Availability products using other methods

This chapter includes the following topics:

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

Installing with JumpStart

These JumpStart instructions assume a working knowledge of JumpStart. See the JumpStart documentation that came with your operating system for details on using JumpStart. Only fresh installations of SFHA are supported using JumpStart. Upgrading is not supported. The following procedure assumes a stand-alone configuration.

For the language pack, you can use JumpStart to install packages. You add the language packages in the script, and put those files in the JumpStart server directory.

Overview of JumpStart installation tasks

Review the summary of tasks before you perform the JumpStart installation.

Summary of tasks

- 1 Add a client (register to the JumpStart server). See the JumpStart documentation that came with your operating system for details.
- 2 Read the JumpStart installation instructions.
- 3 Generate the finish scripts.
- 4 Prepare shared storage installation resources.
- 5 Modify the rules file for JumpStart.
See the JumpStart documentation that came with your operating system for details.
- 6 Run the installer command from the disc or from directory `/opt/VRTS/install` directory to configure the Veritas software.

```
# /opt/VRTS/install/installer -configure
```
- 7 Install the operating system using the JumpStart server.

Generating the finish scripts

Perform these steps to generate the finish script to install SFHA.

To generate the script

- 1 Run the product installer program to generate the scripts.

```
installprod -jumpstart directory_to_generate_scripts
```

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

For example:

```
# ./installsf -jumpstart /js_scripts
```

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

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

4 Specify private region length.

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

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

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

6 JumpStart finish scripts, installer scripts, and encapsulation scripts are generated in the directory you specified in step 1. Output resembles:

```
The finish scripts for SF is generated at /js_scripts/  
jumpstart_sf.fin  
The installer script to configure SF is generated at /js_scripts/  
installsf  
The installer script to uninstall SF is generated at /js_scripts/  
uninstallsf  
The encapsulation boot disk script for VM is generated at  
/js_scripts/encap_bootdisk_vm.fin
```

List the `js_scripts` directory.

```
# ls /js_scripts
```

Output resembles:

```
encap_bootdisk_vm.fin installsf jumpstart_sf51.fin uninstallsf
```

Preparing installation resources

Prepare resources for the JumpStart installation.

To prepare the resources

1 Copy the contents of the installation disc to the shared storage.

```
# cd /cdrom/cdrom0  
# cp -r * BUILDSRC
```

2 Generate the response file with the list of packages.

See “[Installation script options](#)” on page 386.

```
# cd BUILDSRC/pkgs/  
# pkgask -r package_name.response -d /BUILDSRC/pkgs/packages_name.pkg
```

- 3 Create the `adminfile` file under `BUILDSRC/pkgs/` directory.

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

- 4 To configure or uninstall from `/opt/VRTS/install`, copy the `install` and `uninstall` scripts to `BUILDSRC`. You need to configure and uninstall from disc otherwise.
- 5 If you want to encapsulate the root disk automatically when you perform the JumpStart installation, copy the scripts `encap_bootdisk_vm51.fin` generated previously to `ENCAPSRC`.

See [“Generating the finish scripts”](#) on page 74.

Adding language pack information to the finish file

To add the language pack information to the finish file, perform the following procedure.

To add the language pack information to the finish file

- 1 For the language pack, copy the language packages from the language pack installation disc to the shared storage.

```
# cd /cdrom/cdrom0/pkgs  
# cp -r * BUILDSRC/pkgs
```

- 2 In the finish script, copy the product package information and replace the product packages with language packages.

- 3 In the finish script, copy the product patch information, and replace the product patch with language patches.
- 4 The finish script resembles:

```
. . .
for PKG in product_packages
do
...
done. . .
for PATCH in product_patches
do
...
done. . .
for PKG in language_packages
do
...
done. . .
for PATCH in language_patches
do
...
done
```

Installing SFHA using the pkgadd command

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

This procedure describes how to install the software on a stand-alone host. The system can be converted later to a Storage Foundation Manager managed host.

For information about obtaining and installing the SF Manager, refer to the *Veritas Storage Foundation Manager Installation Guide*.

To install SFHA using the pkgadd command

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

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

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

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

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

- `minpkgs`
- `recpkgs`
- `allpkgs`

See “[About installation scripts](#)” on page 385.

See “[Installation script options](#)” on page 386.

- 5 Install the packages listed in step [\[Unresolved xref\]](#).

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

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

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

- `minpkgs`
- `recpkgs`
- `allpkgs`

- 7 Install the patches for Storage Foundation 5.1 SP1 with the `patchadd` command.

```
# patchadd patch-ID
```

- 8 Verify that the packages are installed:

```
# pkginfo -l  
    packagename
```

- 9 If needed, start the VEA server:

```
# /opt/VRTSob/bin/vxsvcctrl start
```

- 10 Start the processes.

See [“Starting and stopping processes for the Veritas products ”](#) on page 317.

Configuration of Storage Foundation and High Availability products

- [Chapter 9. Preparing to configure Storage Foundation and High Availability](#)
- [Chapter 10. Configuring Storage Foundation](#)
- [Chapter 11. Configuring Storage Foundation and High Availability](#)
- [Chapter 12. Configuring Storage Foundation High Availability for data integrity](#)

Preparing to configure Storage Foundation and High Availability

This chapter includes the following topics:

- [Preparing to configure the clusters in secure mode](#)
- [About planning to configure I/O fencing](#)
- [Setting up the CP server](#)

Preparing to configure the clusters in secure mode

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

In a cluster that is online, if you want to enable or disable AT using the `installsfha -security` command, see the *Veritas Cluster Server Administrator's Guide* for instructions.

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

- A system in your enterprise that serves as root broker (RB).
You can either use an external system as root broker, or use one of the cluster nodes as root broker.
- To use an external root broker, identify an existing root broker system in your enterprise or install and configure root broker on a stable system.
See [“Installing the root broker for the security infrastructure”](#) on page 87.

- To use one of the cluster nodes as root broker, the installer does not require you to do any preparatory tasks.

When you configure the cluster in secure mode using the script-based installer, choose the automatic mode and choose one of the nodes for the installer to configure as root broker.

Symantec recommends that you configure a single root broker system for your entire enterprise. If you use different root broker systems, then you must establish trust between the root brokers.

For example, if the management server and the cluster use different root brokers, then you must establish trust.

- For external root broker, an authentication broker (AB) account for each node in the cluster is set up on the root broker system.

See [“Creating authentication broker accounts on root broker system”](#) on page 88.

- The system clocks of the external root broker and authentication brokers must be in sync.

The script-based installer provides the following configuration modes:

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

[Figure 9-1](#) depicts the flow of configuring SFHA cluster in secure mode.

Figure 9-1 Workflow to configure SFHA cluster in secure mode

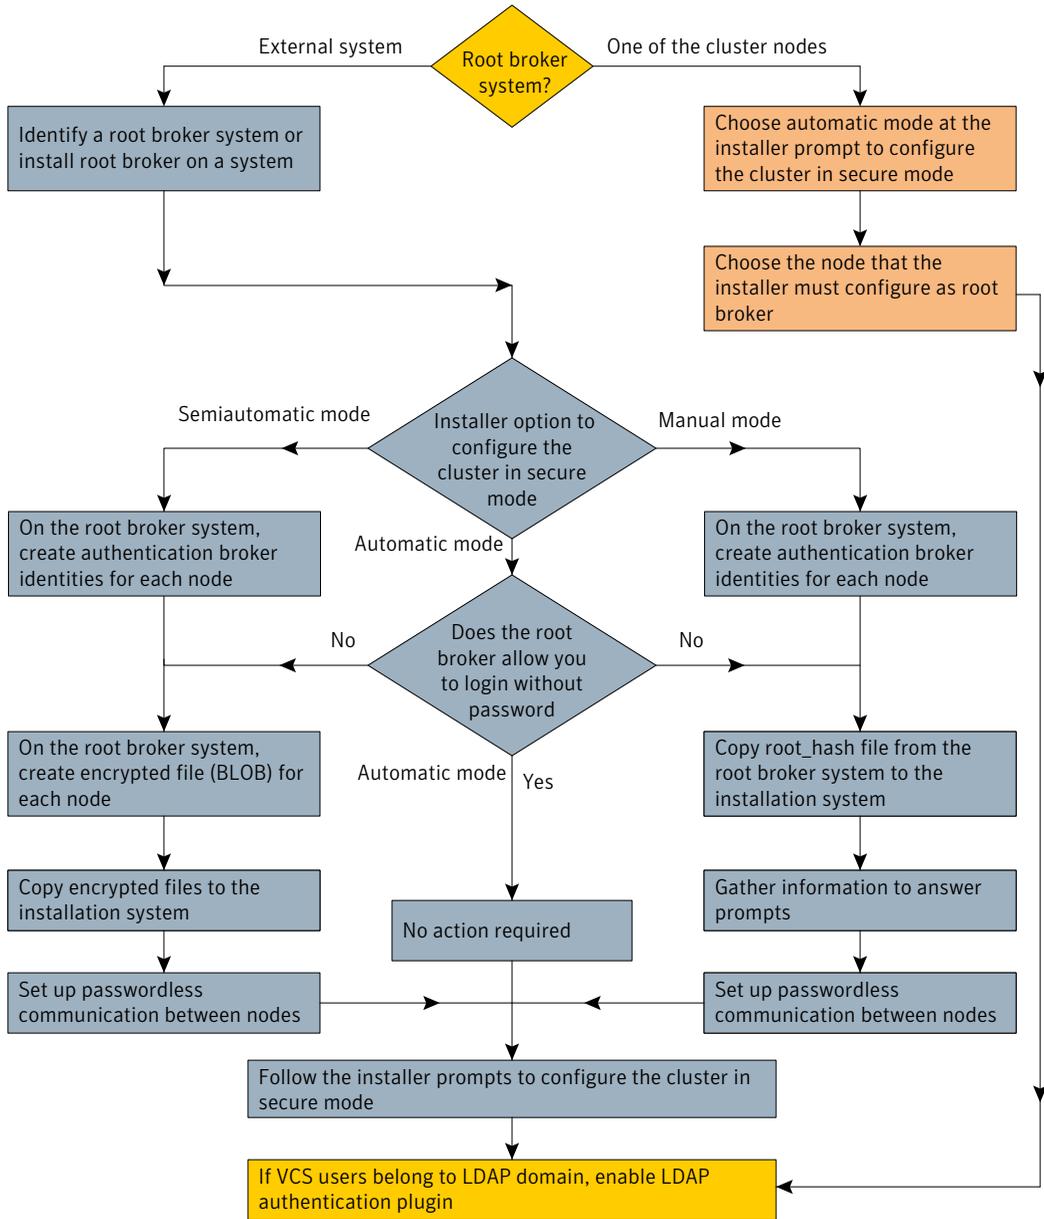


Table 9-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform. These preparatory tasks apply only when you use an external root broker system for the cluster.

Table 9-1 Preparatory tasks to configure a cluster in secure mode (with an external root broker)

Tasks	Who performs this task
<p>Decide one of the following configuration modes to set up a cluster in secure mode:</p> <ul style="list-style-type: none"> ■ Automatic mode ■ Semi-automatic mode ■ Manual mode 	VCS administrator
<p>Install the root broker on a stable system in the enterprise.</p> <p>See “Installing the root broker for the security infrastructure” on page 87.</p>	AT administrator
<p>To use the semi-automatic mode or the manual mode, on the root broker system, create authentication broker accounts for each node in the cluster.</p> <p>See “Creating authentication broker accounts on root broker system” on page 88.</p> <p>The AT administrator requires the following information from the VCS administrator:</p> <ul style="list-style-type: none"> ■ Node names that are designated to serve as authentication brokers ■ Password for each authentication broker 	AT administrator
<p>To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator.</p> <p>See “Creating encrypted files for the security infrastructure” on page 89.</p> <p>The AT administrator requires the following additional information from the VCS administrator:</p> <ul style="list-style-type: none"> ■ Administrator password for each authentication broker Typically, the password is the same for all nodes. 	AT administrator
<p>To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.</p>	AT administrator
<p>Copy the files that are required to configure a cluster in secure mode to the system from where you plan to install and configure SFHA.</p> <p>See “Preparing the installation system for the security infrastructure” on page 91.</p>	VCS administrator

Installing the root broker for the security infrastructure

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

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

See Symantec Product Authentication Service documentation for more information.

To install the root broker

- 1 Mount the product disc and start the installer.

```
# ./installer
```

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

```
Enter the operating_system system names separated by space [q,?]: venus
```

- 6 Review the output as the installer does the following:
 - Checks to make sure that AT supports the operating system
 - Verifies that you install from the global zone
 - Checks if the packages are already on the system.

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

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

- 9 Select a mode to configure the root broker from the three choices that the installer presents:

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

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

```
All AT processes that are currently running must be stopped
```

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

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

Creating authentication broker accounts on root broker system

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

To create authentication broker accounts on root broker system

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

```
venus> # vssat showalltrustedcreds
```

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

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

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

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

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

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

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

```
"Failed To Get Attributes For Principal"
```

Proceed to step 3.

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

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

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

Creating encrypted files for the security infrastructure

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

To create encrypted files

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

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

```
venus> # vssat showbrokerhash
```

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

```
venus> # vssat showalltrustedcreds
```

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

identity	<p>The value for the authentication broker identity, which you provided to create authentication broker principal on the root broker system.</p> <p>This is the value for the <code>--prplname</code> option of the <code>addprpl</code> command.</p> <p>See “Creating authentication broker accounts on root broker system” on page 88.</p>
password	<p>The value for the authentication broker password, which you provided to create authentication broker principal on the root broker system.</p> <p>This is the value for the <code>--password</code> option of the <code>addprpl</code> command.</p> <p>See “Creating authentication broker accounts on root broker system” on page 88.</p>

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

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

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

[configab]
identity=galaxy
password=password
root_domain=root@venus.symanteceexample.com
root_broker=venus.symanteceexample.com:2821
start_broker=false
enable_pbx=false
```

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

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

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

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

For example:

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

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

- 6 After you complete creating the output files for the encrypted file, you must copy these encrypted BLOB files for each node in the cluster.

Preparing the installation system for the security infrastructure

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

To prepare the installation system for the security infrastructure

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

Automatic mode Do the following:

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

Semi-automatic mode Do the following:

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

Manual mode

Do the following:

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

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.

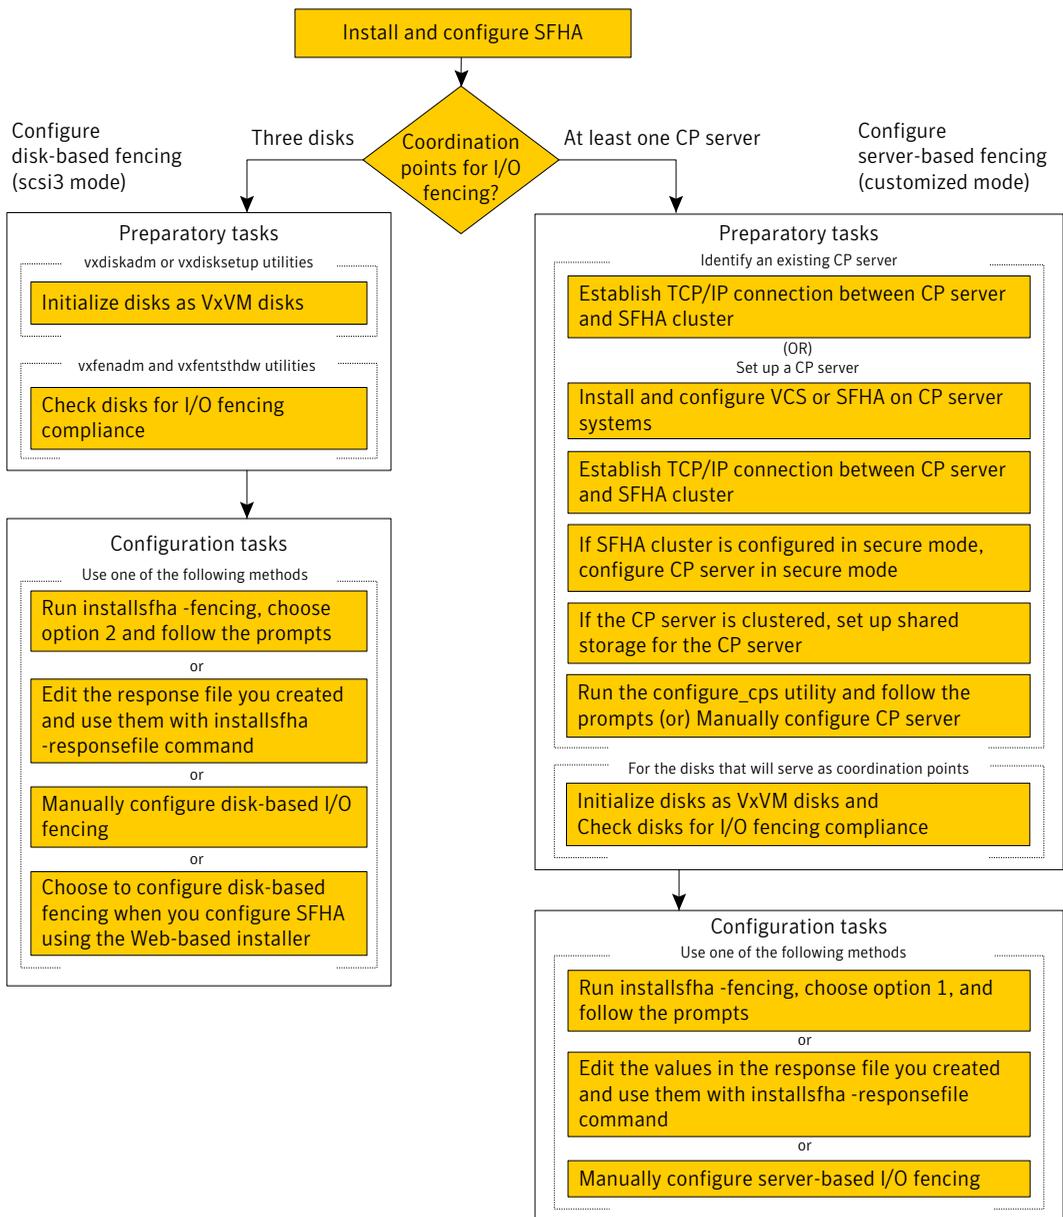
You can configure either disk-based I/O fencing or server-based I/O fencing. If your enterprise setup has multiple clusters that use VCS for clustering, Symantec recommends you to configure server-based I/O fencing.

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.

[Figure 9-2](#) illustrates a high-level flowchart to configure I/O fencing for the SFHA cluster.

Figure 9-2 Workflow to configure I/O fencing



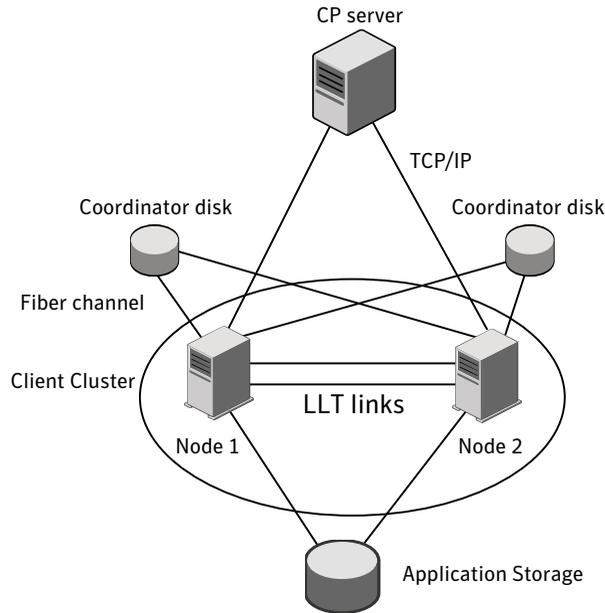
After you perform the preparatory tasks, you can use any of the following methods to configure I/O fencing:

Using the installsfha	See “Setting up disk-based I/O fencing using installsfha” on page 153. See “Setting up server-based I/O fencing using installsfha” on page 166.
Using the Web-based installer	See “Configuring SFHA using the Web-based installer” on page 145. Note: The Web-based installer supports only the disk-based fencing configuration.
Using response files	See “Response file variables to configure disk-based I/O fencing” on page 412. See “Response file variables to configure server-based I/O fencing” on page 414. See “Response file variables to configure server-based I/O fencing” on page 414. See “Configuring I/O fencing using response files” on page 411.
Manually editing configuration files	See “Setting up disk-based I/O fencing manually” on page 161. See “Setting up server-based I/O fencing manually” on page 178.

Typical SF HA cluster configuration with server-based I/O fencing

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

Figure 9-3 CP server, SF HA cluster, and coordinator disks



Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points. See [Figure 9-4](#) on page 96.
- 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 97.
- 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 97.

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 (must be an odd number) coordination points for I/O fencing. 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

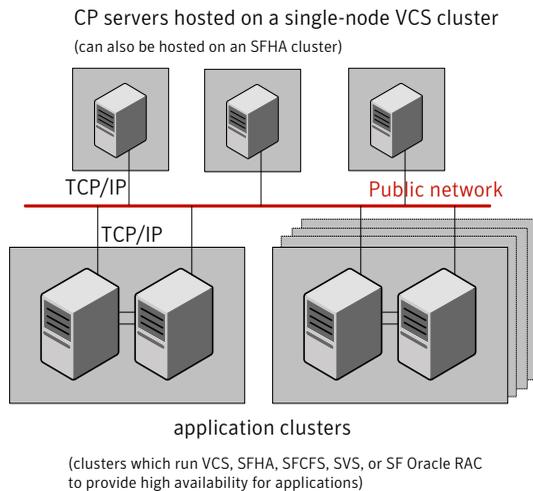


Figure 9-5 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.

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

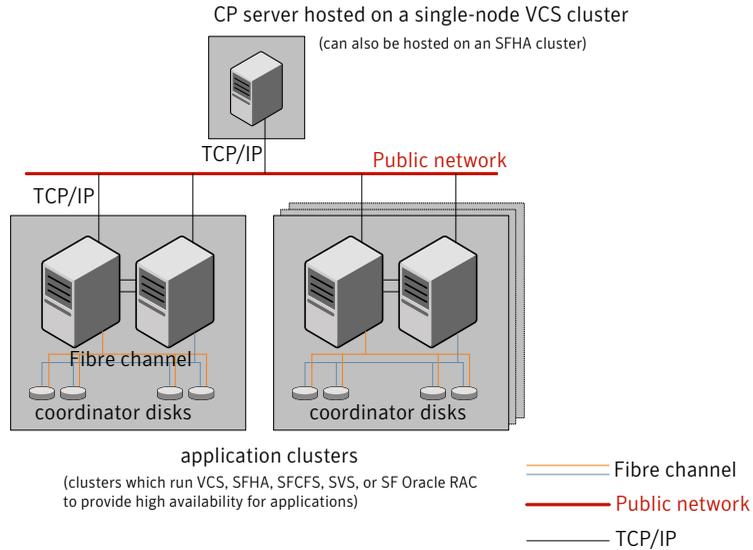
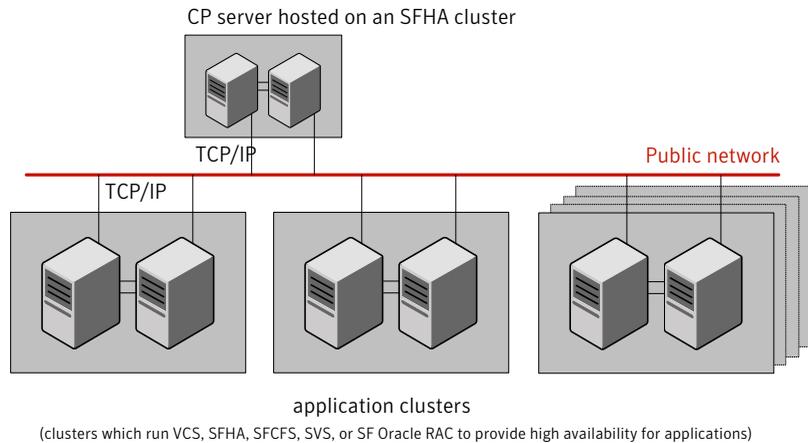


Figure 9-6 displays a configuration using a single CP server that is connected to multiple application clusters.

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



See “[Configuration diagrams for setting up server-based I/O fencing](#)” on page 473.

Setting up the CP server

Table 9-2 lists the tasks to set up the CP server for server-based I/O fencing.

Table 9-2 Tasks to set up CP server for server-based I/O fencing

Task	Reference
Plan your CP server setup	See “Planning your CP server setup” on page 98.
Install the CP server	See “Installing the CP server using the installer” on page 99.
Configure the CP server cluster in secure mode	See “Configuring the CP server cluster in secure mode” on page 100.
Set up shared storage for the CP server database	See “Setting up shared storage for the CP server database” on page 101.
Configure the CP server	See “Configuring the CP server using the configuration utility” on page 102. See “Configuring the CP server manually” on page 110.
Verify the CP server configuration	See “Verifying the CP server configuration” on page 112.

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.
- 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 fencing in enabled mode during the SFHA configuration.
 - You must set up shared storage for the CP server database during your CP server setup.

CP server setup uses a single system

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

Meet the following requirements for CP server:

- During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.
- During configuration, make sure to configure LLT and GAB.
- During configuration, set up the cluster in secure mode if you want secure communication between the CP server and the SF HA cluster (application cluster).

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

Proceed to configure the CP server.

See “[Configuring the CP server using the configuration utility](#)” on page 102.

See “[Configuring the CP server manually](#)” on page 110.

CP server setup uses multiple systems

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

Meet the following requirements for CP server:

- During installation, make sure to select all packages for installation. The VRTScps package is installed only if you select to install all packages.
- During configuration, set up the cluster in secure mode if you want secure communication between the CP server and the SF HA cluster (application cluster). See “[Preparing to configure the clusters in secure mode](#)” on page 83.
- During configuration, configure disk-based fencing (scsi3 mode).

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

Configuring the CP server cluster in secure mode

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

This step secures the HAD communication on the CP server cluster, and makes the authentication broker highly available.

Note: If you already configured Symantec Product Authentication Service (AT) during VCS configuration, you can 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:

```
# installsfha -security
```

See [“Preparing to configure the clusters in secure mode”](#) on page 83.

Setting 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.

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

To set up shared storage for the CP server database

- 1 Create a disk group containing the 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's not already imported.

For example:

```
# vxdg import cps_dg
```

- 3 Create a mirrored volume over the disk group.

For example:

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

- 4 Create a file system over the volume.

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

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

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

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

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

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

Configuring the CP server using the configuration utility

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

Perform one of the following procedures:

For CP servers on single-node VCS cluster: See [“To configure the CP server on a single-node VCS cluster”](#) on page 102.

For CP servers on an SFHA cluster: See [“To configure the CP server on an SFHA cluster”](#) on page 106.

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

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

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

3 Enter **1** at the prompt to configure CP server on a single-node VCS cluster. The configuration utility then runs the following preconfiguration checks:

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

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

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

4 Enter the name of the CP server.

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

5 Enter a valid virtual IP address on which the CP server process should depend on.

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

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 in range [49152, 65535], or
press <enter> for default port (14250):
```

- 7 Choose whether the communication between the CP server and the SF HA 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.symantecexample.com  
(b)CP Server Virtual IP: 10.209.83.85  
(c)CP Server Port: 14250  
(d)CP Server Security : 1  
(e)CP Server Database Dir: /etc/VRTScps/db  
-----
```

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

- 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
-----
```

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

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

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

- 12** Enter networkhosts information for the 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 resource bge0 on
system mycps1? [y/n] : y
Enter a valid IP address to configure NetworkHosts for NIC bge0 on
system mycps1 : 10.209.83.86
Do you want to add another Network Host ?[y/n] : n
```

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

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

- 14** 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.
```

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

For example:

```
# hagrp -state CPSSG
```

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

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

The configuration utility adds the `vxcpsserv` 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 `vxcpsserv` resource and its dependencies.

See [“Sample configuration files for CP server”](#) on page 431.

To configure the CP server on an SFHA cluster

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

```
# /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.symantecexample.com**

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

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

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 in range [49152, 65535], or press <enter> for default port (14250):

8 Choose whether the communication between the CP server and the SF HA 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.symantecexample.com  
(b) CP Server Virtual IP: 10.209.83.85  
(c) CP Server Port: 14250  
(d) CP Server Security : 1  
(e) CP Server Database Dir: /etc/VRTScps/db  
-----
```

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

11 The configuration utility proceeds with the configuration process, 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 Confirm whether you use the same NIC name for the virtual IP on all the systems in the cluster.

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

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

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

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

14 Enter networkhosts information for the 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 resource bge0 on system mycps1? [y/n] : y

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

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

15 Enter the netmask for the virtual IP address.

Enter the netmask for virtual IP 10.209.83.85 :
255.255.252.0

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

Enter the name of diskgroup for cps database :
cps_dg

17 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

- 18** 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.
```

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

For example:

```
# hagrp -state CPSSG
```

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

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

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

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

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

See [“Sample configuration files for CP server”](#) on page 431.

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 431.

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

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

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

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

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

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

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

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

- 5 Start VCS on all the cluster nodes.

```
# hstart
```

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

```
# hagrps -state CPSSG
```

Output similar to the following appears:

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

Verifying the CP server configuration

Perform the following steps to verify the CP server configuration.

To verify the CP server configuration

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

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

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

Configuring Storage Foundation

This chapter includes the following topics:

- [Configuring Storage Foundation using the installer](#)
- [Configuring Storage Foundation manually](#)
- [Configuring the SFDB repository database after installation](#)
- [Veritas Volume Replicator and Volume Manager setup after installation](#)

Configuring Storage Foundation using the installer

Storage Foundation does not require configuration. You need to start it.

To start Storage Foundation

- ◆ Run the installer command with the start option.

```
# ./installer -start
```

Configuring Storage Foundation manually

You can manually configure different components for Storage Foundation.

Configuring Veritas Volume Manager

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

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

To carry out further tasks such as disk encapsulation or initialization, please see the *Veritas Volume Manager Administrator's Guide*.

In releases of VxVM (Volume Manager) prior to 4.0, a system installed with VxVM was configured with a default disk group, `rootdg`, that had to contain at least one disk. By default, operations were directed to the `rootdg` disk group. From release 4.0 onward, VxVM can function without any disk group having been configured. Only when the first disk is placed under control must a disk group be configured. There is no longer a requirement that you name any disk group `rootdg`, and any disk group that is named `rootdg` has no special properties by having this name. During the setup procedures, you will be asked if you want to create a default disk group, and asked to specify its name.

Starting and enabling the configuration daemon

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

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

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

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

```
# vxctl mode
```

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

```
mode: enabled
```

This message indicates that `vxconfigd` is not running:

```
mode: not-running
```

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

```
mode: disabled
```

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

```
# vxconfigd
```

To enable the volume daemon, enter the following command:

```
# vxctl enable
```

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

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

Starting the volume I/O daemon

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

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

```
# vxiod
```

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

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

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

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

```
# vxiod set no_of_daemons
```

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

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

Using `vxinstall` to configure Veritas Volume Manager

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

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

The utility provides the following functions:

- Licensing VxVM.
- Setting up a system-wide default disk group.
- Starting VxVM daemons in case installation of SF has been done manually.

To run the command, enter

```
# vxinstall
```

which will prompt you to enter a license key:

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

If you don't have a license key, refer to the support section.

The presence of certain hardware arrays (for example, A5000) automatically generates a key.

The `vxinstall` program then asks if you want to set up a systemwide default disk group:

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

VxVM will continue with the question:

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

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

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

```
# vxdg defaultdg
```

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

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

Excluding a device that VxVM or Dynamic Multipathing controls

This section describes how to exclude a device that is under VxVM or Dynamic Multipathing control.

To prevent multipathing or suppress devices from being seen by VxVM

- 1 Enter the command

```
# vxdiskadm
```

- 2 Select menu item 17 (Prevent Multipathing/Suppress devices from VxVM's view) from the `vxdiskadm` main menu.

The following message displays:

```
VxVM INFO V-5-2-1239 This operation might lead to some devices
being suppressed from VxVM's view or prevent them from being
multipathed by vxdmp. (This operation can be reversed using the
vxdiskadm command).
```

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

- 3 Enter `y`.
- 4 Select one of the following operations:
 - Suppress all paths through a controller from VxVM's view:
Select Option 1.
Enter a controller name when prompted:

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

- Suppress a path from VxVM's view:
Select Option 2.
Enter a path when prompted.

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

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

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

Enter a VID:PID combination when prompted.

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

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

- Suppress all but one path to a disk:
Select Option 4 and read the messages displayed on the screen before specifying a path.
Enter a path when prompted:

```
Enter pathgroup: [<pattern>,list,list-exclude,q,?]
```

The following options allow you to exclude devices from vxddmp:

- Prevent multipathing of all disks on a controller by VxVM.
Select Option 5 and read the messages displayed on the screen before specifying a controller.
Enter a controller name when prompted. The controller entered is excluded from DMP control.

```
Enter a controller name: [<ctrl-name>,all,list,list-exclude,q,?]
```

- Prevent multipathing of a disk by VxVM.
Select Option 6 to exclude the specified path from multipathing. The corresponding disks are claimed in the OTHER_DISKS category and are not multipathed. Read the messages displayed on the screen before specifying a path.
Enter a path at the prompt:

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

- Prevent multipathing of disks by specifying a VID:PID combination.
Select Option 7 to exclude disks by a VID:PID combination. All disks returning a VID:PID combination are claimed in the OTHER_DISKS category and are not multipathed. Read the messages displayed on the screen before specifying a VID:PID.
Enter the VID:PID combination at the prompt.

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

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

Enabling optional cluster support in VxVM

This release includes an optional cluster feature that enables VxVM to be used in a cluster environment. The cluster functionality in VxVM allows multiple hosts to simultaneously access and manage a set of disks under VxVM control. A cluster is a set of hosts sharing a set of disks; each host is referred to as a node in the cluster.

The VxVM cluster feature requires a license, which can be obtained from your Customer Support channel.

To enable the cluster functionality in VxVM

- 1 Obtain a license for the VxVM cluster feature.
- 2 Install the software packages onto each system (node) to be included in the cluster.
- 3 Create the configuration files required to form a cluster.
- 4 Start the cluster services.
- 5 Configure shared disks.

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

Converting existing VxVM disk groups to shared disk groups

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

To convert existing disk groups to shared disk groups

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

3 Configure the disk groups using the following procedure.

To list all disk groups, use the following command:

```
# vxdg list
```

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

```
# vxdg deport disk_group_name
```

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

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

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

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

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

```
# vxdg list
```

The disk groups are now ready to be shared.

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

Configuring shared disks

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

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

Make sure that the shared disks are not being accessed from another node while you are performing the configuration. If you start the cluster on the node where you perform the configuration only, you can prevent disk accesses from other nodes because the quorum control reserves the disks for the single node.

Also, hot-relocation can be configured.

Verifying existing shared disks

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

To verify that your shared disk groups exist

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

```
# vxdg -s list
```

This displays the existing shared disk groups.

Upgrading in a clustered environment with FastResync set

Upgrading in a clustered environment with FastResync set requires additional steps.

This procedure applies to the following upgrade scenarios:

- Upgrading from VxVM 3.5 to VxVM 5.1 SP1
- Upgrading from VxVM 3.5 Maintenance Pack 4 to VxVM 5.1 SP1

If there are volumes in the shared disk groups with FastResync set (`fastresync=on`), before beginning the upgrade procedure, reattach each snapshot to its data volume, using this procedure:

To upgrade in a clustered environment when FastResync is set

- 1 You should run this procedure from the master node; to find out if you are on the master node, enter the command:

```
# vxdctl -c mode
```

- 2 On the master node, list which disk groups are shared by entering:

```
# vxdg -s list
```

- 3 Using the diskgroup names displayed by the previous command, list the disk groups that have volumes on which FastResync is set:

```
# vxprint -g diskgroup -F "%name" -e "v_fastresync"
```

4 Reattach each snapshot:

```
# vxassist -g diskgroup -o nofmr snapback snapshot_volume
```

5 If you are upgrading from VxVM 3.5 Maintenance Patch 3 or from VxVM 3.2 Maintenance Patch 5, set FastResync to off for each volume:

```
# vxvol -g diskgroup set fastresync=off volume
```

Configuring Veritas File System

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

```
/etc/vfstab
```

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

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

Loading and unloading the file system module

On Solaris 9 and 10, the `vxfs` file system module automatically loads on the first reference to a VxFS file system. This occurs when a user tries to mount a VxFS disk layout. In some instances, you may want to load the file system module manually. To do this, first load `vxfs`, then `vxportal`. `vxportal` is a pseudo device driver that enables VxFS commands to issue ioctls to the VxFS modules even when there are no file systems mounted on the system.

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

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

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

To determine if the modules successfully loaded, enter:

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

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

You can unload the module by entering:

```
# modunload -i portal_module_id
# modunload -i vxfv_module_id
```

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

```
# df -F vxfs
```

vxtunefs command permissions and Cached Quick I/O

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

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

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

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

Configuring the SFDB repository database after installation

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

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

Veritas Volume Replicator and Volume Manager setup after installation

VVR is fully integrated with Veritas Volume Manager (VxVM). Before using VVR, you must have the VxVM volumes set up and initialized.

Refer to the Volume Manager documentation for more information.

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.

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 to use Symantec Security Services.

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.

Before configuring a cluster to operate using Symantec Security Services, another system must already have Symantec Security Services installed and be operating as a Root Broker.

See the *Veritas Cluster Server Installation Guide* for more information on configuring a secure cluster.

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

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

[Overview of tasks to configure SFHA using the script-based installer](#) lists the tasks that are involved in configuring SFHA using the script-based installer.

Table 11-1 Tasks to configure SFHA using the script-based installer

Task	Reference
Start the software configuration	See “Starting the software configuration” on page 127.
Specify the systems where you want to configure SFHA	See “Specifying systems for configuration” on page 128.
Configure the basic cluster	See “Configuring the cluster name and ID” on page 129. See “Configuring private heartbeat links” on page 129.
Configure virtual IP address of the cluster (optional)	See “Configuring the virtual IP of the cluster” on page 132.
Configure the cluster in secure mode (optional)	See “Configuring the cluster in secure mode” on page 134.
Add VCS users (required if you did not configure the cluster in secure mode)	See “Adding VCS users” on page 137.
Configure SMTP email notification (optional)	See “Configuring SMTP email notification” on page 137.
Configure SNMP email notification (optional)	See “Configuring SNMP trap notification” on page 139.
Configure global clusters (optional) Note: You must have enabled Global Cluster Option when you installed SFHA.	See “Configuring global clusters” on page 141.
Complete the software configuration	See “Completing the VCS configuration” on page 142.

Starting the software configuration

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

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:

To configure SFHA using the installsfha program

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

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

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

Specifying systems for configuration

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

To specify system names for configuration

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

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

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

The installer does the following tasks:

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

- Makes sure the systems install from the global zone
 - Checks whether SFHA is installed
 - Exits if Storage Foundation 5.1 SP1 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 92.

Configuring the cluster name and ID

Enter the cluster information when the installer prompts you.

To configure the cluster

- 1** Review the configuration instructions that the installer presents.
- 2** Enter the unique cluster name and cluster ID.

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

Enter a unique Cluster ID number between 0-65535: [b,q,?] **7**

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 489.

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: LLT over Ethernet (allow installer to detect)

Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.

Skip to step 5.

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

The installer discovers and lists the network interface cards. You can use either the standard interfaces or the aggregated interfaces (bonded NICs).

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

- For Solaris SPARC:

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

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

```
[b,q,?] bge0
```

```
Would you like to configure a second private heartbeat link?
```

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

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

```
[b,q,?] bge1
```

```
Would you like to configure a third private heartbeat link?
```

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

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

- For Solaris x64:

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

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

```
[b,q,?] e1000g1
```

```
Would you like to configure a second private heartbeat link?
```

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

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

```
[b,q,?] e1000g2
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
NIC 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
NIC 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 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 to the Cluster Manager (Java Console) 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 To configure virtual IP, enter **y** at the prompt.
- 3 Confirm whether you want to use the discovered public NIC on the first system.

Do one of the following:

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

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

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

Do one of the following:

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

```
Is bge0 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

5 Enter the virtual IP address for the cluster.

You can enter either an IPv4 address or an IPv6 address.

- For IPv4: ■ Enter the virtual IP address.

```
Enter the Virtual IP address for the Cluster:
[b,q,?] 192.168.1.16
```

- Confirm the default netmask or enter another one:

```
Enter the netmask for IP 192.168.1.16: [b,q,?]
(255.255.240.0)
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

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

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

For IPv6

- Enter the virtual IP address.

```
Enter the Virtual IP address for the Cluster:  
[b, q, ?] 2001:454e:205a:110:203:baff:feee:10
```

- Enter the prefix for the virtual IPv6 address you provided. For example:

```
Enter the Prefix for IP  
2001:454e:205a:110:203:baff:feee:10: [b, q, ?] 64
```

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

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

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

Configuring the cluster in secure mode

If you want to configure the cluster in secure mode, make sure that you meet the prerequisites for secure cluster configuration.

The installer provides different configuration modes to configure a secure cluster. Make sure that you completed the pre-configuration tasks for the configuration mode that you want to choose.

See [“Preparing to configure the clusters in secure mode”](#) on page 83.

To configure the cluster in secure mode

- 1 Choose whether to configure SFHA to use Symantec Product Authentication Service.

```
Would you like to configure VCS to use Symantec Security  
Services? [y,n,q] (n) y
```

- If you want to configure the cluster in secure mode, make sure you meet the prerequisites and enter **y**.
- If you do not want to configure the cluster in secure mode, enter **n**. You must add VCS users when the configuration program prompts. See [“Adding VCS users”](#) on page 137.

- 2 Select one of the options to enable security.

Before you choose any of the options, make sure that all the nodes in the cluster can successfully ping the root broker system.

Select the Security option you would like to perform [1-3,b,q,?] (1)

Security Menu

- 1) Configure security completely automatically
- 2) Provide AB credentials using BLOBs
- 3) Provide AB credentials without using BLOBs
- b) Back to previous menu

Review the following configuration modes. Based on the configuration that you want to use, enter one of the following values:

Option 1.
Automatic
configuration

Based on the root broker you want to use, do one of the following:

- To use an external root broker:
Enter the name of the root broker system when prompted.
Requires remote access to the root broker. Make sure that all the nodes in the cluster can successfully ping the root broker system.
Review the output as the installer verifies communication with the root broker system, checks vxatd process and version, and checks security domain.

- To configure one of the nodes as root broker:
■ Press Enter at the following installer prompt:

```
If you already have an external
RB (Root Broker) installed and configured, enter
the RB name, or press Enter to skip: [b]
```

- Choose the node that the installer must configure as root and authentication broker. The installer configures the other nodes as authentication brokers.
At the installer prompt, you can choose the first node in the cluster to configure as RAB, or you can enter n to configure another node as RAB. For example:

```
Do you want to configure <galaxy> as RAB,
and other nodes as AB? [y,n,q,b] (y) n
Enter the node name which you want to
configure as RAB: nebula
```

Option 2.
Semiautomatic
configuration

Enter the path of the encrypted file (BLOB file) for each node when prompted.

Option 3.
Manual
configuration

Enter the following Root Broker information as the installer prompts you:

```
Enter root broker name: [b]
east.symantecexample.com
Enter root broker FQDN: [b]
(symantecexample.com)
symantecexample.com
Enter the root broker domain name for the
Authentication Broker's identity: [b]
root@east.symantecexample.com
Enter root broker port: [b] 2821
Enter path to the locally accessible root hash [b]
(/var/tmp/installvcs-200910221810ROA/root_hash)
/var/tmp/installvcs-200910221810ROA/root_hash
```

Enter the following Authentication Broker information as the installer prompts you for each node:

```
Enter Authentication broker's identity on
galaxy [b]
(galaxy.symantecexample.com)
galaxy.symantecexample.com
Enter the password for the Authentication broker's
identity on galaxy:
Enter Authentication broker's identity on
nebula [b]
(nebula.symantecexample.com)
nebula.symantecexample.com
Enter the password for the Authentication broker's
identity on nebula:
```

- 3 After you provide the required information to configure the cluster in secure mode, the program prompts you to configure SMTP email notification.

Note that the installer does not prompt you to add VCS users if you configured the cluster in secure mode. However, you must add VCS users later.

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

Adding VCS users

If you have enabled Symantec Product Authentication Service, 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 want to set the username and/or password for the Admin user
(default username = 'admin', password='password')? [y,n,q] (n) y
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 139.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

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

- Enter the SMTP server’s host name.

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

- Enter the email address of each recipient.

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

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

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

- 4 Add more SMTP recipients, if necessary.

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

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

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

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

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

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

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

5 Verify and confirm the SMTP notification information.

NIC: bge0

SMTP Address: smtp.example.com

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

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

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

Configuring SNMP trap notification

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

Note that you can also configure the notification after installation.

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

To configure the SNMP trap notification

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

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

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

See [“Configuring global clusters”](#) on page 141.

3 Provide information to configure SNMP trap notification.

Provide the following information:

■ Enter the NIC information.

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

■ Enter the SNMP trap daemon port.

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

■ Enter the SNMP console system name.

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

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

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

4 Add more SNMP consoles, if necessary.

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

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

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

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

5 Verify and confirm the SNMP notification information.

NIC: bge0

SNMP Port: 162

Console: saturn receives SNMP traps for Error or higher events

Console: jupiter receives SNMP traps for SevereError or higher events

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

Configuring global clusters

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster. 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: bge0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

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

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

For IPv6 Global Cluster Option configuration verification:

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

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

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

Completing the VCS configuration

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

If you chose to configure the cluster in secure mode, the installer then does the following before it starts SFHA in secure mode:

- Depending on the security mode you chose to set up Authentication Service, the installer does one of the following:
 - Creates the security principal

- Executes the encrypted file to create security principal on each node in the cluster
- Creates the VxSS service group
- Creates the Authentication Server credentials on each node in the cluster
- Creates the Web credentials for SFHA users
- Sets up trust with the root broker

To complete the VCS configuration

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

```
Do you want to stop VCS 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 395.

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 144.

See [“Updating product licenses using vxlicinst”](#) on page 144.

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              = Solaris
Version               = 5.1 SP1
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 145.

To update product licenses

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

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

Replacing a SFHA demo license with a permanent license

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

To replace a demo key

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

2 Shut down SFHA on all nodes in the cluster:

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

This command does not shut down any running applications.

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

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

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

```
# vxlicrep
```

5 Start SFHA on each node:

```
# hstart
```

Configuring 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.

Note: If you want to configure server-based I/O fencing, you must either use the script-based installer or manually configure.

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 69.
- 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 **Validate**.

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 disk-based I/O fencing, click **Yes**.

If you want to configure server-based I/O fencing, or if you decide to configure I/O fencing later, click **No**. You can either use the `installsfha -fencing` command 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.
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 To configure SFHA on a cluster . 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.
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 In the Confirmation dialog box that appears, choose whether or not to configure the cluster in secure mode using Symantec Product Authentication Service (AT).

To configure the cluster in secure mode, click **Yes**.

If you want to perform this task later, click **No**. You can use the `installsfha-security` command. Go to step [To configure SFHA on a cluster](#).

- 8 On the Security Options page, choose an option to enable security and specify the required information.

Do not configure security services Choose this option if you do not want to enable security. The installer takes you to the next page to configure optional features of SFHA.

Configure security automatically Choose this option to use an external root broker. Enter the name of the root broker that is already configured for your enterprise environment, and click **Validate**. The installer configures the cluster in secure mode.

Configure one node as RAB and the others as AB Select the system that you want to configure as RAB node. The installer configures the cluster in secure mode.

Click **Next**.

- 9 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.
- 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.

Click **Next**.

- 10 On the Stop Processes page, click **Next** after the installer stops all the processes successfully.

- 11 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 [To configure SFHA on a cluster](#), then skip to step [To configure SFHA on a cluster](#). Go to step [To configure SFHA on a cluster](#) to configure fencing.

- 12 On the Select Fencing Type page, specify the following information:

Configure disk based fencing Choose the **Configure disk based fencing** option.

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 mechanism for the disk group. Go to step [To configure SFHA on a cluster](#).
- 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 [To configure SFHA on a cluster](#).

- 13 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 Mechanism Choose the fencing mechanism for the disk group.

- 14 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.

- 15 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 and starting Veritas Enterprise Administrator

Before using the Veritas Enterprise Administrator server or client, start them both.

Optional configuration can also be completed at this time.

Activating, getting status, starting, and stopping the VEA server

After installing the VEA packages, the VEA server may need to be stopped and restarted. The VEA service is automatically started when you reboot your system.

To activate and start the VEA server

- 1 Activate the VEA server.

```
# /opt/VRTSob/bin/vxsvcctl activate
```

- 2 Check the state of the VEA server.

```
# /opt/VRTSob/bin/vxsvcctl status
```

- 3 Start the VEA server.

```
# /opt/VRTSob/bin/vxsvcctl start
```

- 4 Stop the VEA server.

```
# /opt/VRTSob/bin/vxsvcctl stop
```

You can also stop the VEA server manually by killing the `vxsvc` process.

The VEA server is automatically started on a reboot.

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 server-based I/O fencing manually](#)
- [Enabling or disabling the preferred fencing policy](#)

Setting up disk-based I/O fencing using installsfha

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

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

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

```
# devfsadm
```

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

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

```
vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i c2t13d0
```

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

Checking shared disks for I/O fencing

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

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

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

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

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

- Verifying the Array Support Library (ASL)
See [“Verifying Array Support Library \(ASL\)”](#) on page 155.

- Verifying that nodes have access to the same disk
 See “[Verifying that the nodes have access to the same disk](#)” on page 156.
- Testing the shared disks for SCSI-3
 See “[Testing the disks using vxfentsthdw utility](#)” on page 156.

Verifying Array Support Library (ASL)

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

To verify Array Support Library (ASL)

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

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

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

The following output is a sample:

```
# vxddladm listsupport all
```

LIBNAME	VID	PID
libvx3par.so	3PARdata	VV
libvxCLARiiON.so	DGC	All
libvxFJTSYe6k.so	FUJITSU	E6000
libvxFJTSYe8k.so	FUJITSU	All
libvxap.so	SUN	All
libvxatf.so	VERITAS	ATFNODES
libvxcompellent.so	COMPELNT	Compellent Vol
libvxcopan.so	COPANSYS	8814, 8818

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

```
# vxdisk scandisks
```

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

Verifying that the nodes have access to the same disk

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

To verify that the nodes have access to the same disk

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed SFHA.
- 2 Ensure that both nodes are connected to the same disk during the testing. Use the `vxfenadm` command to verify the disk serial number.

```
vxfenadm -i diskpath
```

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

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

From node A, enter:

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

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

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

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

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

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

Testing the disks using vxfcntlsthdw utility

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

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

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

```
The disk /dev/rdisk/ctl1d0s2 is ready to be configured for I/O Fencing on
node galaxy
```

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

To test the disks using `vxfcntlshdw` utility

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

2 From one node, start the utility.

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

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

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

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

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

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

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

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

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

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

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

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

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

- 5 Review the output as the utility performs the checks and report 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/rsk/c1t1d0s2 have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

- 7 Run the vxfcntl 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 5.1 SP1 is configured properly.

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

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

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

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

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

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

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

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

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

See [“Checking shared disks for I/O fencing”](#) on page 154.

- 7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.

- 8 Enter the I/O fencing disk policy that you chose to use. For example:

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

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
 - Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information
- 9 Verify and confirm the I/O fencing configuration information that the installer summarizes.
- 10 Review the output as the configuration program does the following:
- Stops 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 Coordination Point agent to monitor coordination points”](#) on page 182.

Setting up disk-based I/O fencing manually

[Table 12-1](#) lists the tasks that are involved in setting up I/O fencing.

Table 12-1 Tasks to set up I/O fencing manually

Task	Reference
Initializing disks as VxVM disks	See “Initializing disks as VxVM disks” on page 153.
Identifying disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 161.
Checking shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 154.
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 162.
Creating I/O fencing configuration files	See “Creating I/O fencing configuration files” on page 163.
Modifying SFHA configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 164.
Configuring Coordination Point agent to monitor coordination points	See “Configuring Coordination Point agent to monitor coordination points” on page 182.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 166.

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 153.

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 154.

Setting up coordinator disk groups

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

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

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

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

To create the `vxencoorddg` disk group

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

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

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

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

- 1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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

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

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

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

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

Modifying VCS configuration to use I/O fencing

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

To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

- 3 If the I/O fencing driver `vxfen` is already running, stop the I/O fencing driver. Depending on the Solaris version on the cluster nodes, run the following command:

- Solaris 9:

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

- Solaris 10:

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

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

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

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

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

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

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

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

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

For example, on each remaining node, enter:

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

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

- Start the I/O fencing driver.

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

Depending on the Solaris version on the cluster nodes, run the following command:

- Solaris 9:

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

- Solaris 10:

```
# svcadm enable vxfen
```

- Start VCS.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

- 1 On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the 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

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

See [“Verifying the security configuration on the SF HA cluster to use CP server coordination point”](#) on page 167.

See [“Configuring server-based I/O fencing using the installsfha”](#) on page 169.

Verifying the security configuration on the SF HA cluster to use CP server coordination point

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

To verify the security configuration on SF HA cluster to use CP server coordination point

- 1** Run the following command:

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

The following is an example of the command output:

```
Domain(s) Found 1

*****

Domain Name HA_SERVICES@galaxy.symantecexample.com

Expiry Interval 0

*****
```

- 2** There should be a domain name entry with the following format in the command output:

```
HA_SERVICES@hostname.domainname
```

or

```
HA_SERVICES@hostname
```

3 There should not be duplicate entries for HA_SERVICES domain.

The following is an example of an incorrect configuration:

```
showdomains

Domain(s) Found :          3

*****

Domain Name:      HA_SERVICES@galaxy.symantecexample.com

Domain Type:     vx

*****

Domain Name:      broker@galaxy.symantecexample.com

Domain Type:     vx

*****

Domain Name:      HA_SERVICES@galaxy

Domain Type:     vx

*****
```

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

Configuring server-based I/O fencing using the installsfha

You can configure server-based I/O fencing for the SFHA cluster using the `installsfha`.

With server-based fencing, you can have the coordination points in your configuration as follows:

- Combination of CP servers and SCSI-3 compliant coordinator disks
 - CP servers only
- Symantec also supports server-based fencing with a single highly available CP server that acts as a single coordination point.

See [“About planning to configure I/O fencing”](#) on page 92.

See [“Recommended CP server configurations”](#) on page 95.

This section covers the following example procedures:

Mix of CP servers and coordinator disks	See “To configure server-based fencing for the SFHA cluster (one CP server and two coordinator disks)” on page 170.
Single CP server	See “To configure server-based fencing for the SFHA cluster (single CP server)” on page 175.

To configure server-based fencing for the SFHA cluster (one CP server and two coordinator disks)

- 1 Depending on the server-based configuration model in your setup, make sure of the following:

- CP servers are configured and are reachable from the SFHA cluster. The SFHA cluster is also referred to as the application cluster or the client cluster.

See [“Setting up the CP server”](#) on page 98.

- The coordination disks are verified for SCSI3-PR compliance.

See [“Checking shared disks for I/O fencing”](#) on page 154.

- 2 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.

- 3 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 5.1 SP1 is configured properly.

- 4 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 [1-3,b,q] 1
```

- 5 Make sure that the storage supports SCSI3-PR, and answer **y** at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

6 Provide the following details about the coordination points at the installer prompt:

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

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

- Enter the total number of coordinator disks among the coordination points.

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

7 Provide the following CP server details at the installer prompt:

- Enter the virtual IP addresses or host names of the virtual IP address for each of the CP servers. The installer assumes these values to be identical as viewed from all the application cluster nodes.

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

- Enter the port that the CP server would be listening on.

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

8 Provide the following coordinator disks-related details at the installer prompt:

- Enter the I/O fencing disk policy for the coordinator disks.

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

- Choose the coordinator disks from the list of available disks that the installer displays. Ensure that the disk you choose is available from all the SFHA (application cluster) nodes.

The number of times that the installer asks you to choose the disks depends on the information that you provided in step 6. For example, if you had chosen to configure two coordinator disks, the installer asks you to choose the first disk and then the second disk:

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

```
1) c1t1d0s2  
2) c2t1d0s2  
3) c3t1d0s2
```

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

- If you have not already checked the disks for SCSI-3 PR compliance in step 1, check the disks now.
The installer displays a message that recommends you to verify the disks in another window and then return to this configuration procedure.
Press Enter to continue, and confirm your disk selection at the installer prompt.
- Enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):  
[b] (vxsfencoorddg)
```

9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 3  
CP Server (Port):  
  1. 10.209.80.197 (14250)  
SCSI-3 disks:  
  1. c1t1d0s2  
  2. c2t1d0s2  
Disk Group name for the disks in customized fencing: vxsfencoorddg  
Disk mechanism used for customized fencing: raw
```

The installer initializes the disks and the disk group and departs the disk group on the SFHA (application cluster) node.

- ## 10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the SFHA (application cluster):
- Make sure that the security configuration in the application cluster and the CP server is the same. If CP server is configured for security, ensure that the application cluster also runs in secure mode.
 - If the CP server is configured for security, perform the following steps:

- Review the output as the installer verifies if the SFHA (application cluster) nodes have already established trust with an AT root broker.
- If the SFHA (application cluster) nodes and the CP server use different AT root brokers, enter `y` at the installer prompt and provide the following information:
 - Hostname for the authentication broker for any one of the CP servers
 - Port number where the authentication broker for the CP server is listening for establishing trust
 - Hostname for the authentication broker for any one of the SFHA (application cluster) nodes
 - Port number where the authentication broker for the SFHA (application cluster) is listening for establishing trust

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

11 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

- 12 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

```
Updating client cluster information on CP Server 10.210.80.199

Adding the client cluster to the CP Server 10.210.80.199 ..... Done

Registering client node galaxy with CP Server 10.210.80.199..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Registering client node nebula with CP Server 10.210.80.199 ..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Updating /etc/vxfenmode file on galaxy ..... Done
Updating /etc/vxfenmode file on nebula ..... Done
```

See [“About I/O fencing configuration files”](#) on page 429.

- 13 Configure the CP agent on the SFHA (application cluster).

```
Do you want to configure CP Agent on the client cluster? [y,n,q]
(y)

Enter a non-existing name for the service group for CP Agent:
[b] (vxfen)

Adding CP Agent via galaxy ..... Done
```

- 14 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.
- 15 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

To configure server-based fencing for the SFHA cluster (single CP server)

- 1 Make sure that the CP server is configured and is reachable from the SFHA cluster. The SFHA cluster is also referred to as the application cluster or the client cluster.

See “[Setting up the CP server](#)” on page 98.

- 2 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.

- 3 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 5.1 SP1 is configured properly.

- 4 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 [1-3,b,q] 1
```

- 5 Make sure that the storage supports SCSI3-PR, and answer **y** at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

- 6 Enter the total number of coordination points as **1**.

```
Enter the total number of co-ordination points including both
CP servers and disks: [b] (3) 1
```

Read the installer warning carefully before you proceed with the configuration.

- 7 Provide the following CP server details at the installer prompt:
 - Enter the virtual IP address or the host name of the virtual IP address for the CP server. The installer assumes these values to be identical as viewed from all the application cluster nodes.

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

- Enter the port that the CP server would be listening on.

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

8 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 1
CP Server (Port):
  1. 10.209.80.197 (14250)
```

9 If the CP server is configured for security, the installer sets up secure communication between the CP server and the SFHA (application cluster):

- Make sure that the security configuration in the application cluster and the CP server is the same. If CP server is configured for security, ensure that the application cluster also runs in secure mode.
- If the CP server is configured for security, perform the following steps:
 - Review the output as the installer verifies if the SFHA (application cluster) nodes have already established trust with an AT root broker.
 - If the SFHA (application cluster) nodes and the CP server use different AT root brokers, enter `y` at the installer prompt and provide the following information:
 - Hostname for the authentication broker for any one of the CP servers
 - Port number where the authentication broker for the CP server is listening for establishing trust
 - Hostname for the authentication broker for any one of the SFHA (application cluster) nodes
 - Port number where the authentication broker for the SFHA (application cluster) is listening for establishing trust

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

10 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

11 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the /etc/vxfenmode file with the appropriate details in each of the application cluster nodes.

The installer also populates the /etc/vxfenmode file with the entry `single_cp=1` for such single CP server fencing configuration.

```
Updating client cluster information on CP Server 10.210.80.199

Adding the client cluster to the CP Server 10.210.80.199 ..... Done

Registering client node galaxy with CP Server 10.210.80.199..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Registering client node nebula with CP Server 10.210.80.199 ..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Updating /etc/vxfenmode file on galaxy ..... Done
Updating /etc/vxfenmode file on nebula ..... Done
```

See [“About I/O fencing configuration files”](#) on page 429.

12 Configure the CP agent on the SFHA (application cluster).

```
Do you want to configure CP Agent on the client cluster? [y,n,q]
(y)

Enter a non-existing name for the service group for CP Agent:
[b] (vxfen)

Adding CP Agent via galaxy ..... Done
```

- 13 Review the output as the installer stops and restarts VCS with the fencing processes on each application cluster node, and completes the I/O fencing configuration.
- 14 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

Setting up server-based I/O fencing manually

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

Table 12-2 Tasks to set up server-based I/O fencing manually

Action	Description
Preparing the CP servers for use by the SFHA cluster	See “Preparing the CP servers manually for use by the SF HA cluster” on page 178.
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 164.
Configuring Coordination Point agent to monitor coordination points	See “Configuring Coordination Point agent to monitor coordination points” on page 182.
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 184.

Preparing the CP servers manually for use by the SF HA cluster

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

[Table 12-3](#) displays the sample values used in this procedure.

Table 12-3 Sample values in procedure

CP server configuration component	Sample name
CP server	mycps1.symantecexample.com
Node #1 - SF HA cluster	galaxy
Node #2 - SF HA cluster	nebula

Table 12-3 Sample values in procedure (*continued*)

CP server configuration component	Sample name
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the SF HA cluster

- 1 Determine the cluster name and uuid on the SF HA cluster.

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

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster clus1

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2}
```

- 2 Use the `cpsadm` command to check whether the SF HA cluster and nodes are present in the CP server.

For example:

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

ClusName  UUID                               Hostname(Node ID) Registered
clus1     {f0735332-1dd1-11b2} galaxy(0)          0
clus1     {f0735332-1dd1-11b2} nebula(1)         0
```

If the output does not show the cluster and nodes, then add them as described in the next step.

For detailed information about the `cpsadm` command, see the *Veritas Cluster Server Administrator's Guide*.

3 Add the SF HA cluster and nodes to each CP server.

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

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

```
Cluster clus1 added successfully
```

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

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

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

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

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

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

4 If security is to be enabled, check whether the `_HA_VCS_users` are created in the CP server.

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

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

```
Username/Domain Type Cluster Name / UUID Role  
  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantecexample.com/vx  
clus1/{f0735332-1dd1-11b2} Operator  
  
_HA_VCS_nebula@HA_SERVICES@nebula.symantecexample.com/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 `_HA_VCS_users` (for example, cpsclient@galaxy).

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

5 Add the users to the CP server.

First, determine the user@domain to be added on the SF HA cluster (application cluster).

The user for fencing should be of the form `_HA_VCS_short-hostname` and domain name is that of `HA_SERVICES` user in the output of command:

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

Next, issue the following commands on the CP server (mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantecexample.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_galaxy@HA_SERVICES@galaxy.symantecexample.com  
successfully added
```

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_nebula@HA_SERVICES@nebula.symantecexample.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_nebula@HA_SERVICES@nebula.symantecexample.com  
successfully added
```

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

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

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c clus1\  
-u {f0735332-1dd1-11b2}\  
-e _HA_VCS_galaxy@HA_SERVICES@galaxy.symantecexample.com\  
-f cps_operator -g vx  
  
Cluster successfully added to user  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantecexample.com privileges.  
  
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c clus1\  
-u {f0735332-1dd1-11b2}\  
-e _HA_VCS_nebula@HA_SERVICES@nebula.symantecexample.com\  
-f cps_operator -g vx  
  
Cluster successfully added to user  
_HA_VCS_nebula@HA_SERVICES@nebula.symantecexample.com privileges.
```

Configuring Coordination Point agent to monitor coordination points

The following procedure describes how to manually configure the Coordination Point agent to monitor coordination points (CP server or SCSI-3 disks).

To configure Configuration Point agent to monitor coordination points

- 1 Ensure that your SF HA cluster has been properly installed and configured with fencing enabled.
- 2 Create a parallel service group vxfen and add a coordpoint resource to the vxfen service group using the following commands:

```
# haconf -makerw
# 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 SF HA cluster using the `hares` commands. For example:

```
# hares -state coordpoint
```

The following is an example of the command and output:

```
# hares -state coordpoint

# Resource      Attribute  System  Value
coordpoint     State     galaxy  ONLINE
coordpoint     State     nebula  ONLINE
```

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

The agent log contains detailed Coordination Point agent monitoring information; including information about whether the Coordination Point agent is able to access all the coordination points, information to check on which coordination points the Coordination Point agent is reporting missing keys, etc.

To view all such information in the engine log, change the `dbg` level for that node using the following 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
```

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

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.

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 30.

To enable preferred fencing for the I/O fencing configuration

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

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

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

- Make the VCS configuration writable.

```
# haconf -makerw
```

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

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

- Set the value of the system-level attribute FencingWeight for each node in the cluster.
For example, in a two-node cluster, where you want to assign galaxy five times more weight compared to nebula, run the following commands:

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

- Save the VCS configuration.

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

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

- Make the VCS configuration writable.

```
# haconf -makerw
```

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

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

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

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

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

- Save the VCS configuration.

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

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

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

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

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

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

```
# haconf -makerw
```

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

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


Upgrading Storage Foundation and High Availability products

- [Chapter 13. Preparing to upgrade](#)
- [Chapter 14. Upgrading Storage Foundation or Storage Foundation and High Availability](#)
- [Chapter 15. Performing a rolling upgrade](#)
- [Chapter 16. Performing a phased upgrade](#)
- [Chapter 17. Upgrading with Live Upgrade](#)
- [Chapter 18. 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](#)
- [Tasks for upgrading the Storage Foundation for Databases \(SFDB\) tools](#)
- [Preparing to upgrade](#)

About upgrading

You have 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 192.

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

See [“Supported upgrade paths”](#) on page 193.

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

Caution: After you perform an upgrade from 5.1 or 5.1RPx to 5.1 SP1, Symantec recommends that you do not roll-back to 5.1 or 5.1RPx.

If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure you upgraded all application clusters to 5.1 SP1. 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 13-1 Available upgrade methods

Upgrade types and considerations	Methods available for upgrade
Typical upgrades—uses 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—uses 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—uses 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—uses 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>

Note: Script- and Web-based upgrades ask for very similar system information for upgrades.

Supported upgrade paths

The following tables describe upgrading to 5.1 SP1.

Table 13-2 Solaris SPARC upgrades using the script- or Web-based installer

Veritas software versions	Solaris 2.6, 7	Solaris 8	Solaris 9	Solaris 10
3.5 3.5 MP4 4.0 4.0 MP1 4.0 MP2	No upgrade path exists. Uninstall the product and upgrade the operating system to at least Solaris 9. Perform a full 5.1 SP1 installation.	No upgrade path exists. Uninstall the product and upgrade the operating system to at least Solaris 9. Perform a full 5.1 SP1 installation.	No upgrade path exists. Uninstall the product. Perform a full 5.1 SP1 installation.	N/A
4.1 4.1 MP1 4.1 MP2	N/A	No upgrade path exists. Uninstall the product and upgrade the operating system to at least Solaris 9. Perform a full 5.1 SP1 installation.	Upgrade to 4.1 MP2. Upgrade to 5.1 SP1.	Upgrade to 4.1 MP2. Upgrade to 5.1 SP1.
5.0 5.0 MP1 5.0 MP3 5.0 MP3 RP3 5.0 MP3 RP4	N/A	No upgrade path exists. Uninstall the product and upgrade the operating system to at least Solaris 9. Perform a full 5.1 SP1 installation.	Upgrade to 5.1 SP1.	Upgrade to 5.1 SP1.
5.1 5.1 P1 5.1 RPx	N/A	N/A	Upgrade to 5.1 SP1.	Upgrade to 5.1 SP1.

Table 13-3 Solaris x64

Veritas software versions	
4.1 4.1 Phase 2	No upgrade path exists. Uninstall the product. Perform a 5.1 SP1 installation.
5.0 5.0 MP3 5.0 MP3 RP3 5.0 MP3 RP4	Upgrade to 5.1 SP1.
5.1 5.1 P1 5.1 RPx	Upgrade to 5.1 SP1.

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

In prior versions of Storage Foundation and High Availability, when upgrading a system with an encapsulated root disk, you first had to unencapsulate. When upgrading to SFHA 5.1 SP1, that is no longer necessary provided you are upgrading from a supported version, as shown in the table below.

Table 13-4 Upgrading using installer when the root disk is encapsulated

Starting version	Ending version	Action required
4.x or 4.x MPx	5.1 SP1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.0 or 5.0 MPx or 5.0 MP3 RPx	5.1 SP1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 or 5.1 RPx	5.1 SP1	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Tasks for upgrading the Storage Foundation for Databases (SFDB) tools

Tasks for upgrading SFDB tools to version 5.1 SP1:

- Preparing to migrate the repository database before upgrading from 5.0.x or earlier to 5.1 SP1
 See “[Pre-upgrade tasks for migrating the SFDB repository database](#)” on page 198.
- Migrating the repository database after upgrading from 5.0.x or earlier to 5.1 SP1
 See “[Post upgrade tasks for migrating the SFDB repository database](#)” on page 280.

Caution: If you are running Oracle version 11.1.0.6 and upgrading a Storage Foundation product to 5.1 SP1, upgrade the Oracle binaries and database to version 11.1.0.7 before moving to SP1.

Preparing to upgrade

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

Getting ready for the upgrade

Complete the following tasks before you perform the upgrade:

- Review the *Veritas Storage Foundation Release Notes* for any late-breaking information on upgrading your system.
- Review the Symantec Technical Support website for additional information: <http://www.symantec.com/techsupp/>
- For Solaris 10, make sure that all non-global zones are booted and in the running state before you use the Veritas product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must upgrade each package in each non-global zone manually.

For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you reboot the alternative root, you can install `VRTSodm`.

- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the packages, for example `/packages/Veritas` when the root file system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.

Do not put the files under `/tmp`, which is erased during a system reboot. Do not put the files on a file system that is inaccessible prior to running the upgrade script.

You can use a Veritas-supplied disc for the upgrade as long as modifications to the upgrade script are not required. If `/usr/local` was originally created as a slice, modifications are required.
- Unmount all the file systems not on the `root` disk. Comment out their entries in `/etc/vfstab`. Stop the associated volumes and deport the associated disk groups. Any file systems that the Solaris operating system or Storage Foundation assumes should be in `rootdg` but are not, must be unmounted and the associated entry in `/etc/vfstab` commented out.
- For any startup scripts in `/etc/rcS.d`, comment out any application commands or processes that are known to hang if their file systems are not present.
- Make sure that the current operating system supports version 5.1 SP1 of Storage Foundation or Storage Foundation High Availability. If the operating system does not support the Veritas product, plan for a staged upgrade.
- Schedule sufficient outage time for the upgrade. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/vfstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/vfstab` and not mounted during the upgrade. Active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure the file systems are clean before upgrading.
See [“Verifying that the file systems are clean”](#) on page 205.
- If required, upgrade VxFS disk layouts to a supported version.
Some previous layout versions cannot be mounted on VxFS 5.1 SP1. You can upgrade these layout versions online before installing VxFS 5.1 SP1, or upgrade them using `vxfsconvert` after installing VxFS 5.1 SP1.

Note: You must upgrade disk layout version 4 and 5 before you upgrade the product.

- Upgrade arrays (if required)
 - If replication using VVR is configured, Symantec recommends that the disk group version is at least 110 prior to upgrading.
- ```
vxdg list diskgroup
```
- If replication using VVR is configured, make sure the size of the SRL volume is greater than 110 MB.  
Refer to the *Veritas Volume Replicator Administrator's Guide*.
  - If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date on all the hosts.

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

---

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

---

- If VCS is used to manage VVR replication, follow the preparation steps to upgrade VVR and VCS agents.
- Make sure that you have worked out all terminal emulation issues. Make sure that the terminal you use is fully functional for OpenBoot prompts and single-user and multi-user run levels.
- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Veritas products.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. Removing the disk in this manner offers a fallback point.
- To upgrade on a remote host, set up RSH or SSH.  
See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 439.
- Determine if the root disk is encapsulated.  
See [“Determining if the root disk is encapsulated”](#) on page 199.

## Creating backups

Save relevant system information before the upgrade.

### To create backups

- 1 Log in as superuser.
- 2 Before the upgrade, ensure that you have made backups of all data that you want to preserve.

Back up the `/etc/system` file.

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

```
cp /etc/vfstab /etc/vfstab.orig
```

- 4 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 5 If you are installing the high availability version of the Veritas Storage Foundation 5.1 SP1 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.

## Pre-upgrade tasks for migrating the SFDB repository database

If you plan to continue using checkpoints or SmartTier for Oracle policies you created with a 5.0x or earlier version of Storage Foundation for Oracle, you must prepare to migrate the SFDB repository database to 5.1 SP1 before upgrading to Storage Foundation or Storage Foundation for Oracle RAC 5.1 SP1.

If you are upgrading from 5.1 to 5.1 SP1, no upgrade steps are required for the SFDB tools.

---

**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 5.1 SP1.

---

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 5.1 SP1.

---

## Determining if the root disk is encapsulated

Before you upgrade, you need to determine if the root disk is encapsulated by running the following command:

```
mount | grep "/" on"
```

If the output from this command includes a path name that contains `vx` and `rootvol` as in `/dev/vx/dsk/bootdg/rootvol`, then the root disk is encapsulated.

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

## Preupgrade 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.

The following related documents are available:

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

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

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

### Planning an upgrade from the previous VVR version

If you plan to upgrade VVR from the previous VVR version, you can upgrade VVR with reduced application downtime by upgrading the hosts at separate times.

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

See the *Veritas Storage Foundation 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 13-5](#), if either the Primary or Secondary are running a version of VVR prior to 5.1 SP1, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 5.1 SP1, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

**Table 13-5** VVR versions and checksum calculations

| VVR prior to 5.1 SP1<br>(DG version <= 140) | VVR 5.1 SP1<br>(DG version >= 150) | VVR calculates checksum<br>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 High Availability supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

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

## Additional settings for using VVR in a localized environment

If the language packages for VVR are installed, VVR displays localized messages, if the client locale is a supported non-English locale. The client locale is the locale from which you are accessing the VVR command line or GUI. For example, if the Japanese version of VVR is installed, then the messages are displayed in the Japanese locale, if the client locale is Japanese.

Make sure that the appropriate locale has been installed on all the hosts that are intended to be a part of the VVR RDS setup. Otherwise, some VVR error messages will be displayed in English, because it is the default locale. Make sure the following settings are done on all hosts that are intended to be part of the RDS:

- Install the required client locale from the Operating System disc.
- Install the required Volume Manager and VVR localized packages. To use VVR VEA, make sure to install the localized package for the VEA client.
- Set the client locale, before using any of the VVR interfaces:
  - for the VVR command line or VVR VEA, set the locale using the appropriate method for your operating system. When you start VVR VEA, the GUI detects and uses the client locale.

- for VRW, select the locale from the VRW login page.

## Preparing to upgrade VVR when VCS agents are configured

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

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

### Freezing the service groups and stopping all the applications

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

**To freeze the service groups and stop applications 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:** Write down the list of frozen service groups for future use.

---

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

```
haconf -dump -makero
```

---

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

---

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

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

---

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

---

- 9 Repeat step 8 for each cluster.
- 10 For private disk groups, determine and note down the hosts on which the disk groups are imported.
- See “[Determining the nodes on which disk groups are online](#)” on page 203.

### 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 it by typing the following command on the node on which it is imported:

```
vxrecover -bs
```

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

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

---

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

---

## Verifying that the file systems are clean

Verify that all file systems have been cleanly unmounted.

**To make sure the file systems are clean**

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

```
echo "8192B.p S" | /opt/VRTSvxfs/sbin/fsdb filesystem | \
 grep clean
 flags 0 mod 0 clean clean_value
```

A *clean\_value* value of `0x5a` indicates the file system is clean. A value of `0x3c` indicates the file system is dirty. A value of `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.

- 2 If a file system is not clean, enter the following commands for that file system:

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

These commands should complete any extended operations on the file system and unmount the file system cleanly.

A pending large fileset clone removal extended operation might be in progress if the `umount` command fails with the following error:

```
file system device busy
```

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

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

- 3 If an extended operation is in progress, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large fileset clone can take several hours.
- 4 Repeat step 1 to verify that the unclean file system is now clean.

## Upgrading the array support

The Storage Foundation 5.1 SP1 release includes all array support in a single package, `VRTSaslapm`. The array support package includes the array support previously included in the `VRTSvxvm` package. The array support package also includes support previously packaged as external array support libraries (ASLs) and array policy modules (APMs).

See the 5.1 SP1 Hardware Compatibility List for information about supported arrays.

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

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. The installation of the VRTSvxvm package exits with an error if external ASLs or APMs are detected.

After you have installed Storage Foundation 5.1 SP1, Symantec provides support for new disk arrays through updates to the VRTSaslapm package.

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



# Upgrading Storage Foundation or Storage Foundation and High Availability

This chapter includes the following topics:

- [Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required](#)
- [Upgrading Veritas Storage Foundation to 5.1 SP1 using the product installer or manual steps](#)
- [Upgrading SFHA with the Veritas Web-based installer](#)
- [Upgrading the Solaris operating system](#)
- [Upgrading Veritas Volume Replicator](#)
- [Upgrading language packages](#)

## Upgrading Veritas Storage Foundation with the product installer when OS upgrade is not required

This section describes upgrading to the current Veritas Storage Foundation if the root disk is unencapsulated, and you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 5.1 SP1.

Use this procedure to upgrade Veritas Storage Foundation or Veritas Storage Foundation High Availability.

**To upgrade a Veritas Storage Foundation product**

1 Log in as superuser.

2 Unmount any mounted VxFS file systems.

The installer supports the upgrade of multiple hosts, if each host is running the same version of VxVM and VxFS. Hosts must be upgraded separately if they are running different versions.

If any VxFS file systems are mounted with the QuickLog feature, QuickLog must be disabled before upgrading. See the "Veritas QuickLog" chapter of the *Veritas File System Administrator's Guide* for more information.

3 If you are upgrading Storage Foundation and High Availability, 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
```

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

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

5 If your system has separate /opt and /var file systems, make sure they are mounted before proceeding with installation.

6 If replication using VVR is configured, 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.

---

7 Load and mount the disc. If you downloaded the software, navigate to the top level of the download directory.

- 8 From the disc, run the `installer` command. If you downloaded the software, run the `./installer` command.

```
cd /cdrom/cdrom0
./installer
```

- 9 Enter `G` to upgrade and press Return.
- 10 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 system names separated by spaces on which to
install SF: host1 host2
```

Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.

- 11 The installer asks if you agree with the terms of the End User License Agreement. Press `y` to agree and continue.
- 12 The installer lists the packages and any patches to install or to update. You are prompted to confirm that you are ready to upgrade.

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

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

- 13 The installer stops, uninstalls, reinstalls, and starts specified packages.
- 14 The Veritas Storage Foundation software is verified and configured.
- 15 The installer prompts you to provide feedback, and provides the log location for the upgrade.

## Upgrading Veritas Storage Foundation to 5.1 SP1 using the product installer or manual steps

This section describes upgrading a Veritas Storage Foundation product from a prior release to 5.1 SP1. Symantec recommends that you perform this upgrade from single-user mode. No VxFS file systems can be in use at the time of the upgrade.

The following procedures are for Veritas Storage Foundation or Veritas Storage Foundation High Availability.

Choose the appropriate procedure for your situation.

- If the current Storage Foundation product is installed on an operating system supported by 5.1 SP1, you do not need to upgrade the operating system. If you do not plan to upgrade the operating system, use one of the following upgrade procedures:
  - Upgrade SF but not OS with the product installer. This is the recommended upgrade procedure.  
See [“Upgrading Veritas Storage Foundation with the product installer”](#) on page 212.
  - Upgrade SF but not OS with manual steps (pkgadd and patchadd commands).  
See [“Upgrading Veritas Storage Foundation using manual steps”](#) on page 214.
- If you plan to upgrade the operating system, you must perform additional steps to upgrade. If the current Storage Foundation product is installed on an operating system which is no longer supported by 5.1 SP1, you must upgrade the operating system. If you plan to upgrade the operating system, use the following upgrade procedure:  
See [“Upgrading Veritas Storage Foundation to 5.1 SP1 using upgrade scripts \(OS upgrade\)”](#) on page 216.

## Upgrading Veritas Storage Foundation with the product installer

This section describes upgrading to the current Veritas Storage Foundation, and you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 5.1 SP1.

This procedure can be used to upgrade Veritas Storage Foundation or Veritas Storage Foundation High Availability.

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

### To upgrade a Veritas Storage Foundation product

- 1 Log in as superuser.
- 2 Unmount any mounted VxFS file systems.

The installer supports the upgrade of multiple hosts, if each host is running the same version of VxVM and VxFS. Hosts must be upgraded separately if they are running different versions.

If any VxFS file systems are mounted with the QuickLog feature, QuickLog must be disabled before upgrading. See the "Veritas QuickLog" chapter of the *Veritas File System Administrator's Guide* for more information.

- 3 If you are upgrading a high availability (HA) product, take all service groups offline.

List all service groups:

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

For each service group listed, take it offline:

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

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

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

- 5 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation.
- 6 If replication using VVR is configured, 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.

---

- 7 Load and mount the disc.  
See [“Mounting the product disc”](#) on page 55.
- 8 To invoke the common installer, run the `installer` command on the disc as shown in this example:  

```
cd /cdrom/cdrom0
./installer
```
- 9 Enter `G` to upgrade and press Return.

- 10 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 system names separated by spaces on which to
install SF: host1
```

Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.

- 11 Installer asks if you agree with the terms of the End User License Agreement. Press `y` to agree and continue.
- 12 The installer lists the packages and patches to install or update. You are prompted to confirm that you are ready to upgrade.

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

- 13 The installer lists the the patches and the packages to install or upgrade.
- 14 The installer verifies, configures, and starts the Veritas Storage Foundation software.

## Upgrading Veritas Storage Foundation using manual steps

This section describes upgrading from a previous version of Veritas Storage Foundation to the current Veritas Storage Foundation (5.1 SP1) when you do not intend to upgrade your Solaris version. Only use this procedure if you are already running a version of Solaris that is supported with 5.1 SP1.

### To upgrade a Veritas Storage Foundation product

- 1 Stop the VEA service:

```
/opt/VRTS/bin/vxsvcctl stop
```

- 2 Unmount any mounted VxFS file systems.

The installer supports the upgrade of multiple hosts, if each host is running the same version of VxVM and VxFS. Hosts must be upgraded separately if they are running different versions.

If any VxFS file systems are mounted with the QuickLog feature, QuickLog must be disabled before upgrading. See the "Veritas QuickLog" chapter of the *Veritas File System Administrator's Guide* for more information.

- 3 If the VxFS NetBackup libraries package (`VRTSfnsnb1`) is installed, remove it before you install the new packages.

To remove the package, use the `pkgrm` command as follows:

```
pkgrm VRTSfnsb1
```

Respond to any system messages as needed.

The libraries contained in this package are included in the `VRTSVxfs` package in 5.1 SP1.

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

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

---

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

---

- 5 If your system has separate `/opt` and `/var` file systems, make sure they are mounted before proceeding with installation.

- 6 Load and mount the software disc.

See [“Mounting the product disc”](#) on page 55.

- 7 Change to the directory containing the SFHA packages.

```
cd /dvd_mount
```

- 8 If VVR is configured, run the `vvr_upgrade_start` script on all hosts to save the original VVR configuration:

```
./scripts/vvr_upgrade_start
```

- 9 Remove the Veritas packages from your existing installation.

Refer to the *Storage Foundation Installation Guide* for the previous release to obtain the list of packages to remove.

- 10 Run the following command to obtain a list of recommended packages to install:

```
./installsf -recpkgs
```

- 11 Run the following command to obtain a list of recommended patches to install:

```
./installsf -recpkgs -listpatches
```

- 12** Use the `pkgadd` and `patchadd` commands to install the packages from the previous steps.

```
pkgadd -d . package_name.pkg
patchadd patch_id
```

If replication using VVR is configured, ignore the following error messages that appear on the Primary console during the installation process:

```
VxVM VVR vxrlink ERROR V-5-1-3371 Can not recover rlink_name.
rvg_name is in PASSTHRU mode
```

```
VxVM VVR vxrlink ERROR V-5-1-3473 Log header I/O error
```

Also ignore the following error message that appears on the Secondary console:

```
WARNING: VxVM VVR vxio V-5-0-278 Rlink rlink_name is stale and
not replicating
```

- 13** Configure the SF installation using the `installsf -configure` command.

- 14** If VVR is configured, issue the following command on all the hosts to complete the upgrade. If a host contains only Secondary RVGs, we recommend that you first run the following command on that host:

```
/dvd_mount/scripts/vvr_upgrade_finish
```

The `vvr_upgrade_finish` script upgrades only the SRL, after which, the RVG cannot work with the earlier versions of VxVM or VVR.

## Upgrading Veritas Storage Foundation to 5.1 SP1 using upgrade scripts (OS upgrade)

This section describes upgrading to the current Veritas Storage Foundation and need to upgrade the Solaris version. If the operating system is not at a supported Solaris version, you must follow this procedure.

This upgrade procedure allows you to retain existing VxVM and VxFS configurations. After upgrading, you can resume using your file systems and volumes as before (without having to run `vxinstall` again).

It is important that you follow these steps in the specified order.

**To begin the upgrade**

- 1 If VCS agents for VVR are configured, you must perform the pre-upgrade steps before proceeding.  
 See [“Preparing to upgrade VVR when VCS agents are configured”](#) on page 202.

- 2 Load and mount the disc.  
 See [“Mounting the product disc”](#) on page 55.

- 3 Verify that an upgrade is possible on the system. Enter the following command:

```
/dvd_mount/scripts/upgrade_start -check
```

- 4 Run the `upgrade_start` script to preserve the previous configuration of Volume Manager.

```
/dvd_mount/scripts/upgrade_start
```

- 5 If the `upgrade_start` script fails for any reason, run the `upgrade_finish` script to undo any changes already made. Verify that the system is restored by comparing `/etc/system`, `/etc/vfstab`, and the output of the `format` command. Then determine and correct the cause of the `upgrade_start` failure. If you cannot correct the problem in a timely manner, restore the `vfstab` file to the version saved, restore any other applications, and perform an `init 6` to completely restore the system.

- 6 Verify that all the Primary RLINKs are up-to-date on all the hosts.

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

---

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

---

- 7 If VVR is configured, run the `vvr_upgrade_start` script on all hosts to save the original VVR configuration:

```
/dvd_mount/scripts/vvr_upgrade_start
```

- 8 If you have VxFS file systems specified in the `/etc/vfstab` file, comment them out.

- 9 Remove the existing Storage Foundation packages in one of the following ways:

- using the `uninstallsf` script

- using `pkgrm`

For details, refer to the *Storage Foundation Installation Guide* for the existing Storage Foundation version.

After you run the `uninstallsf` script, verify that all VRTS\* packages are removed; otherwise, remove them manually using `pkgrm`.

- 10 If you are upgrading the operating system, do so now.

Refer to the Solaris installation documentation.

- 11 Install the Storage Foundation packages in one of the following ways:

- using the common installer

See [“To upgrade the Veritas Storage Foundation packages with the product installer”](#) on page 218.

- using manual steps

See [“To upgrade the Veritas Storage Foundation packages with manual steps”](#) on page 219.

#### To upgrade the Veritas Storage Foundation packages with the product installer

- 1 Load and mount the disc.

See [“Mounting the product disc”](#) on page 55.

- 2 To invoke the common installer, run the `installer` command on the disc as shown in this example:

```
cd /dvd_mount
./installer
```

- 3 Select **G** to upgrade the product, then select the number for the product you are installing.
- 4 Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.
- 5 If you commented out VxFS File System entries in the `/etc/vfstab` file, uncomment them.
- 6 Complete the upgrade by restoring the configuration.

### To upgrade the Veritas Storage Foundation packages with manual steps

- 1 If you are upgrading from Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, resynchronize all existing snapshots before upgrading.

For Veritas Storage Foundation for DB2:

```
/opt/VRTS/bin/db2ed_vmsnap -D DB2DATABASE -f SNAPPLAN \
-o resync
```

For Veritas Storage Foundation for Oracle:

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

- 2 Load and mount the software disc.
- 3 Change to the directory containing the packages.

```
cd /dvd_mount
```

- 4 Run the following command to obtain a list of recommended packages to install:

```
./installsf -recpkgs
```

Run the following command to obtain a list of all packages to install:

```
./installsf -allpkgs
```

- 5 Add packages with the `pkgadd` command add patches with the `patchadd` command.
- 6 If you commented out VxFS File System entries in the `/etc/vfstab` file, uncomment them.
- 7 Complete the upgrade by restoring the configuration.

### Restoring the configuration and completing the upgrade

- 1 Complete the upgrade using the `upgrade_finish` script.

```
devlinks
/dvd_mount/scripts/upgrade_finish
```

- 2 Configure the product using the following command:

```
/dvd_mount/installer -configure
```

If some Veritas modules fail to unload, perform the following steps:

- Reboot the systems.
- 3 Importing a pre-5.1 SP1 Veritas Volume Manager disk group does not automatically upgrade the disk group version to the VxVM 5.1 SP1 level. You may need to manually upgrade each of your disk groups following a VxVM upgrade.

See [“Upgrading VxVM disk group versions”](#) on page 291.

## Upgrading SFHA with the Veritas Web-based installer

This section describes how to upgrade SFHA with the Veritas Web-based installer. 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.

### To upgrade SFHA

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 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 69.
- 4 On the Select a task and a product page, select **Upgrade a Product**.  
The installer detects the product that is installed on the specified system.
- 5 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Validate**.
- 6 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.
- 7 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.

- 8 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.
- 9 Click **Finish**. The installer prompts you for another task.
- 10 If you want to upgrade VCS or SFHA 5.1 on the CP server systems to version Storage Foundation 5.1 SP1, make sure that you upgraded all application clusters to version Storage Foundation 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems. For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native operating system accounts.

## Upgrading the Solaris operating system

If you are running Storage Foundation 5.1 SP1 with an earlier release of the Solaris operating system, you can upgrade the Solaris operating system using the following procedure.

---

**Warning:** You should only use this procedure to upgrade the Solaris operating system if you are running Storage Foundation 5.1 SP1.

---

The directory `/opt` must exist, be writable, and must not be a symbolic link. This is because the volumes not temporarily converted by the `upgrade_start` are unavailable during the upgrade process. If you have a symbolic link from `/opt` to one of the unconverted volumes, the symbolic link will not function during the upgrade and items in `/opt` will not be installed.

### To upgrade the Solaris operating system only

- 1 Bring the system down to single-user mode using the following command:

```
init s
```

You must mount `/opt` manually if `/opt` is on its own partition.

- 2 Load and mount the software disc from the currently installed version of Storage Foundation.

See ["Mounting the product disc"](#) on page 55.

- 3 Change directory:

```
cd /mount_point/scripts
```

- 4 Run the `upgrade_start` with the `-check` argument to detect any problems that exist which could prevent a successful upgrade. Use the `upgrade_start` script that was supplied with the currently installed SF release. If this command reports success, you can proceed with running the `upgrade_start` script, but if it reports errors, correct the problem(s) and rerun `upgrade_start -check`.

```
./upgrade_start -check
```

- 5 Run the `upgrade_start` script so that the system can come up with partitions. The `upgrade_start` script searches for volumes containing file systems, and if any are found, converts them to partitions:

```
./upgrade_start
```

- 6 Bring the system down to run level 0.

```
init 0
```

- 7 Upgrade the operating system to a supported version of Solaris.

See “[Supported Solaris operating systems](#)” on page 37.

You should boot up the system from run level 0 depending on the Solaris upgrade procedure that you want to follow. Refer to the Solaris installation documentation for instructions on how to upgrade the Solaris operating system.

- 8 After installing the Solaris operating system, install any Solaris patches required by Veritas Storage Foundation 5.1 SP1.

See the *Veritas Storage Foundation Release Notes*.

- 9 After the system is up with the upgraded Solaris operating system, bring the system down to single-user mode by entering:

```
init s
```

- 10 Ensure that `/opt` is mounted.

- 11 Load and mount the software disc from the currently installed version of Storage Foundation.

- 12 If you upgraded to Solaris 10, you must reinstall certain Storage Foundation packages and patches in order to support Solaris 10 functionality.

To reinstall the required packages, follow the steps below:

- Remove the existing packages in the reverse order of their installation. For example, if you chose the installation of all packages and patches then uninstall those in the following order.

For Storage Foundation:

```
pkgrm VRTSsat VRTSodm VRTSdbed
VRTSfssdk VRTSvxfs VRTSsfmh VRTSob VRTSaslapm
VRTSvxvm VRTSspt VRTSperl VRTSvlic
```

For Storage Foundation and High Availability:

```
pkgrm VRTSvxvm VRTSvxfs VRTSvcS
VRTSvxfen VRTSgab VRTSllt
```

For Storage Foundation and High Availability:

```
pkgrm VRTSvxvm VRTSvxfs VRTSvcS
VRTSvxfen VRTSgab VRTSllt
```

- Run the following commands. To obtain a list of recommended packages to install:

```
./installsf -recpkgs
```

Or

To obtain a list of all packages to install:

```
./installsf -allpkgs
```

To obtain a list of all patches to install:

```
./installsf -listpatches
```

- Change to the directory containing the Storage Foundation packages.

```
cd /mount_point/pkgS
```

- Use the `pkgadd` command to install the packages from the list you generated.
- Change to the directory containing the Storage Foundation patches, and use the `patchadd` command to install the patches from the list you generated.

- Reboot the system.

**13** Complete the upgrade from the software disc from the currently installed version of Storage Foundation by entering:

```
devlinks
./upgrade_finish
```

## Upgrading Veritas Volume Replicator

If a previous version of Veritas Volume Replicator (VVR) is configured, the product installer upgrades VVR automatically when you upgrade the Storage Foundation products.

When upgrading from 4.1 MP1 or later, you have the option to upgrade without disrupting replication.

See [“Upgrading VVR without disrupting replication”](#) on page 224.

### 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 199.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade can be performed either on the Primary or on the Secondary. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS. This section includes separate sets of steps, for the Primary upgrade and for the Secondary upgrade.

---

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

---

### Upgrading VVR on the Secondary

Follow these instructions to upgrade the Secondary hosts.

### To upgrade the Secondary

- 1 Stop replication to the Secondary host by initiating a Primary pause using the following command:

```
vradmin -g diskgroup pauserep local_rvgname
```

- 2 Upgrade from VVR after 4.1 MP1 and prior to 5.1 SP1 on the Secondary.
- 3 Resume the replication from the Primary using the following command:

```
vradmin -g diskgroup resumerep local_rvgname sec_hostname
```

### Upgrading VVR on the Primary

After you upgrade the Secondary, use the Veritas product installer to upgrade the Primary.

---

**Note:** Reduce application downtime while upgrading by planning your upgrade.

---

See [“Planning an upgrade from the previous VVR version”](#) on page 199.

## Upgrading language packages

If you are upgrading Veritas products in a language other than English, you must install the required language packages after installing the English packages. Verify that the English installation is correct before proceeding.

Install the language packages as for an initial installation.

See [“Installing language packages”](#) on page 65.



# Performing a rolling upgrade

This chapter includes the following topics:

- [Performing a rolling upgrade using the installer](#)
- [Performing a rolling upgrade of SFHA using the Web-based installer](#)

## Performing a rolling upgrade using the installer

You can use rolling upgrades to upgrade one product from a release to the next with minimal application downtime.

### About rolling upgrades

You can use rolling upgrades to upgrade one product from a release to the next. Rolling upgrades require less downtime. Rolling upgrades are not compatible with phased upgrades. Do not perform "mixed" rolling upgrades with phased upgrades.

Rolling upgrades take two discrete phases. In the first, you upgrade the product kernel packages. In the second, you upgrade the non-kernel packages such as VCS packages and agent packages.

You can perform a rolling upgrade from 5.1, 5.1 P1, 5.1 RP1, or 5.1 RP2 to 5.1 SP1.

### Prerequisites for a rolling upgrade

Meet the following prerequisites before performing a rolling upgrade:

- Make sure the product that you want to upgrade supports rolling upgrades.
- Split your clusters into sub-clusters for the upgrade to keep the service groups available during upgrade.

- Make sure you logged in as superuser and have the media mounted.
- VCS must be running before performing the rolling upgrade.

## Performing a rolling upgrade on kernel packages: phase 1

Note that in the following instructions that a sub-cluster can represent one or more nodes in a full cluster, but is represented by `nodeA`.

### To perform the rolling upgrade on kernel packages: phase 1

- 1 On the first sub-cluster, start the installer for the rolling upgrade with the `-upgrade_kernelpkgs` option.

```
./installer -upgrade_kernelpkgs nodeA
```

- 2 Note that if the boot disk is encapsulated, then you do not need to perform an unencapsulation for upgrades.
- 3 The installer checks system communications, package versions, product versions, and completes prechecks.  
It then upgrades applicable product kernel packages.
- 4 The installer loads new kernel modules.
- 5 The installer starts all the relevant processes and brings all the service groups online.
- 6 If the boot disk is encapsulated, reboot the first sub-cluster's system.  
Otherwise go to 7.
- 7 Before you proceed to phase 2, complete step 1 to 5 on the second subcluster.

## Performing a rolling upgrade on non-kernel packages: phase 2

In this phase installer installs all non-kernel packages on all the nodes in cluster and restarts VCS cluster.

### To perform the rolling upgrade on non-kernel packages: phase 2

- 1 Start the installer for the rolling upgrade with the `-upgrade_nonkernelpkgs` option. Specify all the nodes in the cluster:

```
./installer -upgrade_nonkernelpkgs nodeA nodeB nodeC...
```

- 2 The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1.
- 3 The installer upgrades non-kernel packages.

- 4 The installer checks system communications, package versions, product versions, and completes prechecks. It verifies completion of phase 1. The installer loads the new kernel modules. It then starts all relevant processes and brings all the service groups online.
- 5 Verify the cluster's status:  

```
hastatus -sum
```
- 6 If you want to upgrade VCS or SFHA 5.1 on the CP server systems to version 5.1 SP1 PR1, make sure you upgraded all application clusters to 5.1 SP1 PR1. Then, upgrade VCS or SFHA on the CP server systems.  
  
For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA installation guide.

## 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.

The rolling upgrade is divided into two phases. In the first phase, the installer upgrade kernel packages. In the second phase, it upgrades non-kernel packages. The second phase is required for upgrades that have high-availability components. When you perform a rolling upgrade, you need to divide the number of systems that you plan to upgrade roughly in half. Half of the systems' available capacity is needed to take over processes during the rolling upgrade.

### To start the rolling upgrade—phase 1

- 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 69.

- 3 In the Task pull-down menu, select **Rolling Upgrade**.

In the Product pull-down menu, select the product that you want to upgrade using a rolling upgrade.

Note that the Upgrade Kernel packages for Rolling Upgrade Phase-1 radio button is selected.

Click the **Next** button to proceed.

- 4 In the Systems Names field, enter the sub-cluster's system names. Separate system names with a space.

The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.

- 5 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.
- 6 When the upgrade completes, perform step 3 through step 6 on the second subcluster.

#### To upgrade the non-kernel components—phase 2

- 1 In the Task pull-down menu, make sure that **Rolling Upgrade** and the product are selected.

Note that the Upgrade Non-Kernel packages for Rolling Upgrade Phase-2 radio button is selected.

Click the **Next** button to proceed.

- 2 In the Systems Names field, enter the names of all the systems that you want to upgrade. Separate system names with a space.

The installer validates systems and stops processes. If it throws an error, address the error and return to the installer.

- 3 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.

# 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 the movement of the service groups from one node to another to minimize the downtime for any particular service group.

Some rough guidelines follow:

- Split the cluster into two subclusters 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.

## 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 from any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.

## Phased upgrade example

In this example, you have 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.

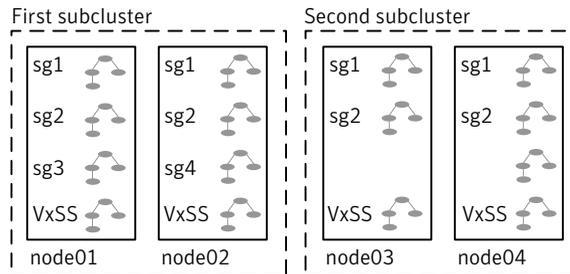
Each service group is running on the nodes as follows:

- 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.
- VxSS service group runs on all nodes (secure mode is enabled)

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.
- VxSS service group runs on all nodes

**Figure 16-1** Example of phased upgrade set up



## Phased upgrade example overview

This example's upgrade path follows:

- Move all the service groups from the first subcluster to the second 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 233.

## 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 and vice versa. 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. The following example illustrates the steps to perform a phased upgrade. The phased upgrade is on a secure cluster.

You can perform a phased upgrade from SFHA 5.1 or other supported previous versions to SFHA 5.1 SP1.

See [“About phased upgrade”](#) on page 231.

See [“Phased upgrade example”](#) on page 232.

## 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.

```
hagrps -state
```

The output resembles the following:

| #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 |
| VxSS   | State     | node01 | ONLINE  |
| VxSS   | State     | node02 | ONLINE  |
| VxSS   | State     | node03 | ONLINE  |
| VxSS   | State     | node04 | ONLINE  |

- 2 Take the parallel service groups (sg1 and sg2) and the VXSS group offline 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).

```
hagrps -offline sg1 -sys node01
hagrps -offline sg2 -sys node01
hagrps -offline sg1 -sys node02
hagrps -offline sg2 -sys node02
hagrps -offline VxSS -sys node01
hagrps -offline VxSS -sys node02
hagrps -switch sg3 -to node03
hagrps -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 -k
```

```
Filesystem kbytes used avail capacity Mounted on
/dev/dsk/c1t0d0s0 66440242 10114415 55661425 16% /
/devices 0 0 0 0% /devices
ctfs 0 0 0 0% /system/contract
proc 0 0 0 0% /proc
mnttab 0 0 0 0% /etc/mnttab
swap 5287408 1400 5286008 1% /etc/svc/volatile
objfs 0 0 0 0% /system/object
sharefs 0 0 0 0% /etc/dfs/sharetab
/platform/sun4u-us3/lib/libc_psr/libc_psr_hwcap1.so.1
66440242 10114415 55661425 16% /platform/sun4u-us3/lib/
libc_psr.so.1
/platform/sun4u-us3/lib/sparcv9/libc_psr/libc_psr_hwcap1.so.1
66440242 10114415 55661425 16% /platform/sun4u-us3/lib/
sparcv9/libc_psr.so.1
fd 0 0 0 0% /dev/fd
swap 5286064 56 5286008 1% /tmp
swap 5286056 48 5286008 1% /var/run
swap 5286008 0 5286008 0% /dev/vx/dmp
swap 5286008 0 5286008 0% /dev/vx/rdmp
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|
VxSS State node01 |OFFLINE|
VxSS State node02 |OFFLINE|
VxSS State node03 |ONLINE|
VxSS 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`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. You want the line in the `main.cf` file to resemble:

```
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 use secure shell or remote shell to connect 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 node1 node2
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 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_5.1SP1.pdf
file present on media? [y,n,q,?] y
```

- 6 Review the available installation options.

- 1 Installs only the minimal required SFHA packages that provides basic functionality of the product.
- 2 Installs the recommended SFHA packages that provides complete functionality of the product.  
Note that this option is the default.
- 3 Installs all the SFHA packages.  
You must choose this option to configure any optional SFHA feature.
- 4 Displays the SFHA packages for each option.

For this example, select **3** for all packages.

```
Select the packages to be installed on all systems? [1-4,q,?]
(2) 3
```

- 7 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility. The installer discovers some warning messages and notes on the systems.
- 8 When you are prompted, reply **y** to continue with the upgrade.

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

- 9 The installer displays the list of packages that get installed or upgraded on the selected systems.
- 10 When you are prompted, reply **y** to stop appropriate processes.

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

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
B VxSS node01 Y N OFFLINE
B VxSS node02 Y N OFFLINE
B VxSS node03 Y N ONLINE
B VxSS node04 Y N ONLINE
```

**2 Unmount all the VxFS file systems that VCS does not manage, for example:**

**# df -k**

```
Filesystem kbytes used avail capacity Mounted on
/dev/dsk/c1t0d0s0 66440242 10114415 55661425 16% /
/devices 0 0 0 0% /devices
ctfs 0 0 0 0% /system/contract
proc 0 0 0 0% /proc
mnttab 0 0 0 0% /etc/mnttab
swap 5287408 1400 5286008 1% /etc/svc/volatile
objfs 0 0 0 0% /system/object
sharefs 0 0 0 0% /etc/dfs/sharetab
/platform/sun4u-us3/lib/libc_psr/libc_psr_hwcap1.so.1
66440242 10114415 55661425 16% /platform/sun4u-us3/
lib/libc_psr.so.1
/platform/sun4u-us3/lib/sparcv9/libc_psr/libc_psr_hwcap1.so.1
66440242 10114415 55661425 16% /platform/sun4u-us3/
lib/sparcv9/libc_psr.so.1
fd 0 0 0 0% /dev/fd
swap 5286064 56 5286008 1% /tmp
swap 5286056 48 5286008 1% /var/run
swap 5286008 0 5286008 0% /dev/vx/dmp
swap 5286008 0 5286008 0% /dev/vx/rdmp
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.

```
hagr -unfreeze sg1 -persistent
hagr -unfreeze sg2 -persistent
hagr -unfreeze sg3 -persistent
hagr -unfreeze sg4 -persistent
hagr -unfreeze VxSS -persistent
```

## 6 Dump the configuration and make it read-only.

```
haconf -dump -makero
```

## 7 Take the service groups offline on node03 and node04.

```
hagr -offline sg1 -sys node03
hagr -offline sg1 -sys node04
hagr -offline sg2 -sys node03
hagr -offline sg2 -sys node04
hagr -offline sg3 -sys node03
hagr -offline sg4 -sys node04
hagr -offline VxSS -sys node03
hagr -offline VxSS -sys node04
```

## 8 Verify the state of the service groups.

```
hagr -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|
VxSS State node01 |OFFLINE|
VxSS State node02 |OFFLINE|
VxSS State node03 |OFFLINE|
VxSS 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`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=disabled
```

- In the `/etc/VRTSvcs/conf/config/main.cf` file, change the value of the `UseFence` attribute from `SCSI3` to `NONE`. You want the line in the `main.cf` file to resemble:

```
UseFence = NONE
```

10 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

- Solaris 9:

```
/opt/VRTSvcs/bin/hastop -local
/etc/init.d/vxfen stop
/etc/init.d/gab stop
/etc/init.d/llt stop
```

- Solaris 10:

```
/opt/VRTSvcs/bin/hastop -local
svcadm disable -t /system/vxfen
svcadm disable -t /system/gab
svcadm disable -t /system/llt
```

11 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 not loaded.

- Solaris 9:

```
/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
```

- Solaris 10:

```
/lib/svc/method/vxfen status
VXFEN module is not loaded

/lib/svc/method/gab status
GAB module is not loaded

/lib/svc/method/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`. You want the line in the `main.cf` file to resemble:

```
UseFence = SCSI3
```

- In the `/etc/vxfenmode` file, change the value of the `vxfen_mode` variable from `disabled` to `scsi3`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=scsi3
```

- 2 Reboot the node01 and node02 in the first subcluster.

```
/usr/sbin/shutdown -y -i6 -g0
```

- 3 Seed node01 and node02 in the first subcluster.

```
gabconfig -xc
```

- 4 Make the configuration writable on the first subcluster.

```
haconf -makerw
```

- 5 Unfreeze the nodes in the first subcluster.

```
hasys -unfreeze -persistent node01
hasys -unfreeze -persistent node02
```

- 6 Dump the configuration and make it read-only.

```
haconf -dump -makero
```

- 7 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
hagrps -online VxSS -sys node01
hagrps -online VxSS -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, disable VCS so that it does not start after reboot. Edit the `vcs` file in `/etc/default`. Open the `vcs` file in an editor, and change the line that reads `VCS_START=1` to `VCS_START=0`. Save and close the file.
- 2 On the second subcluster, disable VXFEN so that it does not start after reboot. Edit the `vxfen` file in `/etc/default`. Open the `vxfen` file in an editor, and change the line that reads `VXFEN_START=1` to `VXFEN_START=0`. Save and close the file.
- 3 On the second subcluster, disable GAB so that it does not start after reboot. Edit the `gab` file in `/etc/default`. Open the `gab` file in an editor, and change the line that reads `GAB_START=1` to `GAB_START=0`. Save and close the file.
- 4 On the second subcluster, disable LLT so that it does not start after reboot. Edit the `llt` file in `/etc/default`. Open the `llt` file in an editor, and change the line that reads `LLT_START=1` to `LLT_START=0`. Save and close the file.

## 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 node3 node4
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 4 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_5.1SP1.pdf
file present on media? [y,n,q,?] y
```

- 5 Review the available installation options.

- 1 Installs only the minimal required SFHA packages that provides basic functionality of the product.

- 2 Installs the recommended SFHA packages that provides complete functionality of the product.

Note that this option is the default.

- 3 Installs all the SFHA packages.

You must choose this option to configure any optional SFHA feature.

- 4 Displays the SFHA packages for each option.

For this example, select **3** for all packages.

```
Select the packages to be installed on all systems? [1-4,q,?]
(2) 3
```

- 6 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility. The installer discovers some warning messages and notes on the systems.
- 7 When you are prompted, reply **y** to continue with the upgrade.  

```
Do you want to continue? [y,n,q] (y)
```
- 8 The installer displays the list of packages to get installed or upgraded on the selected systems.
- 9 When you are prompted, reply **y** to stop appropriate processes.  

```
Do you want to stop SFHA processes now? [y,n,q] (y)
```
- 10 Monitor the installer program answering questions as appropriate until the upgrade completes.

## Finishing the phased upgrade

You now have to reboot the nodes in the second subcluster.

### 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 [-rsh]
-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`. You want the line in the `vxfenmode` file to resemble:

```
vxfen_mode=scsi3
```

- 3 Reboot the node03 and node04 in the second subcluster.

```
/usr/sbin/shutdown -y -i6 -g0
```

The nodes in the second subcluster join the nodes in the first subcluster.

- 4 Check to see if SFHA and High Availabiltiy 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 VxSS node01 Y N ONLINE
B VxSS node02 Y N ONLINE
B VxSS node03 Y N ONLINE
B VxSS node04 Y N ONLINE
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 OFFLINE
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, mount the VxFS file systems and start the VxVM volumes (for each disk group) that VCS does not manage.

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 CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems. For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.

---



# Upgrading with Live Upgrade

This chapter includes the following topics:

- [About Live Upgrade](#)
- [Supported upgrade paths for Live Upgrade](#)
- [Performing Live Upgrade in a Solaris zone environment](#)
- [Before you upgrade SFHA using Solaris Live Upgrade](#)
- [Upgrading SFHA and Solaris using Live Upgrade](#)
- [Upgrading Solaris using Live Upgrade](#)
- [Upgrading SFHA using Live Upgrade](#)
- [Administering boot environments](#)

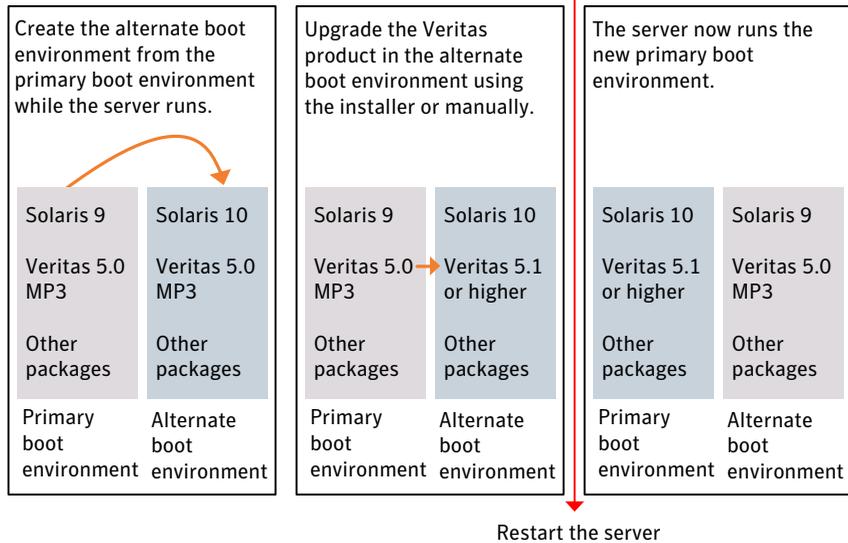
## About Live Upgrade

You can use Live Upgrade to perform the following types of upgrade:

- Upgrade the operating system and SFHA.  
See [“Upgrading SFHA and Solaris using Live Upgrade”](#) on page 261.
- Upgrade the operating system.  
See [“Upgrading Solaris using Live Upgrade”](#) on page 270.
- Upgrade SFHA.  
See [“Upgrading SFHA using Live Upgrade”](#) on page 272.

**Figure 17-1** illustrates an example of an upgrade of Veritas products from 5.0MP3 to 5.1 SP1, and the operating system from Solaris 9 to Solaris 10.

**Figure 17-1** Live Upgrade process



## About Live Upgrade in a Veritas Volume Replicator (VVR) environment

In a SFHA environment that uses Veritas Volume Replicator, the following scripts provide the means to upgrade the VVR configuration:

- `vvr_upgrade_lu_start`
- `vvr_upgrade_lu_finish`

This section provides an overview of the VVR upgrade process. See the Live Upgrade procedures for SFHA for the complete procedure.

See [“Upgrading SFHA and Solaris using Live Upgrade”](#) on page 261.

- Use the `vxlustart` script to perform upgrade steps for SFHA.
- Immediately before rebooting the system to switch over to the alternate boot environment, run the `vvr_upgrade_lu_start` script.

---

**Note:** Use the `vvr_upgrade_lu_start` script only when the applications are stopped and the next step is to switch over to the alternate boot environment.

---

- After the `vvr_upgrade_lu_start` script completes successfully, reboot the system. This reboot results in the system booting from the alternate boot environment.

- After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.

## Supported upgrade paths for Live Upgrade

The systems where you plan to use Live Upgrade must run Solaris 9 or Solaris 10.

For Solaris 10, make sure that all non-global zones are booted and in the running state before you use the Veritas product installer to upgrade the Storage Foundation products in the global zone. If the non-global zones are not mounted and running at the time of the upgrade, you must upgrade each package in each non-global zone manually.

For Live Upgrade, if the alternative root environment also has a zone, you cannot install `VRTSodm`. You must remove the `VRTSodm` package first then install the Storage Foundation product. After you reboot the alternative root, you can install `VRTSodm`.

SFHA version must be at least 4.1.

Symantec requires that both global and non-global zones run the same version of Veritas products.

---

**Note:** If you use Live Upgrade on a system where non-global zones are configured, make sure that all the zones are in the `installed` state before you start Live Upgrade.

---

You can use Live Upgrade in the following virtualized environments:

**Table 17-1** Live Upgrade support in virtualized environments

| Environment          | Procedure                                                                                                                                                                                                                                                                                                                                                                                                             |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Solaris native zones | <p>Perform Live Upgrade to upgrade both global and local zones.</p> <p>If you have a zone root that resides on a VxVM volume, use the following procedure.</p> <p>See <a href="#">“Performing Live Upgrade in a Solaris zone environment”</a> on page 256.</p> <p>Use the standard procedure for the other standby nodes.</p> <p>See <a href="#">“Upgrading SFHA and Solaris using Live Upgrade”</a> on page 261.</p> |

**Table 17-1** Live Upgrade support in virtualized environments (*continued*)

| Environment                    | Procedure                                                                                                                                                                                                                                                                                                                                |
|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Solaris branded zones (BrandZ) | <p>Perform Live Upgrade to upgrade the global zone.<br/>           See <a href="#">“Upgrading SFHA and Solaris using Live Upgrade”</a> on page 261.</p> <p>Manually upgrade the branded zone separately.</p> <p>Note that while you can perform a Live Upgrade in the presence of branded zones, the branded zones are not upgraded.</p> |
| Oracle VM Server for SPARC     | <p>Perform Live Upgrade on the Domain controller only.</p> <p>Perform Live Upgrade on the Guest domain only.</p> <p>Use the standard Live Upgrade procedure for both types of logical domains.</p> <p>See <a href="#">“Upgrading SFHA and Solaris using Live Upgrade”</a> on page 261.</p>                                               |

## Performing Live Upgrade in a Solaris zone environment

If you have a zone root that resides on a VxVM volume, then you must use the following procedure to perform a Live Upgrade on the nodes where zones are online.

Use the standard procedure for the other standby nodes.

See [“Upgrading SFHA and Solaris using Live Upgrade”](#) on page 261.

### To perform a Live Upgrade on a node that has a zone root on a VxVM volume

- 1 Unmount all file systems that do not contain local zone root on shared storage.
- 2 Shut down any application that runs on local zone. Offline its resources and leave only the zone running.

By default, Zone agent `BootState` is set to "multi-user." After you complete the upgrade, you may need to adjust this attribute to the appropriate value before you start your zone through VCS.

---

**Note:** Symantec recommends that you set `BootState` to "multi-user-server" to run applications inside non-global zones.

---

- 3 Freeze the service group that contains the local zone. Note: Make sure that the boot environment disk has enough space for local zone root being copied over during the Live Upgrade.
- 4 Follow the instruction to upgrade using Live Upgrade (which includes vxlustart, the product upgrade, and vxlufinish).

Before rebooting the systems to complete the Live Upgrade, perform the following steps.

- 5 On the system that houses the local zone, copy all files and directories before the upgrade on the local zone root on shared storage to another location.

```
zoneadm list -cv
ID NAME STATUS PATH BRAND IP
 0 global running / native shared
 6 ora-lzone running /oralzones native shared
zoneadm -z ora-lzone halt
cd /oralzones
ls
dev lost+found root SUNWattached.xml
mv dev dev.41
mv root root.41
mv SUNWattached.xml SUNWattached.xml.41
```

- 6 Migrate all files and directories after the upgrade on the local zone root on BE to the shared storage using the tar utility:

```
cd /altroot.5.10/oralzones
ls
dev lost+found lu root SUNWattached.xml
tar cf - . | (cd /oralzones; tar xfbp -)
cd /oralzones
ls
dev .41 lost+found root.41 SUNWattached.xml.41
dev lost+found lu root SUNWattached.xml
```

- 7 Unfreeze the service group that contains the local zone.
- 8 Shut down all systems.

## Before you upgrade SFHA using Solaris Live Upgrade

Before you upgrade, perform the following procedure.

### To prepare for the Live Upgrade

- 1 Make sure that the SFHA installation media and the operating system installation images are available and on hand.
- 2 On the systems to be upgraded, select an alternate boot disk that is at least the same size as the root partition of the primary boot disk. If you are upgrading an SFHA environment, this step is performed on each node in the cluster.

If the primary boot disk is mirrored, you need to break off the mirror for the alternate boot disk.

- 3 Before you perform the Live Upgrade, take offline any services that involve non-root file systems. This prevents file systems from being copied to the alternate boot environment that could potentially cause a root file system to run out of space.
- 4 On the primary boot disk, patch the operating system for Live Upgrade.

---

**Note:** Live Upgrade may fail on a Solaris 9 host if a VxFS file system is in `/etc/vfstab`. On the Solaris 9 host, install the Live Upgrade packages `SUNWlucfg`, `SUNWluu`, and `SUNWlur` from a Solaris 10 image. After you install the packages, install the latest Live Upgrade patch.

---

For more information on required packages and patches, visit the following site and search on "Live Upgrade requirements."

<http://wikis.sun.com>

- 5 The version of the Live Upgrade packages must match the version of the operating system to which you want to upgrade on the alternate boot disk. If you are upgrading the Solaris operating system, do the following steps:
  - Remove the installed Live Upgrade packages for the current operating system version:
    - All Solaris versions: `SUNWluu`, `SUNWlur` packages.
    - Solaris 10 update 5 or later also requires: `SUNWlucfg` package.
    - Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.
  - From the new Solaris installation image, install the new versions of the Live Upgrade packages:
    - All Solaris versions: `SUNWluu`, `SUNWlur` packages.
    - Solaris 10 update 5 or later also requires: `SUNWlucfg` package.
    - Solaris 10 zones or Branded zones also requires: `SUNWluzone` package.

---

**Note:** While you can perform Live Upgrade in the presence of branded zones, they must be halted, and the branded zones themselves are not upgraded.

---

- 6 Symantec provides the `vxlustart` script that runs a series of commands to create the alternate boot disk for the upgrade.

To preview the commands, specify the `vxlustart` script with the `-V` option.

Symantec recommends that you preview the commands to ensure there are no problems before beginning the Live Upgrade process.

The `vxlustart` script is located on the distribution media, in the `scripts` directory.

```
cd /cdrom/scripts

./vxlustart -V -u targetos_version -s osimage_path -d diskname
```

- `-V` Lists the commands to be executed during the upgrade process without executing them and pre-checks the validity of the command.
- If the operating system is being upgraded, the user will be prompted to compare the patches that are installed on the image with the patches installed on the primary boot disk to determine if any critical patches are missing from the new operating system image.
- `-u` Specifies the operating system version for the upgrade on the alternate boot disk. For example, use `5.9` for Solaris 9 and `5.10` for Solaris 10.
- `-U` Specifies that only the Storage Foundation products are upgraded. The operating system is cloned from the primary boot disk.
- `-s` Indicates the path of the operating system image to be installed on the alternate boot disk. If this option is omitted, you are prompted to insert the discs that contain the operating system image.
- If the `-U` option is specified, you can omit the `-s` option. The operating system is cloned from the primary boot disk.
- `-d` Indicates the name of the alternate boot disk on which you intend to upgrade. If you do not specify this option with the script, you are prompted for the disk information.
- `-v` Indicates verbose, the executing commands display before they run.
- `-Y` Indicates a default yes with no questions asked.
- `-D` Prints with debug option on, and is for debugging.
- `-F` Specifies the rootdisk's file system, where the default is `ufs`.
- `-t` Specifies the number of CDs involved in upgrade.

-r Specifies that if the machine crashes or reboots before remounting the alternate disk using this option.

For example, to preview the commands to upgrade only the Veritas product:

```
./vxlustart -V -u 5.10 -U -d disk_name
```

For example, to preview the commands for an upgrade to Solaris 10 update 6:

```
./vxlustart -V -u 5.10 -s /mnt/Solaris_10u6 -d c0t1d0
```

---

**Note:** This command prompts you to compare the patches that are installed on the image with the patches installed on the primary boot disk. If any patches are missing from the new operating system's image, note the patch numbers. To ensure the alternate boot disk is the same as the primary boot disk, you will need to install these patches on the alternate boot disk.

---

- 7 If the specified image is missing patches that are installed on the primary boot disk, note the patch numbers. To ensure that the alternate boot disk is the same as the primary boot disk, you need to install any missing patches on the alternate boot disk.

In the procedure examples, the primary or current boot environment resides on Disk0 (c0t0d0) and the alternate or inactive boot environment resides on Disk1 (c0t1d0).

## Upgrading SFHA and Solaris using Live Upgrade

Perform the Live Upgrade manually or use the installer. For SFHA, the nodes do not form a cluster until all of the nodes are upgraded to Storage Foundation 5.1 SP1. At the end of the Live Upgrade of the last node, all the nodes must boot from the alternate boot environment and join the cluster.

Upgrading SFHA using Live Upgrade involves the following steps:

- Prepare to upgrade using Solaris Live Upgrade.  
See [“Before you upgrade SFHA using Solaris Live Upgrade”](#) on page 257.
- Create a new boot environment on the alternate boot disk.  
See [“Creating a new boot environment on the alternate boot disk”](#) on page 262.
- Upgrade to Storage Foundation 5.1 SP1 on the alternate boot environment manually or using the installer. Refer to one of the following:

To upgrade SFHA manually:

- See “[Upgrading SF manually](#)” on page 265.

To upgrade SFHA using the installer:

- See “[Upgrading SFHA using the installer for a Live Upgrade](#)” on page 263.
- Switch the alternate boot environment to be the new primary.  
See “[Completing the Live Upgrade](#)” on page 268.
- Verify Live Upgrade of SFHA.  
See “[Verifying Live Upgrade of SFHA](#)” on page 269.

## Creating a new boot environment on the alternate boot disk

Run the `vxlustart` command on each system to create a new boot environment on the alternate boot disk. In an HA environment, run the `vxlustart` command on each node in the cluster.

---

**Note:** This step can take several hours to complete. Do not interrupt the session as it may leave the boot environment unstable.

---

At the end of the process:

- The Solaris operating system on the alternate boot disk is upgraded, if you have chosen to upgrade the operating system.
- A new boot environment is created on the alternate boot disk by cloning the primary boot environment.

### To create a new boot environment on the alternate boot disk

- 1 Navigate to the install media for the Symantec products:

```
cd /cdrom/scripts
```

- 2 On each node, run one of the following commands:

To upgrade the operating system, by itself or together with upgrading the Veritas products:

```
./vxlustart -v -u targetos_version \
-s osimage_path -d disk_name
```

To upgrade the Veritas product only:

```
./vxlustart -v -u 5.10 -U -d disk_name
```

The options to the `vxlustart` command are listed in the preupgrade section.

See [“Before you upgrade SFHA using Solaris Live Upgrade”](#) on page 257.

For example, to upgrade to Solaris 10 update 6:

```
./vxlustart -v -u 5.10 -s /mnt/Solaris_10u6
```

- 3 Review the output and note the new mount points. If the system is rebooted before completion of the upgrade or if the mounts become unmounted, you may need to remount the disks.

If you need to remount, run the command:

```
vxlustart -r -u targetos_version -d disk_name
```

- 4 After the alternate boot disk is created, install any operating system patches that are required for the Veritas product installation.

## Upgrading SFHA using the installer for a Live Upgrade

You can use the Veritas product installer to upgrade SFHA as part of the Live Upgrade.

On a node in the cluster, run the installer on the alternate boot disk to upgrade SFHA on all the nodes in the cluster. The program uninstalls the existing version of SFHA on the alternate boot disk during the process.

At the end of the process the following occurs:

- Storage Foundation 5.1 SP1 is installed on the alternate boot disk.

### To perform Live Upgrade of SFHA using the installer

- 1 Insert the product disc with Storage Foundation 5.1 SP1 or access your copy of the software on the network.
- 2 Run the installer script specifying the root path as the alternate boot disk. Do one of the following:
  - For Veritas products that do not have high availability components, enter the following:

```
./installsf -upgrade -rootpath /altroot.5.10
```

- For Veritas products that have high availability components, enter the following:

```
./installsfha -upgrade -rootpath /altroot.5.10
```

If you are upgrading from Solaris 9 to Solaris 10, and you are upgrading SFHA 4.1MP2, 5.0, or 5.0MPx to 5.1SP1, upgrade SFHA using the `installer -upgrade` command.

If you are upgrading from Solaris 9 to Solaris 10 and you are upgrading SFHA from 5.1 or 5.1RPx to 5.1SP1, uninstall and reinstall SFHA.

See [“Removing and reinstalling SFHA using the installer”](#) on page 271.

- 3 Enter the names of the nodes that you want to upgrade to Storage Foundation 5.1 SP1.

---

**Note:** Make sure that the installed version of VxFS uses the disk layout version 6 or later. If you are on a previous disk layout version, upgrade the version before you proceed with the SFHA installation.

---

The installer displays the list of packages to be installed or upgraded on the nodes.

- 4 Press **Return** to continue with the installation.
- 5 Verify that the version of the Veritas packages on the alternate boot disk is 5.1 SP1.

```
pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
pkginfo -R /altroot.5.10 -l VRTSvxvm
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/log`.

## Upgrading SF manually

You can perform a manual upgrade of SFHA using Live Upgrade. On each node, remove and install the appropriate SFHA packages.

At the end of the process the following occurs:

- Storage Foundation 5.1 SP1 is installed on the alternate boot disk.

### To perform Live Upgrade of SF manually

1 Remove the SFHA packages on the alternate boot disk in the following order:

- For Veritas products that do not have high availability components:

```
pkgrm -R /altroot.5.10 \
VRTSmapro VRTSgapms VRTSvxmsa VRTSfasag VRTSfas VRTSvail \
VRTSfsmnd VRTSfssdk VRTSfsman VRTSvrw VRTSweb VRTSjre15 \
VRTSvcsvr VRTSvrpro VRTSddlpr VRTSvdid VRTSalloc VRTSdcli \
VRTSvmpro VRTSvman VRTSfspro VRTSdsa VRTSvxvm VRTSvxfs \
VRTSspt VRTSaa VRTSmh VRTSccg VRTSobgui VRTSob VRTSobc33 \
VRTSat VRTSpbx VRTSicsco VRTSvlic VRTSperl
```

Note that this package list is an example. Full package lists vary from release to release and by product option.

- For Veritas products that have high availability components:

```
pkgrm -R /altroot.5.10 \
VRTSmapro VRTSgapms VRTSvxmsa VRTSfasag VRTSfas VRTSvail \
VRTScmccc VRTScmcs VRTSacclib VRTScssim VRTScscm VRTScscw \
VRTSvcsmn VRTSvcsag VRTSvcsmg VRTSvcs VRTSvxfen VRTSgab \
VRTSllt VRTSfsmnd VRTSfssdk VRTSfsman VRTSvrw VRTSjre15 \
VRTSvcsvr VRTSvrpro VRTSddlpr VRTSvdid VRTSalloc VRTSdcli \
VRTSvmpro VRTSvman VRTSfspro VRTSdsa VRTSvxvm VRTSvxfs \
VRTSspt VRTSaa VRTSmh VRTSccg VRTSobgui VRTSob VRTSobc33 \
VRTSat VRTSpbx VRTSicsco VRTSvlic VRTSperl
```

Note that this package list is an example. Full package lists vary from release to release and by product option.

The `-R` option removes the packages from the root path `/altroot.5.10` on the alternate boot disk.

2 Install the SFHA packages from the 5.1 SP1 pkgs directory. You must install the packages in the following order one at a time to the alternate boot disk using the `pkgadd` command:

- For Veritas products that do not have the high availability components:

```
VRTSvlic.pkg VRTSperl.pkg VRTSspt.pkg VRTSvxvm.pkg VRTSaslapm.pkg
VRTSob.pkg VRTSsfmh.pkg VRTSvxfs.pkg VRTSfssdk.pkg VRTSdbed.pkg
VRTSodm.pkg VRTSat.pkg
```

■ For Veritas products that have high availability components:

```
VRTSvlic.pkg VRTSperl.pkg VRTSspt.pkg VRTSvxvm.pkg VRTSaslapm.pkg
VRTSob.pkg VRTSsfmh.pkg VRTSvxfs.pkg VRTSfssdk.pkg VRTSat.pkg
VRTSllt.pkg VRTSgab.pkg VRTSvxfen.pkg VRTSamf.pkg VRTSvcscs.pkg
VRTScps.pkg VRTSvcscag.pkg VRTSvcsea.pkg VRTSdbed.pkg VRTSodm.pkg
```

For example:

```
pkgadd -R /altroot.5.10 -d package_name.pkg
```

Where you replace *package\_name.pkg* with a package's name, for example VRTSat.pkg.

```
pkgadd -R /altroot.5.10 -d VRTSat.pkg
```

- 3 Copy the SFHA 5.1 SP1 patch files to a temporary location. The patches you copy depend on your Solaris platform. To list the patches for each platform, enter the following:

```
./installer -listpatches
```

To copy the patch files, enter the following:

```
cp patch_name.tar.gz /temporary_location
```

where you replace *patch\_name.tar.gz* with a patch file's tar file name, for example, 143669-01.tar.gz.

```
cp 142629-04.tar.gz /tmp
```

Repeat this step for each patch.

- 4 Unzip and extract the tar files. For example, if you are in the directory where the tar file is located, enter the following:

```
gunzip 142629-04.tar.gz
```

```
tar xvf 142629-04.tar
```

Repeat this step for each patch.

- 5 Install the patches on the alternative boot disk using the `patchadd` command.

```
patchadd -R /altroot.5.10 patch_name
```

For example:

```
patchadd -R /altroot.5.10 142629-04
```

- 6 Verify that the version of the Veritas packages on the alternate boot disk is 5.1 SP1.

```
pkginfo -R /altrootpath -l VRTSpkgname
```

For example:

```
pkginfo -R /altroot.5.10 -l VRTSvxvm
```

- 7 Set the `INSTALL_ROOT_PATH` environment variable to the root path, and then configure a VCS cluster UUID on the alternative root path. Enter the following commands:

```
export INSTALL_ROOT_PATH=/altroot.5.10
/altroot.5.10/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure \
-use_llthost
```

- 8 Confirm that you have created the Universal Unique Identifier for the cluster:

```
/altroot.5.10/opt/VRTSvcs/bin/uuidconfig.pl -clus -display \
-use_llthost
```

- 9 In a zones or branded zones environment, perform the following steps to ensure that all non-global zones contain a universally unique identifier (UUID):

```
zoneadm -z zone1 detach
zoneadm -z zone1 attach
zoneadm -z zone1 boot
zoneadm list -p
0:global:running:/::native:shared
3:zone1:running:/zone1:3770b7b9-f96a-ef34-f4c5-bc125d56ec27:
native:shared
```

For a Solaris environment without zones, run the following command on the alternate root path of any one node in the cluster to configure a unique VCS cluster ID:

```
/mnt/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure -use_llthost
```

The `-use_llthost` option indicates that the `/etc/llthost` file is used to determine the names of the nodes in the cluster. Alternatively, you can specify the node names instead of the file name.

## Completing the Live Upgrade

At the end of the process:

- If the original primary boot disk was encapsulated, the alternate boot disk is encapsulated.
- The alternate boot environment is activated.
- The system is booted from the alternate boot disk.

### To complete the Live Upgrade

- 1 Complete the Live upgrade process using one of the following commands. You must enter the command on all nodes in the cluster.

If the primary root disk is not encapsulated, run the following command:

```
./vxlufinish -u target_os_version
Live Upgrade finish on the Solaris release <5.10>
```

If the primary root disk is encapsulated by VxVM, run the following command:

```
./vxlufinish -u target_os_version -g diskgroup
Live Upgrade finish on the Solaris release <5.10>
```

The Live Upgrade process encapsulates the alternate root disk if the primary root disk was encapsulated.

- 2 If the system crashes or reboots before Live Upgrade completes successfully, you can remount the alternate disk using the following command:

```
./vxlustart -r -u target_os_version
```

Then, rerun the vxlufinish command from step 1

```
./vxlufinish -u target_os_version
```

- 3 If you are upgrading VVR, run the `vvr_upgrade_lu_start` command.

---

**Note:** Only run the `vvr_upgrade_lu_start` command when you are ready to reboot the nodes and switch over to the alternate boot environment.

---

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

---

**Note:** Do not use the `reboot`, `halt`, or `uadmin` commands to reboot the system. Use either the `init` or the `shutdown` commands to enable the system to boot using the alternate boot environment.

---

You can ignore the following error if it appears: ERROR: boot environment <dest.13445> already mounted on </altroot.5.10>.

```
shutdown -g0 -y -i6
```

- 5 After the alternate boot environment is activated, you can switch boot environments. If the root disk is encapsulated, refer to the procedure to switch the boot environments manually.

See “[Administering boot environments](#)” on page 273.

- 6 After the upgrade, perform any required post-upgrade tasks such as upgrading the disk group.
- 7 After the objects are recovered, and the disk group version is upgraded (if desired), run the `vvr_upgrade_lu_finish` script.
- 8 If you want to upgrade CP server systems that use VCS or SFHA to 5.1 SP1, make sure that you upgraded all application clusters to version 5.1 SP1. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the VCS or SFHA Installation Guide.

## Verifying Live Upgrade of SFHA

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

### To verify that Live Upgrade completed successfully

- 1 Verify that the alternate boot environment is active.

```
lustatus
```

If the alternate boot environment is not active, you can revert to the primary boot environment.

See [“Reverting to the primary boot environment”](#) on page 273.

- 2 In a cluster environment, make sure that all the GAB ports are up. Note different ports appear for different products.

```
gabconfig -a
Port a gen c03c01 membership 0
Port h gen c03c03 membership 0
```

- 3 Perform other verification as required to ensure that the new boot environment is configured correctly.
- 4 In a zone environment, verify the zone configuration.

## Upgrading Solaris using Live Upgrade

If you are upgrading Solaris only, you must remove and reinstall SFHA from the alternate boot environment prior to completing the Live Upgrade. You must remove and reinstall because SFHA has kernel components that are specific to Solaris operating system versions. The correct version of the SFHA packages must be installed.

Upgrading Solaris using Live Upgrade involves the following steps:

- Preparing to upgrade using Solaris Live Upgrade.  
See [“Before you upgrade SFHA using Solaris Live Upgrade”](#) on page 257.
- Creating a new boot environment on the alternate boot disk  
See [“Creating a new boot environment on the alternate boot disk”](#) on page 262.
- Removing and reinstalling Storage Foundation 5.1 SP1 on the alternate boot environment, in one of the following ways:  
Using manual steps:  
See [“Upgrading SF manually”](#) on page 265.  
Using the installer:  
See [“Removing and reinstalling SFHA using the installer”](#) on page 271.

---

**Note:** Do NOT configure the Storage Foundation 5.1 SP1

---

- Switching the alternate boot environment to be the new primary  
See “[Completing the Live Upgrade](#)” on page 268.
- Verifying Live Upgrade of SFHA.  
See “[Verifying Live Upgrade of SFHA](#)” on page 269.

## Removing and reinstalling SFHA using the installer

SFHA has kernel components that are specific for Solaris operating system versions. When you use Solaris Live Upgrade to upgrade the Solaris operating system, you must complete these steps to ensure the correct version of SFHA components are installed.

Run the installer on the alternate boot disk to remove and reinstall Storage Foundation 5.1 SP1. In a High Availability environment, you must perform this step on all nodes in the cluster.

At the end of the process the following occurs:

- Storage Foundation 5.1 SP1 is installed on the alternate boot disk, with the correct binaries for the new operating system version

### To remove and reinstall SFHA using the installer

- 1 Insert the product disc with Storage Foundation 5.1 SP1 or access your copy of the software on the network.
- 2 Uninstall using the installer script, specifying the alternate boot disk as the root path:

- For Veritas products that do not have high availability components:

```
/opt/VRTS/install/uninstallsf -rootpath alrootpath
```

- For Veritas products that have high availability components:

```
/opt/VRTS/install/uninstallsfha -rootpath alrootpath
```

- 3 Enter the names of the nodes that you want to uninstall.  
The installer displays the list of packages that will be uninstalled.
- 4 Press **Return** to continue.
- 5 Install using the installer script, specifying the root path as the alternate boot disk as follows:

- For Veritas products that do not have high availability components:

```
/cdrom/storage_foundation/installsf -install \
-rootpath /alrootpath
```

- For Veritas products that have high availability components:

```
/cdrom/storage_foundation_high_availability/installsfha -install \
-rootpath /altrootpath
```

- 6 Press **Return** to continue with the installation.
- 7 Verify that the version of the Veritas packages on the alternate boot disk is 5.1 SP1.

```
pkginfo -R /altroot.5.10 -l VRTSpkgname
```

For example:

```
pkginfo -R /altroot.5.10 -l VRTSvxvm
```

Review the installation logs at `/altroot.5.10/opt/VRTS/install/log`.

## Upgrading SFHA using Live Upgrade

Perform the Live Upgrade manually or use the installer. The nodes will not form a cluster until all of the nodes are upgraded to Storage Foundation 5.1 SP1. At the end of the Live Upgrade of the last node, all the nodes must boot from the alternate boot environment and join the cluster.

Upgrading SFHA using Live Upgrade involves the following steps:

- Prepare to upgrade using Solaris Live Upgrade.  
See [“Before you upgrade SFHA using Solaris Live Upgrade”](#) on page 257.
- Create a new boot environment on the alternate boot disk.  
See [“Creating a new boot environment on the alternate boot disk”](#) on page 262.
- Upgrade to Storage Foundation 5.1 SP1 on the alternate boot environment manually or using the installer. Refer to one of the following:  
To upgrade SFHA manually:
  - See [“Upgrading SF manually”](#) on page 265.To upgrade SFHA using the installer:
  - See [“Upgrading SFHA using the installer for a Live Upgrade”](#) on page 263.
- Switch the alternate boot environment to be the new primary.  
See [“Completing the Live Upgrade ”](#) on page 268.
- Verify Live Upgrade of SFHA.  
See [“Verifying Live Upgrade of SFHA”](#) on page 269.

# Administering boot environments

Use the following procedures to perform relevant administrative tasks for boot environments.

## Reverting to the primary boot environment

If the alternate boot environment fails to start, you can revert to the primary boot environment.

On each node, start the system from the primary boot environment in the PROM monitor mode.

```
ok> boot disk0
```

## Switching the boot environment for Solaris SPARC

You do not have to perform the following procedures to switch the boot environment when you use the `vxlufinish` scripts to process Live Upgrade. You must perform the following procedures when you perform a manual Live Upgrade.

Two different procedures exist to switch the boot environment, choose one of the following procedures based on the encapsulation of the root disk:

- See [“To switch the boot environment if the root disk is not encapsulated”](#) on page 274.
- See [“To switch the boot environment if the root disk is encapsulated”](#) on page 275.

The switching procedures for Solaris SPARC vary, depending on whether VxVM encapsulates the root disk.

**To switch the boot environment if the root disk is not encapsulated**

- 1 Display the status of Live Upgrade boot environments.

```
lustatus

Boot Environment Is Active Active Can Copy
Name Complete Now On Reboot Delete Status

source.2657 yes yes yes no -
dest.2657 yes no no yes -
```

In this example, the primary boot disk is currently (source.2657). You want to activate the alternate boot disk (dest.2657)

- 2 Unmount any file systems that are mounted on the alternate root disk (dest.2657).

```
lufslist dest.2657

boot environment name: dest.2657

Filesystem fstype device size Mounted on Mount Options

/dev/dsk/c0t0d0s1 swap 4298342400 - -
/dev/dsk/c0t0d0s0 ufs 15729328128 / -
/dev/dsk/c0t0d0s5 ufs 8591474688 /var -
/dev/dsk/c0t0d0s3 ufs 5371625472 /vxf -

luumount dest.2657
```

- 3 Activate the Live Upgrade boot environment.

```
luactivate dest.2657
```

- 4 Reboot the system.

```
shutdown -g0 -i6 -y
```

The system automatically selects the boot environment entry that was activated.

### To switch the boot environment if the root disk is encapsulated

#### 1 Display the current boot disk device and device aliases

```
eeprom
boot-device=vx-rootdg vx-int_disk
use-nvramrc?=true
nvramrc=devalias vx-int_disk /pci@1c,600000/scsi@2/disk@0,0:a
devalias vx-rootdg01 /pci@1c,600000/scsi@2/disk@1,0:a
```

#### 2 Set the device from which to boot using the eeprom command. This example shows booting from the primary root disk.

```
eeprom boot-device=vx-rootdg01
```

#### 3 Reboot the system.

```
shutdown -g0 -i6 -y
```

## Switching the boot environment for Solaris x86-64

You do not have to perform the following procedures to switch the boot environment when you use the `vxlufinish` scripts to process Live Upgrade. You must perform the following procedures when you perform a manual Live Upgrade.

Two different procedures exist to switch the boot environment, choose one of the following procedures based on the encapsulation of the root disk:

- See [“To switch the boot environment if root disk is not encapsulated”](#) on page 276.
- See [“To switch the boot environment if root disk is encapsulated”](#) on page 277.

**To switch the boot environment if root disk is not encapsulated**

- 1 Display the status of Live Upgrade boot environments.

```
lustatus

Boot Environment Is Active Active Can Copy
Name Complete Now On Reboot Delete Status

source.2657 yes yes yes no -
dest.2657 yes no no yes -
```

In this example, the primary boot disk is currently (source.2657). You want to activate the alternate boot disk (dest.2657)

- 2 Unmount any file systems that are mounted on the alternate root disk (dest.2657).

```
lufslist dest.2657

boot environment name: dest.2657

Filesystem fstype device size Mounted on Mount Options

/dev/dsk/c0t0d0s1 swap 4298342400 - -
/dev/dsk/c0t0d0s0 ufs 15729328128 / -
/dev/dsk/c0t0d0s5 ufs 8591474688 /var -
/dev/dsk/c0t0d0s3 ufs 5371625472 /vxfst -

luumount dest.2657
```

- 3 Activate the Live Upgrade boot environment.

```
luactivate dest.2657
```

- 4 Reboot the system.

```
shutdown -g0 -i6 -y
```

When the system boots up, the GRUB menu displays the following entries for the Live Upgrade boot environments:

```
source.2657
dest.2657
```

The system automatically selects the boot environment entry that was activated.

### To switch the boot environment if root disk is encapsulated

- ◆ If the root disk is encapsulated, for releases before Solaris 10 update 6 (2.10u6), you can use the `luactivate` method. For Solaris 10 update 6 and subsequent Solaris 10 updates, do one of the following:
  - Select the GRUB entry for the source boot environment or destination boot environment when the system is booted. You can also use the following procedure to manually set the default GRUB menu.lst entry to the source (PBE) or destination (ABE) grub entry:
  - If the system is booted from the alternate boot environment, perform the following steps to switch to the Primary boot environment:

```
mkdir /priroot
mount rootpath /priroot
bootadm list-menu -R /priroot
bootadm list-menu
bootadm set-menu -R /priroot default=PBE_menu_entry
bootadm set-menu default=PBE_menu_entry
shutdown -g0 -i6 -y
```

Where:

**rootpath** is the path to the root device, such as

`/dev/vx/dsk/rootdg/rootvol`

**priroot** is the primary root device

**PBE\_menu\_entry** is the number of the primary boot environment in the GRUB menu.

- If the system is booted from the primary boot environment, perform the following steps to switch to the alternate boot environment:

```
bootadm list-menu
bootadm set-menu default=ABE_menu_entry
ABE booting
```



# Performing post-upgrade tasks

This chapter includes the following topics:

- [Optional configuration steps](#)
- [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](#)
- [Upgrading the Array Support Library](#)
- [Converting from QuickLog to Multi-Volume support](#)
- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [Verifying the Veritas Storage Foundation 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 Volume Manager 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 291.

## Post upgrade tasks for migrating the SFDB repository database

To continue using the checkpoints or tering policies you created with a 5.0x or earlier version of Storage Foundation for Oracle, you must perform one of the following procedures after upgrading SFHA to 5.1 SP1:

- After you upgrade from 5.0.x and before you migrate SFDB:  
See ["After upgrading from 5.0.x and before migrating SFDB"](#) on page 471.
- Migrating from a 5.0 SFDB repository database
- Migrating from a 4.x SFDB repository database
- Upgrading without migrating existing Database Storage Checkpoints and SmartTier parameters

### Migrating from a 5.0 repository database to 5.1 SP1

For clustered environments, perform the following on one node only.

To migrate from a 5.0 repository database to 5.1 SP1

- 1 Rename the startup script NO\_S\*vxdbms3 to S\*vxdbms3.  
See ["After upgrading from 5.0.x and before migrating SFDB"](#) on page 471.
- 2 As root, set the Oracle group permission for various directories used by Oracle.

```
/opt/VRTSdbed/common/bin/sfua_db_config
```

- 3 As root, dump out the old Sybase ASA repository. If you are using SFHA or SF Oracle RAC, you only need to this on one node.

```
/opt/VRTSdbed/migrate/sfua_rept_migrate
```

- 4 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.

For SF, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

For SFHA, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -G \
Oracle_service_group
```

- 5 By default, the repository is created on the filesystem 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.

To create an alternate repository path:

For SF, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -R \
Alternate_path
```

For SFHA, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME \
-G Oracle_service_group -R Alternate_path
```

- 6 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.

If you do not have a repository that you need to migrate from 5.0:

As root, set the Oracle group permission for various directories used by Oracle:

```
/opt/VRTSdbed/common/bin/sfua_db_config
```

- 7** 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.pdx.symantec.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.pdx.symantec.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
```

- 8 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/dbed_vmchecksnap -S $ORACLE_SID \
-H $ORACLE_HOME -f 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*

## Migrating from a 4.x repository database to 5.1 SP1

If you are upgrading Veritas Storage Foundation for Oracle, you can migrate to `/var/vx/vxdba` to save space under the root partition. Migrating to `/var/vx/vxdba` is optional. However, if you do not perform this migration, you cannot remove any file or directory from `/etc/vx/vxdba` to ensure proper operation.

**To migrate from `/etc/vx/vxdba` to `/var/vx/vxdba`**

- 1 Copy the `/etc/vx/vxdba` directory and contents to `/var/vx/vxdba`.

```
cp -rp /etc/vx/vxdba /var/vx/vxdba
```

- 2 Remove `/etc/vx/vxdba`.

```
rm -rf /etc/vx/vxdba
```

- 3 Link the two directories.

```
ln -s /var/vx/vxdba /etc/vx/vxdba
```

### To upgrade the SFDB tools from 4.x to 5.1 SP1

- 1 As root, set Oracle group permission for various directories used by Oracle. For clustered environments, use the following on one node.

```
/opt/VRTSdbed/common/bin/sfua_db_config
```

- 2 On one node, as Oracle user, for each Oracle instance, migrate the old repository data to SQLite repository.

For SF, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME
```

For SFHA, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -G \
Oracle_service_group
```

- 3 By default, the repository is created on the filesystem which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:

- The SFDB 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.

To create an alternate repository path:

For SF, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -R \
Alternate_path
```

For SFHA, on one node, use:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME \
-G Oracle_service_group -R Alternate_path
```

- 4 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=4.0
PRIMARY_HOST=host1
SECONDARY_HOST=host1
PRIMARY_DG=PRODdg
SNAPSHOT_DG=SNAP_PRODdg
ORACLE_SID=PROD
ARCHIVELOG_DEST=/prod_ar
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no

$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=4.0
PRIMARY_HOST=host1
SECONDARY_HOST=host1
PRIMARY_DG=PRODdg
SNAPSHOT_DG_PREFIX=SNAP_PRODdg
ORACLE_SID=PROD
ARCHIVELOG_DEST=/prod_ar
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
```

- 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 steps 1-4 on the secondary host.

If you do not have a repository that you need to migrate from 4.x:

As root, set the Oracle group permission for various directories used by Oracle.

```
/opt/VRTSdbed/common/bin/sfua_db_config
```

- 6 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/dbed_vmchecksnap -S $ORACLE_SID \
-H $ORACLE_HOME -f SNAPPLAN -o validate
```

This completes the migration of the SFDB repository.

To begin using the Storage Foundation for Databases (SFDB) tools:

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

## 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 204.

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.

```
hagrp -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:

```
hagrps -online RVGLogownerGrp -sys masterhost
```

- 7 Start and bring online the failover service groups on the remaining host:

```
hagrps -online groupname -sys nodename
```

- 8 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.

- 9 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 6, Version 7, and Version 8. No prior versions can be created or mounted.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 8.

See the `vxfsconvert` or `vxupgrade` man pages.

For more information about disk layouts, see the *Veritas File System 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 and 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.

After upgrading to Storage Foundation 5.1 SP1, 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.

For 5.1 SP1, the Veritas Volume Manager disk group version is different than in previous VxVM releases. You must upgrade the disk group version if you upgraded from version 5.1 or earlier.

Use the following command to find the version of a disk group:

```
vxdg list diskgroup
```

To upgrade a disk group to the current disk group version, use the following command:

```
vxdg upgrade diskgroup
```

For more information about disk group versions, see the *Veritas Volume Manager 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 Volume Manager Administrator's Guide*.

## Upgrading the Array Support Library

VxVM provides support for new disk arrays in the form of Array Support Library (ASL) software package.

### Adding JBOD support for storage arrays for which there is not an ASL available

If an array is of type A/A-A, A/P or ALUA and a suitable ASL is not available, the array must be claimed as an JBOD of type A/P. This is to prevent path delays and I/O failures arising. As JBODs are assumed to be type A/A by default, you must create appropriate JBOD entries for such arrays.

#### To configure an A/A-A, A/P or ALUA array as a JBOD

- 1 Stop all applications, such as databases, from accessing VxVM volumes that are configured on the array, and unmount all VxFS file systems and checkpoints that are configured on the array.
- 2 Add the array as a JBOD of type A/P:

```
vxddladm addjbod vid=SUN pid=T300 policy=ap
```

- 3 If you have not already done so, upgrade the Storage Foundation or VxVM software to 5.1 SP1. Device discovery will be performed during the upgrade, and the array will be claimed as a JBOD of appropriate type.

If you have already upgraded your system to 5.1 SP1, run the following command to perform device discovery:

```
vxctl enable
```

**4** Verify that the array has been added with the policy set to APdisk:

```
vxddladm listjbod
VID PID Opcode Page Code Page Offset SNO length Policy
=====
SUN T300 18 -1 36 12 APdisk
```

**5** Check that the correct devices are listed for the array:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
APdisk_0 auto:cdsdisk - - online invalid
APdisk_1 auto:cdsdisk - - online invalid
APdisk_2 auto:cdsdisk - - online invalid
...
```

## Unsuppressing DMP for EMC PowerPath disks

This section is only applicable if you are upgrading a system that includes EMC PowerPath disks.

In releases of VxVM before 4.1, a combination of DMP subpaths and the controllers of DMP subpaths were usually suppressed to prevent interference between DMP and the EMC PowerPath multipathing driver. Suppression has the effect of hiding these subpaths and their controllers from DMP, and as a result the disks on these subpaths and controllers cannot be seen by VxVM.

VxVM 4.1 and later releases have the ability to discover EMCpower disks, and configure them as autodiscovered disks that DMP recognizes are under the control of a separate multipathing driver. This has the benefit of allowing such disks to be reconfigured in cluster-shareable disk groups. Before upgrading to VxVM 5.1 SP1, you must remove the suppression of the subpaths and controllers so that DMP can determine the association between EMCpower metadevices and `c#t#d#` disk devices.

In the following scenarios, you may need to unsuppress DMP subpaths and controllers:

- Converting a foreign disk  
See [“Converting a foreign disk to auto:simple”](#) on page 294.
- Converting a defined disk  
See [“Converting a defined disk to auto:simple”](#) on page 296.
- Converting a powervxvm disk  
See [“Converting a powervxvm disk to auto:simple”](#) on page 299.

Because emcpower disks are auto-discovered, the `powervxvm` script should be disabled and removed from the startup script. To remove the `powervxvm` script, use the command:

```
powervxvm remove
```

## Converting a foreign disk to auto:simple

Release 4.0 of VxVM provided the `vxddladm addforeign` command to configure foreign disks with default disk offsets for the private and public regions, and to define them as simple disks. A foreign disk must be manually converted to `auto:simple` format before upgrading to VxVM 5.1 SP1.

If the foreign disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower10c simple fdisk fdg online
...
```

The `vxprint` command is used to display information about the disk group, `fdg`:

```
vxprint
Disk group: fdg
TY NAME ASSOC KSTATE LENGTH PLOFFS STATE TUTILO PUTILO
dg fdg fdg - - - - -
dm fdisk emcpower10c - 17673456 - - - -
...
```

### To convert a foreign disk to `auto:simple` format

- 1 Stop all the volumes in the disk group, and then deport it:

```
vxvol -g fdg stopall
vxdg deport fdg
```

- 2 Use the `vxddladm` command to remove definitions for the foreign devices:

```
vxddladm rmforeign blockpath=/dev/dsk/emcpower10c \
 charpath=/dev/rdisk/emcpower10c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
...
```

- 3 Run the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
/etc/vx/bin/vxprtvtoc -f /tmp/vtoc /dev/rdisk/emcpower10c
```

- 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
/etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/rdisk/emcpower10c
```

```
THE ORIGINAL PARTITIONING IS AS FOLLOWS:
```

```
SLICE TAG FLAGS START SIZE
0 0x0 0x201 0 0
1 0x0 0x200 0 0
2 0x5 0x201 0 17675520
```

```
THE NEW PARTITIONING WILL BE AS FOLLOWS:
```

```
SLICE TAG FLAGS START SIZE
0 0xf 0x201 0 17675520
1 0x0 0x200 0 0
2 0x5 0x201 0 17675520
```

```
DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

- 5 Upgrade to VxVM 5.1 SP1 using the appropriate upgrade procedure.

- 6** After upgrading VxVM, use the `vxdisk list` command to validate the conversion to `auto:simple` format:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower10s2 auto:simple - - online
...
```

To display the physical device that is associated with the metadata, `emcpower10s2`, enter the following command:

```
vxddmpadm getsubpaths dmpnodename=emcpower10s2
```

- 7** Import the disk group and start the volumes:

```
vxdg import fdg
vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower10s2 auto:simple fdisk fdg online
```

## Converting a defined disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (`darec`), and identified as simple disks. If an EMCpower disk is defined with a persistent `darec`, it must be manually converted to `auto:simple` format before upgrading to VxVM 5.1 SP1.

If the defined disk is defined on a slice other than `s2`, you must copy the partition entry for that slice to that for `s0` and change the tag. If the tag of the original slice is changed, the status of the disk is seen as `online:aliased` after the upgrade.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
ls -l /dev/vx/dmp/emcdisk1
lrwxrwxrwx 1 root other 36 Sep 24 17:59 /dev/vx/dmp/emcdisk1->
/dev/dsk/c6t0d11s5
ls -l /dev/vx/rmp/emcdisk1
```

```
lrwxrwxrwx 1 root other 40Sep 24 17:59 /dev/vx/rmdp/emcdisk1->
/dev/dsk/c6t0d11s5
```

Here the fifth partition of `c6t0d11s5` is defined as the persistent disk access record `emcdisk1`.

The `vxdisk list` command can be used to display the EMCpower disks that are known to VxVM:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcdisk1 simple fdisk fdg online
...
```

The `vxprint` command is used to display information about the disk group, `fdg`:

```
vxprint
Disk group: fdg
TY NAME ASSOC KSTATE LENGTH PLOFFS STATE TUTILO PUTILO
dg fdg fdg - - - - -
dm fdisk emcdisk1 - 17673456 - - -
...
```

To convert a disk with a persistent disk access record to `auto:simple` format

- 1 Stop all the volumes in the disk group, and then deport it:

```
vxvol -g fdg stopall
vxdg deport fdg
```

- 2 Use the `vxdisk rm` command to remove the persistent record definitions:

```
vxdisk rm emcdisk1
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
...
```

- 3 Use the `vxprtvtoc` command to retrieve the partition table entry for the device:

```
/etc/vx/bin/vxprtvtoc -f /tmp/hdisk /dev/rdisk/c6t0d11s2
```

- 4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:

```
/etc/vx/bin/vxedvtoc -f /tmp/hdisk /dev/rdisk/c6t0d11s2
```

```
THE ORIGINAL PARTITIONING IS AS FOLLOWS:
```

```
SLICE TAG FLAGS START SIZE
4 0x0 0x200 0 0
5 0x0 0x200 3591000 2100375
6 0x0 0x200 0 0
```

```
THE NEW PARTITIONING WILL BE AS FOLLOWS:
```

```
SLICE TAG FLAGS START SIZE
4 0x0 0x200 0 0
5 0xf 0x200 3591000 2100375
6 0x0 0x200 0 0
```

```
DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :y
```

```
WRITING THE NEW VTOC TO THE DISK #
```

- 5 Upgrade to VxVM 5.1 SP1 using the appropriate upgrade procedure.

- 6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to `auto:simple` format:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower10s2 auto:simple - - online:aliased
...
```

To display the physical device that is associated with the metadvice, `emcpower10s2`, enter the following command:

```
vxddmpadm getsubpaths dmpnodename=emcpower10s2
```

- 7 Import the disk group and start the volumes:

```
vxdg import fdg
vxvol -g fdg startall
```

You can use the `vxdisk list` command to confirm that the disk status is displayed as `online:simple`:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower10s2 auto:simple fdisk fdg online:aliased
```

To allow DMP to receive correct enquiry data, the common Serial Number (C-bit) Symmetrix Director parameter must be set to enabled.

## Converting a powervxvm disk to auto:simple

In VxVM 4.0, and particularly in prior releases, EMCpower disks could be defined by a persistent disk access record (darec) using `powervxvm script`, and identified as simple disks. If an EMCpower disk is used using `powervxvm`, it must be manually converted to `auto:simple` format before upgrading to VxVM 5.1 SP1.

If there are any controllers or devices that are suppressed from VxVM as `powervxvm` requirement, then such controllers/disks must be unsuppressed. This is required for Veritas DMP to determine the association between PowerPath metanodes and their subpaths. After the conversion to `auto:simple` is complete, the `powervxvm script` is no longer useful, and should be disabled from startup script.

The following example is used to illustrate the procedure. The `ls` command shows the mapping of the EMC disks to persistent disk access records:

```
ls -l /dev/vx/rdmp/
crw----- 1 root root 260, 76 Feb 7 02:36 emcpower0c

vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower0c simple ppdisk01 ppdg online

vxprint
Disk group: fdg
TY NAME ASSOC KSTATE LENGTH PLOFFS STATE TUTILO PUTILO
dg ppdg ppdg - - - - - -
dm ppdisk01 emcpower0c - 2094960 - - - -
```

**To convert an EMCpower disk (defined using powervxvm) to auto:simple format**

- 1 Stop all the volumes in the disk group, and then deport it:

```
vxvol -g ppdg stopall
vxdg deport ppdg
```

- 2 Use the `vxdisk rm` command to remove all emcpower disks from VxVM:

```
vxdisk rm emcpower0c
```

If you now run the `vxdisk list` command, the EMCpower disk is no longer displayed:

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
```

- 3 Use the `vxprtvtoc` command to retrieve the partition table entry for this device:

```
/etc/vx/bin/vxprtvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
```

**4 Use the `vxedvtoc` command to modify the partition tag and update the VTOC:**

```
/etc/vx/bin/vxedvtoc -f /tmp/vtoc /dev/vx/rdmp/emcpower0c
THE ORIGINAL PARTITIONING IS AS FOLLOWS:
SLICE TAG FLAGS START SIZE
0 0x0 0x201 0 0
1 0x0 0x200 0 0
2 0x5 0x201 0 17675520

THE NEW PARTITIONING WILL BE AS FOLLOWS:
SLICE TAG FLAGS START SIZE
0 0xf 0x201 0 17675520
1 0x0 0x200 0 0
2 0x5 0x201 0 17675520

DO YOU WANT TO WRITE THIS TO THE DISK ? [Y/N] :Y
WRITING THE NEW VTOC TO THE DISK #
```

**5 Upgrade to VxVM 5.1 SP1 using the appropriate upgrade procedure.**

**6 After upgrading VxVM, use the `vxdisk list` command to validate the conversion to auto:simple format:**

```
vxdisk list
DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower0s2 auto:simple - - online
```

**7 Import the disk group and start the volumes.**

```
vxvg import ppdg
vxvol -g ppdg startall
vxdisk list

DEVICE TYPE DISK GROUP STATUS
c6t0d12s2 auto:sliced - - online
emcpower0s2 auto:simple ppdisk01 ppdg online
```

## Converting from QuickLog to Multi-Volume support

The 4.1 release of the Veritas File System is the last major release to support QuickLog. The Version 6 or Version 7 disk layout does not support QuickLog. The

functionality provided by the Veritas Multi-Volume Support (MVS) feature replaces most of the functionality provided by QuickLog.

The following procedure describes how to convert from QuickLog to MVS. Unlike QuickLog, which allowed logging of up to 31 VxFS file systems to one device, MVS allows intent logging of only one file system per device. Therefore, the following procedure must be performed for each file system that is logged to a QuickLog device if Version 6 or Version 7 disk layout is used.

The QuickLog device did not need to be related to the file system. For MVS, the log volume and the file system volume must be in the same disk group.

### To convert Quicklog to MVS

- 1 Select a QuickLog-enabled file system to convert to MVS and unmount it.

```
umount myfs
```

- 2 Detach one of the QuickLog volumes from the QuickLog device that the file system had been using. This volume will be used as the new intent log volume for the file system.

```
qlodetach -g diskgroup log_vol
```

- 3 Create the volume set.

```
vxvset make myvset myfs_volume
```

- 4 Mount the volume set.

```
mount -F vxfs /dev/vx/dsk/rootdg/myvset /mnt1
```

- 5 Upgrade the volume set's file system to Version 6 or Version 7 disk layout.

For example:

```
vxupgrade -n 6 /mnt1
```

- 6 Add the log volume from step 2 to the volume set.

```
vxvset addvol myvset log_vol
```

- 7 Add the log volume to the file system. The size of the volume must be specified.

```
fsvoladm add /mnt1 log_vol 50m
```

- 8 Move the log to the new volume.

```
fsadm -o logdev=log_vol,logsize=16m /mnt1
```

## About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as Sun Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

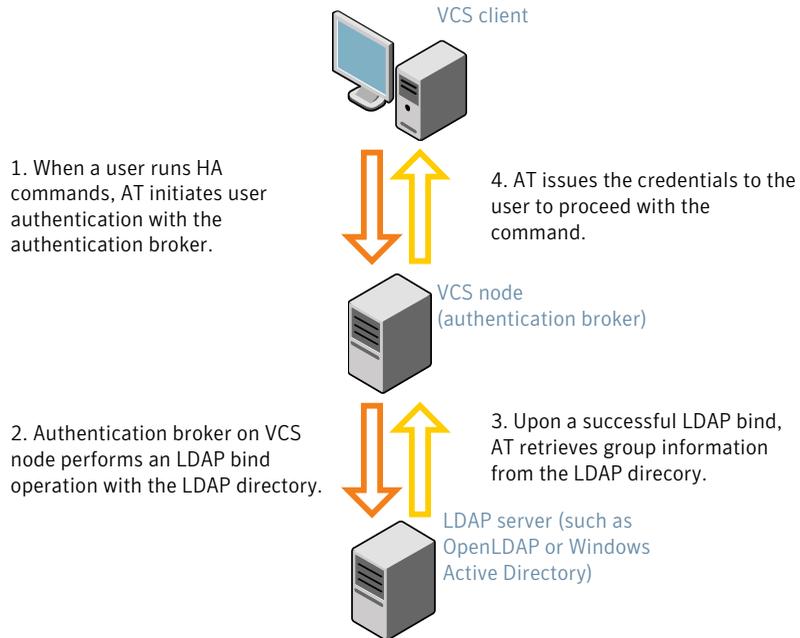
See [“Enabling LDAP authentication for clusters that run in secure mode”](#) on page 305.

If you have not already added VCS users during installation, you can add the users later.

See the *Veritas Cluster Server Administrator's Guide* for instructions to add VCS users.

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

**Figure 18-1** Client communication with LDAP servers



See the *Symantec Product Authentication Service Administrator's Guide*.

The LDAP schema and syntax for LDAP commands (such as `ldapadd`, `ldapmodify`, and `ldapsearch`) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
  - UserObjectClass (the default is `posixAccount`)
  - UserObject Attribute (the default is `uid`)
  - User Group Attribute (the default is `gidNumber`)
  - Group Object Class (the default is `posixGroup`)
  - GroupObject Attribute (the default is `cn`)
  - Group GID Attribute (the default is `gidNumber`)
  - Group Membership Attribute (the default is `memberUid`)
- URL to the LDAP Directory

- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)
- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

## Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 5.0.32.0 or later.

```
/opt/VRTSAt/bin/vssat showversion
vssat version: 5.0.32.0
```

See the `vssat.1m` and the `atldapconf.1m` manual pages.

### To enable OpenLDAP authentication for clusters that run in secure mode

- 1 Add the LDAP domain to the AT configuration using the `vssat` command.

The following example adds the LDAP domain, MYENTERPRISE:

```
/opt/VRTSsat/bin/vssat addldapdomain \
--domainname "MYENTERPRISE.symantecdomain.com"\
--server_url "ldap://my_openldap_host.symantecexample.com"\
--user_base_dn "ou=people,dc=symantecdomain,dc=myenterprise,dc=com"\
--user_attribute "cn" --user_object_class "account"\
--user_gid_attribute "gidNumber"\
--group_base_dn "ou=group,dc=symantecdomain,dc=myenterprise,dc=com"\
--group_attribute "cn" --group_object_class "posixGroup"\
--group_gid_attribute "member"\
--admin_user "cn=manager,dc=symantecdomain,dc=myenterprise,dc=com"\
--admin_user_password "password" --auth_type "FLAT"
```

- 2 Verify that you can successfully authenticate an LDAP user on the 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/VRTSsat/bin/vssat authenticate
--domain ldap:MYENTERPRISE.symantecdomain.com
--prplname vcsadmin1 --broker galaxy:2821
```

```
Enter password for vcsadmin1: #####
```

```
authenticate


```

```
Authenticated User vcsadmin1

```

**3** Add the LDAP user to the main.cf file.

```
haconf makerw
hauser -add "CN=vcsadmin1/CN=people/\
DC=symantecdomain/DC=myenterprise/\
DC=com@myenterprise.symantecdomain.com" -priv Administrator
haconf -dump -makero
```

If you want to enable group-level authentication, you must run the following command:

```
hauser -addpriv \
ldap_group@ldap_domain AdministratorGroup
```

**4** Verify that the main.cf file has the following lines:

```
cat /etc/VRTSvcs/conf/config/main.cf
...
...
cluster clus1 (
 SecureClus = 1
 Administrators = {
 "CN=vcsadmin1/CN=people/DC=symantecdomain/DC=myenterprise/
 DC=com@myenterprise.symantecdomain.com" }
 AdministratorGroups = {
 "CN=symantecusergroups/DC=symantecdomain/DC=myenterprise/
 DC=com@myenterprise.symantecdomain.com " }
)
...
...
```

**5** Set the VCS\_DOMAIN and VCS\_DOMAINTYPE environment variables as follows:

- VCS\_DOMAIN=myenterprise.symantecdomain.com
- VCS\_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
export VCS_DOMAIN=myenterprise.symantecdomain.com
export VCS_DOMAINTYPE=ldap
```

**6** Verify that you can log on to VCS. For example

```
halogin vcsadmin1 password
hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System Attribute Value
galaxy Attribute RUNNING
nebula Attribute RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the 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/VRTSat/bin/atldapconf -d\
-s domain_controller_name_or_ipaddress\
-u domain_user -g domain_group
```

For example:

```
/opt/VRTSat/bin/atldapconf -d -s 192.168.20.32 \
-u Administrator -g "Domain Admins"
Search User provided is invalid or Authentication is required to
proceed further.
Please provide authentication information for LDAP server.
```

```
Username/Common Name: symantecdomain\administrator
Password:
```

Attribute file created.

- 2 Run the LDAP configuration tool `atldapconf` using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
/opt/VRTSat/bin/atldapconf -c -d windows_domain_name
```

For example:

```
/opt/VRTSat/bin/atldapconf -c -d symantecdomain.com
Attribute list file not provided, using default AttributeList.txt.
CLI file name not provided, using default CLI.txt.
```

CLI for `addldapdomain` generated.

- 3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
/opt/VRTSat/bin/atldapconf -x
```

- 4 List the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
/opt/VRTSat/bin/vssat listldapdomains
```

```
Domain Name : symantecdomain.com
Server URL : ldap://192.168.20.32:389
SSL Enabled : No
User Base DN : CN=people,DC=symantecdomain,DC=com
User Object Class : account
User Attribute : cn
User GID Attribute : gidNumber
Group Base DN : CN=group,DC=symantecdomain,DC=com
Group Object Class : group
Group Attribute : cn
Group GID Attribute : cn
Auth Type : FLAT
Admin User :
Admin User Password :
Search Scope : SUB
```

- 5 Set the VCS\_DOMAIN and VCS\_DOMAINTYPE environment variables as follows:

- VCS\_DOMAIN=symantecdomain.com

- VCS\_DOMAINTYPE=ldap

For example, for the Bourne Shell (sh or ksh), run the following commands:

```
export VCS_DOMAIN=symantecdomain.com
export VCS_DOMAINTYPE=ldap
```

- 6 Verify that you can log on to VCS. For example

```
halogin vcsadmin1 password
hasys -state
VCS NOTICE V-16-1-52563 VCS Login:vcsadmin1
#System Attribute Value
galaxy Attribute RUNNING
nebula Attribute RUNNING
```

Similarly, you can use the same LDAP user credentials to log on to the 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 Veritas Storage Foundation upgrade

Refer to the section about verifying the installation to verify the upgrade.

See [“Verifying that the products were installed”](#) on page 316.



# Verification of the installation or the upgrade

- [Chapter 19. Verifying the installation](#)



# Verifying the installation

This chapter includes the following topics:

- [About using the postcheck option](#)
- [Performing a postcheck on a node](#)
- [Verifying that the products were installed](#)
- [Installation log files](#)
- [Starting and stopping processes for the Veritas products](#)
- [Checking Veritas Volume Manager processes](#)
- [Checking Veritas File System installation](#)
- [Verifying the LLT, GAB, and VCS configuration files](#)
- [Verifying LLT, GAB, and cluster operation](#)

## About using the postcheck option

---

**Note:** This command option requires downtime for the node.

---

When you use the `postcheck` option, it returns the results of the following commands for VCS and SFCFS:

- `lltconfig` (to check LLT's status)
- `lltstat -nvv` (to check LLT's status)
- `gabconfig -a` (to check ports a, b, and h)
- `vxfenadm -d` (to check fencing)

- `/opt/VRTS/bin/hasys -state` (to check systems' states)
- `/opt/VRTS/bin/hagrp -state` (to check service groups' states)
- `/opt/VRTS/bin/hares -state` (to check resources' states)

See [“Performing a postcheck on a node”](#) on page 316.

## Performing a postcheck on a node

The installer's `postcheck` command can help you to determine installation-related problems.

See [“About using the postcheck option”](#) on page 315.

---

**Note:** This command option requires downtime for the node.

---

### To run the postcheck command on a node

- ◆ Run the installer with the `-postcheck` option.

```
./installer -postcheck system_name
```

The installer reports some errors or warnings if any of the following issues occur:

- Any processes or drivers do not start
- LLT is not configured
- GAB ports are not started
- Etc.

## Verifying that the products were installed

Verify that the SFHA products are installed.

Use the `pkginfo` command to check which packages have been installed.

```
pkginfo -l VRTSvlic package_name package_name ...
```

Use the following sections to further verify the product installation.

## Installation log files

After every product installation, the installer creates three text files:

- Installation log file
- Response file
- Summary file

The name and location of each file is displayed at the end of a product installation, and are always located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

## Using the installation log file

The installation log file contains all commands executed during the procedure, their output, and errors generated by the commands. This file is for debugging installation problems and can be used for analysis by Veritas Support.

## Using the summary file

The summary file contains the results of the installation by the installer or product installation scripts. The summary includes the list of the packages, and the status (success or failure) of each package. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

# Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

### To stop the processes

- ◆ Use the `-stop` option to stop the product installation script.

For example, to stop the product's processes, enter the following command:

```
./installer -stop
```

### To start the processes

- ◆ Use the `-start` option to start the product installation script.

For example, to start the product's processes, enter the following command:

```
./installer -start
```

## Checking Veritas Volume Manager processes

Use the following procedure to verify that Volume Manager processes are running.

To confirm that key Volume Manager processes are running

- ◆ Type the following command:

```
ps -ef | grep vx
```

Entries for the `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxcached`, and `vxconfigbackupd` processes should appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

## Checking Veritas File System installation

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

### Verifying Veritas File System kernel installation

To ensure that the file system driver is loaded, enter:

```
modinfo | grep vxfs
```

The `modinfo` command displays information about all modules loaded on the system. If the `vxfs` module is loaded, you will see an entry corresponding to `vxfs`. If not, follow the instructions load and then unload the file system module to complete the process.

See “[Loading and unloading the file system module](#)” on page 122.

### Verifying command installation

[Table 19-1](#) lists the directories with Veritas File System commands.

**Table 19-1** VxFS command locations

| Location                          | Contents                                                                                              |
|-----------------------------------|-------------------------------------------------------------------------------------------------------|
| <code>/etc/fs/vxfs</code>         | Contains the Veritas <code>mount</code> command and QuickLog commands required to mount file systems. |
| <code>/usr/lib/fs/vxfs/bin</code> | Contains the VxFS type-specific switch-out commands.                                                  |

**Table 19-1** VxFS command locations (*continued*)

| Location                        | Contents                                                                                            |
|---------------------------------|-----------------------------------------------------------------------------------------------------|
| <code>/opt/VRTSvxfs/sbin</code> | Contains the Veritas-specific commands.                                                             |
| <code>/opt/VRTS/bin</code>      | Contains symbolic links to all Veritas-specific commands installed in the directories listed above. |

Determine whether these subdirectories are present:

```
ls /etc/fs/vxfs
ls /usr/lib/fs/vxfs/bin
ls /opt/VRTSvxfs/sbin
ls /opt/VRTS/bin
```

Make sure you have adjusted the environment variables accordingly.

See [“Setting environment variables”](#) on page 55.

## 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 419.
  - See [“About the VCS configuration files”](#) on page 423.

## 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 On Solaris 9, if you use Sun SCI adapters for your private network, move the scripts `S70llt` and `S92gab` from the directory `/etc/rc2.d` to directory `/etc/rc3.d`, so that they are run after the `S19sci` and `S23scid` scripts.
- 4 Verify LLT operation.  
See “[Verifying LLT](#)” on page 320.
- 5 Verify GAB operation.
- 6 Verify the cluster operation.  
See “[Verifying the cluster](#)” on page 322.

## Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

### To verify LLT

- 1 Log in as superuser on the node galaxy.
- 2 Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node State Links
*0 galaxy OPEN 2
 1 nebula OPEN 2
```

Each node has two links and each node is in the OPEN state. The asterisk (\*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

LLT node information:

| Node       | State | Links |
|------------|-------|-------|
| * 0 galaxy | OPEN  | 2     |
| 1 nebula   | OPEN  | 2     |
| 2 saturn   | OPEN  | 1     |

- 3 Log in as superuser on the node nebula.
- 4 Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

LLT node information:

| Node      | State | Links |
|-----------|-------|-------|
| 0 galaxy  | OPEN  | 2     |
| *1 nebula | OPEN  | 2     |

- 5 To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv active
```

The output on galaxy resembles the following:

■ For Solaris SPARC:

| Node      | State | Link        | Status | Address           |
|-----------|-------|-------------|--------|-------------------|
| *0 galaxy | OPEN  | <i>bge1</i> | UP     | 08:00:20:93:0E:34 |
|           |       | <i>bge2</i> | UP     | 08:00:20:93:0E:34 |
| 1 nebula  | OPEN  | <i>bge1</i> | UP     | 08:00:20:8F:D1:F2 |
|           |       | <i>bge2</i> | DOWN   |                   |

■ For Solaris x64:

| Node      | State | Link            | Status | Address           |
|-----------|-------|-----------------|--------|-------------------|
| *0 galaxy | OPEN  | <i>e1000g:1</i> | UP     | 08:00:20:93:0E:34 |
|           |       | <i>e1000g:2</i> | UP     | 08:00:20:93:0E:38 |
| 1 nebula  | OPEN  |                 |        |                   |

```
e1000g:1 UP 08:00:20:8F:D1:F2
e1000g:2 DOWN
```

The command reports the status on the two active nodes in the cluster, galaxy and nebula.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- A MAC address for each link

However, the output in the example shows different details for the node nebula. The private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

- 6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node galaxy in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
Port Usage Cookie
0 gab 0x0
 opens: 0 2 3 4 5 6 7 8 9 10 11 ... 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

B VxSS galaxy Y N ONLINE
B VxSS nebula Y N ONLINE
```

Note that the VxSS service group is displayed only if you have configured the cluster in secure mode.

- 2 Review the command output for the following information:

- The system state

If the value of the system state is RUNNING, the cluster is successfully started.

## Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for information about the system attributes for VCS.

---

**Note:** The example in the following procedure is for SPARC. x64 clusters have different command output.

---

**To verify the cluster nodes**

- ◆ On one of the nodes, type the `hasys -display` command:

```
hasys -display
```

The example shows the output when the command is run on the node galaxy. The list continues with similar information for nebula (not shown) and any other nodes in the cluster.

| #System | Attribute          | Value                                                                                               |
|---------|--------------------|-----------------------------------------------------------------------------------------------------|
| galaxy  | AgentsStopped      | 0                                                                                                   |
| galaxy  | AvailableCapacity  | 100                                                                                                 |
| galaxy  | CPUBinding         | BindTo None CPUNumber 0                                                                             |
| galaxy  | CPUThresholdLevel  | Critical 90 Warning 80 Note 70<br>Info 60                                                           |
| galaxy  | CPUUsage           | 0                                                                                                   |
| galaxy  | CPUUsageMonitoring | Enabled 0 ActionThreshold 0<br>ActionTimeLimit 0 Action NONE<br>NotifyThreshold 0 NotifyTimeLimit 0 |
| galaxy  | Capacity           | 100                                                                                                 |
| galaxy  | ConfigBlockCount   | 130                                                                                                 |
| galaxy  | ConfigChecksum     | 46688                                                                                               |
| galaxy  | ConfigDiskState    | CURRENT                                                                                             |
| galaxy  | ConfigFile         | /etc/VRTSvcs/conf/config                                                                            |
| galaxy  | ConfigInfoCnt      | 0                                                                                                   |
| galaxy  | ConfigModDate      | Wed 14 Oct 2009 17:22:48                                                                            |
| galaxy  | ConnectorState     | Down                                                                                                |
| galaxy  | CurrentLimits      |                                                                                                     |
| galaxy  | DiskHbStatus       |                                                                                                     |
| galaxy  | DynamicLoad        | 0                                                                                                   |
| galaxy  | EngineRestarted    | 0                                                                                                   |
| galaxy  | EngineVersion      | 5.1.10.0                                                                                            |
| galaxy  | FencingWeight      | 0                                                                                                   |
| galaxy  | Frozen             | 0                                                                                                   |

```

galaxy GUIIPAddr
galaxy HostUtilization CPU 0 Swap 0
galaxy LLTNodeId 0
galaxy LicenseType DEMO
galaxy Limits
galaxy LinkHbStatus
galaxy LoadTimeCounter 0
galaxy LoadTimeThreshold 600
galaxy LoadWarningLevel 80
galaxy NoAutoDisable 0
galaxy NodeId 0
galaxy OnGrpCnt 1
galaxy ShutdownTimeout
galaxy SourceFile ./main.cf
galaxy SwapThresholdLevel Critical 90 Warning 80 Note 70
Info 60
galaxy SysInfo Solaris:galaxy,Generic_
118558-11,5.9,sun4u
galaxy SysName galaxy
galaxy SysState RUNNING
galaxy SystemLocation
galaxy SystemOwner
galaxy TFrozen 0
galaxy TRSE 0
galaxy UpDownState Up
galaxy UserInt 0
galaxy UserStr

```

```
galaxy VCSFeatures DR
galaxy VCMode
```

# Adding and removing nodes

- [Chapter 20. Adding a node to a cluster](#)
- [Chapter 21. 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](#)
- [Preparing to add a node to a cluster](#)
- [Adding a node to a cluster](#)
- [Configuring server-based fencing on the new node](#)
- [After adding the new node](#)
- [Updating the Storage Foundation for Databases \(SFDB\) repository after adding a node](#)

## About adding a node to a cluster

After you install SFHA and create a cluster, you can add and remove nodes from the cluster. You can create clusters of up to 64 nodes.

You can add a node:

- Using the product installer
- Using the Web installer
- Manually

The example procedures describe how to add a node to an existing cluster with two nodes.

## Before adding a node to a cluster

Before preparing to add the node to an existing SFHA cluster, verify the following:

- Hardware and software requirements are met.  
See [“Meeting hardware and software requirements”](#) on page 330.
- Hardware is set up for the new node.  
See [“Setting up the hardware”](#) on page 330.
- The existing cluster is a SFHA cluster and that SFHA is running on the cluster.
- The new system has the same identical operating system versions and patch levels as that of the existing cluster.

## Meeting hardware and software requirements

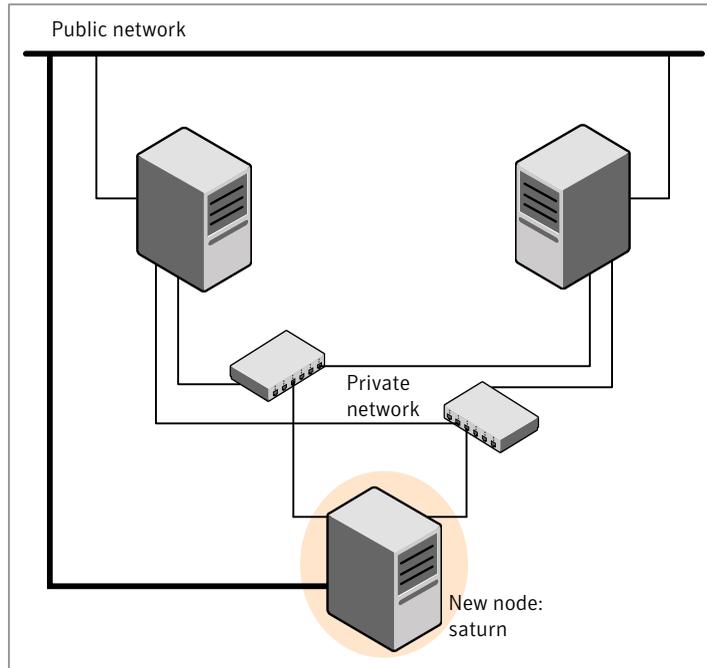
The system you add to the cluster must meet the hardware and software requirements.

See [“Hardware compatibility list \(HCL\)”](#) on page 36.

## Setting up the hardware

[Figure 20-1](#) shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

**Figure 20-1** Adding a node to a two-node cluster using two switches



**To set up the hardware**

- 1 Connect the SFHA private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a 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 20-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.

- 2 Make sure that you meet the following requirements:

- The node must be connected to the same shared storage devices as the existing nodes.
- The node must have private network connections to two independent switches for the cluster.

For more information, see the *Veritas Cluster Server Installation Guide*.

- The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

## Preparing to add a node to a cluster

Complete the following preparatory steps on the new node before you add the node to an existing SFHA cluster.

### To prepare the new node

- 1 Verify that the new node meets installation requirements.

```
./installsfha -precheck saturn
```

- 2 Install SFHA on the new system.

---

**Note:** Use the `-install` option to install SFHA. Do not configure SFHA after the installation.

```
Would you like to configure SFHA on saturn [y, n, q] (n)
```

You can configure the new node later using the configuration from the existing cluster nodes.

---

See [“About installation and configuration methods”](#) on page 32.

## Adding a node to a cluster

You can use one of the following methods to add a node to an existing SFHA cluster:

|                |                                                                                        |
|----------------|----------------------------------------------------------------------------------------|
| SFHA installer | See <a href="#">“Adding a node to a cluster using the SFHA installer”</a> on page 332. |
|                | See <a href="#">“Adding a node using the Web-based installer”</a> on page 336.         |
| Manual         | See <a href="#">“Adding the node to a cluster manually”</a> on page 337.               |

---

**Note:** Before you add the node, make sure that SFHA is not configured on the node.

---

## Adding a node to a cluster using the SFHA installer

You can add a node using the `-addnode` option with the SFHA installer.

The SFHA installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and packages installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:
  - `/etc/llttab`
  - `/etc/VRTSvcs/conf/sysname`
- Updates the following configuration files and copies them on the new node:
  - `/etc/llthosts`
  - `/etc/gabtab`
  - `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files from the existing cluster to the new node
  - `/etc/vxfenmode`
  - `/etc/vxfendg`
  - `/etc/vx/.uuids/clusuid`
  - `/etc/default/llt`
  - `/etc/default/gab`
  - `/etc/default/vxfen`
  - `/etc/`
- Configures security on the new node if the existing cluster is a secure cluster.

---

**Warning:** If the root broker system has failed, then you must recover or reconfigure the root broker system before you add a new node to the cluster.

---

- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.

At the end of the process, the new node joins the SFHA cluster.

---

**Note:** If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

---

### To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the SFHA installer with the `-addnode` option.

```
cd /opt/VRTS/install
./installsfha -addnode
```

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

- 3 Enter the name of a node in the existing SFHA cluster. The installer uses the node information to identify the existing cluster.

```
Enter a node name in the SFHA cluster to which
you want to add a node: galaxy
```

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

```
Enter the system names separated by spaces
to add to the cluster: saturn
```

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

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

---

**Note:** The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

---

```
Enter the NIC for the first private heartbeat
link on saturn: [b,q,?] bge1
```

- 7 Enter **y** to configure a second private heartbeat link.

---

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

---

```
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
```

- 8 Enter the name of the network interface that you want to configure as the second private heartbeat link.

```
Enter the NIC for the second private heartbeat link
on saturn: [b,q,?] bge2
```

- 9 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 10 Review and confirm the information.

- 11 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: bge3
```

- 12** If the existing cluster uses server-based fencing in secure mode, provide responses to the following installer prompts.

If you are using different root brokers for the CP server and the client SFHA cluster, enter **y** to confirm the use of different root brokers. The installer attempts to establish trust between the new node being added to the cluster and the authentication broker of the CP server.

```
Are you using different Root Brokers for the CP Server(s) and the
client cluster? (If so then installer will try to establish trust
between the new node(s) being added and CP Server's
Authentication Broker) [y,n,q] (n) y
```

Enter the host name of the authentication broker used for any one of the CP servers.

```
Enter hostname of the Authentication Broker being used for any one
of the CP Server(s): [b] mycps1.symantecexample.com
```

Enter the port number where the authentication broker for the CP server listens to establish trust with the new node:

```
Enter the port where the Authentication Broker
mycps1.symantecexample.com for the CP Server(s) is listening
for establishing trust: [b] (2821)
```

The installer then starts all the required Veritas processes and joins the new node to cluster.

---

**Note:** Do not quit the installer if you want to perform the Oracle pre-installation tasks using the SFHA installer.

---

- 13** Confirm using `lltstat -n` and `gabconfig -a`.

## 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 In the System Names field enter a name of a node in the cluster where you plan to add the node.  
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.
- 3 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 **Validate** button to check if the system can work in the cluster.  
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.
- 4 From the heartbeat NIC pull-down menus, select the heartbeat NICs for the cluster. Click the **Next** button.
- 5 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 this procedure after you install SFHA only if you plan to add the node to the cluster manually.

### To add the node manually to the cluster

- 1 Start the Volume Manager.  
See [“Starting Volume Manager on the new node”](#) on page 338.
- 2 Configure LLT and GAB.  
See [“Configuring LLT and GAB on the new node”](#) on page 338.
- 3 If the existing cluster is a secure cluster, set up the new node to run in secure mode.  
See [“Setting up the node to run in secure mode”](#) on page 340.

- 4 If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node.  
See “Starting fencing on the new node” on page 343.
- 5 Start VCS.  
See “To start VCS on the new node” on page 347.
- 6 If the ClusterService group is configured on the existing cluster, add the node to the group.  
See “Configuring the ClusterService group for the new node” on page 344.

## Starting Volume Manager on the new node

Volume Manager uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfha` program.

### To start Volume Manager on the new node

- 1 To start Veritas Volume Manager on the new node, use the `vxinstall` utility:

```
vxinstall
```

- 2 Enter **n** when prompted to set up a system wide disk group for the system.  
The installation completes.
- 3 Verify that the daemons are up and running. Enter the command:

```
vxdisk list
```

Make sure the output displays the shared disks without errors.

## Configuring LLT and GAB on the new node

### To configure LLT and GAB on the new node

- 1 Edit the `/etc/llthosts` file on the existing nodes. Using `vi` or another text editor, add the line for the new node to the file. The file resembles:

```
0 galaxy
1 nebula
2 saturn
```

- 2 Copy the `/etc/llthosts` file from one of the existing systems over to the new system. The `/etc/llthosts` file must be identical on all nodes in the cluster.

- 3 Create an `/etc/llttab` file on the new system. For example:

```
set-node saturn
set-cluster 101

link bge1 /dev/bge:1 - ether - -
link bge2 /dev/bge:2 - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 4 Use `vi` or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

```
/sbin/gabconfig -c -nN
```

Where `N` represents the number of systems in the cluster. For a three-system cluster, `N` would equal 3.

- 5 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 6 Use `vi` or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

**7** Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
uuidconfig.pl -rsh -clus -copy \
-from_sys galaxy -to_sys saturn
```

**8** For Solaris 9:

```
/etc/init.d/llt start
/etc/init.d/gab start
/etc/init.d/odm restart
```

For Solaris 10:

```
svcadm enable llt
svcadm enable gab
svcadm restart vxodm
```

## 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 20-1](#) uses the following information for the following command examples.

**Table 20-1** The command examples definitions

| Name   | Fully-qualified host name (FQHN) | Function                                         |
|--------|----------------------------------|--------------------------------------------------|
| saturn | saturn.nodes.example.com         | The new node that you are adding to the cluster. |
| RB1    | RB1.brokers.example.com          | The root broker for the cluster                  |
| RB2    | RB2.brokers.example.com          | Another root broker, not the cluster's RB        |

### To verify the existing security setup on the node

- 1 If node saturn is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Else, proceed to configuring the authentication broker on node saturn.
- 2 Find out the root broker to which the node saturn belongs using the following command.

```
vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
-k "BrokerName"
```

- 3 If the node saturn already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.
- 4 If the node saturn belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node saturn.

- Kill `/opt/VRTSat/bin/vxatd` process.
- Remove the credential that RB2 has given to AB on node saturn.

```
vssat deletecred --domain type:domainname \
--prplname prplname
```

For example:

```
vssat deletecred --domain vx:root@RB2.brokers.example.com \
--prplname saturn.nodes.example.com
```

### Configuring the authentication broker on node saturn

Configure a new authentication broker (AB) on node saturn. This AB belongs to root broker RB1.

**To configure the authentication broker on node saturn**

- 1 Create a principal for node saturn on root broker RB1. Execute the following command on root broker RB1.

```
vssat addprpl --pdrtype root --domain domainname \
 --prplname prplname --password password \
 --prpltype service
```

For example:

```
vssat addprpl --pdrtype root \
 --domain root@RB1.brokers.example.com \
 --prplname saturn.nodes.example.com \
 --password flurbdicate --prpltype service
```

- 2 Ensure that there is no clock skew between the times on node saturn and RB1.
- 3 Copy the `/opt/VRTSat/bin/root_hash` file from RB1 to node saturn.
- 4 Configure AB on node saturn to talk to RB1.

```
vxatd -o -a -n prplname -p password -x vx -y domainname -q \
 rootbroker -z 2821 -h roothash_file_path
```

For example:

```
vxatd -o -a -n saturn.nodes.example.com -p flurbdicate \
 -x vx -y root@RB1.brokers.example.com -q RB1 \
 -z 2821 -h roothash_file_path
```

- 5 Verify that AB is configured properly.

```
vssat showbrokermode
```

The command should return 1, indicating the mode to be AB.

**Setting up SFHA related security configuration**

Perform the following steps to configure SFHA related security settings.

**Setting up SFHA related security configuration**

- 1 Start `/opt/VRTSat/bin/vxatd` process.
- 2 Create `HA_SERVICES` domain for SFHA.

```
vssat createpd --pdrtype ab --domain HA_SERVICES
```

**3** Add SFHA and webserver principal to AB on node saturn.

```
vssat addprpl --pdrttype ab --domain HA_SERVICES --prplname
webserver_VCS_prplname --password new_password --prpltype
service --can_proxy
```

**4** Create `/etc/VRTSvcs/conf/config/.secure` file.

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

## Adding a node in a VxSS group

Perform the following procedure when adding a node in a VxSS group.

### To add a node in the VxSS group using the CLI

**1** Make a backup copy of the `main.cf` file. For example:

```
cd /etc/VRTSvcs/conf/config
cp main.cf main.cf.2node
```

**2** On one of the nodes in the existing cluster, set the cluster configuration to read-write mode:

```
haconf -makerw
```

**3** Add the new node to the VCS configuration:

```
hasys -add saturn
```

**4** Add the node saturn to the existing VxSS group.

```
hagrpl -modify VxSS SystemList -add saturn 2
hagrpl -modify VxSS AutoStartList -add saturn
```

**5** Save the configuration by running the following command from any node in the cluster:

```
haconf -dump -makero
```

## 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 If you are using disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

```
/etc/default/vxfen
/etc/vxfendg
/etc/vxfenmode
```

If you are using pure CP server-based fencing on the existing cluster, then only the `/etc/vxfenmode` file needs to be copied on the new node.

- 2 Start fencing on the new node:

For Solaris 10:

```
svcadm enable vxfen
```

For Solaris 9:

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

### Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.

#### To configure the ClusterService group for the new node

- 1 On an existing node, for example galaxy, write-enable the configuration:

```
haconf -makerw
```

- 2 Add the node saturn to the existing ClusterService group.

```
hagrps -modify ClusterService SystemList -add saturn 2
```

```
hagrps -modify ClusterService AutoStartList -add saturn
```

- 3 Modify the IP address and NIC resource in the existing group for the new node.

```
hares -modify gcoip Device bge0 -sys saturn
hares -modify gconic Device bge0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

```
haconf -dump -makero
```

## Configuring server-based fencing on the new node

Perform this step if your existing cluster uses server-based I/O fencing.

### To configure server-based fencing 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:

```
/opt/VRTScps/bin/cpsadm -s thunderbolt \
-a add_node -c clus1 -u {f0735332-1dd1-11b2} -h saturn -n2
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
/opt/VRTScps/bin/cpsadm -s thunderbolt -a list_nodes
```

The new node must be listed in the output.

- 4 Add the VCS user `cpsclient@saturn` to each CP server:

```
/opt/VRTScps/bin/cpsadm -s thunderbolt \
-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 As the root user, create the VCS user and the domain on the new node:
  - Create a dummy configuration file `/etc/VRTSvcscs/conf/config/main.cf` that resembles the following example:

```
cat main.cf
include "types.cf"
```

**Configuring server-based fencing on the new node**

```
cluster clus1 {
 SecureClus = 1
}
system saturn {
}
```

- Start VCS in one node mode on the new node:

```
/opt/VRTSvcs/bin/hastart -onenode
```

- 2 Verify that the VCS user and the domain are created on the new node:

```
/opt/VRTSvcs/bin/cpsat showcred | grep _HA_VCS_
/opt/VRTSvcs/bin/cpsat listpd -t local | grep HA_SERVICES
```

- 3 Stop VCS if the VCS user and domain are created successfully on the new node:

```
/opt/VRTSvcs/bin/hastop
```

- 4 If the root broker for the CP server and the new node are different, run the following command to establish trust between the authentication broker of the CP Server and the new node:

```
/usr/bin/echo y | /opt/VRTSvcs/bin/cpsat setuptrust \
-b thunderbolt -s high
```

- 5 Log in to each CP server as the root user.
- 6 Update each CP server configuration with the new node information:

```
/opt/VRTSvcs/bin/cpsadm -s thunderbolt \
-a add_node -c clus1 -u {f0735332-1dd1-11b2} -h saturn -n2
Node 2 (saturn) successfully added
```

- 7 Verify that the new node is added to the CP server configuration:

```
/opt/VRTScps/bin/cpsadm -s thunderbolt -a list_nodes
```

The new node must be listed in the output.

- 8 Add the VCS user `_HA_VCS_saturn@HA_SERVICES@saturn.veritas.com` to each CP server:

```
/opt/VRTScps/bin/cpsadm -s thunderbolt \
-a add_user -e _HA_VCS_saturn@HA_SERVICES@saturn.veritas.com \
-f cps_operator -g vx
User _HA_VCS_saturn@HA_SERVICES@saturn.veritas.com successfully added
```

## 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 Checkpoints, Database Flashsnap, or Adding a Node in your configuration, update the SFDB repository to enable access for the new node after it is added to the cluster.

**To update the SFDB repository after adding a node**

- 1 Run the following to change permission, owner, group of various SFDB directories on the newly added node:

```
sfua_db_config
```

- 2 Run the `dbed_update` command on any one node in the cluster. For example:

```
$ dbed_update -S $ORACLE_SID -H $ORACLE_HOME -G $ORACLE_SERVICE_GROUP
```

This completes the addition of the node to the SFDB repository.

For information on using SFDB tools features:

See the Storage Foundation guide: *Storage Foundation: Storage and Availability Management for Oracle Databases*.

# 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 21-1](#) 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 21-1** Tasks that are involved in removing a node

| Task                                                                                                                                                                                | Reference                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------|
| <ul style="list-style-type: none"> <li>■ Back up the configuration file.</li> <li>■ Check the status of the nodes and the service groups.</li> </ul>                                | See <a href="#">“Verifying the status of nodes and service groups”</a> on page 350.     |
| <ul style="list-style-type: none"> <li>■ Switch or remove any SFHA service groups on the node departing the cluster.</li> <li>■ Delete the node from SFHA configuration.</li> </ul> | See <a href="#">“Deleting the departing node from SFHA configuration”</a> on page 351.  |
| Modify the llthosts and gabtab files to reflect the change.                                                                                                                         | See <a href="#">“Modifying configuration files on each remaining node”</a> on page 354. |
| If the existing cluster is configured to use server-based I/O fencing, remove the node configuration from the CP server.                                                            | See <a href="#">“Removing the node configuration from the CP server”</a> on page 354.   |

**Table 21-1** 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 <a href="#">“Removing security credentials from the leaving node”</a> on page 355.          |
| On the node departing the cluster: <ul style="list-style-type: none"><li>■ Modify startup scripts for LLT, GAB, and SFHA to allow reboot of the node without affecting the cluster.</li><li>■ Unconfigure and unload the LLT and GAB utilities.</li><li>■ Remove the SFHA packages.</li></ul> | See <a href="#">“Unloading LLT and GAB and removing VCS on the departing node”</a> on page 355. |

## Verifying the status of nodes and service groups

Start by issuing the following commands from one of the nodes to remain, node galaxy or node nebula.

## 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 the `/etc/llthosts` file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 galaxy
1 nebula
2 saturn
```

To:

```
0 galaxy
1 nebula
```

## Removing the node configuration from the CP server

After removing a node from a SF HA 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 secure mode:

```
cpsadm -s cp_server -a rm_user \
-e _HA_VCS_saturn@HA_SERVICES@saturn.nodes.example.com \
-f cps_operator -g vx
```

For CP server in non-secure mode:

```
cpsadm -s cp_server -a rm_user \
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
cpsadm -s cp_server -a rm_node -h saturn -c clus1 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
cpsadm -s cp_server -a list_nodes
```

## Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node saturn. Perform the following steps.

### To remove the security credentials

- 1 Kill the `/opt/VRTSat/bin/vxatd` process.
- 2 Remove the root credentials on node saturn.

```
vssat deletecred --domain type:domainname --prplname prplname
```

## Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node that is departing the cluster.

If you have configured Storage Foundation and High Availability as part of the Storage Foundation and High Availability products, you may have to delete other dependent packages before you can delete all of the following ones.

### To unconfigure and unload LLT and GAB and remove SFHA

- 1 If you had configured I/O fencing in enabled mode, then stop I/O fencing.

On Solaris 9:

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

On Solaris 10:

```
/lib/svc/method/vxfen stop
```

- 2 Unconfigure GAB and LLT:

```
/sbin/gabconfig -U
```

```
/sbin/lltconfig -U
```

- 3 Unload the GAB and LLT modules from the kernel.

- Determine the kernel module IDs:

```
modinfo | grep gab
```

```
modinfo | grep llt
```

The module IDs are in the left-hand column of the output.

- Unload the module from the kernel:

```
modunload -i gab_id
```

```
modunload -i llt_id
```

- 4 Disable the startup files to prevent LLT, GAB, or SFHA from starting up:

- Solaris 9:

```
/etc/init.d/llt stop
```

```
/etc/init.d/gab stop
```

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

```
/opt/VRTSvcs/bin/hastop
```

- Solaris 10:

```
/usr/sbin/svcadm disable -t llt
/usr/sbin/svcadm disable -t gab
/usr/sbin/svcadm disable -t vcs
```

- 5 To determine the packages to remove, enter:

```
pkginfo | grep VRTS
```

- 6 To permanently remove the SFHA packages from the system, use the `pkgrm` command. Start by removing the following packages, which may have been optionally installed, in the order shown:

```
pkgrm VRTSvcsea
pkgrm VRTSat
pkgrm VRTSvcsag
pkgrm VRTScps
pkgrm VRTSvcs
pkgrm VRTSamf
pkgrm VRTSvxfen
pkgrm VRTSgab
pkgrm VRTSllt
pkgrm VRTSspt
pkgrm VRTSperl
pkgrm VRTSvlic
```

- 7 Remove the LLT and GAB configuration files.

```
rm /etc/llttab
rm /etc/gabtab
rm /etc/llthosts
```

- 8 Remove the language packages and patches.



# Uninstallation of Storage Foundation and High Availability products

- [Chapter 22. Uninstalling Storage Foundation and High Availability products](#)



# Uninstalling Storage Foundation and High Availability products

This chapter includes the following topics:

- [About removing Veritas Storage Foundation](#)
- [Preparing to uninstall](#)
- [Disabling VCS agents for VVR the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling SFHA packages using the script-based installer](#)
- [Uninstalling SFHA with the Veritas Web-based installer](#)
- [Uninstalling Storage Foundation using the pkgrm command](#)
- [Removing the CP server configuration using the removal script](#)
- [Removing the Storage Foundation for Databases \(SFDB\) repository after removing the product](#)

## About removing Veritas Storage Foundation

This section covers uninstallation requirements and steps to uninstall the Veritas software.

Only users with superuser privileges can uninstall Veritas Storage Foundation.

---

**Warning:** Failure to follow the instructions in the following sections may result in unexpected behavior.

---

## Preparing to uninstall

Review the following removing the Veritas software.

### Preparing to remove Veritas Volume Manager

This section describes the steps you need to take before removing Veritas Volume Manager (VxVM) to preserve the contents of the volumes.

---

**Warning:** Failure to follow the preparations in this section might result in unexpected behavior.

---

#### **Moving volumes from an encapsulated root disk**

Use the following procedure to move volumes from an encapsulated root disk.

### To uninstall VxVM if root, swap, usr, or var is a volume under Volume Manager control

- 1 Ensure that the `rootvol`, `swapvol`, `usr`, and `var` volumes have only one associated plex each.

The plex must be contiguous, non-striped, non-spanned, and non-sparse. To obtain this information, enter the following:

```
vxprint -ht rootvol swapvol usr var
```

If any of these volumes have more than one associated plex, remove the unnecessary plexes using the following command:

```
vxplex -o rm dis plex_name
```

- 2 Run the `vxunroot` command:

```
/etc/vx/bin/vxunroot
```

The `vxunroot` command changes the volume entries in `/etc/vfstab` to the underlying disk partitions for `rootvol`, `swapvol`, `usr`, and `var`. It also modifies `/etc/system` and prompts for a reboot so that disk partitions are mounted instead of volumes for `root`, `swap`, `usr`, and `var`.

- 3 Once you have changed the `root`, `swap`, `usr`, and `var` volumes, move all remaining volumes to disk partitions.

You can do this using one of the following procedures:

- Back up the entire system to tape and then recover from tape.
- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Move volumes incrementally to disk partitions.  
See [“Moving volumes to disk partitions”](#) on page 363.  
Otherwise, shut down VxVM.

### Moving volumes to disk partitions

Use the following procedure to move volumes incrementally to disk partitions.

### To move volumes incrementally to disk partitions

- 1 Evacuate disks using `vxdiskadm`, the VEA GUI, or the `vxevac` utility.

Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.

- 2 Remove the evacuated disks from VxVM control by entering:

```
vxdg rmdisk diskname
vxdisk rm devname
```

- 3 Decide which volume to move first, and if the volume is mounted, unmount it.
- 4 If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that you have applied the `sync` command to the data on the volume.
- 5 Create a partition on free disk space of the same size as the volume using the `format` command.

If there is not enough free space for the partition, add a new disk to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this first volume.

- 6 Copy the data on the volume onto the newly created disk partition using a command such as `dd`.

```
dd if=/dev/vx/dsk/diskgroup/lhome of=/dev/dsk/c2t2d2s7
```

where `c2t2d2` is the disk outside of Volume Manager and `s7` is the newly created partition.

- 7 Replace the entry for that volume (if present) in `/etc/vfstab` with an entry for the newly created partition.
- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Stop and remove the volume from VxVM using the commands.

```
vxvol -g diskgroup stop volume_name
vxedit -rf rm volume_name
```

- 10** Remove any free disks (those having no subdisks defined on them) by removing the volumes from VxVM control.

To check if there are still some subdisks remaining on a particular disk, use the `vxprint` command.

```
vxprint -F '%snum' diskname
```

If the output is not 0, there are still some subdisks on this disk that you need to remove. If the output is 0, remove the disk from VxVM control.

```
vxdg rmdisk diskname
vxdisk rm devname
```

Use the free space created for adding the data from the next volume you want to remove.

- 11** After you successfully convert all volumes into disk partitions, reboot the system.
- 12** After the reboot, make sure none of the volumes are open by using the `vxprint` command.

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

- 13** If any volumes remain open, repeat the steps listed above.

## Example of moving volumes to disk partitions on Solaris

This example shows how to move the data on a volume to a disk partition. In the example, there are three disks: `disk1` and `disk2` are subdisks on volume `vol101` and `disk3` is a free disk. The data on `vol101` is copied to `disk3` using `vxevac`.

These are the contents of the disk group `voldg` before the data on `vol101` is copied to `disk3`.

```
vxprint -g voldg -ht
DG NAME NCONFIG NLOG MINORS GROUP-ID
DM NAME DEVICE TYPE PRIVLEN PUBLN STATE
RV NAME RLINK_CNT KSTATE STATE PRIMARY DATAVOLS SRL
RL NAME RVG KSTATE STATE REM_HOST REM_DG REM_RLNK
V NAME RVG KSTATE STATE LENGTH READPOL PREFPLEX UTYPE
PL NAME VOLUME KSTATE STATE LENGTH LAYOUT NCOL/WID MODE
SD NAME PLEX DISK DISKOFFS LENGTH [COL/]OFF DEVICE MODE
SV NAME PLEX VOLNAME NVOLLAYR LENGTH [COL/]OFF AM/NM MODE
DC NAME PARENTVOL LOGVOL
SP NAME SNAPVOL DCO
```

```
dg voldg default default 115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2 sliced 2591 17900352 -
dm disk2 c1t14d0s2 sliced 2591 17899056 -
dm disk3 c1t3d0s2 sliced 2591 17899056 -
```

```
v vol1 - ENABLED ACTIVE 4196448 ROUND - fsgen
pl pl1 vol1 ENABLED ACTIVE 4196448 CONCAT - RW
sd sd1 pl1 disk1 0 2098224 0 c1t12d0 ENA
sd sd2 pl1 disk2 0 2098224 2098224 c1t14d0 ENA
```

Evacuate disk1 to disk3.

```
/etc/vx/bin/vxevac -g voldg disk1 disk3
vxprint -g voldg -ht
```

| DG NAME | NCONFIG   | NLOG    | MINORS   | GROUP-ID |           |          |       |  |
|---------|-----------|---------|----------|----------|-----------|----------|-------|--|
| DM NAME | DEVICE    | TYPE    | PRIVLEN  | PUBLEN   | STATE     |          |       |  |
| RV NAME | RLINK_CNT | KSTATE  | STATE    | PRIMARY  | DATAVOLS  | SRL      |       |  |
| RL NAME | RVG       | KSTATE  | STATE    | REM_HOST | REM_DG    | REM_RLNK |       |  |
| V NAME  | RVG       | KSTATE  | STATE    | LENGTH   | READPOL   | PREFPLEX | UTYPE |  |
| PL NAME | VOLUME    | KSTATE  | STATE    | LENGTH   | LAYOUT    | NCOL/WID | MODE  |  |
| SD NAME | PLEX      | DISK    | DISKOFFS | LENGTH   | [COL/]OFF | DEVICE   | MODE  |  |
| SV NAME | PLEX      | VOLNAME | NVOLLAYR | LENGTH   | [COL/]OFF | AM/NM    | MODE  |  |
| DC NAME | PARENTVOL | LOGVOL  |          |          |           |          |       |  |
| SP NAME | SNAPVOL   | DCO     |          |          |           |          |       |  |

```
dg voldg default default 115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2 sliced 2591 17900352 -
dm disk2 c1t14d0s2 sliced 2591 17899056 -
dm disk3 c1t3d0s2 sliced 2591 17899056 -
```

```
v vol1 - ENABLED ACTIVE 4196448 ROUND - fsgen
pl pl1 vol1 ENABLED ACTIVE 4196448 CONCAT - RW
sd disk3-0111 disk3 0 2098224 0 c1t3d0 ENA
sd sd2 pl1 disk2 0 2098224 2098224 c1t14d0 ENA
```

Evacuate disk2 to disk3.

```
/etc/vx/bin/vxevac -g voldg disk2 disk3
vxprint -g voldg -ht
```

```
DG NAME NCONFIG NLOG MINORS GROUP-ID
DM NAME DEVICE TYPE PRIVLEN PUBLLEN STATE
RV NAME RLINK_CNT KSTATE STATE PRIMARY DATAVOL SRL
RL NAME RVG KSTATE STATE REM_HOST REM_DG REM_RLNK
V NAME RVG KSTATE STATE LENGTH READPOL PREFPLEX UTYPE
PL NAME VOLUME KSTATE STATE LENGTH LAYOUT NCOL/WID MODE
SD NAME PLEX DISK DISKOFFS LENGTH [COL/]OFF DEVICE MODE
SV NAME PLEX VOLNAME NVOLLAYR LENGTH [COL/]OFF AM/NM MODE
DC NAME PARENTVOL LOGVOL
SP NAME SNAPVOL DCO
```

```
dg voldg default default 115000
1017856044.1141.hostname.veritas.com
```

```
dm disk1 c1t12d0s2 sliced 2591 17900352 -
dm disk2 c1t14d0s2 sliced 2591 17899056 -
dm disk3 c1t3d0s2 sliced 2591 17899056 -
```

```
v vol1 - ENABLED ACTIVE 4196448 ROUND - fsgen
pl pl1 vol1 ENABLED ACTIVE 4196448 CONCAT - RW
sd disk3-01 pl1 disk3 0 2098224 0 c1t3d0 ENA
sd disk3-02 pl1 disk3 2098224 2098224 2098224 c1t3d0 ENA
```

Remove the evacuated disks from VxVM control.

```
vxdisk -g voldg list
```

```
DEVICE TYPE DISK GROUP STATUS
c1t3d0s2 sliced disk3 voldg online
c1t12d0s2 sliced disk1 voldg online
c1t14d0s2 sliced disk2 voldg online
```

```
vxdg rmdisk disk1
vxdg rmdisk disk2
vxdisk rm c1t12d0
vxdisk rm c1t14d0
```

Verify that the evacuated disks have been removed from VxVM control.

```
vxdisk -g voldg list
```

```
DEVICE TYPE DISK GROUP STATUS
c1t3d0s2 sliced disk3 voldg online
```

Check to see whether the volume you want to move first is mounted.

```
mount | grep voll
/voll on /dev/vx/dsk/voldg/voll
read/write/setuid/log/nolargefiles/dev=12dc138 on Wed Apr
3 10:13:11 2002
```

Create a partition on free disk space of the same size as the volume. In this example, a 2G partition is created on `disk1 (clt12d0s1)`.

```
format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
 0. c0t0d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
 /sbus@1f,0/SUNW,fas@e,8800000/sd@a,0
 1. clt3d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@3,0
 2. clt9d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@9,0
 3. clt10d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@a,0
 4. clt11d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@b,0
 5. clt12d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@c,0
 6. clt14d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@e,0
 7. clt15d0 <QUANTUM-ATLASIV9SCA-0808 cyl 13814 alt 2 hd 4 sec 324>
 /sbus@1f,0/SUNW,fas@2,8800000/sd@f,0

Specify disk (enter its number): 5
selecting clt12d0
[disk formatted]

FORMAT MENU:
disk - select a disk
type - select (define) a disk type
partition - select (define) a partition table
current - describe the current disk
format - format and analyze the disk
repair - repair a defective sector
label - write label to the disk
analyze - surface analysis
defect - defect list management
```

```

 backup - search for backup labels
 verify - read and display labels
 save - save new disk/partition definitions
 inquiry - show vendor, product and revision
 volname - set 8-character volume name
 !<cmd> - execute <cmd>, then return
 quit
format> p

PARTITION MENU:
 0 - change '0' partition
 1 - change '1' partition
 2 - change '2' partition
 3 - change '3' partition
 4 - change '4' partition
 5 - change '5' partition
 6 - change '6' partition
 7 - change '7' partition
select - select a predefined table
modify - modify a predefined partition table
name - name the current table
print - display the current table
label - write partition map and label to the disk
!<cmd> - execute <cmd>, then return
quit

partition> 1
Part Tag Flag Cylinders Size Blocks
 1 unassigned wm 0 0 (0/0/0) 0
Enter partition id tag[unassigned]:
Enter partition permission flags[wm]:
Enter new starting cyl[0]:
Enter partition size[0b, 0c, 0.00mb, 0.00gb]: 2.00gb
partition> 1
Ready to label disk, continue? y

partition> p
Current partition table (unnamed):
Total disk cylinders available: 13814 + 2 (reserved cylinders)
Part Tag Flag Cylinders Size Blocks
 0 unassigned wm 0 0 (0/0/0) 0
 1 unassigned wm 0 - 3236 2.00GB (3237/0/0) 4195152
partition> q

```

Copy the data on `vol01` to the newly created disk partition.

```
dd if=/dev/vx/dsk/voldg/vol01 of=/dev/dsk/c1t12d0s1
```

In the `/etc/vfstab` file, remove the following entry.

```
/dev/vx/dsk/voldg/vol1 /dev/vx/rdisk/voldg/vol1 /vol1 vxfs 4 yes rw
```

Replace it with an entry for the newly created partition.

```
/dev/dsk/c1t12d0s1 /dev/rdsk/c1t12d0s1 /vol01 vxfs 4 yes rw
```

Mount the disk partition.

```
mount -F vxfs /dev/dsk/c1t12d0s1 /vol01
```

Remove `vol01` from VxVM.

```
vxedit -rf rm /dev/vx/dsk/voldg/vol01
```

To complete the procedure, follow the remaining steps.

## Preparing to remove Veritas File System

The `VRTSvxfs` package cannot be removed if there are any mounted VxFS file systems or Storage Checkpoints. Unmount the VxFS file systems and Storage Checkpoints before uninstalling Veritas Storage Foundation. After you remove the `VRTSvxfs` package, VxFS file systems are not mountable or accessible until another `VRTSvxfs` package is installed.

### To unmount a file system

- 1 Check if any VxFS file systems are mounted.

```
cat /etc/mnttab | grep vxfs
```

- 2 Unmount any file systems.

```
umount special | mount_point
```

Specify the file system to be unmounted as a *mount\_point* or *special* (the device on which the file system resides). See the `umount_vxfs(1M)` manual page for more information about this command and its available options.

You can use the `-a` option to unmount all file systems except `/`, `/usr`, `/usr/kvm`, `/var`, `/proc`, `/dev/fd`, and `/tmp`.

**To unmount a Storage Checkpoint**

- 1 Check if any Storage Checkpoints are mounted.

```
cat /etc/mnttab | grep vxfs
```

- 2 Unmount any Storage Checkpoints.

```
umount /checkpoint_name
```

## Disabling VCS agents for VVR the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

**To disable the agents**

- 1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
hagr -state service_group -sys system_name
```

If none of the service groups is online, skip to [3](#).

- 2 If the service group is online, take it offline.

To take the service group offline without bringing it online on any other system in the cluster, enter:

```
hagr -offline service_group -sys system_name
```

- 3 Stop the agent on the system by entering:

```
haagent -stop agent_name -sys system_name
```

When you get the message Please look for messages in the log file, check the file `/var/VRTSvcs/log/engine_A.log` for a message confirming that each agent has stopped.

You can also use the `ps` command to confirm that the agent is stopped.

- 4 Remove the system from the `SystemList` of the service group. If you disable the agent on all the systems in the `SystemList`, you can also remove the service groups and resource types from the VCS configuration.

Read information on administering VCS from the command line.

Refer to the *Veritas Cluster Server 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.

---

**Note:** If you are upgrading Veritas Volume Replicator, do not remove the Replicated Data Set.

---

## To remove the Replicated Data Set

- 1 Verify that all RLINKs are up-to-date:

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

If the Secondary is not required to be up-to-date, proceed to [2](#) and stop replication using the `-f` option with the `vradmin stoprep` command.

- 2 Stop replication to the Secondary by issuing the following command on any host in the RDS:

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

```
vradmin -g diskgroup stoprep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 3 Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

```
vradmin -g diskgroup delsec local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

- 4 Remove the Primary from the RDS by issuing the following command on the Primary:

```
vradmin -g diskgroup delpri local_rvgname
```

When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

The RDS is removed.

- 5 If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

```
vxedit -r -g diskgroup rm srl_name
```

# Uninstalling SFHA packages using the script-based installer

Use the following procedure to remove SFHA products.

Not all packages may be installed on your system depending on the choices that you made when you installed the software.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 439.

Language packages are uninstalled when you uninstall the English language packages.

## To shut down and remove the installed SFHA packages

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems.

```
umount /mount_point
```

- 3 If the VxVM package (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

See [“Preparing to remove Veritas Volume Manager”](#) on page 362.

- 4 Stop the VEA Service.

```
/opt/VRTS/bin/vxsvcctl stop
```

- 5 Make sure you have performed all of the prerequisite steps.

- 6 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
hastop -local
```

To stop VCS processes on all systems:

```
hastop -all
```

- 7 Move to the `/opt/VRTS/install` directory and run the uninstall script.

```
cd /opt/VRTS/install
```

For Veritas Storage Foundation

```
./uninstallsf
```

For Veritas Storage Foundation High Availability

```
./uninstallsfha
```

- 8 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 from which to
uninstall Storage Foundation: host1
```

- 9 The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the packages are uninstalled.

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

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

- 11 To verify the removal of the packages, use the `pkginfo` command.

```
pkginfo | grep VRTS
```

## Uninstalling SFHA with the Veritas Web-based installer

This section describes how to uninstall Storage Foundation or Storage Foundation High Availability with the Veritas Web-based installer.

### 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 69.
- 3 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.

- 4 Select Storage Foundation or 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 **Validate**.
- 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**.

The Web-based installer prompts you for another task.

## Uninstalling Storage Foundation using the `pkgrm` command

Use the following procedure to uninstall Storage Foundation using the `pkgrm` command.

If you are uninstalling Veritas Storage Foundation using the `pkgrm` command, the packages must be removed in a specific order, or else the uninstallation will fail. Removing the packages out of order will result in some errors, including possible core dumps, although the packages will still be removed.

### To uninstall Storage Foundation

- 1 Unmount all VxFS file systems and Storage Checkpoints, and close all VxVM volumes.

Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems and Storage Checkpoints.

```
umount /mount_point
```

- 3 Stop all applications from accessing VxVM volumes, and close all VxVM volumes.
- 4 Stop various daemons, if applicable.

```
/opt/VRTS/bin/vxsvctrl stop
```

- 5 Remove the packages in the following order:

- For Storage Foundation:

```
pkgrm VRTSsat VRTSodm \
VRTSdbed VRTSfssdk VRTSvxfs VRTSsfmh VRTSob \
VRTSaslapm VRTSvxvm VRTSspt VRTSperl VRTSvlic
```

- For Storage Foundation and High Availability:

```
pkgrm VRTSvlic VRTSperl VRTSspt VRTSob \
VRTSvxvm VRTSaslapm VRTSsfmh VRTSvxfs VRTSfssdk VRTSdbed \
VRTSodm VRTSsat VRTSamf VRTSvcssea VRTSvcsag VRTSgms VRTSglm \
VRTSdbac VRTScps VRTScavf
```

## Uninstalling the language packages using the `pkgrm` command

If you would like to remove only the language packages, you can do so with the `pkgrm` command.

If you use the product installer menu or the uninstallation script, you can remove the language packages along with the English packages.

### To remove the language packages

- 1 Stop the VEA service on each system using the `vxsvcctl stop` command.

```
/opt/VRTS/bin/vxsvcctl stop
```

- 2 Use the `pkgrm` command to remove the appropriate packages.

See “[Chinese language packages](#)” on page 450.

See “[Japanese language packages](#)” on page 451.

```
pkgrm package_name package_name ...
```

Because the packages do not contain any dependencies, you can remove them in any order.

- 3 After removing the appropriate packages, restart the VEA service on each system using the `vxsvcctl start` command.

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

## Removing the CP server configuration using the removal script

This section describes how to remove the CP server configuration from a node or cluster hosting the CP server.

---

**Warning:** Ensure that no SF HA cluster is using the CP server that will have its CP server configuration removed.

---

A configuration utility that is part of VRTScps package is used to remove the CP server configuration. When using the configuration utility, a configuration removal script is run and the following tasks are performed:

- All CP server configuration files are removed
- The VCS configuration for CP server is removed

After running the utility and script, you can then uninstall VCS from the node or cluster.

---

**Note:** The configuration script has to run only once per CP server (which can be on a single node or SFHA cluster), when removing the CP server configuration.

---

The configuration utility performs the following steps to remove the CP server configuration:

- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration

The following procedure describes how to remove the CP server configuration.

**To remove the CP server configuration**

- 1 To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2 The Veritas Coordination Point Server Configuration utility appears with an option menu.

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

Select one of the following:

```
[1] Configure Coordination Point Server on single node VCS system
[2] Configure Coordination Point Server on SFHA cluster
[3] Unconfigure Coordination Point Server
```

- 3 Select option 3 to unconfigure the Coordination Point Server.
- 4 A warning appears and prompts you to confirm the action to unconfigure the Coordination Point Server.

Enter "y" to proceed.

```
WARNING: Unconfiguring Coordination Point Server stops the
vxcpsserv process. VCS clusters using this server for
coordination purpose will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)
(Default:n) :y
```

- 5 After entering "y" to proceed, messages appear informing you of the progress in removing the CP server configuration.

When the CP server configuration has been unconfigured, a success message appears.

For an example of the messages from a single node VCS cluster:

```
A single node VCS cluster is currently configured.
Stopping the CP server ...

Removing the CP Server from VCS configuration..

Removing resource dependencies...
Deleting the resources configured under CPSSG service group...
Deleting the CPSSG service group...

Successfully unconfigured the Veritas Coordination Point Server.
```

For an example of the messages from a CP server on an SFHA cluster:

```
A multinode CP Server cluster is currently configured.
Stopping the CP server ...

Removing the CP Server from VCS configuration..

Removing resource dependencies...
Deleting the resources configured under CPSSG service group...
Deleting the CPSSG service group...

Successfully unconfigured the Veritas Coordination Point Server.
```

- 6 You are then prompted to delete the CP server database. Enter "y" to delete the database. For example:

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

- 7 Enter "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) :
```

- 8 Enter "y" to delete the CP server configuration file and log files. For example:

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 9 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 will disable the SFDB tools.

### To remove the SFDB repository

- 1 Change directories to the location of the local lookup information for the Oracle SID.

For example:

```
cd /var/vx/vxdba/$ORACLE_SID
```

- 2 Identify the SFDB repository file and any associated links:

For example:

```
ls -al
```

```
lrwxrwxrwx 1 oracle oinstall 26 Jul 21 13:58 .sfdb_rept -> \
/ora_data1/TEST/.sfdb_rept
```

```
cd /ora_data1/TEST
```

Follow the symlink of `.sfdb_rept`.

**Removing the Storage Foundation for Databases (SFDB) repository after removing the product**

- 3 Remove the repository directory containing the repository file and all backups.

For example:

```
rm -rf .sfdb_rept
```

- 4 Remove the local lookup directory for the Oracle SID:

```
cd /var/vx/vxdba
```

```
rm -rf $ORACLE_SID
```

This completes the removal of the SFDB repository.

# Installation reference

- [Appendix A. Installation scripts](#)
- [Appendix B. Response files](#)
- [Appendix C. Configuring I/O fencing using a response file](#)
- [Appendix D. Configuration files](#)
- [Appendix E. Configuring the secure shell or the remote shell for communications](#)
- [Appendix F. Storage Foundation and High Availability components](#)
- [Appendix G. Troubleshooting installation issues](#)
- [Appendix H. Troubleshooting cluster installation](#)
- [Appendix I. Sample SF HA cluster setup diagrams for CP server-based I/O fencing](#)
- [Appendix J. Reconciling major/minor numbers for NFS shared disks](#)
- [Appendix K. Configuring LLT over UDP using IPv4](#)
- [Appendix L. Configuring LLT over UDP using IPv6](#)



# Installation scripts

This appendix includes the following topics:

- [About installation scripts](#)
- [Installation script options](#)

## About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.1 SP1 provides several installation scripts.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from an electronic download site, which does not include the installer, use the appropriate product installation script.

The following product installation scripts are available:

|                                                           |                           |
|-----------------------------------------------------------|---------------------------|
| Veritas Cluster Server (VCS)                              | <code>installvcs</code>   |
| Veritas Storage Foundation (SF)                           | <code>installsf</code>    |
| Veritas Storage Foundation and High Availability (SFHA)   | <code>installsfha</code>  |
| Veritas Storage Foundation Cluster File System (SFCFS)    | <code>installsfcfs</code> |
| Veritas Storage Foundation for Oracle RAC (SF Oracle RAC) | <code>installsfrac</code> |
| Symantec Product Authentication Service (AT)              | <code>installat</code>    |
| Veritas Dynamic Multi-pathing                             | <code>installdmp</code>   |

Symantec VirtualStore

`installsvs`

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

## Installation script options

[Table A-1](#) shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts, except where otherwise noted.

See [“About installation scripts”](#) on page 385.

**Table A-1** Available command line options

| Command Line Option       | Function                                                                                                                                                                                                                                          |
|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>system1 system2...</i> | Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.                                                                               |
| <code>-addnode</code>     | Adds a node to a high availability cluster.                                                                                                                                                                                                       |
| <code>-allpkgs</code>     | Displays all packages and patches required for the specified product. The packages and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network. |
| <code>-comcleanup</code>  | The <code>-comcleanup</code> option removes the ssh or rsh configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of ssh or rsh are abruptly terminated.       |
| <code>-configure</code>   | Configures the product after installation.                                                                                                                                                                                                        |

**Table A-1** Available command line options (*continued*)

| Command Line Option                     | Function                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><code>-copyinstallscripts</code></p> | <p>Use this option when you manually install products and want to use the installation scripts that are stored on the system to perform product configuration, uninstallation, and licensing tasks without the product media.</p> <p>Use this option to copy the installation scripts to an alternate rootpath when you use it with the <code>-rootpath</code> option.</p> <p>The following examples demonstrate the usage for this option:</p> <ul style="list-style-type: none"> <li>■ <code>./installer -copyinstallscripts</code><br/> Copies the installation and uninstallation scripts for all products in the release to <code>/opt/VRTS/install</code>. It also copies the installation Perl libraries to <code>/opt/VRTSperl/lib/site_perl/release_name</code>.<br/> .</li> <li>■ <code>./installproduct_name -copyinstallscripts</code><br/> Copies the installation and uninstallation scripts for the specified product and any subset products for the product to <code>/opt/VRTS/install</code>. It also copies the installation Perl libraries to <code>/opt/VRTSperl/lib/site_perl/release_name</code>.<br/> .</li> <li>■ <code>./installer -copyinstallscripts -rootpath alt_root_path</code><br/> The path <i>alt_root_path</i> can be a directory like <code>/rdisk2</code>. In that case, this command copies installation and uninstallation scripts for all the products in the release to <code>/rdisk2/opt/VRTS/install</code>. CPI perl libraries are copied to <code>/rdisk2/opt/VRTSperl/lib/site_perl/release_name</code>, where the <i>release_name</i> is a string that starts with UXRT and includes the release version with no punctuation.</li> </ul> |
| <p><code>-fencing</code></p>            | <p>Configures I/O fencing in a running cluster.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

**Table A-1** Available command line options (*continued*)

| Command Line Option                | Function                                                                                                                                                                                                                |
|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -hostfile <i>full_path_to_file</i> | Specifies the location of a file that contains a list of hostnames on which to install.                                                                                                                                 |
| -ignorepatchreqs                   | The -ignorepatchreqs option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.                                                                           |
| -install                           | The -install option is used to install products on systems.                                                                                                                                                             |
| -installallpkgs                    | Specifies that all packages are installed.                                                                                                                                                                              |
| -installminpkgs                    | Specifies that the minimum package set is installed.                                                                                                                                                                    |
| -installrecpkgs                    | Specifies that the required package set is installed.                                                                                                                                                                   |
| -jumpstart <i>dir_path</i>         | Produces a sample finish file for Solaris JumpStart installation. The <i>dir_path</i> indicates the path to the directory in which to create the finish file.                                                           |
| -keyfile <i>ssh_key_file</i>       | Specifies a key file for secure shell (SSH) installs. This option passes -i <i>ssh_key_file</i> to every SSH invocation.                                                                                                |
| -license                           | Registers or updates product licenses on the specified systems.                                                                                                                                                         |
| -listpatches                       | The -listpatches option displays product patches in correct installation order.                                                                                                                                         |
| -logpath <i>log_path</i>           | Specifies a directory other than /opt/VRTS/install/logs as the location where installer log files, summary files, and response files are saved.                                                                         |
| -makeresponsefile                  | The -makeresponsefile generates a response file without doing an actual installation. Text displaying install, uninstall, start, and stop actions are simulations. These actions are not being performed on the system. |

**Table A-1** Available command line options (*continued*)

| Command Line Option          | Function                                                                                                                                                                                                                                                                                                                     |
|------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -minpkgs                     | Displays the minimal packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option. |
| -patchpath <i>patch_path</i> | Designates the path of a directory that contains all patches to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.                                                                                                                                    |
| -pkginfo                     | Displays a list of packages and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the <code>-pkginfo</code> option with the <code>installvcs</code> script to display VCS packages.                                            |
| -pkgpath <i>package_path</i> | Designates the path of a directory that contains all packages to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.                                                                                                                                   |
| -pkgset                      | Discovers and displays the package group (minimum, recommended, all) and packages that are installed on the specified systems.                                                                                                                                                                                               |
| -pkgtable                    | Displays product's packages in correct installation order by group.                                                                                                                                                                                                                                                          |
| -postcheck                   | Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.                                                                                                                                                                      |
| -precheck                    | Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.                                                                                                                                                               |

**Table A-1** Available command line options (*continued*)

| Command Line Option                | Function                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -recpkgs                           | Displays the recommended packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <a href="#">allpkgs option</a> .       |
| -redirect                          | Displays progress details without showing the progress bar.                                                                                                                                                                                                                                                                                |
| -requirements                      | The <code>-requirements</code> option displays required OS version, required patches, file system space, and other system requirements in order to install the product.                                                                                                                                                                    |
| -responsefile <i>response_file</i> | Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file. |
| -rootpath <i>root_path</i>         | Specifies an alternative root directory on which to install packages.<br><br>On Solaris operating systems, <code>-rootpath</code> passes <code>-R path</code> to <code>pkgadd</code> command.                                                                                                                                              |
| -rsh                               | Specify this option when you want to use <code>rsh</code> and <code>rcp</code> for communication between systems instead of the default <code>ssh</code> and <code>scp</code> .<br><br>See <a href="#">“About configuring secure shell or remote shell communication modes before installing products”</a> on page 439.                    |

**Table A-1** Available command line options (*continued*)

| Command Line Option      | Function                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -security                | <p>Enable or disable Symantec Product Authentication Service in a VCS cluster that is running.</p> <p>You can specify this option with the <code>installvcs</code>, <code>installsfha</code> or <code>installsfdfs</code> scripts.</p> <p>For more information about Symantec Product Authentication Service in a VCS cluster, see the <i>Veritas Cluster Server Installation Guide</i>.</p> |
| -serial                  | <p>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.</p>                                                                                                                                                     |
| -start                   | <p>Starts the daemons and processes for the specified product.</p>                                                                                                                                                                                                                                                                                                                           |
| -stop                    | <p>Stops the daemons and processes for the specified product.</p>                                                                                                                                                                                                                                                                                                                            |
| -tmppath <i>tmp_path</i> | <p>Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.</p>                                                                                                                                             |
| -uninstall               | <p>The <code>-uninstall</code> option is used to uninstall products from systems.</p>                                                                                                                                                                                                                                                                                                        |
| -upgrade                 | <p>Specifies that an existing version of the product exists and you plan to upgrade it.</p>                                                                                                                                                                                                                                                                                                  |
| -upgrade_kernelpkgs      | <p>The <code>-upgrade_kernelpkgs</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel packages get upgraded to the latest version</p>                                                                                                                                                                                                                  |
| -upgrade_nonkernelpkgs   | <p>The <code>-upgrade_nonkernelpkgs</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent packages upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version."</p>                                                                                                                                    |

**Table A-1** Available command line options (*continued*)

| Command Line Option | Function                                                                                                                                                                                                                                                                     |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| -version            | Checks and reports the installed products and their versions. Identifies the installed and missing packages and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing packages and patches where applicable. |

# Response files

This appendix includes the following topics:

- [About response files](#)
- [Installing Storage Foundation or Storage Foundation and High Availability using response files](#)
- [Configuring SFHA using response files](#)
- [Upgrading Storage Foundation or Storage Foundation and High Availability using response files](#)
- [Uninstalling Storage Foundation or Storage Foundation and High Availability using response files](#)
- [Syntax in the response file](#)
- [Response file variables to install, upgrade, or uninstall Storage Foundation or Storage Foundation and High Availability](#)
- [Response file variables to configure SFHA](#)
- [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 386.

## Installing Storage Foundation or Storage Foundation and High Availability 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 Storage Foundation or Storage Foundation and High Availability 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
./installsf -responsefile /tmp/response_file
./installsfha -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

## Configuring SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA configuration on one cluster to configure SFHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

### To configure SFHA using response files

- 1 Make sure the SFHA packages are installed on the systems where you want to configure SFHA.
- 2 Copy the response file to one of the cluster systems where you want to configure SFHA.
- 3 Edit the values of the response file variables as necessary.

To configure optional features, you must define appropriate values for all the response file variables that are related to the optional feature.

See “[Response file variables to configure SFHA](#)” on page 400.

- 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 Storage Foundation or Storage Foundation and High Availability using response files

Typically, you can use the response file that the installer generates after you perform SFHA upgrade on one cluster to upgrade SFHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

### To perform automated Storage Foundation or Storage Foundation High Availability 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 cluster systems where you want to upgrade SFHA.
- 4 Edit the values of the response file variables as necessary.

- 5 Mount the product disk, 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
./installsf -responsefile /tmp/response_file
./installsfha -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

## Uninstalling Storage Foundation or Storage Foundation and High Availability 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 one of the cluster systems 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/uninstallsf -responsefile /tmp/response_file
/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 or Storage Foundation and High Availability

[Table B-1](#) lists the response file variables that you can define to configure Storage Foundation or Storage Foundation and High Availability.

**Table B-1** Response file variables to specific installing, upgrading, or uninstalling Storage Foundation or Storage Foundation and High Availability

| Variable              | Description                                                                                                                                                                          |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{install}     | Installs SFHA packages. Configuration can be performed at a later time using the <code>-configure</code> option.<br><br>List or scalar: scalar<br><br>Optional or required: optional |
| CFG{accepteula}       | Specifies whether you agree with the <code>EULA.pdf</code> file on the media.<br><br>List or scalar: scalar<br><br>Optional or required: required                                    |
| \$CFG{opt}{vxkeyless} | Installs the product with keyless license.<br><br>List of scalar: scalar<br><br>Optional or required: optional                                                                       |
| CFG{systems}          | List of systems on which the product is to be installed or uninstalled.<br><br>List or scalar: list<br><br>Optional or required: required                                            |

**Table B-1** Response file variables to specific installing, upgrading, or uninstalling Storage Foundation or Storage Foundation and High Availability (*continued*)

| Variable            | Description                                                                                                                                                                                                                                 |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{systemscfs}     | <p>List of systems for configuration if secure environment prevents the installer to install SFHA on all systems at once.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>                                             |
| CFG{prod}           | <p>Defines the product to be installed or uninstalled.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>                                                                                                              |
| CFG{opt}{keyfile}   | <p>Defines the location of an ssh keyfile that is used to communicate with all remote systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                      |
| CFG{at_rootdomain}  | <p>Defines the name of the system where the root broker is installed.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                                               |
| CFG{opt}{patchpath} | <p>Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>  |
| CFG{opt}{pkgpath}   | <p>Defines a location, typically an NFS mount, from which all remote systems can install product packages. The location must be accessible from all target systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p> |

**Table B-1** Response file variables to specific installing, upgrading, or uninstalling Storage Foundation or Storage Foundation and High Availability (*continued*)

| Variable                    | Description                                                                                                                                                                                                                                                       |
|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{tmppath}           | <p>Defines the location where a working directory is created to store temporary files and the packages that are needed during the install. The default location is <code>/var/tmp</code>.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p> |
| CFG{opt}{rsh}               | <p>Defines that <code>rsh</code> must be used instead of <code>ssh</code> as the communication method between systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>                                                                    |
| CFG{donotinstall} {package} | <p>Instructs the installation to not install the optional packages in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>                                                                                                             |
| CFG{donotremove} {package}  | <p>Instructs the uninstallation to not remove the optional packages in the list.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>                                                                                                            |
| CFG{opt}{logpath}           | <p>Mentions the location where the log files are to be copied. The default location is <code>/opt/VRTS/install/logs</code>.</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>                                                                                                                                                                |
| CFG{opt}{upgrade}           | <p>Upgrades all packages installed, without configuration.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>                                                                                                                                  |

**Table B-1** Response file variables to specific installing, upgrading, or uninstalling Storage Foundation or Storage Foundation and High Availability (*continued*)

| Variable            | Description                                                                           |
|---------------------|---------------------------------------------------------------------------------------|
| CFG{opt}{uninstall} | Uninstalls SFHA packages.<br>List or scalar: scalar<br>Optional or required: optional |

## Response file variables to configure SFHA

[Table B-2](#) lists the response file variables that you can define to configure SFHA.

**Table B-2** Response file variables specific to configuring SFHA

| Variable            | List or Scalar | Description                                                                                                                       |
|---------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{configure} | Scalar         | Performs the configuration if the packages are already installed.<br>(Required)                                                   |
| CFG{accepteula}     | Scalar         | Specifies whether you agree with <code>EULA.pdf</code> on the media.<br>(Required)                                                |
| CFG{systems}        | List           | List of systems on which the product is to be configured.<br>(Required)                                                           |
| CFG{prod}           | Scalar         | Defines the product to be configured.<br>(Required)                                                                               |
| CFG{opt}{keyfile}   | Scalar         | Defines the location of an ssh keyfile that is used to communicate with all remote systems.<br>(Optional)                         |
| CFG{opt}{rsh}       | Scalar         | Defines that <code>rsh</code> must be used instead of <code>ssh</code> as the communication method between systems.<br>(Optional) |

**Table B-2** Response file variables specific to configuring SFHA (*continued*)

| Variable          | List or Scalar | Description                                                                                                                                                                                                                                                              |
|-------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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> .<br><br><b>Note:</b> The installer copies the response files and summary files also to the specified <i>logpath</i> location.<br><br>(Optional) |
| \$CFG{uploadlogs} | Scalar         | Defines Boolean value 0 or 1.<br><br>The value 1 indicates that the installation logs are uploaded to the Symantec Web site.<br><br>The value 0 indicates that the installation logs are not uploaded to the Symantec Web site.<br><br>(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 (the `csgnic`, `csgvip`, and `csgnetmask` variables) must be defined if any are defined. The same is true for the SMTP notification (the `smtserver`, `smtprec`, and `smtprsev` variables), the SNMP trap notification (the `snmpport`, `snmpcons`, and `snmpcsev` variables), and the Global Cluster Option (the `gconic`, `gcovip`, and `gconetmask` variables).

[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.<br><br>(Required) |
| CFG{vcs_clustername} | Scalar         | Defines the name of the cluster.<br><br>(Required)                                     |

**Table B-3** Response file variables specific to configuring a basic SFHA cluster  
(continued)

| Variable              | List or Scalar | Description                                                                                                                                            |
|-----------------------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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).<br><br>(Required) |
| \$CFG{fencingenabled} | Scalar         | In a SFHA configuration, defines if fencing is enabled.<br><br>Valid values are 0 or 1.<br><br>(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#}<br>{"system"} | Scalar         | Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links.<br><br>You must enclose the system name within double quotes.<br><br>(Required) |

**Table B-4** Response file variables specific to configuring private LLT over Ethernet (*continued*)

| Variable                              | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vcs_lltlinklowpri#}<br>{"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{lltoverudp}=1                          | Scalar         | <p>Indicates whether to configure heartbeat link using LLT over UDP.</p> <p>(Required)</p>                                                                                                                                                             |
| CFG{vcs_udplink<n>_address}<br>{<system1>} | Scalar         | <p>Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1.</p> <p>You can have four heartbeat links and &lt;n&gt; for this response file variable can take values 1 to 4 for the respective heartbeat links.</p> <p>(Required)</p> |

**Table B-5** Response file variables specific to configuring LLT over UDP  
(continued)

| Variable                                             | List or Scalar | Description                                                                                                                                                                                                                                                                            |
|------------------------------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG<br>{vcs_udplinklowpri<n>_address}<br>{<system1>} | Scalar         | Stores the IP address (IPv4 or IPv6) that the low-priority heartbeat link uses on node1.<br><br>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.<br><br>(Required)       |
| CFG{vcs_udplink<n>_port}<br>{<system1>}              | Scalar         | Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1.<br><br>You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links.<br><br>(Required)                                        |
| CFG{vcs_udplinklowpri<n>_port}<br>{<system1>}        | Scalar         | Stores the UDP port (16-bit integer value) that the low-priority heartbeat link uses on node1.<br><br>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.<br><br>(Required) |
| CFG{vcs_udplink<n>_netmask}<br>{<system1>}           | Scalar         | Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1.<br><br>You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links.<br><br>(Required)                                              |
| CFG{vcs_udplinklowpri<n>_netmask}<br>{<system1>}     | Scalar         | Stores the netmask (prefix for IPv6) that the low-priority heartbeat link uses on node1.<br><br>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.<br><br>(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}<br>{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.<br><br>(Optional) |
| CFG{vcs_csgvip}             | Scalar         | Defines the virtual IP address for the cluster.<br><br>(Optional)                                                                          |
| CFG{vcs_csgnetmask}         | Scalar         | Defines the Netmask of the virtual IP address for the cluster.<br><br>(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{at_rootdomain}       | Scalar         | Defines the name of the system where the root broker is installed.<br><br>(Optional)                                                                                                                              |
| CFG{at_rootbroker}       | Scalar         | Defines the root broker's name.                                                                                                                                                                                   |
| CFG{vcs_securitymenuopt} | Scalar         | Specifies the menu option to choose to configure the cluster in secure mode.<br><br><ul style="list-style-type: none"> <li>■ 1—Automatic</li> <li>■ 2—Semi-automatic</li> <li>■ 3—Manual</li> </ul><br>(Optional) |

**Table B-7** Response file variables specific to configuring SFHA cluster in secure mode (*continued*)

| Variable                         | List or Scalar | Description                                                                       |
|----------------------------------|----------------|-----------------------------------------------------------------------------------|
| CFG{vcs_vssdefport}              | Scalar         | Specifies the default port address of the root broker.<br><br>(Optional)          |
| CFG{vcs_roothashpath}            | Scalar         | Specifies the path of the root hash file.<br><br>(Optional)                       |
| CFG{vcs_ab_prplname}<br>{system} | Scalar         | Specifies the authentication broker's principal name on system.<br><br>(Optional) |
| CFG{vcs_ab_password}<br>{system} | Scalar         | Specifies the authentication broker's password on system.<br><br>(Optional)       |
| CFG{vcs_blobpath}<br>{system}    | Scalar         | Specifies the path of the encrypted BLOB file for system.<br><br>(Optional)       |

[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.<br><br>The value in the list can be "Administrators Operators Guests."<br><br><b>Note:</b> The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list.<br><br>(Optional) |
| CFG{vcs_username} | List           | List of names of VCS users.<br><br>(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.<br><br><b>Note:</b> The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list.<br><br>(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.<br><br>(Optional)                                                                                                                |
| CFG{vcs_smtprecip}  | List           | List of full email addresses (example: user@symantecexample.com) of SMTP recipients.<br><br>(Optional)                                                                                                                                                     |
| CFG{vcs_smtprsev}   | List           | Defines the minimum severity level of messages (Information, Warning, Error, and SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients.<br><br>(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).<br><br>(Optional)                                                                                                                                                                                    |
| CFG{vcs_snmpcons} | List           | List of SNMP console system names.<br><br>(Optional)                                                                                                                                                                                                  |
| CFG{vcs_snmpcsev} | List           | Defines the minimum severity level of messages (Information, Warning, Error, and SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names.<br><br>(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}<br>{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.<br><br>(Optional) |
| CFG{vcs_gcovip}             | Scalar         | Defines the virtual IP address to that the Global Cluster Option uses.<br><br>(Optional)                                                                                |
| CFG{vcs_gconetmask}         | Scalar         | Defines the Netmask of the virtual IP address that the Global Cluster Option uses.<br><br>(Optional)                                                                    |

## Sample response file for SFHA configuration

The following example shows a response file for configuring Storage Foundation High Availability.

```

#Auto generated sfha responsefile #

our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{rsh}=1;
$CFG{vcs_allowcomms}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
$CFG{opt}{prodmode}="SF Enterprise HA";
$CFG{opt}{configure}=1;
$CFG{prod}="SFHA51";
$CFG{systems}=[qw(system01 system02)];
$CFG{vm_restore_cfg}{system01}=0;
$CFG{vm_restore_cfg}{system02}=0;
$CFG{vcs_clusterid}=127;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_username}=[qw(admin operator)];
$CFG{vcs_userenpw}=[qw(JlmElgLimHmmKumGlj bQOsOUUnVQoOUUnTQsOSnUQuOUUnPQtOS)];
$CFG{vcs_userpriv}=[qw(Administrators Operators)];
$CFG{vcs_lltlink1}{"system01"}="bge1";
$CFG{vcs_lltlink2}{"system01"}="bge2";
$CFG{vcs_lltlink1}{"system02"}="bge1";
$CFG{vcs_lltlink2}{"system02"}="bge2";
$CFG{opt}{uuid}=normC;
$CFG{opt}{logpath}="/opt/VRTS/install/logs/installsf-xxxxxx/installsf-xxxxxx.response";

1;
```

## Sample response file for SFHA install

The following example shows a response file for installing Storage Foundation High Availability.

```

#Auto generated sfha responsefile #
#####
```

**Sample response file for SF upgrade**

```

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{opt}{upgrade}=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.

# 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](#)

## 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 92.
- 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 413.  
See [“Sample response file for configuring server-based I/O fencing”](#) on page 416.

- 4 Edit the values of the response file variables as necessary.  
 See “[Response file variables to configure disk-based I/O fencing](#)” on page 412.  
 See “[Response file variables to configure server-based I/O fencing](#)” on page 414.
- 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 C-1](#) lists the response file variables that specify the required information to configure disk-based I/O fencing for SFHA.

**Table C-1** Response file variables specific to configuring disk-based I/O fencing

| Variable                             | List or Scalar | Description                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{opt}{fencing}                    | Scalar         | Performs the I/O fencing configuration.<br>(Required)                                                                                                                                                                                                                                                                        |
| CFG{vxfen_config_fencing_option}     | Scalar         | Specifies the I/O fencing configuration mode.<br><ul style="list-style-type: none"> <li>■ 1—Coordination Point Server-based I/O fencing</li> <li>■ 2—Coordinator disk-based I/O fencing</li> <li>■ 3—Disabled mode</li> </ul> (Required)                                                                                     |
| CFG {vxfen_config_fencing_mechanism} | Scalar         | Specifies the I/O fencing mechanism.<br>This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the vxfen_config_fencing_mechanism variable and either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.<br>(Optional) |

**Table C-1** Response file variables specific to configuring disk-based I/O fencing  
*(continued)*

| Variable                              | List or Scalar | Description                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CFG{vxfen_config_fencing_dg}          | Scalar         | Specifies the disk group for I/O fencing.<br><br>(Optional)<br><br><b>Note:</b> You must define the vxfen_config_fencing_dg variable to use an existing disk group. If you want to create a new disk group, you must use both the vxfen_config_fencing_dg variable and the vxfen_config_fencing_newdg_disks variable.                              |
| CFG{vxfen_config_fencing_newdg_disks} | List           | Specifies the disks to use to create a new disk group for I/O fencing.<br><br>(Optional)<br><br><b>Note:</b> You must define the vxfen_config_fencing_dg variable to use an existing disk group. If you want to create a new disk group, you must use both the vxfen_config_fencing_dg variable and the vxfen_config_fencing_newdg_disks variable. |

## Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions. See [“Response file variables to configure disk-based I/O fencing”](#) on page 412.

```
#
Configuration Values:
#
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
```

```
$CFG{prod}="SFHA51";

$CFG{systems}=[qw(galaxy nebula)];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{vxfen_config_fencing_dg}="fendg";
$CFG{vxfen_config_fencing_mechanism}="dmp";
$CFG{vxfen_config_fencing_newdg_disks}=
 [qw(c1t1d0s2 c2t1d0s2 c3t1d0s2)];
$CFG{vxfen_config_fencing_option}=2;
```

## Response file variables to configure server-based I/O fencing

You can use a CP server response file to configure server-based customized I/O fencing. The installer uses the CP server response file for the following types of I/O fencing configurations:

- **Client cluster fencing (server-based I/O fencing configuration itself)**  
The installer configures server-based customized I/O fencing on the SF HA cluster without prompting for user input.
- **Disk-based fencing with the disk group already created**  
The installer configures fencing in disk-based mode on the SF HA cluster without prompting for user input.  
Disk-based fencing configuration is one in which SCSI-3 disks are used as the only coordination points.  
Disk-based fencing with the disk group already created means that the disk group consisting of the coordinating disks already exists on the SF HA cluster nodes.
- **Disk-based fencing with the disk group to be created**  
The installer creates the disk group and configures fencing properly on all the nodes in the SF HA cluster without user intervention.  
Disk-based fencing with the disk group to be created means that the disk group does not exist yet, but will be created with the disks mentioned as coordination point.

[Table C-2](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

**Table C-2** CP server response file definitions

| Response file field        | Definition                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencing_cpc_config_cpagent | <p>Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not.</p> <p>Enter "0" if you do not want to configure the Coordination Point agent using the installer.</p> <p>Enter "1" if you want to use the installer to configure the Coordination Point agent.</p>                                                                                                                                                                                                                                                    |
| fencing_cpc_cpagentgrp     | <p>Name of the service group which will have the Coordination Point agent resource as part of it.</p> <p><b>Note:</b> This field is obsolete if the <code>fencing_cpc_config_cpagent</code> field is given a value of '0'.</p>                                                                                                                                                                                                                                                                                                                                                   |
| fencing_cpc_cps            | <p>Virtual IP address or Virtual hostname of the CP servers.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| fencing_cpc_reusedg        | <p>This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks).</p> <p>Enter either a "1" or "0".</p> <p>Entering a "1" indicates reuse, and entering a "0" indicates do not reuse.</p> <p>When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text , such as<br/> <code>"\$CFG{fencing_cpc_reusedg}=0"</code> or<br/> <code>"\$CFG{fencing_cpc_reusedg}=1"</code> before proceeding with a silent installation.</p> |
| fencing_cpc_dgname         | <p>The name of the disk group to be used in the customized fencing, where at least one disk is being used.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| fencing_cpc_diffab         | <p>This response field indicates whether the CP servers and SF HA clusters use different root brokers.</p> <p>Entering a "1" indicates that they are using different root brokers.</p> <p>Entering a "0" indicates that they are not using different root brokers.</p>                                                                                                                                                                                                                                                                                                           |

**Table C-2** CP server response file definitions (*continued*)

| Response file field   | Definition                                                                                                                                                                        |
|-----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| fencing_cpc_disks     | The disks being used as coordination points if any.                                                                                                                               |
| fencing_cpc_ncps      | Total number of coordination points being used, including both CP servers and disks.                                                                                              |
| fencing_cpc_ndisks    | The number of disks being used.                                                                                                                                                   |
| fencing_cpc_ports     | The port of the CP server that is denoted by <i>cps</i> .                                                                                                                         |
| fencing_cpc_ccab      | The name of the authentication broker (AB) for any one of the SF HA cluster nodes.                                                                                                |
| fencing_cpc_cpsabport | The port at which the authentication broker (AB) mentioned above listens for authentication..                                                                                     |
| fencing_cpc_ccabport  | The port at which the authentication broker (AB) mentioned above listens for authentication.                                                                                      |
| fencing_cpc_mechanism | The disk mechanism that is used by customized fencing.<br><br>The value for this field is either "raw" or "dmp"                                                                   |
| fencing_cpc_cpsab     | The name of the authentication broker (AB) for any one of the CP servers.                                                                                                         |
| fencing_cpc_security  | This field indicates whether security is enabled or not<br><br>Entering a "1" indicates that security is enabled.<br>Entering a "0" indicates that security has not been enabled. |

## Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing :

```

$CFG{fencing_cpc_config_cpagent}=0;
$CFG{fencing_cpc_cps}=[qw(10.200.117.145)];
$CFG{fencing_cpc_dgname}="vxfencoorddg";
$CFG{fencing_cpc_diffab}=0;
$CFG{fencing_cpc_disks}=[qw(emc_clariion0_37 emc_clariion0_13)];
$CFG{fencing_cpc_mechanism}="raw";
$CFG{fencing_cpc_ncps}=3;
$CFG{fencing_cpc_ndisks}=2;
$CFG{fencing_cpc_ports}{"10.200.117.145"}=14250;

```

```
$CFG{fencing_cpc_reusedg}=1;
$CFG{fencing_cpc_security}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SF51";
$CFG{systems}=[qw(galaxy nebula)];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{vxfen_config_fencing_option}=1;
```



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

**Table D-1** LLT configuration files

| File             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/default/llt | <p>This file stores the start and stop environment variables for LLT:</p> <ul style="list-style-type: none"> <li>■ <b>LLT_START</b>—Defines the startup behavior for the LLT module after a system reboot. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that LLT is enabled to start up.</li> <li>0—Indicates that LLT is disabled to start up.</li> </ul> </li> <li>■ <b>LLT_STOP</b>—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that LLT is enabled to shut down.</li> <li>0—Indicates that LLT is disabled to shut down.</li> </ul> </li> </ul> <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 data-bbox="346 869 514 921"> 0      galaxy 1      nebula </pre>                                                                                                                                                                                                                                                                                             |

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

| File        | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/llttab | <p>The file <code>llttab</code> contains the information that is derived during installation and used by the utility <code>lltconfig(1M)</code>. After installation, this file lists the private network links that correspond to the specific system. For example, the file <code>/etc/llttab</code> contains the entries that resemble the following:</p> <ul style="list-style-type: none"> <li>■ For Solaris SPARC: <pre style="margin-left: 40px;">set-node galaxy set-cluster 2 link bge0 /dev/bge0 - ether - - link bge0 /dev/bge1 - ether - -</pre> </li> <li>■ For Solaris x64: <pre style="margin-left: 40px;">set-node galaxy set-cluster 2 link e1000g1 /dev/e1000g:1 - ether - - link e1000g2 /dev/e1000g:2 - ether - -</pre> </li> </ul> <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 D-2** lists the GAB configuration files and the information that these files contain.

**Table D-2** GAB configuration files

| File             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/default/gab | <p>This file stores the start and stop environment variables for GAB:</p> <ul style="list-style-type: none"> <li>■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that GAB is enabled to start up.</li> <li>0—Indicates that GAB is disabled to start up.</li> </ul> </li> <li>■ GAB_STOP—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that GAB is enabled to shut down.</li> <li>0—Indicates that GAB is disabled to shut down.</li> </ul> </li> </ul> <p>The installer sets the value of these variables to 1 at the end of SFHA configuration.</p> |
| /etc/gabtab      | <p>After you install SFHA, the file /etc/gabtab contains a <code>gabconfig(1)</code> command that configures the GAB driver for use.</p> <p>The file /etc/gabtab contains a line that resembles:</p> <pre data-bbox="588 838 870 864">/sbin/gabconfig -c -nN</pre> <p>The <code>-c</code> option configures the driver for use. The <code>-nN</code> specifies that the cluster is not formed until at least <i>N</i> nodes are ready to form the cluster. Symantec recommends that you set <i>N</i> to be the total number of nodes in the cluster.</p> <p><b>Note:</b> 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 D-3](#) lists the AMF configuration files.

**Table D-3** AMF configuration files

| File                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>/etc/default/amf</code> | <p>This file stores the start and stop environment variables for AMF:</p> <ul style="list-style-type: none"> <li>■ <b>AMF_START</b>—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"> <li>1—Indicates that AMF is enabled to start up.</li> <li>0—Indicates that AMF is disabled to start up. (default)</li> </ul> </li> <li>■ <b>AMF_STOP</b>—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"> <li>1—Indicates that AMF is enabled to shut down. (default)</li> <li>0—Indicates that AMF is disabled to shut down.</li> </ul> </li> </ul> |
| <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 style="margin-left: 20px;"><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 425.

See [“Sample main.cf file for global clusters”](#) on page 426.
- `types.cf`  

The file `types.cf`, which is listed in the include statement in the `main.cf` file, defines the VCS bundled types for VCS resources. The file `types.cf` is also located in the folder `/etc/VRTSvcs/conf/config`.

Additional files similar to `types.cf` may be present if agents have been added, such as `OracleTypes.cf`.

- `/etc/default/vcs`

This file stores the start and stop environment variables for VCS engine:

- **VCS\_START**—Defines the startup behavior for VCS engine after a system reboot. Valid values include:
  - 1—Indicates that VCS engine is enabled to start up.
  - 0—Indicates that VCS engine is disabled to start up.
- **VCS\_STOP**—Defines the shutdown behavior for VCS engine during a system shutdown. Valid values include:
  - 1—Indicates that VCS engine is enabled to shut down.
  - 0—Indicates that VCS engine is disabled to shut down.

The installer sets the value of these variables to 1 at the end of SFHA configuration.

Note the following information about the VCS configuration file after installing and configuring VCS:

- The cluster definition includes the cluster information that you provided during the configuration. This definition includes the cluster name, cluster address, and the names of users and administrators of the cluster. Notice that the cluster has an attribute `UserNames`. The `installsfha` creates a user "admin" whose password is encrypted; the word "password" is the default password.
- If you set up the optional I/O fencing feature for VCS, then the `UseFence = SCSI3` attribute is present.
- If you configured the cluster in secure mode, the `main.cf` includes the `VxSS` service group and "`SecureClus = 1`" cluster attribute.
- The `installsfha` creates the `ClusterService` service group if you configured the virtual IP, SMTP, SNMP, or global cluster options.

The service group also has the following characteristics:

- The group includes the IP and NIC resources.
- The service group also includes the notifier resource configuration, which is based on your input to `installsfha` prompts about notification.
- The `installsfha` also creates a resource dependency tree.
- If you set up global clusters, the `ClusterService` service group contains an Application resource, `wac` (wide-area connector). This resource's attributes contain definitions for controlling the cluster in a global cluster environment.

Refer to the *Veritas Cluster Server Administrator's Guide* for information about managing VCS global clusters.

Refer to the *Veritas Cluster Server Administrator's Guide* to review the configuration concepts, and descriptions of `main.cf` and `types.cf` files for Solaris systems.

## Sample `main.cf` file for VCS clusters

The following sample `main.cf` file is for a three-node cluster in secure mode.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"
include "Db2udbTypes.cf"
include "SybaseTypes.cf"

cluster vcs02 (
 SecureClus = 1
)

system sysA (
)

system sysB (
)

system sysC (
)

group ClusterService (
 SystemList = { sysA = 0, sysB = 1, sysC = 2 }
 AutoStartList = { sysA, sysB, sysC }
 OnlineRetryLimit = 3
 OnlineRetryInterval = 120
)

NIC csgnic (
 Device = bge0
)

NotifierMngr ntfr (
 SnmpConsoles = { jupiter" = SevereError }
 SntpServer = "smtp.example.com"
 SntpRecipients = { "ozzie@example.com" = SevereError }
```

```

)

ntfr requires csgnic

// resource dependency tree
//
// group ClusterService
// {
// NotifierMngr ntfr
// {
// NIC csgnic
// }
// }

group VxSS (
 SystemList = { sysA = 0, sysB = 1, sysC = 2 }
 Parallel = 1
 AutoStartList = { sysA, sysB, sysC }
 OnlineRetryLimit = 3
 OnlineRetryInterval = 120
)

Phantom phantom_vxss (
)

ProcessOnOnly vxatd (
 IgnoreArgs = 1
 PathName = "/opt/VRTSat/bin/vxatd"
)

// resource dependency tree
//
// group VxSS
// {
// Phantom phantom_vxss
// ProcessOnOnly vxatd
// }

```

## Sample main.cf file for global clusters

If you installed SFHA with the Global Cluster option, note that the ClusterService group also contains the Application resource, wac. The wac resource is required to control the cluster in a global cluster environment.

In the following main.cf file example, bold text highlights global cluster specific entries.

```
include "types.cf"

cluster vcs03 (
 ClusterAddress = "10.182.13.50"
 SecureClus = 1
)

system sysA (
)

system sysB (
)

system sysC (
)

group ClusterService (
 SystemList = { sysA = 0, sysB = 1, sysC = 2 }
 AutoStartList = { sysA, sysB, sysC }
 OnlineRetryLimit = 3
 OnlineRetryInterval = 120
)

Application wac (
 StartProgram = "/opt/VRTSvcs/bin/wacstart"
 StopProgram = "/opt/VRTSvcs/bin/wacstop"
 MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
 RestartLimit = 3
)

IP gcoip (
 Device = bge0
 Address = "10.182.13.50"
 NetMask = "255.255.240.0"
)

NIC csgnic (
 Device = bge0
)
```

```
NotifierMngr ntfr (
 SnmpConsoles = { jupiter" = SevereError }
 SmtServer = "smtp.example.com"
 SmtpRecipients = { "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
// }
// }
// }
// }

group VxSS (
 SystemList = { sysA = 0, sysB = 1, sysC = 2 }
 Parallel = 1
 AutoStartList = { sysA, sysB, sysC }
 OnlineRetryLimit = 3
 OnlineRetryInterval = 120
)

Phantom phantom_vxss (
)

ProcessOnOnly vxatd (
 IgnoreArgs = 1
 PathName = "/opt/VRTSat/bin/vxatd"
)
```

```
// resource dependency tree
//
// group VxSS
// {
// Phantom phantom_vxss
// ProcessOnOnly vxatd
// }
```

## About I/O fencing configuration files

[Table D-4](#) lists the I/O fencing configuration files.

**Table D-4** I/O fencing configuration files

| File               | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| /etc/default/vxfen | <p>This file stores the start and stop environment variables for I/O fencing:</p> <ul style="list-style-type: none"> <li>■ <b>VXFEN_START</b>—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that I/O fencing is enabled to start up.</li> <li>0—Indicates that I/O fencing is disabled to start up.</li> </ul> </li> <li>■ <b>VXFEN_STOP</b>—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: <ul style="list-style-type: none"> <li>1—Indicates that I/O fencing is enabled to shut down.</li> <li>0—Indicates that I/O fencing is disabled to shut down.</li> </ul> </li> </ul> <p>The installer sets the value of these variables to 1 at the end of 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 D-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"> <li>■ <b>vxfen_mode</b> <ul style="list-style-type: none"> <li>■ <b>scsi3</b>—For disk-based fencing</li> <li>■ <b>customized</b>—For server-based fencing</li> <li>■ <b>disabled</b>—To run the I/O fencing driver but not do any fencing operations.</li> </ul> </li> <li>■ <b>vxfen_mechanism</b><br/>This parameter is applicable only for server-based fencing. Set the value as cps.</li> <li>■ <b>scsi3_disk_policy</b> <ul style="list-style-type: none"> <li>■ <b>dmp</b>—Configure the vxfen module to use DMP devices<br/>The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp.</li> <li>■ <b>raw</b>—Configure the vxfen module to use the underlying raw character devices</li> </ul> </li> </ul> <p><b>Note:</b> You must use the same SCSI-3 disk policy on all the nodes.</p> <ul style="list-style-type: none"> <li>■ <b>security</b><br/>This parameter is applicable only for server-based fencing.<br/>1—Indicates that Symantec Product Authentication Service is used for CP server communications. This setting is the default.<br/>0—Indicates that communication with the CP server is in non-secure mode.<br/><b>Note:</b> The CP server and the SFHA clusters must have the same security setting.</li> <li>■ <b>List of coordination points</b><br/>This list is required only for server-based fencing configuration.<br/>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.<br/>Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points.</li> <li>■ <b>single_cp</b><br/>This parameter is applicable only for server-based fencing which uses a single highly available CP server as its coordination point.</li> </ul> |

**Table D-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 <code>vxfentab</code> file with the current list of all the coordinator points.</p> <p><b>Note:</b> 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"> <li>■ Raw disk:           <ul style="list-style-type: none"> <li><code>/dev/rdisk/c1t1d0s2</code></li> <li><code>/dev/rdisk/c2t1d0s2</code></li> <li><code>/dev/rdisk/c3t1d2s2</code></li> </ul> </li> <li>■ DMP disk:           <ul style="list-style-type: none"> <li><code>/dev/vx/rdmp/c1t1d0s2</code></li> <li><code>/dev/vx/rdmp/c2t1d0s2</code></li> <li><code>/dev/vx/rdmp/c3t1d0s2</code></li> </ul> </li> </ul> <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 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 432.
- 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 434.

---

**Note:** The CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with SF HA clusters. The following 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"

// cluster name: cps1
// CP server: mycps1

cluster cps1 (
 UserNames = { admin = bMNFmHmJNiNNlVnHMK, haris = fopKojNvpHouNn }

 Administrators = { admin, haris }

 SecureClus = 1
 HacliUserLevel = COMMANDROOT
)

system mycps1 (
)

group CPSSG (
 SystemList = { mycps1 = 0 }
 AutoStartList = { mycps1 }
)

IP cpsvip (
 Device @mycps1 = bge0
 Address = "10.209.3.1"
 NetMask = "255.255.252.0"
)

NIC cpsnic (
```

```

 Device @mycps1 = bge0
)

 Process vxcpserv (
 PathName = "/opt/VRTScps/bin/vxcpserv"
 ConfInterval = 30
 RestartLimit = 3
)

cpsvip requires cpsnic
vxcpserv requires cpsvip

// resource dependency tree
//
// group CPSSG
// {
// Process vxcpserv
// {
// IP cpsvip
// {
// NIC cpsnic
// }
// }
// }

group VxSS (
 SystemList = { mycps1 = 0 }
 Parallel = 1
 AutoStartList = { mycps1 }
 OnlineRetryLimit = 3
 OnlineRetryInterval = 120
)

Phantom phantom_vxss (
)

ProcessOnOnly vxatd (
 IgnoreArgs = 1
 PathName = "/opt/VRTSat/bin/vxatd"
)

```

```
// resource dependency tree
//
// group VxSS
// {
// Phantom phantom_vxss
// ProcessOnOnly vxatd
// }
```

## 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"

// cluster: cps1
// CP servers:
// mycps1
// mycps2

cluster cps1 (
 UserNames = { admin = ajkCjeJgkFkkIskEjh
 }
 Administrators = { admin }
 SecureClus = 1
)

system mycps1 (
)

system mycps2 (
)

group CPSSG (
```

```

SystemList = { mycps1 = 0, mycps2 = 1 }
AutoStartList = { mycps1, mycps2 }

DiskGroup cpsdg (
 DiskGroup = cps_dg
)

IP cpsvip (
 Device @mycps1 = bge0
 Device @mycps2 = bge0
 Address = "10.209.81.88"
 NetMask = "255.255.252.0"
)

Mount cpsmount (
 MountPoint = "/etc/VRTScps/db"
 BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
 FSType = vxfs
 FsckOpt = "-y"
)

NIC cpsnic (
 Device @mycps1 = bge0
 Device @mycps2 = bge0
)

Process vxcpserv (
 PathName = "/opt/VRTScps/bin/vxcpserv"
)

Volume cpsvol (
 Volume = cps_volume
 DiskGroup = cps_dg
)

```

```

cpsmount requires cpsvol
cpsvip requires cpsnic
cpsvol requires cpsdg
vxcpserv requires cpsmount
vxcpserv requires cpsvip

```

```
// resource dependency tree
```

```
//
// group CPSSG
// {
// Process vxcpserv
// {
// Mount cpsmount
// {
// Volume cpsvol
// {
// DiskGroup cpsdg
// }
// }
// IP cpsvip
// {
// NIC cpsnic
// }
// }
// }

group VxSS (
 SystemList = { mycps1 = 0, mycps2 = 1 }
 Parallel = 1
 AutoStartList = { mycps1, mycps2 }
 OnlineRetryLimit = 3
 OnlineRetryInterval = 120
)

Phantom phantom_vxss (
)

ProcessOnOnly vxatd (
 IgnoreArgs = 1
 PathName = "/opt/VRTSat/bin/vxatd"
)

// resource dependency tree
//
// group VxSS
// {
// Phantom phantom_vxss
```

```
// ProcessOnOnly vxatd
// }
```



# 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](#)
- [Configuring and enabling ssh](#)
- [Restarting the ssh session](#)
- [Enabling and disabling rsh for Solaris](#)

## About configuring secure shell or remote shell communication modes before installing products

Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installer is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install Veritas software.

You can install products to remote systems using either secure shell (`ssh`) or remote shell (`rsh`). Symantec recommends that you use `ssh` as it is more secure than `rsh`.

This section contains an example of how to set up `ssh` password free communication. The example sets up `ssh` between a source system (`system1`) that

contains the installation directories, and a target system (system2). This procedure also applies to multiple target systems.

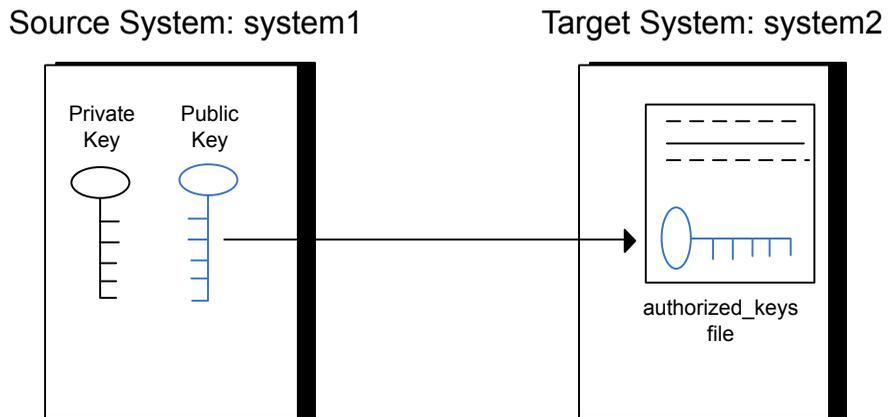
## Configuring and enabling ssh

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

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

Figure E-1 illustrates this procedure.

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



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

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

### To create the DSA key pair

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

```
system1 # cd /
```

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

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

System output similar to the following is displayed:

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

- 3 Press Enter to accept the default location of `/.ssh/id_dsa`.
- 4 When the program asks you to enter the passphrase, press the Enter key twice.

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

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

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

```
system2 # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

```
system2 # chmod go-w /.ssh
```

**To append the public key from the source system to the authorized\_keys file on the target system, using secure file transfer**

- 1 Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems (system2 in this example).

To enable SFTP, the `/etc/ssh/sshd_config` file must contain the following two lines:

```
PermitRootLogin yes
Subsystem sftp /usr/lib/ssh/sftp-server
```

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

To restart ssh on Solaris 10, type the following command:

```
system1 # svcadm restart ssh
```

To restart on Solaris 9, type the following commands:

```
system1 # /etc/init.d/sshd stop
system1 # /etc/init.d/sshd start
```

- 3 From the source system (system1), move the public key to a temporary file on the target system (system2).

Use the secure file transfer program.

In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

Use the following command for secure file transfer:

```
system1 # sftp system2
```

If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

```
Connecting to system2 ...
The authenticity of host 'system2 (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
Are you sure you want to continue connecting (yes/no)?
```

- 4 Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'system2,10.182.00.00'
(DSA) to the list of known hosts.
root@system2 password:
```

- 5 Enter the root password of `system2`.

- 6 At the `sftp` prompt, type the following command:

```
sftp> put /.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /.ssh/id_dsa.pub to /id_dsa.pub
```

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

```
sftp> quit
```

- 8 To begin the `ssh` session on the target system (`system2` in this example), type the following command on `system1`:

```
system1 # ssh system2
```

Enter the root password of `system2` at the prompt:

```
password:
```

- 9 After you log in to `system2`, enter the following command to append the `id_dsa.pub` file to the `authorized_keys` file:

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

- 10 After the `id_dsa.pub` public key file is copied to the target system (`system2`), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, enter the following command on `system2`:

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

- 11 To log out of the `ssh` session, enter the following command:

```
system2 # exit
```

- 12 When you install from a source system that is also an installation target, also add the local system `id_dsa.pub` key to the local `authorized_keys` file. The installation can fail if the installation source system is not authenticated.

To add the local system `id_dsa.pub` key to the local `authorized_keys` file, enter the following command:

```
system1 # cat /.ssh/id_dsa.pub >> /.ssh/authorized_keys
```

- 13 Run the following commands on the source installation system. If your ssh session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and make the key globally available to the user `root`:

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

Identity added: //./ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

#### To verify that you can connect to a target system

- 1 On the source system (system1), enter the following command:

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

where `system2` is the name of the target system.

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

## Restarting the ssh session

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

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

### To restart ssh

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

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

- 2 Make the key globally available for the user `root`

```
system1 # ssh-add
```

## Enabling and disabling rsh for Solaris

The following section describes how to enable remote shell on Solaris system.

Veritas recommends configuring a secure shell environment for Veritas product installations.

See [“Configuring and enabling ssh”](#) on page 440.

See the operating system documentation for more information on configuring remote shell.

### To enable rsh

- 1 To determine the current status of `rsh` and `rlogin`, type the following command:

```
inetadm | grep -i login
```

If the service is enabled, the following line is displayed:

```
enabled online svc:/network/login:rlogin
```

If the service is not enabled, the following line is displayed:

```
disabled disabled svc:/network/login:rlogin
```

- 2 To enable a disabled `rsh/rlogin` service, type the following command:

```
inetadm -e rlogin
```

- 3 To disable an enabled `rsh/rlogin` service, type the following command:

```
inetadm -d rlogin
```

- 4 Modify the `.rhosts` file. A separate `.rhosts` file is in the `$HOME` directory of each user. This file must be modified for each user who remotely accesses the system using `rsh`. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system having access to the local system. For example, if the root user must remotely access `system1` from `system2`, you must add an entry for `system2.companyname.com` in the `.rhosts` file on `system1`.

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

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

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

# Storage Foundation and High Availability components

This appendix includes the following topics:

- [Storage Foundation and High Availability installation packages](#)
- [Veritas Cluster Server installation packages](#)
- [Chinese language packages](#)
- [Japanese language packages](#)
- [Veritas Storage Foundation obsolete and reorganized installation packages](#)

## Storage Foundation and High Availability installation packages

[Table F-1](#) shows the package name and contents for each English language package for Storage Foundation and High Availability. The table also gives you guidelines for which packages to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and High Availability and Veritas Cluster Server (VCS) packages, the combined functionality is called Storage Foundation and High Availability and High Availability.

See [“Veritas Cluster Server installation packages”](#) on page 449.

**Table F-1** Storage Foundation and High Availability packages

| packages   | Contents                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Configuration |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| VRTSaslapm | Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries<br><br>Required for the support and compatibility of various storage arrays.                                                                                                                                                                                                                                                                                                         | Minimum       |
| VRTSat     | Symantec Product Authentication Service<br><br>Installs the Symantec Product Authentication Service, which provides authentication services to other Symantec products.<br><br>This package contains a server and client component. The server provides services for a root broker, authentication broker, or both.<br><br>The client allows Symantec products to communicate with the brokers.<br><br>Required to use Symantec Product Authentication Service. | All           |
| VRTSperl   | Perl 5.10.0 for Veritas                                                                                                                                                                                                                                                                                                                                                                                                                                         | Minimum       |
| VRTSvlic   | Veritas License Utilities<br><br>Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.                                                                                                                                                                                                                                                  | Minimum       |
| VRTSvxfs   | Veritas File System binaries<br><br>Required for VxFS file system support.                                                                                                                                                                                                                                                                                                                                                                                      | Minimum       |
| VRTSvxvm   | Veritas Volume Manager binaries                                                                                                                                                                                                                                                                                                                                                                                                                                 | Minimum       |
| VRTSdbed   | Veritas Storage Foundation for Oracle                                                                                                                                                                                                                                                                                                                                                                                                                           | Recommended   |
| VRTSob     | Veritas Enterprise Administrator                                                                                                                                                                                                                                                                                                                                                                                                                                | Recommended   |

**Table F-1** Storage Foundation and High Availability packages (*continued*)

| packages  | Contents                                                                                                                                                                                                                                                                                                                                                                                                 | Configuration |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| VRTSodm   | ODM Driver for VxFS<br><br>Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager allows Oracle 9i and 10g to improve performance and manage system bandwidth.                                                                                                                                                          | Recommended   |
| VRTSsfmh  | Veritas Storage Foundation Managed Host<br><br>Discovers configuration information on a Storage Foundation managed host. This information is stored on a central database, which is not part of this release. You must download the database separately at:<br><br><a href="http://www.symantec.com/business/storage-foundation-manager">http://www.symantec.com/business/storage-foundation-manager</a> | Recommended   |
| VRTSspt   | Veritas Software Support Tools                                                                                                                                                                                                                                                                                                                                                                           | Recommended   |
| VRTSfssdk | Veritas File System Software Developer Kit<br><br>For VxFS APIs, the package contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.                                                                                                                                                                                       | All           |

## Veritas Cluster Server installation packages

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

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

See “[Storage Foundation and High Availability installation packages](#)” on page 447.

**Table F-2** VCS installation packages

| package    | Contents                                                                                                                                                                                                                                | Configuration |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|
| VRTSgab    | Veritas Cluster Server group membership and atomic broadcast services                                                                                                                                                                   | Minimum       |
| VRTSllt    | Veritas Cluster Server low-latency transport                                                                                                                                                                                            | Minimum       |
| VRTSamf    | Veritas Cluster Server Asynchronous Monitoring Framework                                                                                                                                                                                | Minimum       |
| VRTSvcsc   | Veritas Cluster Server                                                                                                                                                                                                                  | Minimum       |
| VRTSvcscag | Veritas Cluster Server Bundled Agents                                                                                                                                                                                                   | Minimum       |
| VRTSvcxfen | Veritas I/O Fencing                                                                                                                                                                                                                     | Minimum       |
| VRTSvcsea  | Consolidated database and enterprise agent packages                                                                                                                                                                                     | Recommended   |
| VRTScps    | Veritas Coordination Point Server<br><br>The Coordination Point Server is an alternate mechanism for I/O fencing. It implements I/O fencing through a client/server architecture and can provide I/O fencing for multiple VCS clusters. | All           |

## Chinese language packages

The following table shows the package name and contents for each Chinese language package.

**Table F-3** Chinese language packages

| package  | Contents                                                                       |
|----------|--------------------------------------------------------------------------------|
| VRTSatZH | Symantec Product Authentication Service Software Chinese Language Kit          |
| VRTSzhvm | Chinese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages |

## Japanese language packages

The following table show the package name and contents for each Japanese language package.

**Table F-4** Japanese language packages

| package   | Contents                                                                                                                          |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------|
| VRTSatJA  | Symantec Product Authentication Service Software Japanese Language Kit                                                            |
| VRTSjacav | Japanese Veritas Cluster Server Agents for Storage Foundation Cluster File System – Manual Pages and Message Catalogs by Symantec |
| VRTSjacs  | Veritas Cluster Server Japanese Message Catalogs by Symantec                                                                      |
| VRTSjacse | Japanese Veritas High Availability Enterprise Agents by Symantec                                                                  |
| VRTSjadba | Japanese Veritas Oracle Real Application Cluster Support package by Symantec                                                      |
| VRTSjadbe | Japanese Veritas Storage Foundation for Oracle from Symantec – Message Catalogs                                                   |
| VRTSjafs  | Japanese Veritas File System – Message Catalog and Manual Pages                                                                   |
| VRTSjaodm | Veritas Oracle Disk Manager Japanese Message Catalog and Manual Pages by Symantec                                                 |
| VRTSjavm  | Japanese Veritas Volume Manager by Symantec – Message Catalogs and Manual Pages                                                   |
| VRTSmulic | Multi-language Symantec License Utilities                                                                                         |

## Veritas Storage Foundation obsolete and reorganized installation packages

**Table F-5** lists the packages that are obsolete or reorganized for Storage Foundation and Storage Foundation High Availability.

**Table F-5** Veritas Storage Foundation obsolete and reorganized packages

| package        | Description |
|----------------|-------------|
| Infrastructure |             |
| SYMClma        | Obsolete    |

**Table F-5** Veritas Storage Foundation obsolete and reorganized packages  
*(continued)*

| package          | Description                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| VRTSaa           | Included in VRTSsfmh                                                                                                                                                                                                                                                                                                                                                                                            |
| VRTSccg          | Included in VRTSsfmh                                                                                                                                                                                                                                                                                                                                                                                            |
| VRTSdbms3        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSicsco        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSjre          | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSjre15        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSmh           | Included in VRTSsfmh                                                                                                                                                                                                                                                                                                                                                                                            |
| VRTSobc33        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSobweb        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSobgui        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSpbx          | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSsfm          | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTSweb          | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| Product packages |                                                                                                                                                                                                                                                                                                                                                                                                                 |
| VRTSacclib       | <p>Obsolete</p> <p>The following information is for installations, upgrades, and uninstalls using the script- or Web-based installer.</p> <ul style="list-style-type: none"> <li>■ For fresh installations VRTSacclib is not installed.</li> <li>■ For upgrades, the existing VRTSacclib is uninstalled and a new VRTSacclib is installed.</li> <li>■ For uninstalls, VRTSacclib is not uninstalled.</li> </ul> |
| VRTSalloc        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTScmccc        | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |
| VRTScmcm         | Obsolete                                                                                                                                                                                                                                                                                                                                                                                                        |

**Table F-5** Veritas Storage Foundation obsolete and reorganized packages  
*(continued)*

| package   | Description            |
|-----------|------------------------|
| VRTScmcs  | Obsolete               |
| VRTScscm  | Obsolete               |
| VRTScscw  | Obsolete               |
| VRTScsocw | Obsolete               |
| VRTScssim | Obsolete               |
| VRTScutil | Obsolete               |
| VRTSd2gui | Included in VRTSdbed   |
| VRTSdb2ed | Included in VRTSdbed   |
| VRTSdbcom | Included in VRTSdbed   |
| VRTSdbed  | Included in VRTSdbed   |
| VRTSdcli  | Obsolete               |
| VRTSddlpr | Obsolete               |
| VRTSdsa   | Obsolete               |
| VRTSfas   | Obsolete               |
| VRTSfasag | Obsolete               |
| VRTSfsman | Included in mainpkg    |
| VRTSfsmnd | Included in mainpkg    |
| VRTSfspro | Included in VRTSsfmh   |
| VRTSgapms | Obsolete               |
| VRTSmapro | Included in VRTSsfmh   |
| VRTSorgui | Obsolete               |
| VRTSsybed | Included in VRTSdbed   |
| VRTSvail  | Obsolete               |
| VRTSvcsdb | Included in VRTSvcssea |

**Table F-5** Veritas Storage Foundation obsolete and reorganized packages  
*(continued)*

| <b>package</b> | <b>Description</b>                  |
|----------------|-------------------------------------|
| VRTSvcsmn      | Included in VRTSvc                  |
| VRTSvcSor      | Included in VRTSvcsea               |
| VRTSvcSSy      | Included in VRTSvcsea               |
| VRTSvcSvr      | Included in VRTSvc                  |
| VRTSvDid       | Obsolete                            |
| VRTSvMman      | Included in mainpkg                 |
| VRTSvMpro      | Included in VRTSsfmh                |
| VRTSvrpro      | Included in VRTSob                  |
| VRTSvrw        | Obsolete                            |
| VRTSvxmsa      | Obsolete                            |
| Documentation  | All Documentation packages obsolete |

# Troubleshooting installation issues

This appendix includes the following topics:

- [Restarting the installer after a failed connection](#)
- [What to do if you see a licensing reminder](#)
- [Troubleshooting information](#)
- [Incorrect permissions for root on remote system](#)
- [Inaccessible system](#)
- [Upgrading Veritas Storage Foundation for Databases \(SFDB\) tools from 5.0.x to 5.1SP1 \(2184482\)](#)

## 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. See “[Installing Veritas product license keys](#)” on page 46. After you install the license key, you must validate the license key using the following command:

```
vxkeyless display
```

- 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 VRTSspt package provides a group of tools for troubleshooting a system and collecting information on its configuration. The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. Although the tools are not required for the operation of any Veritas product, Symantec recommends installing them should a support case be needed to be opened with Symantec Support. If you are unfamiliar with their use and purpose, use caution when using them or use them in concert with Symantec Support.

## Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Failed to setup rsh communication on 10.198.89.241:
'rsh 10.198.89.241 <command>' failed
Trying to setup ssh communication on 10.198.89.241.
Failed to setup ssh communication on 10.198.89.241:
Login denied
```

```
Failed to login to remote system(s) 10.198.89.241.
Please make sure the password(s) are correct and superuser(root)
can login to the remote system(s) with the password(s).
If you want to setup rsh on remote system(s), please make sure
rsh with command argument ('rsh <host> <command>') is not
denied by remote system(s).
```

```
Either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

```
Would you like the installer to setup ssh/rsh communication
automatically between the nodes?
Superuser passwords for the systems will be asked. [y,n,q] (y) n
```

```
System verification did not complete successfully
```

```
The following errors were discovered on the systems:
```

```
The ssh permission denied on 10.198.89.241
rsh exited 1 on 10.198.89.241
either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

**Suggested solution:** You need to set up the systems to allow remote access using ssh or rsh.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 439.

---

**Note:** Remove remote shell permissions after completing the SFHA installation and configuration.

---

## Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
Verifying systems: 12%
Estimated time remaining: 0:10 1 of 8
Checking system communication Done
System verification did not complete successfully
The following errors were discovered on the systems:
cannot resolve hostname host1
Enter the system names separated by spaces: q,? (host1)
```

Suggested solution: Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

## Upgrading Veritas Storage Foundation for Databases (SFDB) tools from 5.0.x to 5.1SP1 (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 5.1SP1. The error message is:

When upgrading from SFHA version 5.0 to SFHA 5.1SP1 the `S*vxdbms3` startup script is renamed to `NO_S*vxdbms3`. The `S*vxdbms3` startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` 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.
```

### Workaround

Before running `sfua_rept_migrate`, rename the startup script `NO_S*vxdbms3` to `S*vxdbms3`.

# 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 server-based I/O fencing](#)
- [Troubleshooting server-based fencing on the SF HA cluster nodes](#)
- [Troubleshooting server-based I/O fencing in mixed mode](#)
- [After upgrading from 5.0.x and before migrating SFDB](#)

## 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 55.

- 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.

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

## The `vxfcntlsthdw` utility fails when SCSI TEST UNIT READY command fails

While running the `vxfcntlsthdw` utility, you may see a message that resembles as follows:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node
FAILED.
Contact the storage provider to have the hardware configuration
fixed.
```

The disk array does not support returning success for a `SCSI TEST UNIT READY` command when another host has the disk reserved using SCSI-3 persistent reservations. This happens with the Hitachi Data Systems 99XX arrays if bit 186 of the system mode option is not enabled.

# Troubleshooting server-based I/O fencing

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/cpserver_[ABC].log`
- `/var/VRTSat/vrtsat_broker.txt` (Security related)
- If the `vxcperv` process fails on the CP server, then review the following diagnostic files:
  - `/var/VRTScps/diag/FFDC_CPS_pid_vxcperv.log`
  - `/var/VRTScps/diag/stack_pid_vxcperv.txt`

---

**Note:** If the `vxcperv` process fails on the CP server, these files are present in addition to a core file. VCS restarts `vxcperv` process automatically in such situations.

---

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs and text files that may be useful in understanding and troubleshooting fencing-related issues on a SF HA cluster (client cluster) node.

See [“Troubleshooting issues related to the CP server service group”](#) on page 461.

See [“Checking the connectivity of CP server”](#) on page 462.

See [“Issues during fencing startup on SF HA cluster nodes set up for server-based fencing”](#) on page 463.

See [“Issues during online migration of coordination points”](#) on page 465.

See [“Troubleshooting server-based I/O fencing in mixed mode”](#) on page 466.

See [“Checking keys on coordination points when `vxfen\_mechanism` value is set to `cps`”](#) on page 470.

## Troubleshooting issues related to the CP server service group

If you cannot bring up the CPSSG service group after the CP server configuration, perform the following steps:

- Verify that the CPSSG service group and its resources are valid and properly configured in the VCS configuration.
- Check the VCS engine log (`/var/VRTSvcs/log/engine_[ABC].log`) to see if any of the CPSSG service group resources are FAULTED.
- Review the sample dependency graphs to make sure the required resources are configured correctly.

## Checking the connectivity of CP server

You can test the connectivity of CP server using the `cpsadm` command.

You must have set the environment variables `CPS_USERNAME` and `CPS_DOMAINTYPE` to run the `cpsadm` command on the SF HA cluster (client cluster) nodes.

### To check the connectivity of CP server

- ◆ Run the following command to check whether a CP server is up and running at a process level:

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

where `cp_server` is the virtual IP address or virtual hostname on which the CP server is listening.

## Troubleshooting server-based fencing on the SF HA cluster nodes

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs and text files that may be useful in understanding and troubleshooting fencing-related issues on a SF HA cluster (client cluster) node.

## Issues during fencing startup on SF HA cluster nodes set up for server-based fencing

**Table H-1** Fencing startup issues on SF HA cluster (client cluster) nodes

| Issue                                                      | Description and resolution                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cpsadm command on the SF HA cluster gives connection error | <p>If you receive a connection error message after issuing the <code>cpsadm</code> command on the SF HA cluster, perform the following actions:</p> <ul style="list-style-type: none"> <li>■ Ensure that the CP server is reachable from all the SF HA cluster nodes.</li> <li>■ Check that the SF HA cluster nodes use the correct CP server virtual IP or virtual hostname and the correct port number.<br/>Check the <code>/etc/vxfenmode</code> file.</li> <li>■ Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.</li> </ul>         |
| Authorization failure                                      | <p>Authorization failure occurs when the CP server's nodes or users are not added in the CP server configuration. Therefore, fencing on the SF HA cluster (client cluster) node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points.</p> <p>To resolve this issue, add the CP server node and user in the CP server configuration and restart fencing.</p> <p>See <a href="#">“Preparing the CP servers manually for use by the SF HA cluster”</a> on page 178.</p> |

**Table H-1** Fencing startup issues on SF HA cluster (client cluster) nodes  
(continued)

| Issue                  | Description and resolution                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Authentication failure | <p>If you had configured secure communication between the CP server and the SF HA cluster (client cluster) nodes, authentication failure can occur due to the following causes:</p> <ul style="list-style-type: none"> <li>■ Symantec Product Authentication Services (AT) is not properly configured on the CP server and/or the SF HA cluster.</li> <li>■ The CP server and the SF HA cluster nodes use the same root broker but the certificate hash of the root broker is not same on the SF HA cluster and the CP server. Run the following command on both the CP server and the SF HA cluster to see the certificate hash:           <pre># cpsat showalltrustedcreds</pre> </li> <li>■ The CP server and the SF HA cluster nodes use different root brokers, and trust is not established between the authentication brokers:</li> <li>■ The hostname of the SF HA cluster nodes is not the same hostname used when configuring AT.           <p>The hostname of the SF HA cluster nodes must be set to the hostname used when configuring AT. You can view the fully qualified hostname registered with AT using the <code>cpsat showcred</code> command. After entering this command, the hostname appears in the User Name field.</p> </li> <li>■ The CP server and SF HA cluster do not have the same security setting.           <p>In order to configure secure communication, both the CP server and the SF HA cluster must have same security setting.</p> <p>In order to have the same security setting, the security parameter must have same value in the <code>/etc/vxcps.conf</code> file on CP server and in the <code>/etc/vxfenmode</code> file on the SF HA cluster (client cluster) nodes.</p> </li> </ul> |

**Table H-1** Fencing startup issues on SF HA cluster (client cluster) nodes  
*(continued)*

| Issue                   | Description and resolution                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Preexisting split-brain | <p>Assume the following situations to understand preexisting split-brain in server-based fencing:</p> <ul style="list-style-type: none"> <li>■ There are three CP servers acting as coordination points. One of the three CP servers then becomes inaccessible. While in this state, also one client node leaves the cluster. When the inaccessible CP server restarts, it has a stale registration from the node which left the SF HA cluster. In this case, no new nodes can join the cluster. Each node that attempts to join the cluster gets a list of registrations from the CP server. One CP server includes an extra registration (of the node which left earlier). This makes the joiner node conclude that there exists a preexisting split-brain between the joiner node and the node which is represented by the stale registration.</li> <li>■ All the client nodes have crashed simultaneously, due to which fencing keys are not cleared from the CP servers. Consequently, when the nodes restart, the vxfen configuration fails reporting preexisting split brain.</li> </ul> <p>These situations are similar to that of preexisting split-brain with coordinator disks, where the problem is solved by the administrator running the <code>vxfenclearpre</code> command. A similar solution is required in server-based fencing using the <code>cpsadm</code> command.</p> <p>Run the <code>cpsadm</code> command to clear a registration on a CP server:</p> <pre># cpsadm -s cp_server -a unreg_node -c cluster_name -n nodeid</pre> <p>where <code>cp_server</code> is the virtual IP address or virtual hostname on which the CP server is listening, <code>cluster_name</code> is the VCS name for the SF HA cluster, and <code>nodeid</code> specifies the node id of SF HA cluster node. Ensure that fencing is not already running on a node before clearing its registration on the CP server.</p> <p>After removing all stale registrations, the joiner node will be able to join the cluster.</p> |

## Issues during online migration of coordination points

During online migration of coordination points using the `vxfenswap` utility, the operation is automatically rolled back if a failure is encountered during validation of coordination points from all the cluster nodes.

Validation failure of the new set of coordination points can occur in the following circumstances:

- The `/etc/vxfenmode` file is not updated on all the SF HA cluster nodes, because new coordination points on the node were being picked up from an old `/etc/vxfenmode` file.

- The coordination points listed in the `/etc/vxfenmode` file on the different SF HA cluster nodes are not the same. If different coordination points are listed in the `/etc/vxfenmode` file on the cluster nodes, then the operation fails due to failure during the coordination point snapshot check.
- There is no network connectivity from one or more SF HA cluster nodes to the CP server(s).
- Cluster, nodes, or users for the SF HA cluster nodes have not been added on the new CP servers, thereby causing authorization failure.

### Vxfen service group activity after issuing the `vxfenswap` command

After issuing the `vxfenswap` command, the Coordination Point agent reads the details of coordination points from the `vxfenconfig -l` output and starts monitoring the registrations on them.

During `vxfenswap`, when the `vxfenmode` file is being changed by the user, the Coordination Point agent does not move to FAULTED state but continues monitoring the old set of coordination points.

As long as the changes to `vxfenmode` file are not committed or the new set of coordination points are not re-elected in `vxfenconfig -l` output, the Coordination Point agent continues monitoring the old set of coordination points it read from `vxfenconfig -l` output in every monitor cycle.

The status of the Coordination Point agent (either ONLINE or FAULTED) depends upon the accessibility of the coordination points, the registrations on these coordination points, and the fault tolerance value.

When the changes to `vxfenmode` file are committed and reflected in the `vxfenconfig -l` output, then the Coordination Point agent reads the new set of coordination points and proceeds to monitor them in its new monitor cycle.

## Troubleshooting server-based I/O fencing in mixed mode

Use the following procedure to troubleshoot a mixed I/O fencing configuration (configuration which uses both coordinator disks and CP server for I/O fencing).

This procedure uses the following commands to obtain I/O fencing information:

- To obtain I/O fencing cluster information on the coordinator disks, run the following command on one of the cluster nodes:

```
vxfenadm -s diskname
```

Any keys other than the valid keys used by the cluster nodes that appear in the command output are spurious keys.

- To obtain I/O fencing cluster information on the CP server, run the following command on one of the cluster nodes:

```
cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster name* is the VCS name for the SF HA cluster. Nodes which are not in GAB membership, but registered with CP server indicate a pre-existing network partition.

Note that when running this command on the SF HA cluster nodes, you need to first export the CPS\_USERNAME and CPS\_DOMAINTYPE variables.

The CPS\_USERNAME value is the user name which is added for this node on the CP server.

- To obtain the user name, run the following command on the CP server:

```
cpsadm -s cp_server -a list_users
```

where *cp server* is the virtual IP address or virtual hostname on which the CP server is listening.

The CPS\_DOMAINTYPE value is vx.

The following are export variable command examples:

```
export CPS_USERNAME=_HA_VCS_test-system@HA_SERVICES@test-system.symantec.com
```

```
export CPS_DOMAINTYPE=vx
```

Once a pre-existing network partition is detected using the above commands, all spurious keys on the coordinator disks or CP server must be removed by the administrator.

### To troubleshoot server-based I/O fencing configuration in mixed mode

- 1 Review the current I/O fencing configuration by accessing and viewing the information in the `vxfenmode` file.

Enter the following command on one of the SF HA cluster nodes:

```
cat /etc/vxfenmode

vxfen_mode=customized
vxfen_mechanism=cps
scsi3_disk_policy=dmp
security=0
cps1=[10.140.94.101]:14250
vxfendg=vxfencoordg
```

- 2 Review the I/O fencing cluster information.

Enter the `vxfenadm -d` command on one of the cluster nodes:

```
vxfenadm -d

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

Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:

 * 0 (galaxy)
 1 (nebula)

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

**3 Review the SCSI registration keys for the coordinator disks used in the I/O fencing configuration.**

The variables *disk\_7* and *disk\_8* in the following commands represent the disk names in your setup.

Enter the `vxfenadm -s` command on each of the SF HA cluster nodes.

```
vxfenadm -s /dev/vx/rdmp/disk_7
```

```
Device Name: /dev/vx/rdmp/disk_7
Total Number Of Keys: 2
key[0]:
 [Numeric Format]: 86,70,66,69,65,68,48,48
 [Character Format]: VFBEAD00
 [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
 [Numeric Format]: 86,70,66,69,65,68,48,49
 [Character Format]: VFBEAD01
* [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

Run the command on the other node:

```
vxfenadm -s /dev/vx/rdmp/disk_8
```

```
Device Name: /dev/vx/rdmp/disk_8
Total Number Of Keys: 2
key[0]:
 [Numeric Format]: 86,70,66,69,65,68,48,48
 [Character Format]: VFBEAD00
 [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
 [Numeric Format]: 86,70,66,69,65,68,48,49
 [Character Format]: VFBEAD01
* [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

- 4 Review the CP server information about the cluster nodes. On the CP server, run the `cpsadm list nodes` command to review a list of nodes in the cluster.

```
cpsadm -s cp_server -a list_nodes
```

where *cp\_server* is the virtual IP address or virtual hostname on which the CP server is listening.

- 5 Review the CP server list membership. On the CP server, run the following command to review the list membership.

```
cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp\_server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster\_name* is the VCS name for the SF HA cluster.

For example:

```
cpsadm -s 10.140.94.101 -a list_membership -c gl-ss2
```

```
List of registered nodes: 0 1
```

## Checking keys on coordination points when `vxfen_mechanism` value is set to `cps`

When I/O fencing is configured in customized mode and the `vxfen_mechanism` value is set to `cps`, the recommended way of reading keys from the coordination points (coordinator disks and CP servers) is as follows:

- For coordinator disks, the disks can be put in a file and then information about them supplied to the `vxfenadm` command.

For example:

```
vxfenadm -s all -f file_name
```

- For CP servers, the `cpsadm` command can be used to obtain the membership of the SF HA cluster.

For example:

```
cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp\_server* is the virtual IP address or virtual hostname on which CP server is configured, and *cluster\_name* is the VCS name for the SF HA cluster.

## After upgrading from 5.0.x and before migrating SFDB

When upgrading from SFHA version 5.0 to SFHA 5.1 SP1 the S\*vxdms3 startup script is renamed to NO\_S\*vxdms3. The S\*vxdms3 startup script is required by sfua\_rept\_migrate. Thus when sfua\_rept\_migrate is run, it is unable to find the S\*vxdms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdms3 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\*vxdms3 startup script error

- ◆ Rename the startup script NO\_S\*vxdms3 to S\*vxdms3.



## Sample SF HA cluster setup diagrams for CP server-based I/O fencing

This appendix includes the following topics:

- [Configuration diagrams for setting up server-based I/O fencing](#)

### Configuration diagrams for setting up server-based I/O fencing

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:

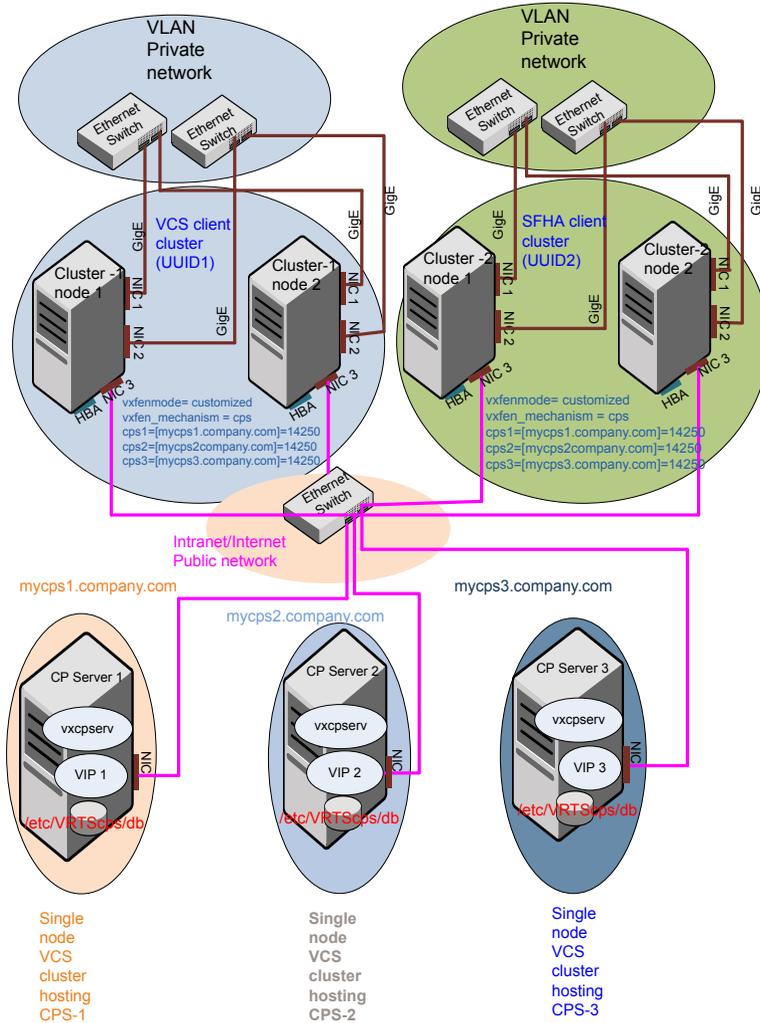
- Two unique client clusters that are served by 3 CP servers:  
See [Figure I-1](#) on page 474.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:

#### Two unique client clusters served by 3 CP servers

[Figure I-1](#) displays a configuration where two unique client clusters are being served by 3 CP servers (coordination points). Each client cluster has its own unique user ID (UUID1 and UUID2).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

**Figure I-1** Two unique client clusters served by 3 CP servers



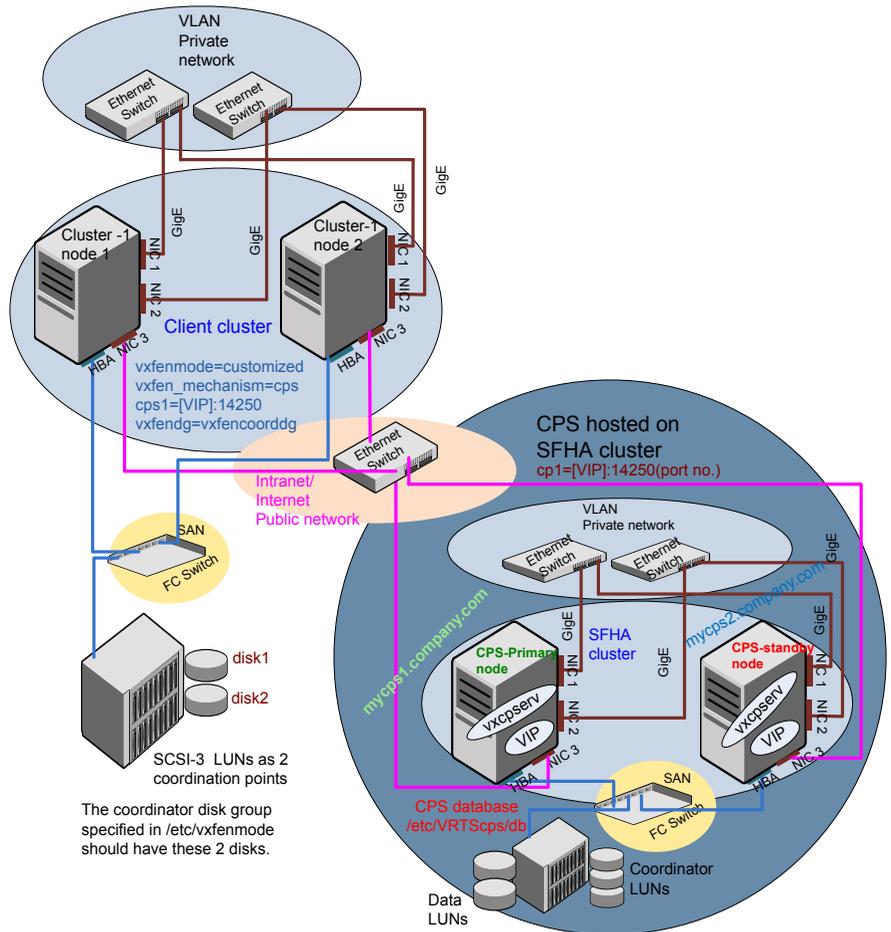
## Client cluster served by highly available CPS and 2 SCSI-3 disks

Figure I-2 displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to customized with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfencoorddg`. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

**Figure I-2** Client cluster served by highly available CP server and 2 SCSI-3 disks



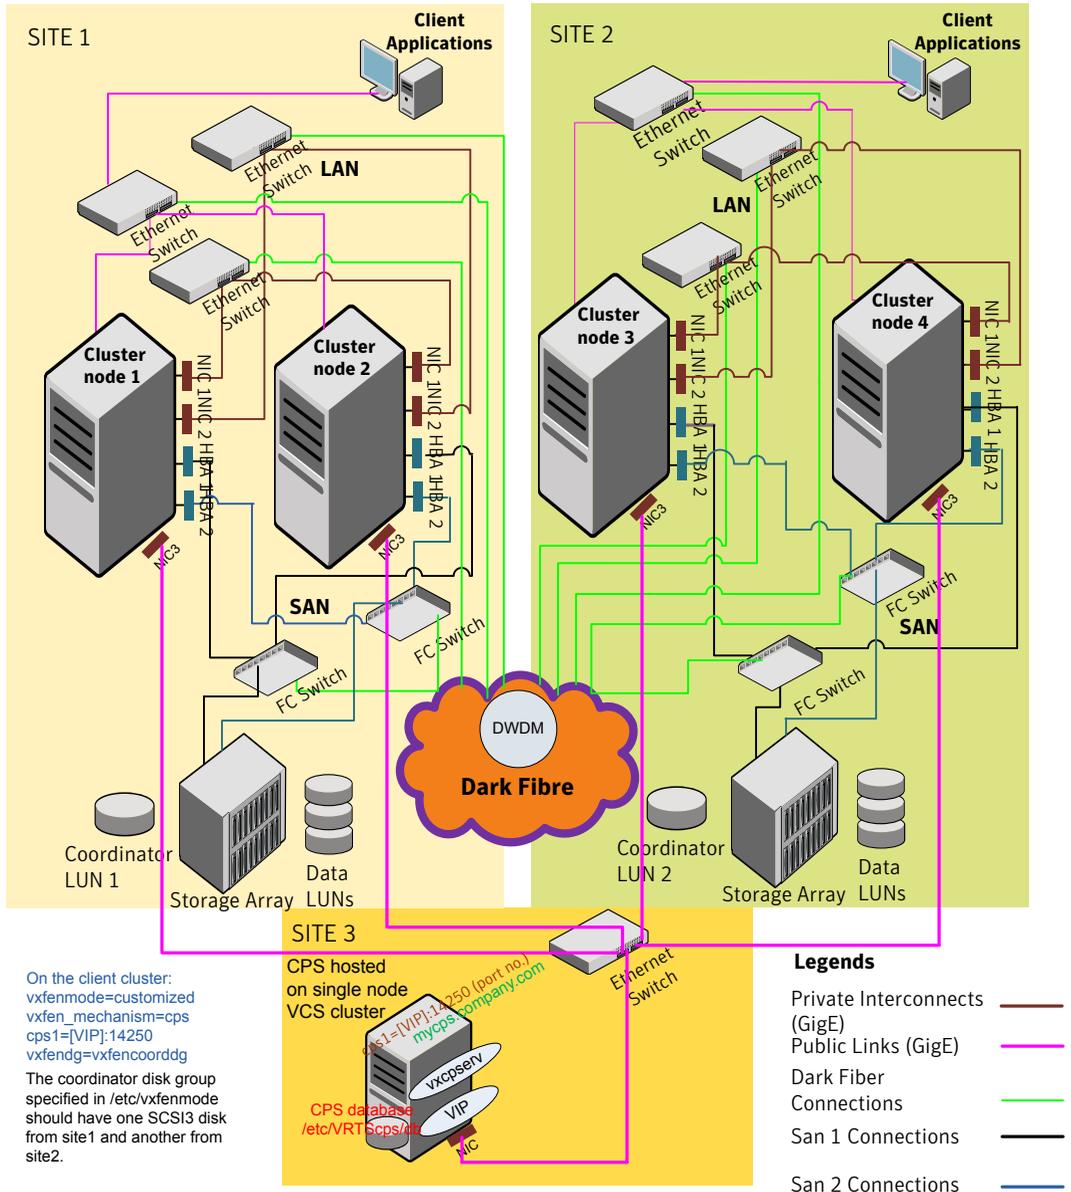
## Two node campus cluster served by remote CP server and 2 SCSI-3 disks

[Figure I-3](#) displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks (one from each site) are part of disk group `vxfencoorddg`. The third coordination point is a CP server on a single node VCS cluster.

**Figure I-3** Two node campus cluster served by remote CP server and 2 SCSI-3



On the client cluster:  
 vxfenmode=customized  
 vxfen\_mechanism=cps  
 cps1=[VIP]:14250  
 vxfendg=vxfencoordg  
 The coordinator disk group specified in /etc/vxfermode should have one SCSI3 disk from site1 and another from site2.

CPS hosted on single node  
 VCS cluster  
 cps1=[VIP]:14250 (port no.)  
 mycps.company.com  
 vxcpssrv  
 VIP  
 CPS database /etc/VRTScps/01  
 NIC

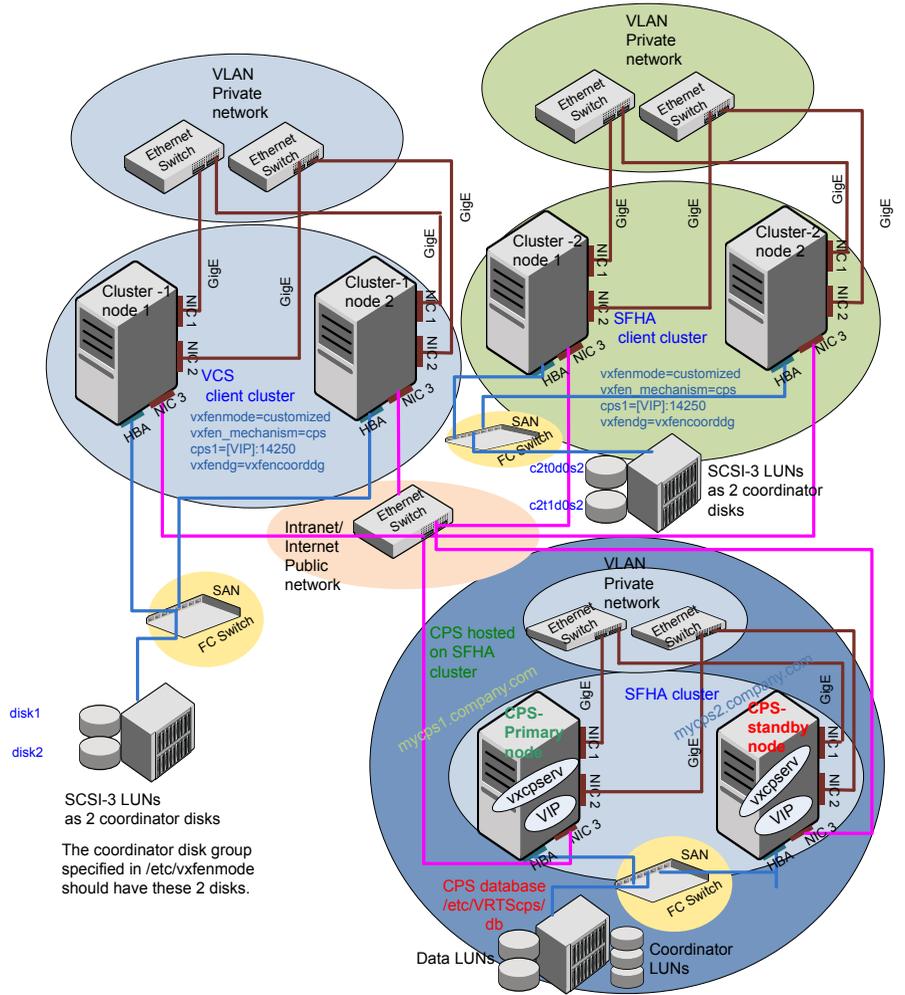
## Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

[Figure I-4](#) displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNS (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfencoordg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

**Figure I-4** Multiple client clusters served by highly available CP server and 2 SCSI-3 disks





# Reconciling major/minor numbers for NFS shared disks

This appendix includes the following topics:

- [Reconciling major/minor numbers for NFS shared disks](#)

## Reconciling major/minor numbers for NFS shared disks

Your configuration may include disks on the shared bus that support NFS. You can configure the NFS file systems that you export on disk partitions or on Veritas Volume Manager volumes. An example disk partition name is `/dev/dsk/c1t1d0`. An example volume name is `/dev/vx/dsk/shreddg/vol3`. Each name represents the block device on which the file system is to be mounted.

In a VCS cluster, block devices providing NFS service must have the same major and minor numbers on each cluster node. Major numbers identify required device drivers (such as a Solaris partition or a VxVM volume). Minor numbers identify the specific devices themselves. NFS also uses major and minor numbers to identify the exported file system.

Major and minor numbers must be verified to ensure that the NFS identity for the file system is the same when exported from each node.

## Checking major and minor numbers for disk partitions

The following sections describe checking and changing, if necessary, the major and minor numbers for disk partitions used by cluster nodes.

### To check major and minor numbers on disk partitions

- ◆ Use the following command on all nodes exporting an NFS file system. This command displays the major and minor numbers for the block device.

```
ls -lL block_device
```

The variable *block\_device* refers to a partition where a file system is mounted for export by NFS. Use this command on each NFS file system. For example, type:

```
ls -lL /dev/dsk/c1t1d0
```

Output on Node A resembles:

```
crw-r----- 1 root sys 32,1 Dec 3 11:50 /dev/dsk/c1t1d0
```

Output on Node B resembles:

```
crw-r----- 1 root sys 32,1 Dec 3 11:55 /dev/dsk/c1t1d0
```

Note that the major numbers (32) and the minor numbers (1) match, satisfactorily meeting the requirement for NFS file systems.

### To reconcile the major numbers that do not match on disk partitions

- 1 Reconcile the major and minor numbers, if required. For example, if the output in the previous section resembles the following, perform the instructions beginning step 2:

Output on Node A:

```
crw-r----- 1 root sys 32,1 Dec 3 11:50 /dev/dsk/c1t1d0
```

Output on Node B:

```
crw-r----- 1 root sys 36,1 Dec 3 11:55 /dev/dsk/c1t1d0
```

- 2 Place the VCS command directory in your path. For example:

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

- 3 Attempt to change the major number on System B (now 36) to match that of System A (32). Use the command:

```
haremajor -sd major_number
```

For example, on Node B, enter:

```
haremajor -sd 32
```

- 4 If the command succeeds, go to step 8.
- 5 If the command fails, you may see a message resembling:

```
Error: Preexisting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

- 6 Notice that the number 36 (the major number on Node A) is not available on Node B. Run the `haremajor` command on Node B and change it to 128,

```
haremajor -sd 128
```

- 7 Run the same command on Node A. If the command fails on Node A, the output lists the available numbers. Rerun the command on both nodes, setting the major number to one available to both.
- 8 Reboot each system on which the command succeeds.
- 9 Proceed to reconcile the major numbers for your next partition.

#### To reconcile the minor numbers that do not match on disk partitions

- 1 In the example, the minor numbers are 1 and 3 and are reconciled by setting to 30 on each node.
- 2 Type the following command on both nodes using the name of the block device:

```
ls -l /dev/dsk/c1t1d0
```

Output from this command resembles the following on Node A:

```
lrwxrwxrwx 1 root root 83 Dec 3 11:50
/dev/dsk/c1t1d0 -> ../../
devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:d,raw
```

The device name (in bold) includes the slash following the word `devices`, and continues to, but does not include, the colon.

- 3 Type the following command on both nodes to determine the instance numbers that the SCSI driver uses:

```
grep sd /etc/path_to_inst | sort -n -k 2,2
```

Output from this command resembles the following on Node A:

```
"/sbus@1f,0/QLGC,isp@0,10000/sd@0,0" 0 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@1,0" 1 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@2,0" 2 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@3,0" 3 "sd"
.
.
"/sbus@1f,0/SUNW,fas@e,8800000/sd@d,0" 27 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@e,0" 28 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@f,0" 29 "sd"
```

In the output, the instance numbers are in the second field.

The instance number that is associated with the device name that matches the name for Node A displayed in step 2, is "1."

- 4 Compare instance numbers for the device in the output on each node.

After you review the instance numbers, perform one of the following tasks:

- If the instance number from one node is unused on the other— it does not appear in the output of step 3—edit `/etc/path_to_inst`. You edit this file to make the second node's instance number similar to the number of the first node.
- If the instance numbers in use on both nodes, edit `/etc/path_to_inst` on both nodes. Change the instance number that is associated with the device name to an unused number. The number needs to be greater than the highest number that other devices use. For example, the output of step 3 shows the instance numbers that all devices use (from 0 to 29). You edit the file `/etc/path_to_inst` on each node and reset the instance numbers to 30.

- 5 Type the following command to reboot each node on which `/etc/path_to_inst` was modified:

```
reboot -- -rv
```

## Checking the major and minor number for VxVM volumes

The following sections describe checking and changing, if necessary, the major and minor numbers for the VxVM volumes that cluster systems use.

### To check major and minor numbers on VxVM volumes

- 1 Place the VCS command directory in your path. For example:

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

- 2 To list the devices, use the `ls -lL block_device` command on each node:

```
ls -lL /dev/vx/dsk/shareddg/vol3
```

On Node A, the output may resemble:

```
brw----- 1 root root 32,43000 Mar 22 16:4 1
/dev/vx/dsk/shareddg/vol3
```

On Node B, the output may resemble:

```
brw----- 1 root root 36,43000 Mar 22 16:4 1
/dev/vx/dsk/shareddg/vol3
```

- 3 Import the associated shared disk group on each node.

- 4 Use the following command on each node exporting an NFS file system. The command displays the major numbers for `vxio` and `vxspec` that Veritas Volume Manager uses. Note that other major numbers are also displayed, but only `vxio` and `vxspec` are of concern for reconciliation:

```
grep vx /etc/name_to_major
```

Output on Node A:

```
vxdump 30
vxio 32
vxspec 33
vxfen 87
vxg1m 91
```

Output on Node B:

```
vxdump 30
vxio 36
vxspec 37
vxfen 87
vxg1m 91
```

- 5 To change Node B's major numbers for `vxio` and `vxspec` to match those of Node A, use the command:

```
haremajor -vx major_number_vxio major_number_vxspec
```

For example, enter:

```
haremajor -vx 32 33
```

If the command succeeds, proceed to step 8. If this command fails, you receive a report similar to the following:

```
Error: Preexisting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

- 6 If you receive this report, use the `haremajor` command on Node A to change the major number (32/33) to match that of Node B (36/37). For example, enter:

```
haremajor -vx 36 37
```

If the command fails again, you receive a report similar to the following:

```
Error: Preexisting major number 36
These are available numbers on this node: 126...
Check /etc/name_to_major on all systems for
available numbers.
```

- 7 If you receive the second report, choose the larger of the two available numbers (in this example, 128). Use this number in the `haremajor` command to reconcile the major numbers. Type the following command on both nodes:

```
haremajor -vx 128 129
```

- 8 Reboot each node on which `haremajor` was successful.
- 9 If the minor numbers match, proceed to reconcile the major and minor numbers of your next NFS block device.
- 10 If the block device on which the minor number does not match is a volume, consult the `vxvg(1M)` manual page. The manual page provides instructions on reconciling the Veritas Volume Manager minor numbers, and gives specific reference to the `reminor` option.

Node where the `vxio` driver number have been changed require rebooting.



# Configuring LLT over UDP using IPv4

This appendix includes the following topics:

- [Using the UDP layer for LLT](#)
- [Manually configuring LLT over UDP using IPv4](#)

## Using the UDP layer for LLT

Storage Foundation 5.1 SP1 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 different physical networks.

If the LLT private links are not on different physical networks, then make sure that the links are on separate subnets. Set the broadcast address in `/etc/llttab` explicitly depending on the subnet for each link.

See [“Broadcast address in the `/etc/llttab` file”](#) on page 490.

- Make sure that each NIC has an IP address that is configured before configuring LLT.
- Make sure the IP addresses in the `/etc/llttab` files are consistent with the IP addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.  
See [“Selecting UDP ports”](#) on page 492.
- Set the broadcast address correctly for direct-attached (non-routed) links.  
See [“Sample configuration: direct-attached links”](#) on page 494.
- For the links that cross an IP router, disable broadcast features and specify the IP address of each link manually in the `/etc/llttab` file.  
See [“Sample configuration: links crossing IP routers”](#) on page 496.

## Broadcast address in the `/etc/llttab` file

The broadcast address is set explicitly for each link in the following example.

- Display the content of the `/etc/llttab` file on the first node galaxy:

```
galaxy # cat /etc/llttab

set-node galaxy
set-cluster 1
link link1 /dev/udp - udp 50000 - 192.168.9.1 192.168.9.255
link link2 /dev/udp - udp 50001 - 192.168.10.1 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

- Display the content of the `/etc/llttab` file on the second node nebula:

```
nebula # cat /etc/llttab

set-node nebula
set-cluster 1
link link1 /dev/udp - udp 50000 - 192.168.9.2 192.168.9.255
link link2 /dev/udp - udp 50001 - 192.168.10.2 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

## The link command in the /etc/llttab file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See “[Sample configuration: direct-attached links](#)” on page 494.
- See “[Sample configuration: links crossing IP routers](#)” on page 496.

[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 <code>/dev/udp</code> .                                                                                                                                  |
| <i>node-range</i>    | Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.                                                                                                                 |
| <i>link-type</i>     | Type of link; must be "udp" for LLT over UDP.                                                                                                                                                             |
| <i>udp-port</i>      | Unique UDP port in the range of 49152-65535 for the link.<br>See “ <a href="#">Selecting UDP ports</a> ” on page 492.                                                                                     |
| <i>MTU</i>           | "-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.              |
| <i>IP address</i>    | IP address of the link on the local node.                                                                                                                                                                 |
| <i>bcast-address</i> | <ul style="list-style-type: none"> <li>■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address.</li> <li>■ "-" is the default for clusters spanning routers.</li> </ul> |

## The set-addr command in the /etc/llttab file

The `set-addr` command in the `/etc/llttab` file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 496.

[Table K-2](#) describes the fields of the `set-addr` command.

**Table K-2** Field description for set-addr command in /etc/llttab

| Field                | Description                                                                  |
|----------------------|------------------------------------------------------------------------------|
| <i>node-id</i>       | The ID of the cluster node; for example, 0.                                  |
| <i>link tag-name</i> | The string that LLT uses to identify the link; for example link1, link2,.... |
| <i>address</i>       | IP address assigned to the link for the peer node.                           |

## Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
  - Ports from the range of well-known ports, 0 to 1023
  - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
netstat -a | more
UDP
 Local Address Remote Address State

 *.sunrpc Idle
 *. Unbound
 *.32771 Idle
 *.32776 Idle
 *.32777 Idle
 *.name Idle
 *.biff Idle
 *.talk Idle
 *.32779 Idle
 .
 .
 .
 *.55098 Idle
 *.syslog Idle
```

```
*.58702 Idle
. Unbound
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

## Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
ifconfig interface_name netmask netmask
```

For example:

- For the first network interface on the node galaxy:

```
IP address=192.168.9.1, Broadcast address=192.168.9.255,
Netmask=255.255.255.0
```

For the first network interface on the node nebula:

```
IP address=192.168.9.2, Broadcast address=192.168.9.255,
Netmask=255.255.255.0
```

- For the second network interface on the node galaxy:

```
IP address=192.168.10.1, Broadcast address=192.168.10.255,
Netmask=255.255.255.0
```

For the second network interface on the node nebula:

```
IP address=192.168.10.2, Broadcast address=192.168.10.255,
Netmask=255.255.255.0
```

## Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.

An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

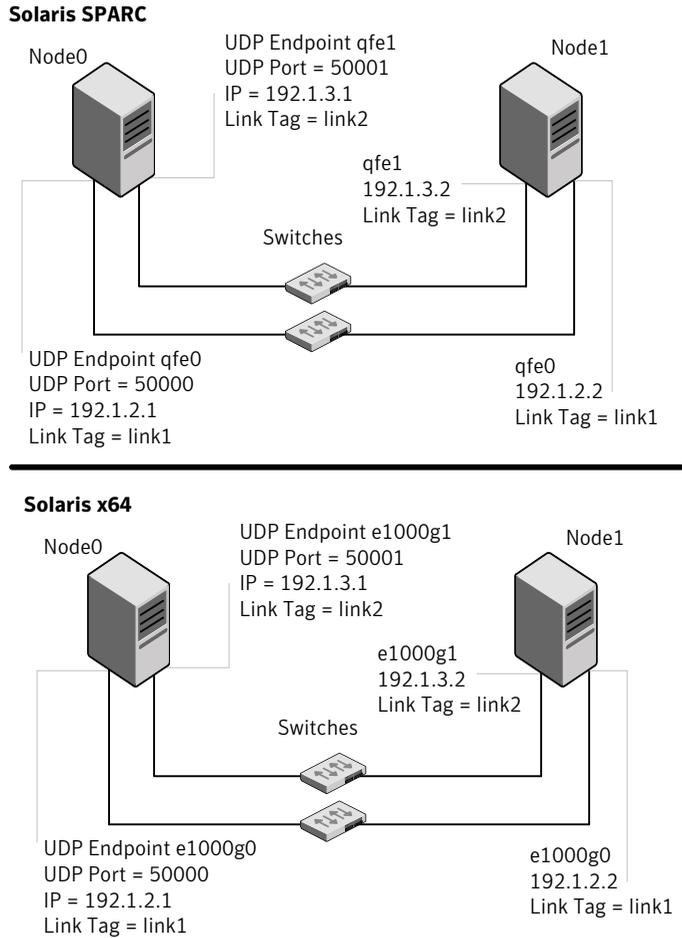
```
cat /etc/llttab
set-node nodexyz
set-cluster 100

link link1 /dev/udp - udp 50000 - 192.168.30.1 192.168.30.255
link link2 /dev/udp - udp 50001 - 192.168.31.1 192.168.31.255
```

## Sample configuration: direct-attached links

[Figure 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 /dev/udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 /dev/udp - udp 50001 - 192.1.3.1 192.1.3.255
```

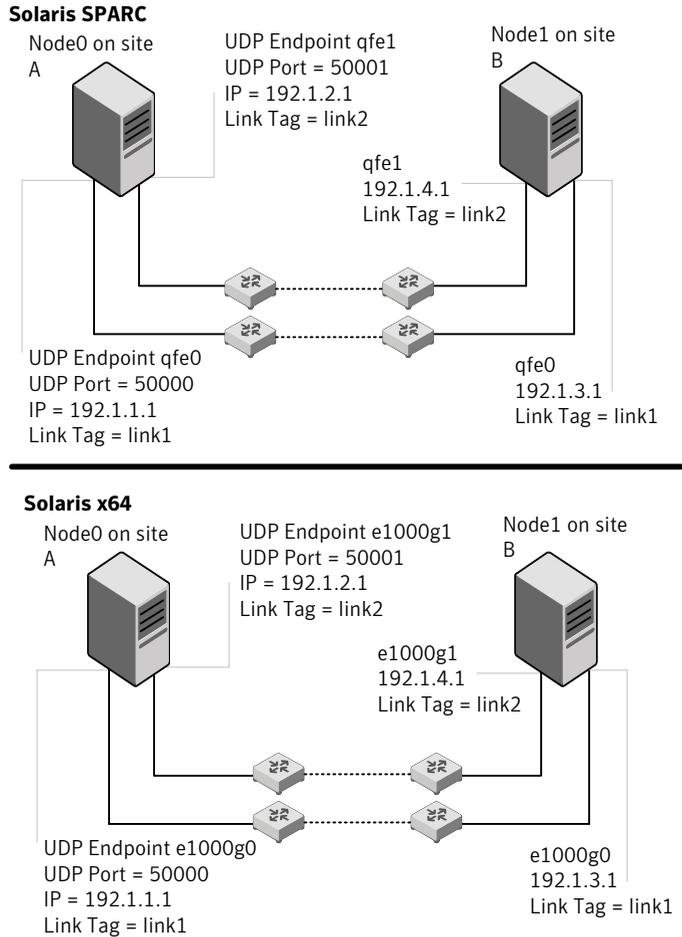
The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /dev/udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 /dev/udp - udp 50001 - 192.1.3.2 192.1.3.255
```

## Sample configuration: links crossing IP routers

[Figure K-2](#) depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

**Figure K-2** A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1
```

```
link link1 /dev/udp - udp 50000 - 192.1.3.1 -
link link2 /dev/udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 0 link1 192.1.1.1
set-addr 0 link2 192.1.2.1
set-addr 2 link1 192.1.5.2
set-addr 2 link2 192.1.6.2
set-addr 3 link1 192.1.7.3
set-addr 3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb 0
set-arp 0
```

**The /etc/llttab file on Node 0 resembles:**

```
set-node Node0
set-cluster 1

link link1 /dev/udp - udp 50000 - 192.1.1.1 -
link link2 /dev/udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 1 link1 192.1.3.1
set-addr 1 link2 192.1.4.1
set-addr 2 link1 192.1.5.2
set-addr 2 link2 192.1.6.2
set-addr 3 link1 192.1.7.3
set-addr 3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb 0
set-arp 0
```

# Configuring LLT over UDP using IPv6

This appendix includes the following topics:

- [Using the UDP layer of IPv6 for LLT](#)
- [Manually configuring LLT over UDP using IPv6](#)

## Using the UDP layer of IPv6 for LLT

Storage Foundation 5.1 SP1 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:

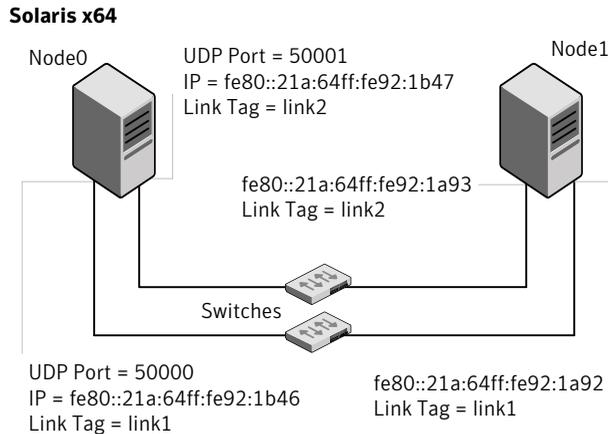
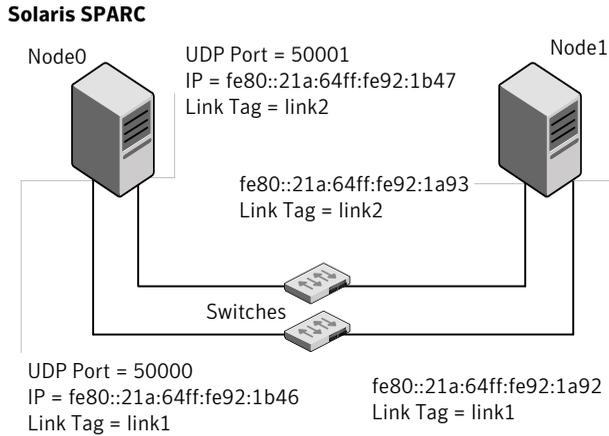
- For UDP6, the multicast address is set to "-".
- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.

- 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 502.

## Sample configuration: direct-attached links

[Figure L-1](#) depicts a typical configuration of direct-attached links employing LLT over UDP.

**Figure L-1** A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
```

```
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```

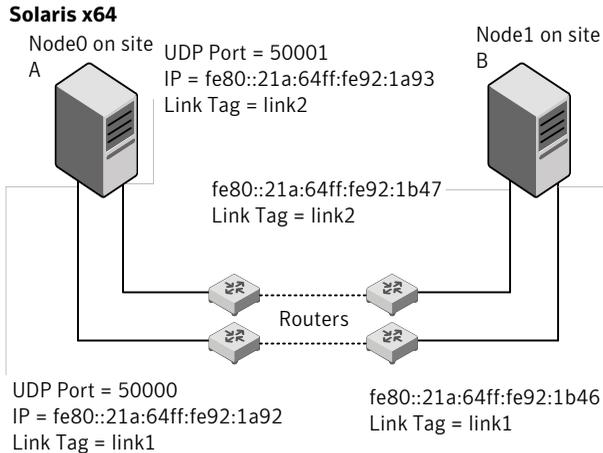
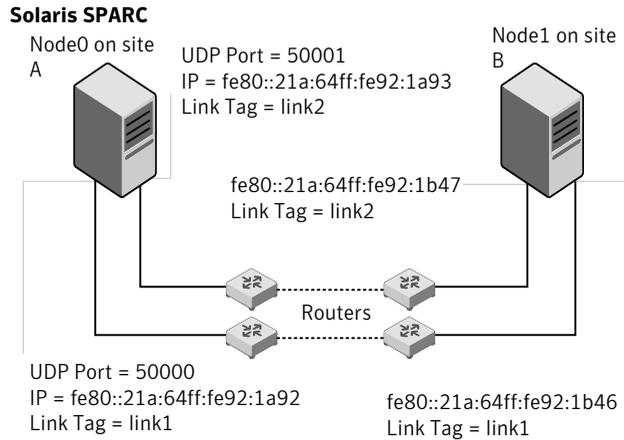
The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

## Sample configuration: links crossing IP routers

**Figure L-2** depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

**Figure L-2** A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

#set address of each link for all peer nodes in the cluster
```

```
#format: set-addr node-id link tag-name address
set-addr 0 link1 fe80::21a:64ff:fe92:1b46
set-addr 0 link2 fe80::21a:64ff:fe92:1b47
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

**The /etc/llttab file on Node 0 resembles:**

```
set-node Node0
set-cluster 1

link link1 /dev/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 /dev/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 1 link1 fe80::21a:64ff:fe92:1a92
set-addr 1 link2 fe80::21a:64ff:fe92:1a93
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

# Index

## A

- about
  - global clusters 30
- adding
  - users 137
- agents
  - disabling 371
- applications, stopping 204
- attributes
  - UseFence 164

## B

- block device
  - partitions
    - example file name 481
  - volumes
    - example file name 481
- bootdg 117

## C

- cables
  - cross-over Ethernet 331
- cluster
  - removing a node from 349
  - verifying operation 322
- cluster functionality
  - enabling 119
  - environment requirements 37
  - shared disks 120
- command failures 460
- commands
  - format 54
  - hastatus 322
  - hasys 323
  - lltconfig 419
  - lltstat 320
  - vxdisksetup (initializing disks) 153
  - vxlicinst 144-145
  - vxlicrep 144

- configuration
  - restoring the original 288
- configuration daemon (vxconfigd)
  - starting 114
- configuring
  - rsh 50
  - shared disks 120
  - ssh 50
- configuring SFHA
  - script-based installer 126
- configuring VCS
  - adding users 137
  - event notification 137, 139
  - global clusters 141
  - secure mode 134
  - starting 127
- controllers
  - SCSI 51
- coordinator disks
  - DMP devices 28
  - for I/O fencing 28
  - setting up 162

## D

- data disks
  - for I/O fencing 28
- default disk group 116
- defaultdg 116
- devices
  - suppress devices 117
- disabling the agents 371
- disk groups
  - bootdg 117
  - default 116
  - nodg 117
  - root 116
  - rootdg 114, 116
- disks
  - adding and initializing 153
  - coordinator 162
  - testing with vxfentsthdw 154

disks (*continued*)

- verifying node access 156

## DMP

- prevent multipathing 117

**E**

- Ethernet controllers 331

**F**

- FC-AL controllers 54

- freezing service groups 204

**G**

## GAB

- description 26

## gabtab file

- verifying after installation 419

## global clusters 30

- configuration 141

**H**

- hastatus -summary command 322

- hasys -display command 323

## hubs

- independent 331

**I**

## I/O daemon (vxiod)

- starting 115

## I/O fencing

- checking disks 154

- setting up 161

- shared storage 154

## Installing

- SFHA with the Web-based installer 70

## installing

- JumpStart 73

- post 142

- Root Broker 87

**J**

## JumpStart

- installing 73

**L**

## language packages

- removal 377

## license keys

- adding with vxlicinst 144

- replacing demo key 145

## licenses

- information about 144

## links

- private network 419

## Live Upgrade

- preparing 257

- upgrade paths 253

- upgrading Solaris on alternate boot disk 262

## LLT

- description 26

- interconnects 27

- verifying 320

## lltconfig command 419

## llthosts file

- verifying after installation 419

## lltstat command 320

## llttab file

- verifying after installation 419

## localized environment settings for using VVR

- settings for using VVR in a localized

- environment 201

## log files 461

**M**

## main.cf file

- contents after installation 425

## main.cf files 431

## major and minor numbers

- checking 482, 485

- shared devices 481

## manual pages

- potential problems 459

- troubleshooting 459

## media speed 27

- optimizing 26

## mounting

- software disc 55

**N**

## NFS services

- shared storage 481

## nodes

- adding application nodes
- configuring GAB 338
- configuring LLT 338
- configuring VXFEN 338
- starting Volume Manager 338

nodg 117

**O**

## optimizing

- media speed 26

## original configuration

- restoring the 288

**P**

## PATH variable

- VCS commands 319

## persistent reservations

- SCSI-3 51

## phased 231

## phased upgrade 231

- example 232

## planning an upgrade from

- previous VVR version 199

## planning to upgrade VVR 199

## preinstallation 199

## preparing

- Live Upgrade 257

## preparing to upgrade VVR 204

## Prevent Multipathing/Suppress Devices from

- VxVMbsxd5 s view 117

## previous VVR version

- planning an upgrade from 199

## problems

- accessing manual pages 459
- executing file system commands 460

**R**

## removing

- the Replicated Data Set 372

## removing a system from a cluster 349

## Replicated Data Set

- removing the 372

## restoring the original configuration 288

## Root Broker

- installing 87

## root disk group 114, 116

## rootdg 116

## rsh 128

- configuration 50

**S**

## script-based installer

- SFHA configuration overview 126

## SCSI driver

- determining instance numbers 483

## SCSI-3

- persistent reservations 51

## SCSI-3 persistent reservations

- verifying 161

## service groups

- freezing 204

- unfreezing 288

## settings for using VVR in a localized environment

- localized environment settings for using VVR 201

## SFHA

- configuring 126

- coordinator disks 162

## SFHA installation

- verifying

- cluster operations 319

- GAB operations 319

- LLT operations 319

## shared disks, configuring 120

## shared storage

- Fibre Channel

- setting up 54

- NFS services 481

## SMTP email notification 137

## SNMP trap notification 139

## ssh 128

- configuration 50

## starting configuration

- installvcs program 128

- Veritas product installer 128

## starting vxconfigd configuration daemon 114

## starting vxiod daemon 115

## stopping

- applications 204

## storage

- setting up shared fibre 54

## suppress devices 117

## Symantec Product Authentication Service 87, 134

## system state attribute value 322

**T**

- troubleshooting
  - accessing manual pages 459
  - executing file system commands 460

**U**

- unfreezing service groups 288
- upgrade
  - phased 231
- upgrade paths
  - Live Upgrade 253
- upgrading
  - clustered environment 121
  - phased 231
- upgrading VVR
  - planning 199
  - preparing 204
- using Live Upgrade 253

**V**

- VCS
  - command directory path variable 319
  - configuration files
    - main.cf 423
- verifying installation
  - kernel component 318
- Veritas Operations Manager 25
- Volume Manager
  - Fibre Channel 54
- vradmin
  - delpri 373
  - stoprep 373
- vvr\_upgrade\_finish script 290
- vxconfigd configuration daemon
  - starting 114
- vxctl mode command 114
- vxdisksetup command 153
- vxinstall program 115-117
- vxinstall program, running 115
- vxiod I/O daemon
  - starting 115
- vxlicinst command 144
- vxlicrep command 144
- vxplex
  - used to remove mirrors of root disk volumes 212

**W**

- Web-based installer 70