

Veritas Storage Foundation™ and High Availability Installation Guide

AIX

5.1

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

Document version: 5.1.0

Legal Notice

Copyright © 2009 Symantec Corporation. All rights reserved.

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

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

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

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

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

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

<http://www.symantec.com>

Technical Support

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

Symantec's maintenance offerings include the following:

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

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

www.symantec.com/techsupp/

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information
- Operating system

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

Licensing and registration

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

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

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

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

Maintenance agreement resources

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

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

Additional enterprise services

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

Enterprise services that are available include the following:

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

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

www.symantec.com

Select your country or language from the site index.

Contents

Technical Support	4	
Chapter 1	About Storage Foundation and High-Availability Solutions	17
	Veritas Storage Foundation product suites	17
	About I/O fencing	19
	About Veritas product licensing	19
Chapter 2	Planning to install the Storage Foundation and High Availability products	23
	About planning for a Storage Foundation installation	23
	About installation and configuration methods	24
	Assessing your system preparedness	25
	Veritas Operations Services	25
	Preinstallation or upgrade planning for Veritas Volume Replicator	25
	Planning an upgrade from the previous VVR version	26
	Downloading the Storage Foundation and High Availability software	27
Chapter 3	System requirements	29
	Hardware and software requirements	29
	I/O fencing requirements	29
	Coordinator disk requirements for I/O fencing	30
	CP server requirements	30
	Release notes	32
	Supported AIX operating systems	32
	Database requirements	33
	Disk space requirements	33

Chapter 4	Preparing your systems for installation	35
	Configuring secure shell (ssh) or remote shell before installing products	35
	Configuring and enabling ssh	36
	Restarting ssh	40
	Enabling rsh for AIX	40
	Setting up shared storage	41
	Setting the SCSI identifier value	41
	Setting up Fiber Channel	42
	Setting environment variables	43
	Mounting a software disc	44
Chapter 5	Installing Storage Foundation and High Availability Solutions using the common product installer	45
	Installation quick reference	45
	About the common product installer	46
	Installing Storage Foundation using the common product installer	47
	Installing Storage Foundation and High Availability Solutions using the common product installer	50
Chapter 6	Installing Storage Foundation and High Availability Solutions using the web-based installer	55
	About the Web-based installer	55
	Features supported with Web-based installer	56
	Features not supported with Web-based installer	57
	Before using the Veritas Web-based installer	57
	Starting the Veritas Web-based installer	57
	Obtaining a security exception on Mozilla Firefox	58
	Performing a pre-installation check with the Veritas Web-based installer	58
	Installing Storage Foundation with the Veritas Web-based installer	59
Chapter 7	Installing Storage Foundation, other methods	61
	Installing Storage Foundation using NIM and the installer	61
	Preparing the bundle and script resources on NIM server	61
	Installing Storage Foundation on the NIM client using SMIT	63

	Installing Storage Foundation and the operating system on the NIM client using SMIT	63
	Installing Storage Foundation on an alternate disk	64
	Preparing to install Storage Foundation on an alternate disk ... 6	5
	Installing Storage Foundation on an alternate disk	65
	Verifying the installation	68
	Installing SF using the mksysb utility	69
	Creating the mksysb backup image	69
	Installing mksysb image on alternate disk	70
	Verifying the installation	71
Chapter 8	Preparing to configure Storage Foundation and High Availability	73
	Preparing to configure the clusters in secure mode	73
	Installing the root broker for the security infrastructure	77
	Creating authentication broker accounts on root broker system	78
	Creating encrypted files for the security infrastructure	79
	Preparing the installation system for the security infrastructure	81
	About configuring Storage Foundation clusters for data integrity	82
	About I/O fencing components	83
	About data disks	83
	About coordination points	83
	About I/O fencing configuration files	84
	About planning to configure I/O fencing	87
	Typical SF HA cluster configuration with server-based I/O fencing	89
	Recommended CP server configurations	89
	About configuring server-based I/O fencing	91
	Setting up the CP server	92
	Installing the CP server using the installer	93
	Configuring security on the CP server	94
	Setting up shared storage for the CP server database	95
	Configuring the CP server using the configuration utility	96
	Configuring the CP server manually	102
	Verifying the CP server configuration	103

Chapter 9	Configuring Storage Foundation and High Availability products	105
	Configuring Storage Foundation and High Availability Solutions	105
	Required information for configuring Storage Foundation and High Availability Solutions	106
	Configuring Storage Foundation High Availability using installsf	106
	Configuring Storage Foundation High Availability using the web-based installer	123
	Configuring Veritas Volume Manager	126
	Configuring DMP support for booting over a SAN	127
	Configuring Veritas File System	128
	Configuring the SFDB repository database	129
	Setting or changing the product level for keyless licensing	129
	Installing Veritas product license keys	130
Chapter 10	Configuring Storage Foundation High Availability for data integrity	131
	Setting up disk-based I/O fencing using installsf	131
	Initializing disks as VxVM disks	131
	Checking shared disks for I/O fencing	133
	Configuring disk-based I/O fencing using installsf	137
	Setting up disk-based I/O fencing manually	139
	Removing permissions for communication	140
	Identifying disks to use as coordinator disks	140
	Setting up coordinator disk groups	140
	Creating I/O fencing configuration files	141
	Modifying VCS configuration to use I/O fencing	142
	Verifying I/O fencing configuration	144
	Setting up server-based I/O fencing using installsf	144
	Verifying security configuration on SF HA cluster to use CP server coordination point	145
	Configuring server-based I/O fencing	147
	Setting up server-based I/O fencing manually	155
	Preparing the CP servers manually for use by the SF HA cluster	155
	Configuring server-based fencing on the SF HA cluster manually	159
	Configuring Coordination Point agent to monitor coordination points	163
	Verifying server-based I/O fencing configuration	165

Chapter 11	Upgrading Storage Foundation	167
	Upgrading Storage Foundation products or the operating system	167
	Planning the upgrade	168
	Saving system information before upgrade	168
	About upgrading the Storage Foundation for Databases (SFDB) tools to 5.1	169
	Upgrade paths	169
	Performing pre-installation checks and configuration	170
	Pre-upgrade tasks for migrating the SFDB repository database	171
	Preparing for an upgrade of Storage Foundation	172
	Preparing for upgrade of VVR in the presence of VCS agents	173
	Verifying that the file systems are clean	177
	Upgrading the array support	178
	Upgrading Storage Foundation products to 5.1	179
	Upgrading Veritas Storage Foundation with the product installer	179
	Upgrading Storage Foundation with the Veritas Web-based installer	181
	Upgrading Storage Foundation and AIX on a DMP-enabled rootvg	182
	Upgrading using SMIT	183
	Upgrading the AIX operating system	184
	Upgrading Veritas Volume Replicator	186
	Upgrading VVR without disrupting replication	186
	Post-upgrade tasks	187
	Optional configuration steps	187
	Recovering VVR if automatic upgrade fails	188
	Post-upgrade tasks when VCS Agents for VVR are configured	188
	Upgrading disk layout versions	191
	Post upgrade tasks for migrating the SFDB repository database	192
	Upgrading VxVM disk group versions	198
	Updating variables	199
	Setting the default disk group	199
	Verifying the Veritas Storage Foundation upgrade	200

Chapter 12	Upgrading Storage Foundation using an alternate disk	201
	About upgrading Storage Foundation using an alternate disk	201
	Supported upgrade scenarios	202
	Supported upgrade paths	202
	Preparing to upgrade Storage Foundation on an alternate disk	203
	Upgrading Storage Foundation on an alternate disk	204
	Upgrading a cluster that is in secure mode	214
	Upgrading Storage Foundation and AIX on an alternate rootvg that is enabled for DMP on an LVM root disk	215
	Configuring fencing for an ADI upgrade	215
	Configuring fencing in disabled mode for an ADI upgrade	215
	Configuring fencing in SCSI-3 mode for an ADI upgrade	216
	Verifying the upgrade	217
Chapter 13	Performing a phased upgrade of Storage Foundation High Availability	219
	About phased upgrade	219
	Prerequisites for a phased upgrade	220
	Planning for a phased upgrade	220
	Phased upgrade limitations	220
	Phased upgrade example	220
	Phased upgrade example overview	221
	Performing a phased upgrade from Storage Foundation 5.0 MP3	222
	Moving the service groups to the second subcluster	222
	Upgrading the operating system on the first subcluster	226
	Upgrading the first subcluster	226
	Preparing the second subcluster	228
	Activating the first subcluster	232
	Upgrading the operating system on the second subcluster	234
	Upgrading the second subcluster	234
	Finishing the phased upgrade	235
Chapter 14	Verifying the Storage Foundation installation	239
	Verifying that the products were installed	239
	Installation log files	240
	Using the installation log file	240
	Using the summary file	240
	Starting and stopping processes for the Veritas products	240
	Checking Volume Manager processes	241
	Checking Veritas File System installation	241

About enabling LDAP authentication for clusters that run in secure mode	241
Enabling LDAP authentication for clusters that run in secure mode	243
About the LLT and GAB configuration files	249
Verifying the LLT, GAB, and VCS configuration files	251
Verifying LLT, GAB, and cluster operation	251
Verifying LLT	252
Verifying the cluster	255
Verifying the cluster nodes	256

Chapter 15

Adding and removing nodes in Storage Foundation and High Availability clusters	259
About adding and removing nodes	259
Adding nodes using the Storage Foundation installer	259
Manually adding a node to a cluster	262
Setting up the hardware	263
Installing the Storage Foundation software manually when adding a node	264
Setting up the node to run in secure mode	265
Configuring LLT and GAB	267
Configuring I/O fencing on the new node	269
Adding the node to the existing cluster	274
Starting Storage Foundation and verifying the cluster	275
Removing a node from a cluster	275
Verifying the status of nodes and service groups	276
Deleting the departing node from Storage Foundation configuration	277
Modifying configuration files on each remaining node	280
Removing the node configuration from the CP server	280
Removing security credentials from the leaving node	281
Unloading LLT and GAB and removing VCS on the departing node	281
Adding a node to a single-node cluster	282
Setting up a node to join the single-node cluster	283
Installing and configuring Ethernet cards for private network	283
Configuring the shared storage	284
Bringing up the existing node	284
Installing the Storage Foundation software manually when adding a node to a single node cluster	285
Creating configuration files	285

	Starting LLT and GAB	286
	Reconfiguring Storage Foundation on the existing node	286
	Verifying configuration on both nodes	287
Chapter 16	Uninstalling Storage Foundation	289
	About removing Veritas Storage Foundation	289
	Preparing to uninstall a Storage Foundation product	290
	Moving volumes to physical disks	291
	Disabling the agents on a system	292
	Removing the Replicated Data Set	293
	Uninstalling Storage Foundation with the Veritas Web-based installer	295
	Uninstalling Storage Foundation filesets using the script-based installer	296
	Removing Storage Foundation products using SMIT	297
	Removing the CP server configuration using the removal script	298
	Removing the Storage Foundation for Databases (SFDB) repository after removing the product	301
Appendix A	Installation scripts	303
	About installation scripts	303
	Installation script options	304
Appendix B	Response files	309
	About response files	309
	About the installation simulator	310
	Installing Storage Foundation using response files	311
	Configuring Storage Foundation using response files	311
	Upgrading Storage Foundation using response files	312
	Uninstalling Storage Foundation using response files	313
	Syntax in the response file	313
	Response file variable definitions	314
	Sample response file for SFHA configuration	320
	Sample response file for SFHA install	320
	Sample response file for SF upgrade	321
	Sample response file for SFHA upgrade	321
Appendix C	Configuring I/O fencing using a response file	323
	Response file variables to configure disk-based I/O fencing	323
	Sample response file for configuring disk-based I/O fencing	324
	Configuring I/O fencing using response files	325

	Response file variables to configure server-based I/O fencing	326
	Sample response file for configuring server-based I/O fencing	328
Appendix D	Storage Foundation and High Availability components	329
	Veritas Storage Foundation installation filesets	329
	Veritas Cluster Server installation filesets	331
	Veritas Storage Foundation obsolete and reorganized installation filesets	332
Appendix E	Troubleshooting installation issues	335
	Restarting the installer after a failed connection	335
	What to do if you see a licensing reminder	335
	Troubleshooting an installation on AIX	336
	Incorrect permissions for root on remote system	336
	Resource temporarily unavailable	337
	Inaccessible system	338
	Storage Foundation Cluster File System problems	338
	High availability issues	338
Appendix F	Troubleshooting cluster installation	341
	Unmount failures	341
	Command failures	341
	Installer cannot create UUID for the cluster	342
	The vxfsentsthdw utility fails when SCSI TEST UNIT READY command fails	342
	Troubleshooting on the CP server	343
	CP server service group issues	343
	Testing the connectivity of the CP server	344
	Troubleshooting server-based I/O fencing on the SF HA cluster	344
	Issues during server-based fencing start up on SF HA cluster node	344
	Issues during online migration of coordination points	345
	Troubleshooting server-based I/O fencing in mixed mode	346
	Checking keys on coordination points when vxfsen_mechanism value is set to cps	351

Appendix G	Sample SF HA cluster setup diagrams for CP server-based I/O fencing	353
	Configuration diagrams for setting up server-based I/O fencing	353
	Two unique client clusters served by 3 CP servers	353
	Client cluster served by highly available CPS and 2 SCSI-3 disks	354
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	356
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	357
Appendix H	Changing NFS server major numbers for VxVM volumes	359
	Changing NFS server major numbers for VxVM volumes	359
Appendix I	Configuring LLT over UDP using IPv4	361
	Using the UDP layer for LLT	361
	When to use LLT over UDP	361
	Configuring LLT over UDP	361
	Broadcast address in the /etc/llttab file	362
	The link command in the /etc/llttab file	363
	The set-addr command in the /etc/llttab file	363
	Selecting UDP ports	364
	Configuring the netmask for LLT	365
	Configuring the broadcast address for LLT	365
	Sample configuration: direct-attached links	366
	Sample configuration: links crossing IP routers	367
Appendix J	Configuring LLT over UDP using IPv6	369
	Using the UDP layer of IPv6 for LLT	369
	When to use LLT over UDP	369
	Configuring LLT over UDP using IPv6	372
	Sample configuration: direct-attached links	370
	Sample configuration: links crossing IP routers	371
	Configuring LLT over UDP using IPv6	372
	The link command in the /etc/llttab file	373
	The set-addr command in the /etc/llttab file	374
	Selecting UDP ports	374
Index	377

About Storage Foundation and High-Availability Solutions

This chapter includes the following topics:

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

Veritas Storage Foundation product suites

The following table lists the Symantec products and optionally licensed features available with each Veritas Storage Foundation product suite.

In the 5.1 release of Storage Foundation, the database utilities are included in the Storage Foundation release, rather than as options.

Table 1-1 Contents of Veritas Storage Foundation products

Storage Foundation version	Products and features
Storage Foundation Basic	Veritas File System Veritas Volume Manager

Table 1-1 Contents of Veritas Storage Foundation products (*continued*)

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

About I/O fencing

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

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

The fencing operation determines the following:

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

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

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

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

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

About Veritas product licensing

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

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

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

- Install a license key for the product and features that you want to install.
When you purchase a Symantec product, you receive a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased.
- Continue to install without a license key.
The installer prompts for the product modes and options that you want to install, and then sets the required product level.
Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled or continue with keyless licensing by managing the server or cluster with a management server. If you do not comply with the above terms, continuing to use the Veritas product is a violation of your end user license agreement, and results in warning messages. For more information about keyless licensing, see the following URL:
<http://go.symantec.com/sfhakeyless>

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

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

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.
See “[Setting or changing the product level for keyless licensing](#)” on page 129.
See the `vxkeyless(1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the 5.1 products you have purchased.
See “[Installing Veritas product license keys](#)” on page 130.
See the `vxlicinst(1m)` manual page.

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

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

We recommend updating to keyless licensing for the following reasons:

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

Planning to install the Storage Foundation and High Availability products

This chapter includes the following topics:

- [About planning for a Storage Foundation installation](#)
- [About installation and configuration methods](#)
- [Assessing your system preparedness](#)
- [Preinstallation or upgrade planning for Veritas Volume Replicator](#)
- [Downloading the Storage Foundation and High Availability software](#)

About planning for a Storage Foundation installation

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

http://sfdoccentral.symantec.com/sf/5.1/aix/sf_install.pdf

This document is version 5.1.1.

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 Storage Foundation will be installed.

Follow the preinstallation instructions if you are installing one of the Veritas Storage Foundation 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 Storage Foundation products.

See “[Veritas Storage Foundation product suites](#)” on page 17.

About installation and configuration methods

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

Use one of the following methods to install and configure Storage Foundation:

- The Veritas product installer (Recommended)
The common product installer displays a menu that simplifies the selection of installation options.
See “[About the common product installer](#)” on page 46.
- The product-specific installation scripts
The installation scripts provide a command-line interface to installing a specific product. The product-specific scripts enable you to specify some additional command-line options. Otherwise, installing with the installation script is identical to specifying Storage Foundation from the common product installer menu.
- The Web-based Veritas installer
The installer provides an interface to manage the installation from a remote site using a standard Web browser.
In this release, there are some limitations in the Web-based installer.
See “[About the Web-based installer](#)” on page 55.
- Silent installation with response files
You can use any of the above options to generate a response file. You can then customize the response file for another system. Run the product installation script with the response file to install silently on one or more other systems.
See “[About response files](#)” on page 309.
- Network Installation Manager (NIM)

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

Assessing your system preparedness

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

Veritas Operations Services

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

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

Simulation option of the Veritas product installer

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

See “[About the installation simulator](#)” on page 310.

Veritas Operations Services

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

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

To access VOS, go to:

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

Preinstallation or upgrade planning for Veritas Volume Replicator

Before installing or upgrading VVR:

- Confirm that your system has enough free disk space to install VVR.
- Make sure you have root permissions. You must have root permissions to perform the install and upgrade procedures.

The following related documents are available:

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

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

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

Planning an upgrade from the previous VVR version

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

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

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

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

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

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

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

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.

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

See [“About installation scripts”](#) on page 303.

To download the software

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

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

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

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

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

```
# df -k filesystem
```

Caution: When you select a location to download files, do not select a directory that contains Veritas products from a previous release or maintenance pack. You must download the Veritas 5.0 software and the Veritas 5.1 software into separate directories.

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

System requirements

This chapter includes the following topics:

- [Hardware and software requirements](#)
- [I/O fencing requirements](#)
- [Release notes](#)
- [Supported AIX operating systems](#)
- [Database requirements](#)
- [Disk space requirements](#)

Hardware and software requirements

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

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

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

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

I/O fencing requirements

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

- Coordinator disks

See “Coordinator disk requirements for I/O fencing” on page 30.

- CP servers
See “CP server requirements” on page 30.

Coordinator disk requirements for I/O fencing

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

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

CP server requirements

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

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

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

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

Table 3-1 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: <ul style="list-style-type: none"> ■ 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 server(s) to the SF HA cluster.

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

Table 3-2 CP server supported operating systems and versions

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

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

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

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

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

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

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

For additional information, see *Veritas Cluster Server User's Guide*.

Release notes

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

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

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

Supported AIX operating systems

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

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

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

The minimum system requirements for this release are as follows:

AIX 5.3 at one of the following levels:

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

AIX 6.1 at one of the following levels:

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

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

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

Database requirements

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

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

Note: Storage Foundation supports running Oracle, DB2, and Sybase on VxFS and VxVM.

Storage Foundation does not support running SFDB tools with DB2 and Sybase.

Disk space requirements

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

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

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


Preparing your systems for installation

This chapter includes the following topics:

- [Configuring secure shell \(ssh\) or remote shell before installing products](#)
- [Setting up shared storage](#)
- [Setting environment variables](#)
- [Mounting a software disc](#)

Configuring secure shell (ssh) or remote shell before installing products

Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. These utilities must run as `root` on all cluster nodes or remote systems.

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

This section contains an example of how to set up ssh password free communication. The example sets up ssh between a source system (system1) that contains the installation directories, and a target system (system2). This procedure also applies to multiple target systems.

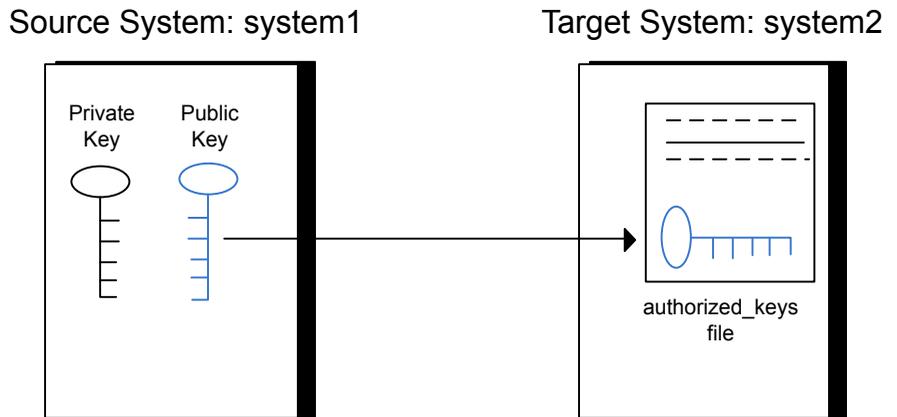
Configuring and enabling ssh

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

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

Figure 4-1 illustrates this procedure.

Figure 4-1 Creating the DSA key pair and appending it to target systems



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

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

To create the DSA key pair

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

```
system1 # cd /
```

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

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

System output similar to the following is displayed:

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

- 3** Press Enter to accept the default location of `/.ssh/id_dsa`.
- 4** When the program asks you to enter the passphrase, press the Enter key twice.

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

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

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

```
system2 # cd /  
system2 # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

```
system2 # chmod go-w /.ssh
```

To append the public key from the source system to the authorized_keys file on the target system, using secure file transfer

- 1 From the source system (system1), move the public key to a temporary file on the target system (system2).

Use the secure file transfer program.

In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

Use the following command for secure file transfer:

```
system1 # sftp system2
```

If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

```
Connecting to system2 ...
The authenticity of host 'system2 (10.182.00.00)'
can't be established. DSA key fingerprint is
fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
Are you sure you want to continue connecting (yes/no)?
```

- 2 Enter `yes`.

Output similar to the following is displayed:

```
Warning: Permanently added 'system2,10.182.00.00'
(DSA) to the list of known hosts.
root@system2 password:
```

- 3 Enter the root password of system2.
- 4 At the `sftp` prompt, type the following command:

```
sftp> put /.ssh/id_dsa.pub
```

The following output is displayed:

```
Uploading /.ssh/id_dsa.pub to /id_dsa.pub
```

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

```
sftp> quit
```

- 6** To begin the `ssh` session on the target system (system2 in this example), type the following command on system1:

```
system1 # ssh system2
```

Enter the root password of system2 at the prompt:

```
password:
```

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

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

- 8** After the `id_dsa.pub` public key file is copied to the target system (system2), and added to the authorized keys file, delete it. To delete the `id_dsa.pub` public key file, type the following command on system2:

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

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

```
system2 # exit
```

- 10** When you install from a source system that is also an installation target, also add the local system `id_dsa.pub` key to the local `authorized_keys` file. The installation can fail if the installation source system is not authenticated.

To add the local system `id_dsa.pub` key to the local `authorized_keys` file, enter the following command:

```
system1 # cat /.ssh/id_dsa.pub >> /.ssh/authorized_keys
```

- 11** Run the following commands on the source installation system. If your `ssh` session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and make the key globally available for the user `root`:

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

```
Identity added: //./ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

To verify that you can connect to a target system

- 1 On the source system (system1), type the following command:

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

where system2 is the name of the target system.

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

Restarting ssh

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

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

To restart ssh

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

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

- 2 Make the key globally available for the user root

```
system1 # ssh-add
```

Enabling rsh for AIX

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

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

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

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

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

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

Setting up shared storage

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

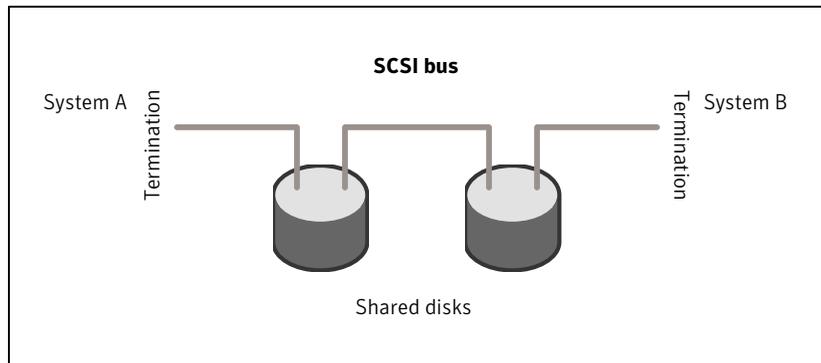
See also the *Veritas Cluster Server User's Guide* for a description of I/O fencing.

Setting the SCSI identifier value

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

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

Figure 4-2 Cabling the shared storage



To set the SCSI identifier value

- 1 Determine the SCSI adapters on each system:

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

- 2 Verify the SCSI ID of each adapter:

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

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

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

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

Setting up Fiber Channel

Perform the following steps to set up fiber channel.

To set up fiber channel

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

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

- 2 Reboot each system:

```
shutdown -Fr
```

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

Setting environment variables

Most of the commands used in the installation are in the `/sbin` or `/usr/sbin` directory. Add these directories to your `PATH` environment variable as necessary.

After installation, Veritas Storage Foundation commands are stored in `/opt/VRTS/bin` and HA commands are stored in `/opt/VRTSvcs/bin`. Storage Foundation HA manual pages are stored in `/opt/VRTS/man`.

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

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

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

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

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

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

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

Mounting a software disc

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

To mount the software disc

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

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

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

Installing Storage Foundation and High Availability Solutions using the common product installer

This chapter includes the following topics:

- [Installation quick reference](#)
- [About the common product installer](#)
- [Installing Storage Foundation using the common product installer](#)
- [Installing Storage Foundation and High Availability Solutions using the common product installer](#)

Installation quick reference

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

[Table 5-1](#) provides a quick overview of a stand-alone installation using the product installer.

Table 5-1 Installation overview

Installation task	For more information, refer to the following section:
Obtain product licenses.	
Download the software, or insert the product DVD.	See “Downloading the Storage Foundation and High Availability software” on page 27. See “Mounting a software disc” on page 44.
Set environment variables.	See “Setting environment variables” on page 43.
Configure the secure shell (SSH) on all nodes.	See “Configuring secure shell (ssh) or remote shell before installing products” on page 35.
Verify that hardware, software, and operating system requirements are met.	
Check that sufficient disk space is available.	See “Disk space requirements” on page 33.
Use the installer to install the products.	See “Installing Storage Foundation using the common product installer” on page 47.

About the common product installer

The product installer is the recommended method to license and install the Veritas products. The installer also enables you to configure the product, verify preinstallation requirements, and view the product’s description.

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

See [“About installation scripts”](#) on page 303.

At most points during an installation, you can type `b` (back) to return to a previous section of the installation procedure. The back feature of the installation scripts is context-sensitive, so it returns to the beginning of a grouped section of questions. If an installation procedure hangs, use `Control-C` to stop and exit the program. After a short delay, the script exits. You can also enter `q` to quit the installer or `?` to display help information.

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

Additional options are available for the common product installer.

See [“Installation script options”](#) on page 304.

Installing Storage Foundation using the common product installer

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

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

To install Storage Foundation

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

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 35.

- 2 Load and mount the software disc.

See [“Mounting a software disc”](#) on page 44.

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

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

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

```
# ./installer
```

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

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

- 8 Select from one of the following install options:

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

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

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

```
1) Install minimal Storage Foundation filesets -  
   235 MB required  
2) Install recommended Storage Foundation filesets -  
   450 MB required  
3) Install all Storage Foundation filesets -  
   596 MB required  
4) Display filesets to be installed for each option  
  
Select the filesets to be installed on all systems?  
[1-4,q,?] (2)
```

- 9** You are prompted to enter the system names (in the following example, "host1") on which the software is to be installed. Enter the system name or names and then press Return.

```
Enter the platform system names separated by spaces: host1
```

Where *platform* indicates the operating system.

10 You are prompted to choose your licensing method.

To comply with the terms of Symantec's End User License Agreement, you have 60 days to either:

- * Enter a valid license key matching the functionality in use on the systems
- * Enable keyless licensing and manage the systems with a Management Server. For more details visit <http://go.symantec.com/sfhakeyless>. The product is fully functional during these 60 days.

- 1) Enter a valid license key
- 2) Enable keyless licensing and complete system licensing later

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

If you have a valid license key, select 1 and enter the license key at the prompt. Skip to step 13.

To install using keyless licensing, select 2. You are prompted for the product modes and the options that you want to install and license.

Note: The keyless license option enables you to install without entering a key. However, 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>

11 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,?] (1) **1**

12 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**

- 13** 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-****
```

- 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
```

Installing Storage Foundation and High Availability Solutions using the common product 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 [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 35.

- 2** Load and mount the software disc.

See [“Mounting a software disc”](#) on page 44.

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

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

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

```
# ./installer
```

- 5** Enter `␣` to install and press Return.

- 6 When the list of available products is displayed, select Veritas Storage Foundation High Availability (SF HA), 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 EULA.pdf file present on the media? [y,n,q,?] **y**

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

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

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

```
1) Install minimal Storage Foundation HA filesets -
   719 MB required
2) Install recommended Storage Foundation HA filesets -
   944 MB required
3) Install all Storage Foundation HA filesets -
   993 MB required
4) Display filesets to be installed for each option
```

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

- 9 You are prompted to enter the system names (in the following example, "host1" and "host2") on which the software is to be installed. Enter the system name or names and then press Return.

Enter the *platform* system names separated by spaces: **host1 host2**

Where *platform* indicates the operating system.

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

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

See “[Configuring secure shell \(ssh\) or remote shell before installing products](#)” on page 35.

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

12 You are prompted to choose your licensing method.

To comply with the terms of Symantec's End User License Agreement, you have 60 days to either:

- * Enter a valid license key matching the functionality in use on the systems
 - * Enable keyless licensing and manage the systems with a Management Server. For more details visit <http://go.symantec.com/sfhakeyless>. The product is fully functional during these 60 days.
- 1) Enter a valid license key
 - 2) Enable keyless licensing and complete system licensing later

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

If you have a valid license key, select 1 and enter the license key at the prompt. Skip to step 16 .

To install using keyless licensing, select 2. You are prompted for the product modes and the options that you want to install and license.

Note: The keyless license option enables you to install without entering a key. However, you must still have a valid license to install and use Veritas products.

Keyless licensing requires that you manage the systems with a Management Server.

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

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

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

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.

See [“Configuring Storage Foundation and High Availability Solutions”](#) on page 105.

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

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

17 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-****
```

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

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

y

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

This chapter includes the following topics:

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

About the Web-based installer

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

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

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

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

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

Features supported with Web-based installer

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

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

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

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:

- Simulating any of the previously listed tasks
- Configuring Authentication (AT)
- Configuring VxSS security for VCS
- Adding a node to a cluster
- Configuring the I/O fencing feature
- Uninstalling or configuring from the installation server rather than from the media.
- Installing SFRAC
- Configuring SFRAC
- Upgrading VCS
- Upgrading SFHA

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 6-1 Web-based installer requirements

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

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

- 1 Start the Veritas XPortal Server process `xprt1d`, on the installation server:

```
# ./webinstaller start
```

The webinstaller script displays a URL.

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

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

Obtaining a security exception on Mozilla Firefox

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

To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **Add Exception** button.
- 3 Click **Get Certificate** button.
- 4 Uncheck **Permanently Store this exception checkbox (recommended)**.
- 5 Click **Confirm Security Exception** button.
- 6 Enter `root` in User Name field and root password of the web server in the Password field.

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

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

To perform a pre-installation check

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

Installing Storage Foundation with the Veritas Web-based installer

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

To install Storage Foundation

- 1 Perform preliminary steps.
- 2 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 57.
- 3 On the Select a task and a product page, select **Install a Product** from the **Task** drop-down list.
- 4 Select Storage Foundation or Storage Foundation High Availability from the Product drop-down list, and click Next.
- 5 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 6 Choose minimal, recommended, or all packages. Click **Next**.
- 7 Indicate the systems on which to install. Enter one or more system names, separated by spaces. Click **Validate**.

- 8 After the validation completes successfully, click **Next** to install Storage Foundation 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.

For Storage Foundation High Availability, choose whether you want to enable Global Cluster option.

Click **Register**.

- Enter license key

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

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

For Storage Foundation High Availability, the installer prompts you to configure the cluster.

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

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

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

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

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

Installing Storage Foundation, other methods

This chapter includes the following topics:

- [Installing Storage Foundation using NIM and the installer](#)
- [Installing Storage Foundation on an alternate disk](#)
- [Installing SF using the mksysb utility](#)

Installing Storage Foundation using NIM and the installer

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

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

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

Preparing the bundle and script resources on NIM server

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

To prepare the bundle and script resources

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

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

- 3 Navigate to the cluster_server directory on the disc and run the `installsf` command to prepare the bundle and script resources:

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

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

Output resembles:

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

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

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

Output resembles:

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

```
# lsnim -l copy_cpi
```

Output resembles:

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

Installing Storage Foundation on the NIM client using SMIT

You can install Storage Foundation on the NIM client using the SMIT tool.

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

To install Storage Foundation

- 1 On the NIM client, start smitty.

```
# smitty install
```

- 2 In the menu, select **Network Installation Management**.
- 3 In the menu, select **Install and Update Software**.
- 4 In the menu, select **Install Software Bundle**.
- 5 In the menu, select the LPP_SOURCE. In this example, specify **LPP-6100-up2date**.
- 6 In the menu, select the bundle. In this example, specify the **SF51_bundle**.
- 7 For the customization script that you want to run after installation, specify **copy_cpi**.
- 8 For the installp flags, specify that the ACCEPT new license agreements flag has a **yes** value.
- 9 Press the Enter key to start the installation. Note that it may take some time to finish.

Installing Storage Foundation and the operating system on the NIM client using SMIT

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

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

To install Storage Foundation and the operating system

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

```
# smitty nim_bosinst
```

- 2 In the menu, select the stand-alone target.
- 3 In the menu, select **rte- Install from installation images**.

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

Installing Storage Foundation on an alternate disk

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

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

The installation process involves the following steps:

Preparing to install Storage Foundation on an alternate disk	See “Preparing to install Storage Foundation on an alternate disk” on page 65.
Installing Storage Foundation on an alternate disk	See “Installing Storage Foundation on an alternate disk” on page 65.
Verifying the installation	See “Verifying the installation” on page 68.

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

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

Preparing to install Storage Foundation on an alternate disk

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

To prepare to install Storage Foundation on an alternate disk

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

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

```
# lspv
```

Output similar to the following displays:

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

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

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

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

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

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

Installing Storage Foundation on an alternate disk

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

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

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

To install Storage Foundation on an alternate disk using the SMIT interface

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

Start the SMIT menu :

```
# smit alt_clone
```

Provide information for the following fields in the SMIT menu.

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

```
hdisk1
```

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

```
MyBundle
```

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

```
/usr/
```

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

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

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

The installation process takes some time.

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

```
# lspv
Output similar to the following displays:
hdisk0          0009710fa9c79877    rootvg
hdisk1          0009710f0b90db93    altinst_rootvg
```

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

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

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

```
# shutdown -r
```

- 6 Verify the installation.

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

- 7 Configure Storage Foundation.

To install Storage Foundation on an alternate disk manually

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

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

Where:

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

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

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

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

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

```
# shutdown -r
```

- 5 Configure Storage Foundation.

Verifying the installation

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

To verify the installation

- 1 Verify that the alternate boot environment is active:

```
# lspv  
Output similar to the following displays:  
hdisk0          0009710fa9c79877    old_rootvg  
hdisk1          0009710f0b90db93    rootvg
```

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

```
# lspp -Lb MyBundle
```

- 3 In a cluster environment, make sure that all the GAB ports are up.

```
# gabconfig -a  
Port a gen 39d901 membership 01  
Port b gen 39d905 membership 01  
Port d gen 39d904 membership 01  
Port f gen 39d90f membership 01  
Port h gen 39d909 membership 01  
Port o gen 39d903 membership 01  
Port v gen 39d90b membership 01  
Port w gen 39d90d membership 01
```

Installing SF using the mksysb utility

On AIX, one can use the `mksysb` utility to back up the system image. This image can be installed on another host. For example, you can use this utility to set up a disaster recovery site. Storage Foundation can be installed through `mksysb` image.

You can install the `mksysb` image on the same machine or on any NIM client through a NIM server. This procedure assumes working knowledge of `mksysb`. See your operating system installation guide for more details about `mksysb`.

The installation process involves the following steps:

- Creating the `mksysb` image.
- Installing the SF stack through `mksysb` image on a machine.
- Verifying the installation.

Creating the mksysb backup image

You can create the `mksysb` backup image with the SMIT interface or with manual steps.

Before you begin, make sure that the SF installation media is available.

To create an `mksysb` image using SMIT interface

- 1 Check maximum file size limit with `ulimit`. It should be sufficient for creating backup image
- 2 Check that all the required filesets are installed for a particular product stack. You can obtain the list of packages from the `pkginfo.txt` file that is available on the product disc in the corresponding Veritas product stack directory.

The recommended approach is to install all of the packages but do not configure product stack before taking `mksysb` image if the image is to be installed on a different machine.

- 3 Enter fast path `smitty mksysb` and enter the required values.
- 4 Press enter to start the backup image creation.

To create an `mksysb` image using commands manually

- 1 Check maximum file size limit with `ulimit`. It should be sufficient for creating backup image
- 2 Check that all the required file sets are installed for a particular product stack. You can obtain the list of packages from the `pkginfo.txt` file that is available on the product disc in the corresponding Veritas product stack directory.

The recommended approach is to install all of the packages but do not configure product stack before taking `mksysb` image if the image is to be installed on a different machine.

- 3 Enter command

```
/usr/bin/mksysb '-i' '-X' backup file name
```

Installing `mksysb` image on alternate disk

You can install the `mksysb` image on the same machine or on any NIM client through a NIM server.

On same machine

To install SF with `mksysb` on an alternate disk of the same machine using SMIT

- 1 Type `smitty` and then select Software Installation and Maintenance -> Alternate Disk Installation -> Install `mksysb` on an Alternate Disk
- 2 Select target disks
- 3 Select `mksysb` image to be installed
- 4 Select appropriate values for remaining options
- 5 Press enter to start the `mksysb` image installation.
- 6 After installation is complete reboot from the alternate disk.
- 7 If SF was not configured in the `mksysb` image then run `installsf -configure` after reboot.

To install SF with `mksysb` on an alternate disk of the same machine using commands manually

- ◆ To install SF with `mksysb` on an alternate disk of the same machine using commands manually

```
/usr/sbin/alt_disk_mksysb -m mksysb_image -P "all" -d "disk_name"
```

On NIM client

To install SF with mksysb on an alternate disk of the same machine using SMIT

- 1 Create an `mksysb` resource from the `mksysb` image created above on NIM server.
- 2 Set up the machine on which you want to install SF as NIM client.
- 3 Type `smitty nim` then select Perform NIM Software Installation and Maintenance Tasks -> Alternate Disk Installation -> Install `mksysb` on an Alternate Disk
- 4 Select target machine.
- 5 Select target disks.
- 6 Select `mksysb` image to be installed.
- 7 Select appropriate values for remaining options.
- 8 Press enter to start the `mksysb` image installation.
- 9 If SF was not configured in the `mksysb` image then run `installsf -configure` after rebooting NIM client.

To install SF with mksysb on an alternate disk of a NIM client using commands manually

- 1 Create an `mksysb` resource from the `mksysb` image created above on NIM server.
- 2 Set up the machine on which you want to install SF as NIM client.
- 3 To install SF with `mksysb` on an alternate disk of a NIM client using commands manually

```
/usr/sbin/nim -o alt_disk_install \  
-a source=mksysb -a mksysb=mksysb_resource -a disk=hdisk_name machine_
```

- 4 If SF was not configured in the `mksysb` image then run `installsf -configure` after rebooting NIM client.

Verifying the installation

After the installation is finished, verify the installation using the following command:

```
lspp -l
```

All the packages should be installed properly.

Preparing to configure Storage Foundation and High Availability

This chapter includes the following topics:

- [Preparing to configure the clusters in secure mode](#)
- [About configuring Storage Foundation clusters for data integrity](#)
- [About I/O fencing components](#)
- [About I/O fencing configuration files](#)
- [About planning to configure I/O fencing](#)
- [About configuring server-based 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 Storage Foundation configuration.

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

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

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

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

- A system in your enterprise that serves as root broker (RB).
You can either use an external system as root broker, or use one of the cluster nodes as root broker.
- To use an external root broker, identify an existing root broker system in your enterprise or install and configure root broker on a stable system.
See [“Installing the root broker for the security infrastructure”](#) on page 77.
- 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 `installsf`, 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 78.
- The system clocks of the external root broker and authentication brokers must be in sync.

The `installsf` provides the following configuration modes:

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

[Figure 8-1](#) depicts the flow of configuring Storage Foundation cluster in secure mode.

Figure 8-1 Workflow to configure Storage Foundation cluster in secure mode

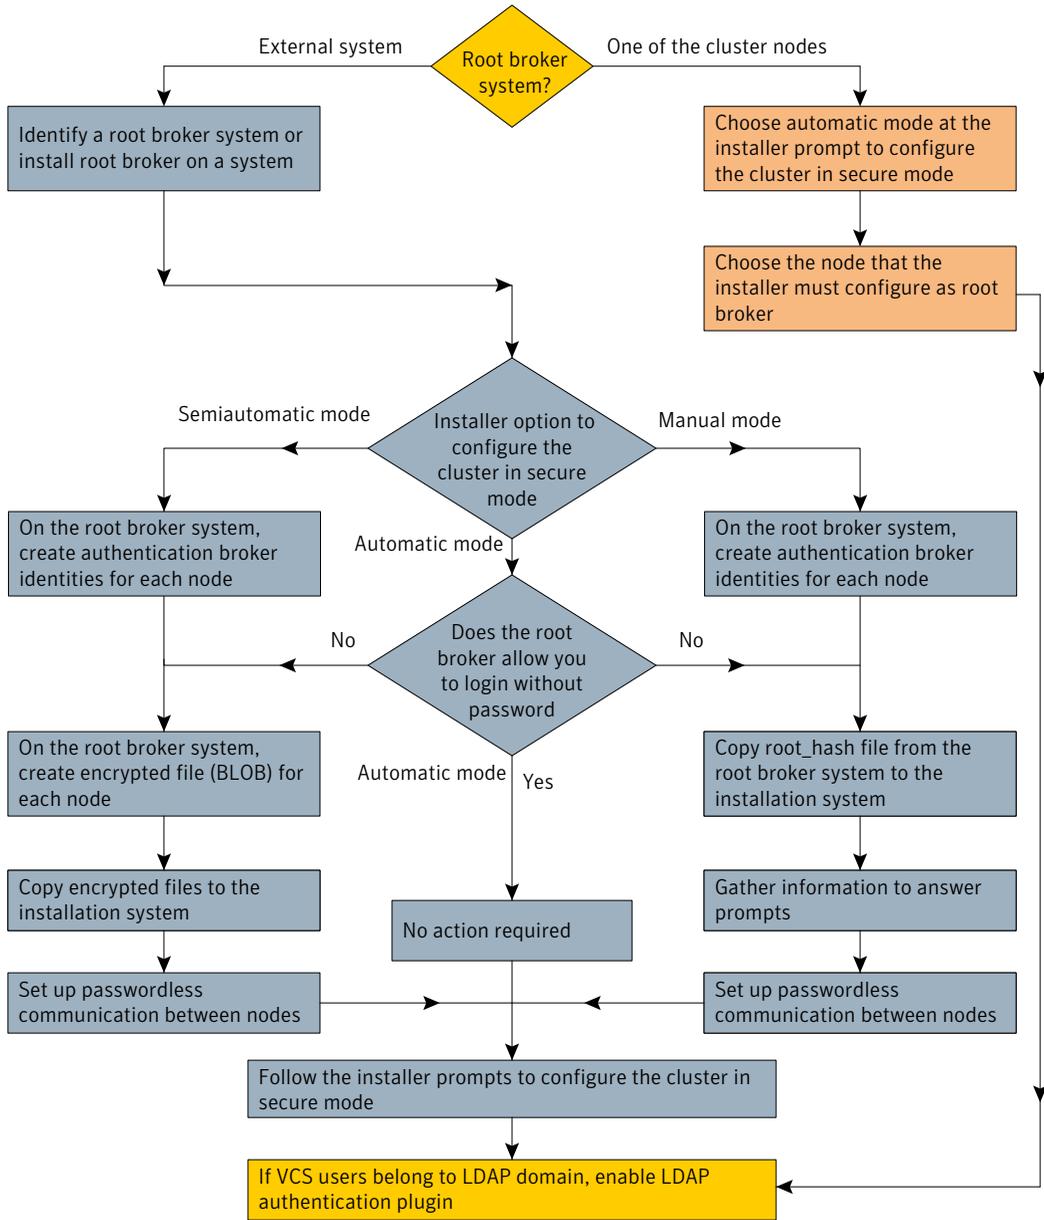


Table 8-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform. These preparatory tasks apply only when you use an external root broker system for the cluster.

Table 8-1 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 77.</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 78.</p> <p>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 79.</p> <p>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

Table 8-1 Preparatory tasks to configure a cluster in secure mode (with an external root broker) (*continued*)

Tasks	Who performs this task
<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 Storage Foundation.</p> <p>See “Preparing the installation system for the security infrastructure” on page 81.</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 Storage Foundation. 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.
- 2 From the Task Menu, choose I for "Install a Product."
- 3 From the displayed list of products to install, choose: Symantec Product Authentication Service (AT).
- 4 Enter **y** to agree to the End User License Agreement (EULA).
- 5 Enter **2** to install the recommended packages.
- 6 Enter the name of the system where you want to install the Root Broker.

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

- 7 Review the output as the installer does the following:

- Checks to make sure that Storage Foundation supports the operating system
- Checks if the filesets are already on the system.

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

- 8 Review the output as the installer installs the root broker on the system.
- 9 After the installation, configure the root broker.
- 10 Select to configure the root broker from the three choices that the installer presents:

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

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

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

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

Creating authentication broker accounts on root broker system

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

To create authentication broker accounts on root broker system

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

```
venus> # vssat showalltrustedcreds
```

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

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

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

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

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

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

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

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

Proceed to step 3.

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

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

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

Creating encrypted files for the security infrastructure

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

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 The value for the authentication broker identity, which you provided to create authentication broker principal on the root broker system.

This is the value for the `--prplname` option of the `addprpl` command.

See [“Creating authentication broker accounts on root broker system”](#) on page 78.

password The value for the authentication broker password, which you provided to create authentication broker principal on the root broker system.

This is the value for the `--password` option of the `addprpl` command.

See [“Creating authentication broker accounts on root broker system”](#) on page 78.

broker_admin_password The value for the authentication broker password for Administrator account on the node. This password must be at least five characters.

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

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

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

```
[configab]
identity=galaxy
password=password
```

```
root_domain=vx:root@venus.symantecexample.com
root_broker=venus.symantecexample.com:2821

start_broker=false
enable_pbx=false
```

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

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

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

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

For example:

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

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

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

Preparing the installation system for the security infrastructure

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

To prepare the installation system for the security infrastructure

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

Automatic mode Do the following:

- Gather the root broker system name from the AT administrator.
- During Storage Foundation configuration, choose the configuration option 1 when the installsf prompts.

- | | |
|---------------------|--|
| Semi-automatic mode | <p>Do the following:</p> <ul style="list-style-type: none">■ 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 Storage Foundation configuration, choose the configuration option 2 when the installsf prompts. |
| Manual mode | <p>Do the following:</p> <ul style="list-style-type: none">■ 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 Storage Foundation configuration, choose the configuration option 3 when the installsf prompts. |

About configuring Storage Foundation clusters for data integrity

When a node fails, Storage Foundation takes corrective action and configures its components to reflect the altered membership. If an actual node failure did not occur and if the symptoms were identical to those of a failed node, then such corrective action would cause a split-brain situation.

Some example scenarios that can cause such split-brain situations are as follows:

- **Broken set of private networks**
If a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects. The remaining node then takes corrective action. The failure of the private interconnects, instead of the actual nodes, presents identical symptoms and causes each node to determine its peer has departed. This situation typically results in data corruption because both nodes try to take control of data storage in an uncoordinated manner
- **System that appears to have a system-hang**
If a system is so busy that it appears to stop responding, the other nodes could declare it as dead. This declaration may also occur for the nodes that use the

hardware that supports a "break" and "resume" function. When a node drops to PROM level with a break and subsequently resumes operations, the other nodes may declare the system dead. They can declare it dead even if the system later returns and begins write operations.

I/O fencing is a feature that prevents data corruption in the event of a communication breakdown in a cluster. Storage Foundation 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 Storage Foundation, you must configure I/O fencing in Storage Foundation to ensure data integrity.

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

About I/O fencing components

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

- Data disks—Store shared data
See [“About data disks”](#) on page 83.
- Coordination points—Act as a global lock during membership changes
See [“About coordination points”](#) on page 83.

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 are part of standard VxVM or CVM disk groups.

CVM is responsible for fencing data disks on a disk group basis. Disks that are added to a disk group and new paths that are discovered for a device are automatically fenced.

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. Racing for control of the coordination points to fence data disks is the key to understand how fencing prevents split-brain.

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

■ **Coordinator disks**

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

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

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

■ **Coordination point servers**

The coordination point server (CP server) is a software solution which runs on a remote system or cluster. CP server provides arbitration functionality by allowing the 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 activeSF HA cluster
- Self-unregister from this activeSF 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 I/O fencing configuration files

[Table 8-2](#) lists the I/O fencing configuration files.

Table 8-2 I/O fencing configuration files

File	Description
/etc/default/vxfen	<p>This file stores the start and stop environment variables for I/O fencing:</p> <ul style="list-style-type: none"> ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: <ul style="list-style-type: none"> 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: <ul style="list-style-type: none"> 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of Storage Foundation HA configuration.</p>
/etc/vxfendg	<p>This file includes the coordinator disk group information.</p> <p>This file is not applicable for server-based fencing.</p>

Table 8-2 I/O fencing configuration files (*continued*)

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

Table 8-2 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	<p>When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.</p> <p>Note: The /etc/vxfentab file is a generated file; do not modify this file.</p> <p>For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:</p> <ul style="list-style-type: none"> ■ Raw disk: <pre data-bbox="373 656 541 736">/dev/rhdisk75 /dev/rhdisk76 /dev/rhdisk77</pre> ■ DMP disk: <pre data-bbox="373 821 642 900">/dev/vx/rdmp/rhdisk75 /dev/vx/rdmp/rhdisk76 /dev/vx/rdmp/rhdisk77</pre> <p>For server-based fencing, the /etc/vxfentab file also includes the security settings information.</p>

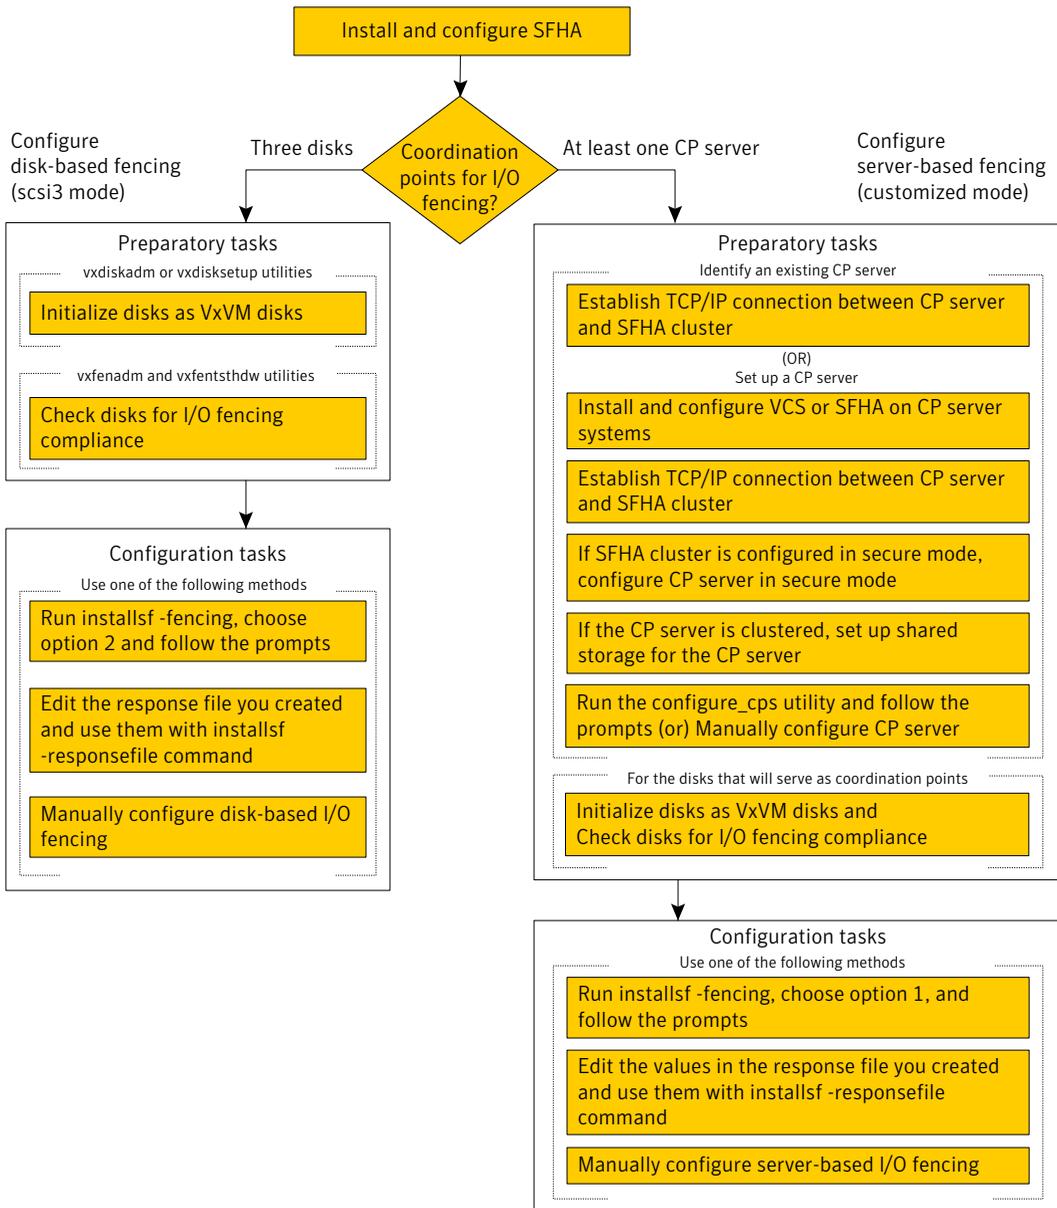
About planning to configure I/O fencing

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

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

Figure 8-2 illustrates a high-level flowchart to configure I/O fencing for the Storage Foundation HA cluster.

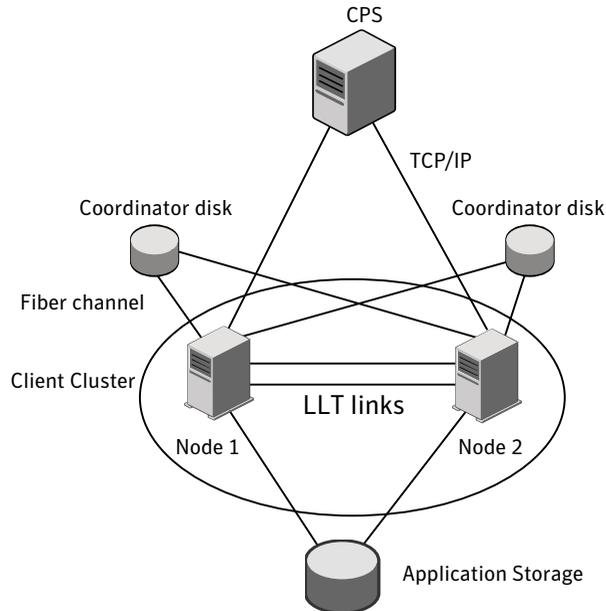
Figure 8-2 Workflow to configure I/O fencing



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

Figure 8-3 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 8-3 CP server, SF HA cluster, and coordinator disks



Recommended CP server configurations

This section discusses the following recommended CP server configurations:

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

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

Figure 8-4 displays a configuration using a single CP server that is connected to multiple SF HA clusters with each SF HA cluster also using two coordinator disks.

Figure 8-4 Single CP server connecting to multiple SF HA clusters

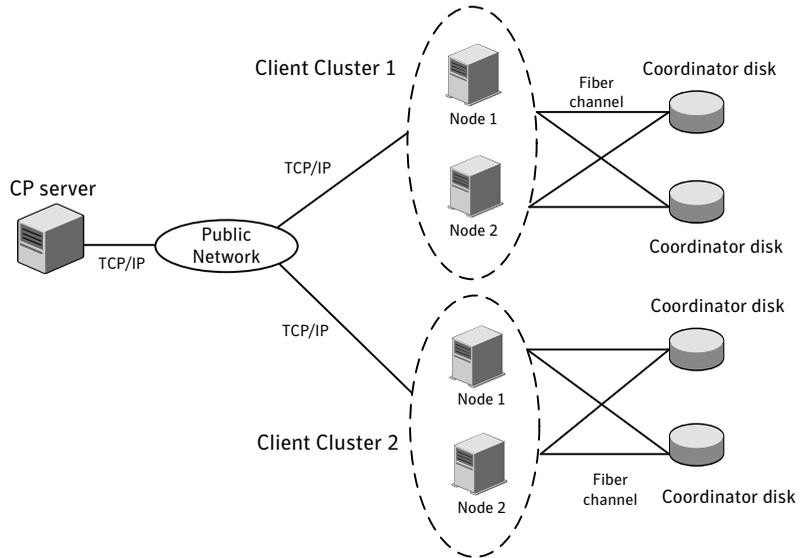
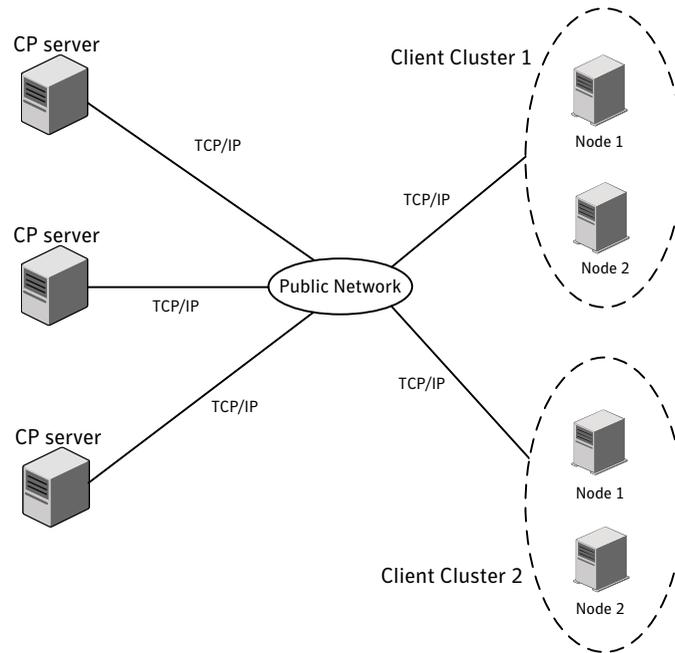


Figure 8-5 displays a configuration using 3 CP servers that are connected to multiple SF HA clusters.

Figure 8-5 Three CP servers connecting to multiple SF HA clusters



For detailed deployment diagrams for server-based fencing:

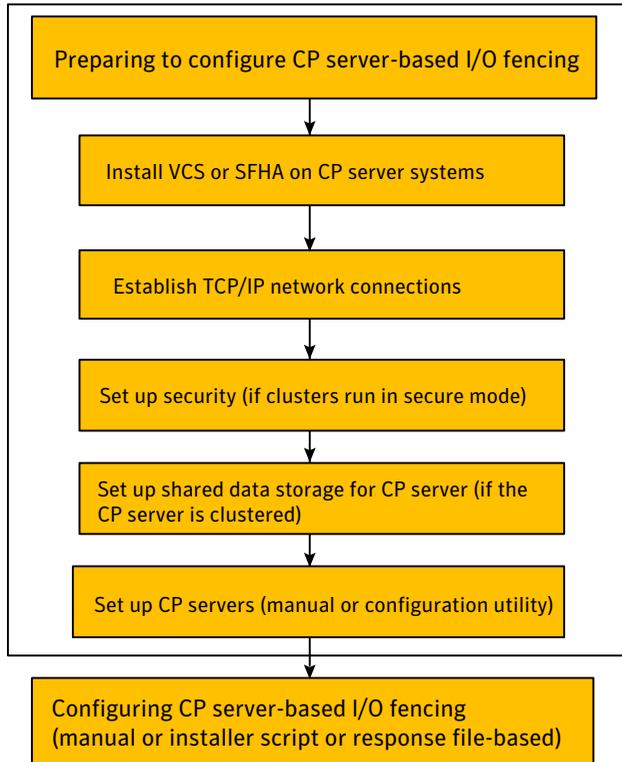
See [“Configuration diagrams for setting up server-based I/O fencing”](#) on page 353.

About configuring server-based I/O fencing

You can configure the CP server using the CP server configuration utility. Alternatively, you can configure the CP server manually.

[Figure 8-6](#) displays the steps to be performed to configure CP server using the configuration utility or manually.

Figure 8-6 Configuring CP server using the configuration utility or manually



See: [Setting up the CP server](#)

See: [Setting up server-based I/O fencing using installsf](#)

See: [Setting up server-based I/O fencing manually](#)

See: [Configuring I/O fencing using response files](#)

Setting up the CP server

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

To prepare to configure the CP server

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

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

- 2 If the CP server is hosted on an SFHA cluster, configure fencing in enabled mode during the SFHA configuration using either the installer or manually.
- 3 Decide if you want to secure the communication between the CP server and SF HA clusters using the Symantec Product Authentication Service (AT).

Symantec recommends setting up security for the CP server and SF HA cluster communications.

For information about configuring security on the CP server:

See [“Configuring security on the CP server”](#) on page 94.

- 4 Choose a name for the CP server.

The CP server name should not contain any special characters.

- 5 Choose a port number for the CP server.

Allocate a TCP/IP port for use by the CP server.

The default port number is 14250. Alternatively, the user can specify any other valid port from the following valid port range: 49152-65535.

- 6 If CP server is hosted on an SFHA cluster, then set up shared storage for the CP server database.

For information about setting up shared storage for the CP server database:

See [“Setting up shared storage for the CP server database”](#) on page 95.

- 7 Choose a valid virtual IP address, network interface, and netmask for the CP server.

Installing the CP server using the installer

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

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

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

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

- 1 Review the CP server hardware and networking requirements, and set up the CP server hardware and network.
- 2 Establish network connections between the CP server(s) and the SF HA clusters through the TCP/IP network. This step requires that you have valid IP addresses, hostnames, and netmasks set up for the CP servers.
- 3 For installing CP server on a single node:
 - Install VCS 5.1 onto the system where you are installing the CP server. Installing VCS 5.1 also installs CP server on the system. Refer to the *Veritas™ Cluster Server Installation Guide, Version 5.1* for instructions on installing VCS 5.1.

When installing VCS 5.1, be sure to select the complete installation option and not the minimum package installation option. The VRTScps package is only part of the complete installation.

- 4 For installing CP server to be hosted on an SFHA cluster:
 - Install SFHA 5.1 onto each system where you are installing CP server to be hosted on a cluster. Installing SFHA 5.1 also installs CP server on the system. Refer to the *Veritas Storage Foundation™ and High Availability Installation Guide* for instructions on installing SFHA 5.1.

When installing SFHA 5.1, be sure to select the complete installation option and not the minimum package installation option. The VRTScps package is only part of the complete installation.

- 5 Proceed to configure the single node or SFHA cluster for CP server.

Configuring security on the CP server

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

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

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

For additional information:

See [“Preparing to configure the clusters in secure mode”](#) on page 73.

Setting up shared storage for the CP server database

To set up shared storage for the CP server database

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

For a command example:

```
# vxvg init cps_dg disk1 disk2
```

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

For a command example:

```
# vxvg import cps_dg
```

- 3 Create a mirrored volume over the disk group.

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

For a command example:

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

- 4 Create a file system over the volume.

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

Symantec recommends the vxfs file system type.

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

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

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

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

Configuring the CP server using the configuration utility

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

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

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

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

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

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

See [“Setting up the CP server”](#) on page 92.

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

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

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

Use the -n option for using rsh communication.

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

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

```
=====
```

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

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

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

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

```
Enter the option:
```

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

4 Depending upon your configuration, select either option 1 or option 2.

The configuration utility then runs the following preconfiguration checks:

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

5 Enter the name of the CP server.

For example:

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

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

For example:

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

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

For example:

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

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

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

For example:

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

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

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

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

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

For information about configuring security on the CP server:

See [“Configuring security on the CP server”](#) on page 94.

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

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

For a single node VCS configuration for CP server example:

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

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

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

For configuring CP server on an SFHA cluster example:

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

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

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

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

- 10** Review the displayed CP server configuration information.

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

For example:

Following is the CP Server configuration information:

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

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

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

For a single node VCS configuration for CP server example:

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

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

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

For configuring CP server on an SFHA cluster example:

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

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

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

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

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

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

For a single node VCS configuration for CP server example:

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

For configuring CP server on an SFHA cluster example:

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

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

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

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

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

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

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

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

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

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

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

For example:

```
# hagrp -state CPSSG
```

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

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

The configuration utility adds the `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 User's Guide*.

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

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

To manually configure the CP server

- 1 Ensure that the CP server preparation procedures have been performed:
- 2 Stop VCS on each node by using the following command:

```
# hastop -local
```

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

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

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

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

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

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

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

For more information about configuring security on the CP server:

See [“Configuring security on the CP server”](#) on page 94.

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

- 6 Start VCS on all the cluster nodes.

Enter the following command:

```
# hastart
```

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

Enter the following command:

```
# hagrps -state CPSSG
```

Output similar to the following should appear:

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

Verifying the CP server configuration

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

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

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

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

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

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

Configuring Storage Foundation and High Availability products

This chapter includes the following topics:

- [Configuring Storage Foundation and High Availability Solutions](#)
- [Configuring Veritas Volume Manager](#)
- [Configuring DMP support for booting over a SAN](#)
- [Configuring Veritas File System](#)
- [Configuring the SFDB repository database](#)
- [Setting or changing the product level for keyless licensing](#)
- [Installing Veritas product license keys](#)

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 common product installer. Use the same procedures to configure Storage Foundation for Oracle High Availability.

Required information for configuring Storage Foundation and High Availability Solutions

To configure Storage Foundation High Availability or Storage Foundation for Oracle 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
 - One heartbeat link may be configured as a low priority link

Veritas Storage Foundation can be configured to use Symantec Security Services.

Running Storage Foundation in Secure Mode guarantees that all inter-system communication is encrypted and that users are verified with security credentials. When running Storage Foundation in Secure Mode, NIS and system usernames and passwords are used to verify identity. Storage Foundation 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 installsf

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

Overview of tasks for Storage Foundation HA configuration using installsf

Storage Foundation HA configuration requires configuring the HA (VCS) cluster.

Tasks involved in configuring Storage Foundation HA are as follows:

- Start the software configuration
See [“Starting the software configuration”](#) on page 107.
- Specify the systems where you want to configure VCS
See [“Specifying systems for configuration”](#) on page 108.
- Configure the basic cluster
See [“Configuring the basic cluster”](#) on page 109.
- Configure virtual IP address of the cluster (optional)
See [“Configuring the virtual IP of the cluster”](#) on page 110.
- Configure the cluster in secure mode (optional)
See [“Configuring the cluster in secure mode”](#) on page 113.
- Add VCS users (required if you did not configure the cluster in secure mode)
See [“Adding VCS users”](#) on page 117.
- Configure SMTP email notification (optional)
See [“Configuring SMTP email notification”](#) on page 117.
- Configure SNMP email notification (optional)
See [“Configuring SNMP trap notification”](#) on page 119.
- Configure global clusters (optional)
You must have enabled Global Cluster Option when you installed VCS.
See [“Configuring global clusters”](#) on page 121.
- Complete the software configuration
See [“Completing the VCS configuration”](#) on page 122.

Starting the software configuration

You can configure Storage Foundation HA using the Veritas product installer or the installvcs program.

To configure Storage Foundation HA 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: Veritas Storage Foundation.

To configure Storage Foundation HA using the installsf

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

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

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

Specifying systems for configuration

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

To specify system names for installation

- 1 Enter the names of the systems where you want to configure Storage Foundation HA.

```
Enter the 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 the systems use the proper operating system
- Checks whether Storage Foundation HA is installed
- Exits if Veritas Storage Foundation 5.1 is not installed

Configuring the basic cluster

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,?] (0) 7
```

- 3 Review the NICs available on the first system as the installer discovers and reports them.

The private heartbeats can either use NIC or aggregated interfaces. To use aggregated interfaces for private heartbeat, enter the name of the aggregated interface. To use a NIC for private heartbeat, enter a NIC which is not part of an aggregated interface.

4 Enter the network interface card details for the private heartbeat links.

You must choose the network interface cards or the aggregated interfaces that the installer discovers and reports. To use any aggregated interfaces that the installer has not discovered, you must manually add the aggregated interfaces to use as private links after you configure Storage Foundation HA.

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

You must not enter the network interface card that is used for the public network (typically en5.)

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

```
[b,q,?] en2
```

```
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,?] en3
```

```
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)
```

5 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.

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.

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: en5
Enter the NIC for Virtual IP of the Cluster to use on galaxy:
[b,q,?] (en5)
```

- 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 en5 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

- 5** Enter the virtual IP address for the cluster.

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

For IPv4: ■ Enter the virtual IP address.

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

■ Confirm the default netmask or enter another one:

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

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

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

■ Verify and confirm the Cluster Virtual IP information.

Cluster Virtual IP verification:

NIC: *en5*
IP: 192.168.1.16
Netmask: 255.255.240.0

NetworkHosts: 192.168.1.17

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

For IPv6

- Enter the virtual IP address.

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

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

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

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

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

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

```
NIC: en5
IP: 2001:454e:205a:110:203:baff:feee:10
Prefix: 64
```

```
NetworkHosts: 192.168.1.17
```

```
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 `installsf` 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 73.

To configure the cluster in secure mode

- 1 Choose whether to configure Storage Foundation HA 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.
- 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. Enter the following Root Broker information as the installer prompts you:
Manual configuration

```
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): [?] 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 Storage Foundation HA 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 User'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.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: en5
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (en5)
Is en5 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,?] hazriet@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: en5
```

```
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 Storage Foundation HA 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 User's Guide* for more information.

To configure the SNMP trap notification

- 1 Review the required information to configure the SNMP notification feature of Storage Foundation HA.
- 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 Storage Foundation HA based on the configuration details you provided.

- 3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on galaxy: en5
Enter the NIC for the VCS Notifier to use on galaxy:
[b,q,?] (en5)
Is en5 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: en5

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

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 User's Guide* for instructions to set up VCS 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.

```
Global Cluster Option configuration verification:
```

```
NIC: en5  
IP: 192.168.1.16  
Netmask: 255.255.240.0
```

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

Completing the VCS configuration

After you enter the Storage Foundation HA 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 Storage Foundation HA.

If you chose to configure the cluster in secure mode, the installer then does the following before it starts Storage Foundation HA 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 Storage Foundation HA users
- Sets up trust with the root broker

To complete the VCS configuration

- 1 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 Storage Foundation HA.
- 3 Enter y at the prompt to send the installation information to Symantec.

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

- 4 After the installer configures Storage Foundation HA 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 Storage Foundation using response files” on page 311.

Configuring Storage Foundation High Availability using the web-based installer

This section describes the procedure to configure Storage Foundation High Availability using the web-based installer. Before you begin with the procedure, review the requirements for configuring Storage Foundation High Availability.

To configure Storage Foundation High Availability on a cluster

- 1 Start the web-based installer.
- 2 Select the following on the **Select Product/Task** screen:
 - From the list of tasks, select **Configure a Product**.
 - From the list of products, select **Storage Foundation High Availability**.

- 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.
- Click **Next**.

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

- 3 Select the following on the **Select Systems** screen:
 - Enter the system names on which VCS is to be configured, and then click **Validate**. System names are separated by spaces.
Example: **system01 system02**
The installer performs the initial system verification. It checks that communication between systems has been set up. 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 Select the following on the **Set Cluster Name/ID** screen.
 - Enter the unique cluster name and Cluster ID number.
 - Select the number of heartbeat links.
 - Select **Low priority heartbeat** if you want to configure one heartbeat link as a low priority link.
 - Select **Unique NICs per system** if you do not want to use the same NIC details to configure private heartbeat links on other systems.
 - Click **Next**.
- 5 Select the following on the **Set Cluster Heartbeat** screen.
 - If you are using the same NICs to configure private heartbeat links on all the systems, select the NIC for the first private heartbeat NIC on each system.
Select the NIC for the second private heartbeat NIC on each system.
 - If you have selected **Unique NICs per system** in the previous screen, provide the NIC details for each system.
 - Click **Next**.

- 6** In the **Storage Foundation High Availability Optional Configure** screen, select the Storage Foundation High Availability options that you want to configure, namely Virtual IP, VCS Users SMTP, SNMP, and GCO. Depending on the options that you select, you can enter the details regarding each option.
- To configure the virtual IP, do the following:
 - Select **Configure Virtual IP**.
 - If each system uses a separate NIC, select **Configure NICs for every system separately**.
 - Select the interface on which you want to configure the virtual IP.
 - Enter a virtual IP address and value for the netmask.
 - To configure the Storage Foundation High Availability users, enter the following information:
 - 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.
 - To configure SMTP notification, enter the following information:
 - If all the systems use the same NIC, select the NIC for the Storage Foundation High Availability Notifier to be used on all systems. If not, select the NIC to be used by each system.
 - Enter the domain-based hostname of the SMTP server. Example: `smtp.yourcompany.com`
 - Enter the full email address of the SMTP recipient. Example: `user@yourcompany.com`.
 - Select the minimum security level of messages to be sent to each recipient.
 - Click **Add** to add more SMTP recipients, if necessary.
 - To configure SNMP notification, enter the following information.
 - If all the systems use the same NIC, select the NIC for the Storage Foundation High Availability Notifier to be used on all systems. If not, select the NIC to be used by each system.
 - Enter the SNMP trap daemon port: (162).
 - Enter the SNMP console system name.
 - Select the minimum security level of messages to be sent to each console.

- Click **Add** to add more SNMP consoles, if necessary.
 - If you installed a valid HA/DR license, you can select the gco option to configure this cluster as a global cluster.
See *Veritas Cluster Server User's Guide* for instructions to set up VCS global clusters.
 - Select a NIC.
 - Enter a virtual IP address and value for the netmask.
 - Click **Next**.
The installer proceeds to configure Storage Foundation High Availability based on the configuration details you provided.
- 7 In the **Starting Processes** screen, the installer completes the Storage Foundation High Availability configuration.
The installer starts Storage Foundation High Availability and its components on each system.
After the startup process is complete, click **Next** to move to the next screen.
- 8 Click **Next** to complete the process of configuring Storage Foundation High Availability.
View the summary file, log file, or response file, if needed, to confirm the configuration.
- 9 Select the checkbox to specify whether you want to send your installation information to Symantec.

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

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

Configuring Veritas Volume Manager

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

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

Configuring DMP support for booting over a SAN

Configuring DMP to work with an LVM root disk over a SAN requires that the system is correctly configured to use the boot device over all possible paths.

To configure DMP support for booting over a SAN

- 1 The PVID and volume group entries in the second and third columns should be identical for all the paths. Use the `lspv` command for the root volume group to verify that the PVID and volume group entries are set correctly.

In this example, the LVM root disk is multipathed with 4 paths. The output from the `lspv` command for the root volume group (`rootvg`) is as follows:

```
# lspv | grep rootvg
hdisk374 00cbf5ce56def54d rootvg active
hdisk375 00cbf5ce56def54d rootvg
hdisk376 00cbf5ce56def54d rootvg
hdisk377 00cbf5ce56def54d rootvg
```

- 2 If the PVID and volume group entries are not set correctly on any of the paths, use the `chdev` command to set the correct value.

For example, the following output shows that the `hdisk377` path is not set correctly:

```
# lspv | grep rootvg
hdisk374 00cbf5ce56def54d rootvg active
hdisk375 00cbf5ce56def54d rootvg
hdisk376 00cbf5ce56def54d rootvg
hdisk377 none None
```

To correct the setting for the path, use the following command:

```
# chdev -l hdisk377 -a pv=yes
hdisk377 changed
```

The output of the `lspv` command now shows the correct values:

```
# lspv | grep rootvg
hdisk374 00cbf5ce56def54d rootvg active
hdisk375 00cbf5ce56def54d rootvg
hdisk376 00cbf5ce56def54d rootvg
hdisk377 00cbf5ce56def54d rootvg
```

- 3 Check that the output from the `bootlist` command is correct:

```
# bootlist -m normal -o
hdisk374 blv=hd5
hdisk375 blv=hd5
hdisk376 blv=hd5
hdisk377 blv=hd5
```

In this example, the output should show the default boot volume, `hd5`, for each path.

- 4 If the `blv` option is not set for a path to the disk, use the `bootlist` command to set it:

```
# bootlist -m normal hdisk374 hdisk375 hdisk376 hdisk377 blv=hd5
```

- 5 Run the following command to configure DMP on the root disk:

```
# vxddmpadm native enable vgroupname=rootvg
```

- 6 Reboot the system. DMP takes control of the SAN boot device to perform load balancing and failover.

- 7 Verify whether DMP controls the root disk.

```
# vxddmpadm native list vgroupname=rootvg
```

Configuring Veritas File System

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

```
/etc/filesystems
```

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

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

Configuring the SFDB repository database

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

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

Setting or changing the product level for keyless licensing

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

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

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

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

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

To set or change the product level

- 1 View the current setting for the product level.

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

- 2 View the possible settings for the product level.

```
# vxkeyless displayall
```

- 3 Set the desired product level.

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

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

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas product license keys

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

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

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

To install a new license

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

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

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

Configuring Storage Foundation High Availability for data integrity

This chapter includes the following topics:

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

Setting up disk-based I/O fencing using installsfc

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

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

- 1 Scan for the new hdisk devices.

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

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

```
# lsdev -Cc disk
```

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

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

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

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

```
# vxdisk list EMC0_17
```

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

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

```
vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i EMC0_17
```

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

Checking shared disks for I/O fencing

Make sure that the shared storage you set up while preparing to configure Storage Foundation meets the I/O fencing requirements. You can test the shared disks using the `vxfcntlsthaw` 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 `vxfcntlsthaw` command with the `-i` option. This command option verifies that the same serial number for the LUN is returned on all paths to the LUN.

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

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

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

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

The `vxfcntlsthaw` 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 User'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 133.
- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 134.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfcntlsthaw utility”](#) on page 135.

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
libvxFUJTSYe6k.so FUJITSU     E6000
libvxFUJTSYe8k.so FUJITSU     All
libvxcompellent.so COMPELNT    Compellent Vol
libvxcopan.so   COPANSYS    8814, 8818
libvxddns2a.so  DDN         S2A 9550, S2A 9900,
                S2A 9700
```

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

```
# vxdisk scandisks
```

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

Verifying that the nodes have access to the same disk

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

To verify that the nodes have access to the same disk

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

```
vxfenadm -i diskpath
```

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

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

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

From node A, enter:

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

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

The same serial number information should appear when you enter the equivalent command on node B using the `/dev/rhdisk76` path.

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

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

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

Testing the disks using vxfentsthdw utility

This procedure uses the `/dev/rhdisk75` disk in the steps.

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

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

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

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

To test the disks using `vxfcntlsthdw` utility

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

2 From one node, start the utility.

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

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

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

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

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

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

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

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

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

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

5 Review the output as the utility performs the checks and report its activities.

6 If a disk is ready for I/O fencing on each node, the utility reports success:

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

```
ALL tests on the disk /dev/rhdisk75 have PASSED
```

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

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

Configuring disk-based I/O fencing using `installsf`

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

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

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

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

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

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

The program checks that the local node running the script can communicate with remote nodes and checks whether Storage Foundation HA 5.1 is configured properly.

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

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

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

- If the check fails, configure and enable VxVM before you repeat this procedure.

- If the check passes, then the program prompts you for the coordinator disk group information.
- 5 Choose whether to use an existing disk group or create a new disk group to configure as the coordinator disk group.

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

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.
The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
 - To create a new disk group, perform the following steps:
 - Enter the number corresponding to the **Create a new disk group** option. The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks. Symantec recommends to use three disks as coordination points for disk-based I/O fencing.
 - Enter the numbers corresponding to the disks that you want to use as coordinator disks.
 - Enter the disk group name.
- 6 Verify that the coordinator disks you chose meet the I/O fencing requirements. You must verify that the disks are SCSI-3 PR compatible using the `vxfsntsthdw` utility and then return to this configuration program.
- 7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.
- 8 Enter the I/O fencing disk policy that you chose to use. For example:

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

The program also does the following:

- Populates the `/etc/vxfendg` file with this disk group information
 - Populates the `/etc/vxfenmode` file on each cluster node with the I/O fencing mode information and with the SCSI-3 disk policy information
- 9 Verify and confirm the I/O fencing configuration information that the installer summarizes.

- 10 Review the output as the configuration program does the following:
 - Stops Storage Foundation HA and I/O fencing on each node.
 - Configures disk-based I/O fencing and starts the I/O fencing process.
 - Updates the VCS configuration file `main.cf` if necessary.
 - Copies the `/etc/vxfenmode` file to a date and time suffixed file `/etc/vxfenmode-date-time`. This backup file is useful if any future fencing configuration fails.
 - Starts Storage Foundation HA on each node to make sure that the Storage Foundation HA 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 163.

Setting up disk-based I/O fencing manually

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

Table 10-1 Tasks to set up I/O fencing manually

Action	Description
Initializing disks as VxVM disks	See “Initializing disks as VxVM disks” on page 131.
Identifying disks to use as coordinator disks	See “Identifying disks to use as coordinator disks” on page 140.
Checking shared disks for I/O fencing	See “Checking shared disks for I/O fencing” on page 133.
Setting up coordinator disk groups	See “Setting up coordinator disk groups” on page 140.
Creating I/O fencing configuration files	See “Creating I/O fencing configuration files” on page 141.
Modifying Storage Foundation configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 142.

Table 10-1 Tasks to set up I/O fencing manually (*continued*)

Action	Description
Configuring Coordination Point agent to monitor coordination points	See “Configuring Coordination Point agent to monitor coordination points” on page 163.
Verifying I/O fencing configuration	See “Verifying I/O fencing configuration” on page 144.

Removing permissions for communication

Make sure you completed the installation of Storage Foundation 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

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

See [“Initializing disks as VxVM disks”](#) on page 131.

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 133.

Setting up coordinator disk groups

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

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

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

The following example procedure assumes that the disks have the device names EMC0_12, EMC0_16, and EMC0_17.

To create the `vxfencoorddg` disk group

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

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

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

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

- 3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

- 1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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

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

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

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

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

Modifying VCS configuration to use I/O fencing

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

To modify VCS configuration to enable I/O fencing

- 1 Save the existing configuration:

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

- 2 Stop VCS on all nodes:

```
# hastop -all
```

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

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

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

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

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

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

- 6 Save and close the file.

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

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

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

For example, on each remaining node, enter:

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

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

- Start the I/O fencing driver.

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

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

- Start VCS.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

- ◆ On one of the nodes, type:

```
# vxfenadm -d

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

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

* 0 (galaxy)
  1 (nebula)

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

Setting up server-based I/O fencing using installsf

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

See [“Verifying security configuration on SF HA cluster to use CP server coordination point”](#) on page 145.

See [“Configuring server-based I/O fencing”](#) on page 147.

Verifying security configuration on SF HA cluster to use CP server coordination point

After configuring security using the `installsf -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
```

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

```
*****
```

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

```
Expiry Interval 0
```

```
*****
```

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

```
HA_SERVICES@hostname.domainname
```

or

```
HA_SERVICES@hostname
```

3 There should not be duplicate entries for HA_SERVICES domain.

An example of incorrect configuration is given below.

```
showdomains

Domain(s) Found :          3

*****

Domain Name:      HA_SERVICES@galaxy.symantec.com

Domain Type:     vx

*****

Domain Name:      broker@galaxy.symantec.com

Domain Type:     vx

*****

Domain Name:      HA_SERVICES@galaxy

Domain Type:     vx

*****
```

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

Configuring server-based I/O fencing

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

To configure the SF HA cluster with server-based I/O fencing

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

For example:

```
/opt/VRTS/install/installsf -fencing
```

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

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

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

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

```
/opt/VRTS/install/installsf -fencing
```

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

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

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

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

```
Cluster information verification:
```

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

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

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

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

```
Fencing configuration
```

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

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

Select the first option for CP client-based fencing.

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

For example:

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

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

For example:

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

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

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

For example:

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

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

For example:

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

9 Enter the fencing mechanism for the disk or disks.

For example:

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

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

Select disk number 1 for co-ordination point

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

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

Select a disk from the displayed list.

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

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

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

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

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

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

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

Select a disk from the displayed list for the second coordinator point.

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

- 13** Proceed to read the displayed recommendation from the installer to verify the disks prior to proceeding.

Press Enter to continue.

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

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

- 15** Proceed to enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):
[b] (vxfencoorddg)
```

- 16** The installer now begins verification of the coordination points. At the end of the verification process, the following information is displayed:

- Total number of coordination points being used
- CP Server Virtual IP/hostname and port number
- SCSI-3 disks

- Disk Group name for the disks in customized fencing
- Disk mechanism used for customized fencing

For example:

```
Total number of coordination points being used: 3
```

```
CP Server (Port):
```

```
1. 10.209.80.197 (14250)
```

```
SCSI-3 disks:
```

```
1. c3t0d0s2
```

```
2. c3t1d0s3
```

```
Disk Group name for the disks in customized fencing: vxsfencorddg
```

```
Disk mechanism used for customized fencing: raw
```

You are then prompted to accept the above information. Press Enter to accept the default (y) and continue.

The disks and disk group are initialized and the disk group deported on the SF HA cluster node.

17 The installer now automatically determines the security configuration of the CP server's side and takes the appropriate action:

- If the CP server's side is configured for security, then the SF HA cluster's side will be configured for security.
- If the CP server's side is not configured for security, then the SF HA cluster's side will not be configured for security.

For example:

```
While it is recommended to have secure communication configured between CP Servers and CP client cluster, the client cluster must be in the same mode (secure or non-secure) as the CP servers are.
```

```
Since the CP servers are configured in secure mode, the installer will configure the client cluster also as a secure cluster.
```

```
Press [Enter] to continue:
```

```
Trying to configure Security on the cluster:
```

```
All systems already have established trust within the
```

```
Symantec Product Authentication Service domain  
root@galaxy.symantec.com
```

- 18** Enter whether you are using different root brokers for the CP servers and SF HA clusters.

If you are using different root brokers, then the installer tries to establish trust between the authentication brokers of the CP servers and the SF HA cluster nodes for their communication.

After entering "y" for yes or "n" for no, press Enter to continue.

- 19** If you entered "y" for yes in step 18, then you are also prompted for the following information:
- Hostname for the authentication broker for any one of the CP servers
 - Port number where the authentication broker for the CP server is listening for establishing trust
 - Hostname for the authentication broker for any one of the SF HA cluster nodes
 - Port number where the authentication broker for the SF HA cluster is listening for establishing trust

Press Enter to continue.

- 20** The installer then displays your I/O fencing configuration and prompts you to indicate whether the displayed I/O fencing configuration information is correct.

If the information is correct, enter "y" for yes.

For example:

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm  
Cluster ID: 2122  
Cluster Name: clus1  
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

- 21 The installer then updates the SF HA cluster information on each of the CP Servers to ensure connectivity between them.

The installer then populates the file `/etc/vxfenmode` with the above details in each of the CP SF HA cluster nodes.

For example:

```
Updating client cluster information on CP Server 10.210.80.199

Adding the client cluster to the CP Server 10.210.80.199 ..... Done

Registering client node galaxy with CP Server 10.210.80.199..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Registering client node nebula with CP Server 10.210.80.199 ..... Done
Adding CPClient user for communicating to CP Server 10.210.80.199 ..... Done
Adding cluster clus1 to the CPClient user on CP Server 10.210.80.199 ... Done

Updating /etc/vxfenmode file on galaxy ..... Done
Updating /etc/vxfenmode file on nebula ..... Done
```

For additional information about the `vxfenmode` file in mixed disk and CP server mode, or pure server-based mode:

See [“About I/O fencing configuration files”](#) on page 84.

- 22 You are then prompted to configure the CP agent on the client cluster.

```
Do you want to configure CP Agent on the client cluster? [y,n,q]
(y)

Enter a non-existing name for the service group for CP Agent:
[b] (vxfen)

Adding CP Agent via galaxy ..... Done
```

- 23** The VCS and the fencing process are then stopped and restarted on each SF HA cluster node, and the I/O configuration process then finished.

```
Stopping VCS on galaxy ..... Done
Stopping Fencing on galaxy ..... Done
Stopping VCS on nebula ..... Done
Stopping Fencing on nebula ..... Done
```

- 24** At the end of this process, the installer then displays the location of the configuration log files, summary files, and response files.

Setting up server-based I/O fencing manually

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

Table 10-2 Tasks to set up server-based I/O fencing manually

Action	Description
Preparing the CP servers for use by the Storage Foundation HA cluster	See “Preparing the CP servers manually for use by the SF HA cluster” on page 155.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See “Configuring server-based fencing on the SF HA cluster manually” on page 159.
Configuring Coordination Point agent to monitor coordination points	See “Configuring Coordination Point agent to monitor coordination points” on page 163.
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 165.

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 10-3](#) displays the sample values used in this procedure.

Table 10-3 Sample values in procedure

CP server configuration component	Sample name
CP server	mycps1.symantecexample.com

Table 10-3 Sample values in procedure (*continued*)

CP server configuration component	Sample name
Node #1 - SF HA cluster	galaxy
Node #2 - SF HA cluster	nebula
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the SF HA cluster

- 1 Determine the cluster name and uuid on the SF HA cluster.

For example, issue the following commands on one of the SF HA cluster nodes (galaxy):

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster clus1

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2}
```

- 2 Check whether the SF HA cluster and nodes are present in the CP server.

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

ClusName  UUID                               Hostname(Node ID)  Registered
clus1     {f0735332-1dd1-11b2} galaxy(0)           0
clus1     {f0735332-1dd1-11b2} nebula(1)          0
```

If the output does not show the cluster and nodes, then add them as described in the next step.

3 Add the 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
<code>_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com/vx</code>	<code>clus1/{f0735332-1dd1-11b2}</code>	Operator
<code>_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com/vx</code>	<code>clus1/{f0735332-1dd1-11b2}</code>	Operator

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

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

5 Add the users to the CP server.

First, determine the user@domain to be added.

The user for fencing should be of the form `_HA_VCS_`*short-hostname* and domain name is that of HA_SERVICES user in the output of command:

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

Next, issue the following commands on the CP server (mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com  
successfully added
```

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com\  
-f cps_operator -g vx
```

```
User _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com  
successfully added
```

- 6 Authorize the CP server user to administer the 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.symantec.com\  
-f cps_operator -g vx
```

```
Cluster successfully added to user  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com privileges.
```

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

```
Cluster successfully added to user  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com privileges.
```

Configuring server-based fencing on the SF HA cluster manually

The configuration process for the client or SF HA cluster to use CP server as a coordination point requires editing the `/etc/vxfenmode` file. You need to edit this file to specify the following information for your configuration:

- Fencing mode
- Fencing mechanism
- Fencing disk policy (if applicable to your I/O fencing configuration)
- Appropriate value for the security configuration
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)

Whenever coordinator disks are used as coordination points in your I/O fencing configuration, a disk group (vxfendg) has to be created. This disk group has to be

specified in the `/etc/vxfenmode` file. For information about creating the disk group, see the *Veritas™ Cluster Server Installation Guide*.

The customized fencing framework also generates the `/etc/vxfentab` file which has security setting and the coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

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

`/etc/default/vxfen`

Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

The following file output provides an example of what the `/etc/vxfenmode` file contains:

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

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

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
```

```
scsi3_disk_policy=dmp

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

#
# Specify 3 or more odd number of coordination points in this file,
# one in each row. They can be all-CP servers, all-SCSI-3 compliant
# coordinator disks, or a combination of CP servers and SCSI-3
# compliant coordinator disks. Please ensure that the CP server
# coordination points are numbered sequentially and in the same
# order on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as:
#
# cps<number>=<Virtual IP/ Virtual hostname of cp server> in
# square brackets ([]), followed by ":" and CPS port number.
#
# Examples:
# cps1=[192.168.0.23]:14250
# cps2=[mycps.company.com]:14250
#
# SCSI-3 compliant coordinator disks are specified as:
#
# vxfendg=<coordinator disk group name>
# Example:
# vxfendg=vxfencoordg
#
# Examples of different configurations:
# 1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
```

```
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
#
```

Table 10-4 defines the vxfenmode parameters that must be edited.

Table 10-4 vxfenmode file parameters

vxfenmode File Parameter	Description
vxfen_mode	Fencing mode of operation. This parameter must be set to "customized".
vxfen_mechanism	Fencing mechanism. This parameter defines the mechanism that is used for fencing. If one of the three coordination points is a CP server, then this parameter must be set to "cps".
scsi3_disk_policy	Configure the vxfen module to use either DMP devices, "dmp" or the underlying raw character devices, "raw". Note: The configured disk policy is applied on all the nodes.
security	Security parameter 1 indicates that Symantec Product Authentication Service is used for CP server communications. Security parameter 0 indicates that communication with the CP server is made in non-secure mode. The default security value is 1. Note: Symantec only supports a configuration where both the CP server and client sides have the same security setting. The security setting on both sides must be either enabled or disabled.

Table 10-4 vxfenmode file parameters (*continued*)

vxfenmode File Parameter	Description
cps1, cps2, cps3, or vxfendg	<p>Coordination point parameters.</p> <p>Enter either the Virtual IP address or FQHN (whichever is accessible) of the CP server.</p> <p>Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfendg) and specified in the /etc/vxfenmode file. Additionally, the customized fencing framework also generates the /etc/vxfentab file which specifies the security setting and the coordination points (all the CP servers and the disks from disk group specified in /etc/vxfenmode file).</p>

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

For example:

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

Configuring Coordination Point agent to monitor coordination points

The following procedure describes how to manually configure the Coordination Point agent to monitor coordination points (CP server or SCSI-3 disks).

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

To configure Configuration Point agent to monitor coordination points

- 1 Ensure that your 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 as follows:

```
# haconf -makerw
# hagr -add vxfen
# hagr -modify vxfen SystemList system1 0 system2 1
# hagr -modify vxfen AutoFailOver 0
# hagr -modify vxfen Parallel 1
# hagr -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 1
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

- 3 Verify the status of the agent on the SF HA cluster using the `hares` commands.

For example:

```
# hares -state coordpoint
```

The following is an example of the command and output::

```
# hares -state

# Resource      Attribute      System      Value
coordpoint     State          galaxy      ONLINE
```

- 4 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed Coordination Point agent monitoring information; including information about whether the Coordination Point agent is able to access all the coordination points, information to check on which coordination points the Coordination Point agent is reporting missing keys, etc.

To view all such information in the engine log, change the `dbg` level for that node using the following command:

```
# hatype -modify coordpoint LogDbg 10
```

The agent log can now be viewed at the following location:

```
/var/VRTSvcs/log/engine_A.log
```

Verifying server-based I/O fencing configuration

During the SF HA cluster installation, the installer populates the following files based on inputs that are received during the configuration phase:

- `/etc/vxfenmode` (edited for CP server)
- `/etc/vxfentab` (edited for CP server)

Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

For troubleshooting server-based I/O fencing configuration issues, refer to the *Veritas Cluster Server User's Guide*.

Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```


Upgrading Storage Foundation

This chapter includes the following topics:

- [Upgrading Storage Foundation products or the operating system](#)
- [Planning the upgrade](#)
- [Upgrading Storage Foundation products to 5.1](#)
- [Upgrading Storage Foundation with the Veritas Web-based installer](#)
- [Upgrading Storage Foundation and AIX on a DMP-enabled rootvg](#)
- [Upgrading using SMIT](#)
- [Upgrading the AIX operating system](#)
- [Upgrading Veritas Volume Replicator](#)
- [Post-upgrade tasks](#)
- [Verifying the Veritas Storage Foundation upgrade](#)

Upgrading Storage Foundation products or the operating system

If your system is already running a previous release of a Storage Foundation (or Foundation Suite) product, this section describes how to upgrade it to Veritas Storage Foundation 5.1. The operating system must be at a supported level for this upgrade. Perform the procedures in the following sections to upgrade Storage Foundation or your operating system, or both. You can perform an upgrade to

Storage Foundation using the Veritas product installer or product installation script if you already have Storage Foundation installed.

This section describes how to upgrade Veritas Storage Foundation, Veritas Storage Foundation for DB2, and Veritas Storage Foundation for Oracle.

Caution: Make sure that supported combinations of Storage Foundation and the operating system are present on your system during the upgrades. Do not upgrade to a version of Storage Foundation that is not supported with the current operating system.

Planning the upgrade

Be sure that the administrator doing the upgrade has root access and a working knowledge of UNIX operating system administration.

Complete the following tasks in advance of upgrading:

- Check the latest *Storage Foundation Release Notes* to verify that the system meets all the requirements for software and hardware, including any required operating system patches.
- Schedule sufficient outage time for the upgrade, and downtime for any applications using the VxFS file systems or VxVM volumes.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. (This may not be practical, but if done, offers a failback point.)
- To upgrade on a remote host, rsh or ssh must be set up.
See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 35.
- Select the method to upgrade.
See [“Upgrade paths”](#) on page 169.

Saving system information before upgrade

Use the following procedure to save system information before an upgrade.

To save system information before an upgrade

- 1 Log in as superuser.
- 2 Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/filesystems` file. You will need to recreate these entries in the `/etc/filesystems` file on the freshly upgraded system.

- 3 Before upgrading, ensure that you have made backups of all data that you want to preserve.
- 4 Copy `filesystems` to `filesystems.orig`:

```
# cp /etc/filesystems /etc/filesystems.orig
```
- 5 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 6 If you are installing the HA version of the Veritas Storage Foundation 5.1 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.

About upgrading the Storage Foundation for Databases (SFDB) tools to 5.1

If you plan to continue using the Storage Foundation for Databases (SFDB) tools you are using with Storage Foundation for Oracle 5.0 or 4.x, you must migrate the SFDB repository database to 5.1.

Tasks for upgrading SFDB tools to version 5.1

- Preparing to migrate the repository database before upgrading Storage Foundation to 5.1
See [“Pre-upgrade tasks for migrating the SFDB repository database”](#) on page 171.
- Migrating the repository database after upgrading to Storage Foundation 5.1.
See [“Post upgrade tasks for migrating the SFDB repository database”](#) on page 192.

Upgrade paths

Verify that the upgrade path is supported. Make sure that the release version of Storage Foundation is supported on the AIX operating system you are running, at each point in the upgrade. Refer to the *Veritas Storage Foundation Release Notes*.

Before upgrading the operating system, you may need to shut down Storage Foundation processes.

Table 11-1

AIX release	Storage Foundation version	Upgrade procedure
5.2, 5.3	4.0, 4.0MP1, 4.0MP2, 4.0MP3	Upgrade Storage Foundation to 4.0MP4. Upgrade OS to 5.3 TL7/SP6. Upgrade Storage Foundation using the installer script.
5.2, 5.3	4.0MP4, 5.0, 5.0MP1	Upgrade OS to 5.3 TL7/SP6. Upgrade Storage Foundation using the installer script.
5.3	5.0MP3	Upgrade OS to 5.3 TL7/SP6. Upgrade Storage Foundation using the installer script.
6.1	5.0MP1, 5.0MP3	Upgrade OS to 6.1 TL0/SP6. Upgrade Storage Foundation using the installer script.

Performing pre-installation checks and configuration

Only users with superuser privileges can upgrade and initialize Storage Foundation and AIX.

Complete the following tasks in advance of upgrading:

- Ensure that you have created a valid backup.
- 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/>

- Be sure that the administrator doing the upgrade has root access and a working knowledge of the specific operating system administration.
- Upgrade the operating system to the required level for the upgraded version of Storage Foundation. Make sure that your current version of Storage Foundation is supported with the upgraded operating system. If it is not supported, you need to plan a staged upgrade. At all stages of the upgrade, make sure the combination of the operating system and Storage Foundation is supported.
- Schedule sufficient outage time for the upgrade of Veritas Storage Foundation, Veritas Storage Foundation for DB2, or Veritas Storage Foundation for Oracle. Depending on your configuration, the outage may take several hours.
- Make sure the file systems are clean before upgrading.
See “[Verifying that the file systems are clean](#)” on page 177.

To determine which Storage Foundation product is installed on your system, run the following commands:

```
# ls1pp -L VRTSvxfs
# ls1pp -L VRTSvxvm
```

If the versions of `VRTSvxfs` and `VRTSvxvm` are 3.4.6.0 and 3.2.3.0, respectively, you have a 1.0 MP3 product installed.

If the versions of `VRTSvxfs` and `VRTSvxvm` are both 4.0.3.0, you have a 4.0 MP3 product installed.

Pre-upgrade tasks for migrating the SFDB repository database

If you plan to continue using the Storage Foundation for Databases (SFDB) tools you are using with Storage Foundation for Oracle 5.0 or 4.x, you must prepare to migrate the SFDB repository database to 5.1 before upgrading Storage Foundation for Oracle to 5.1.

Note: When using the CPI to install Storage Foundation 5.1, the `VRTSdbms3` package will not be removed.

Note: For clustered products, the `Sfua_Base` repository resource group will be removed from the `main.cf` file. It is not required as a separate service group for Storage Foundation 5.1 products.

Perform the following before upgrading Storage Foundation.

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 carried over after upgrading. You will have to create a new Database Flashsnap clone database after upgrading to 5.1

Preparing for an upgrade of Storage Foundation

Before the upgrade of Storage Foundation to a new release, shut down processes and synchronize snapshots.

To prepare for an upgrade of Storage Foundation

- 1 Before the upgrade of Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle, log in as `oracle` or `db2` and then shut down the database before proceeding.
- 2 Log in as `root`.
- 3 Resynchronize all existing snapshots before upgrading. In particular, resynchronize snapshots for Veritas Storage Foundation for DB2 or Veritas Storage Foundation for Oracle.

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
```

- 4 Stop activity to all file systems and raw volumes, for example by unmounting any file systems that have been created on volumes.

```
# umount mnt_point
```

- 5 Stop all the volumes by entering the following command for each disk group:

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

- 6 Before the upgrade of 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
```

- 7 Stop the VEA backend service by entering the following command:

```
# /opt/VRTS/bin/vxsvcctrl stop
```

- 8 Upgrade AIX on your system to the required levels if applicable.

- 9 Before the upgrade of Veritas Storage Foundation 1.0 for DB2 MP3 or Veritas Storage Foundation 1.0 for Oracle MP3, uninstall the 1.0 MP3 software.

To list all installed filesets:

```
# lsllp -a|grep VRTS
```

To uninstall the filesets, run the `installp -u` command for all filesets:

```
# installp -u fileset_name  
# installp -u fileset_name  
...
```

Then, reboot the system.

Preparing for upgrade of VVR in the presence of VCS agents

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

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

Freezing the service groups and stopping all the applications

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

To freeze the service groups and stop applications

Perform the following steps for the Primary and Secondary clusters:

- 1 Log in as the superuser.
- 2 Make sure that `/opt/VRTS/bin` is in your PATH so that you can execute all the product commands.
- 3 Before the upgrade, cleanly shut down all applications.
 - 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 175.

Determining the nodes on which disk groups are online

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

To determine the online disk groups

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

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

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

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

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

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

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

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

Preparing for the upgrade when VCS agents are configured

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

To prepare a configuration with VCS agents for an upgrade

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

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

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

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

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

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

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

```
# vxrecover -bs
```

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

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

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

Verifying that the file systems are clean

Prior to upgrading to release 5.1, 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/VRTSvxfss/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 -V vxfs filesystem
# mount -V vxfs filesystem mountpoint
# umount mountpoint
```

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

The clone removal of a large filesset may be an extended operation in progress, if the unmount 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 filesset 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 release includes all array support in a single filesset, VRTSaslapm. The array support filesset includes the array support previously included in the VRTSvxvm filesset. The array support filesset also includes support previously packaged as external array support libraries (ASLs) and array policy modules (APMs).

See the 5.1 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. Installing the VRTSvxvm fileset exits with an error if external ASLs or APMs are detected.

After you have installed Storage Foundation 5.1, 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 products to 5.1

If your system is already running a previous release of Storage Foundation (or Foundation Suite), this section describes how to upgrade it to Veritas Storage Foundation 5.1. The operating system must be at a supported level for this upgrade.

Before the upgrade of Storage Foundation to a new release, shut down processes and synchronize snapshots.

See “[Preparing for an upgrade of Storage Foundation](#)” on page 172.

Upgrading Veritas Storage Foundation with the product installer

This section describes upgrading from Veritas Storage Foundation products to 5.1.

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 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.

- 6 Load and mount the disc.
See [“Mounting a software disc”](#) on page 44.
- 7 To invoke the common installer, run the `installer` command on the disc as shown in this example:

```
# cd /cdrom/cdrom0  
# ./installer
```

- 8 Enter `G` to upgrade and press Return.

- 9 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.

- 10 The installer lists the filesets that will be installed or updated. You are prompted to confirm that you are ready to upgrade.

```
Are you sure you want to upgrade Storage Foundation? [y,n,q] (y) y
```

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

- 11 The installer uninstalls and reinstalls the listed filesets.
12 The Veritas Storage Foundation software is verified and configured.

Start the Veritas Storage Foundation processes.

```
Do you want to start Veritas Storage Foundation processes
now? [y,n,q] (y) y
```

Upgrading Storage Foundation with the Veritas Web-based installer

This section describes upgrading Storage Foundation 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 Storage Foundation

- 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 57.
- 4 Select **Upgrade**.
The installer detects the product that is installed on the specified system.
- 5 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.
- 6 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Validate**.
- 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 Click **Finish**. The installer prompts you for another task.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native operating system accounts.

Upgrading Storage Foundation and AIX on a DMP-enabled rootvg

If the rootvg is enabled for DMP, use these steps to upgrade Storage Foundation, AIX or both.

To upgrade Storage Foundation and AIX on a DMP-enabled rootvg

- 1 Disable DMP control on the root disk. For a Storage Foundation 5.1 installation, run the following command to disable DMP support for the root disk:

```
# vxddmpadm native disable vgroupname=rootvg
```

For a Storage Foundation release 5.0MP3 installation, run the following command to disable DMP support for the root disk:

```
# /usr/sbin/vxddmroot uninstall
```

- 2 Reboot the system.

```
# reboot
```

- 3 Upgrade AIX and Storage Foundation according to the instructions in the *Veritas Storage Foundation High Availability Installation Guide*.
- 4 Enable DMP on the root disk.
See [“Configuring DMP support for booting over a SAN”](#) on page 127.

Upgrading using SMIT

Veritas provides the `gunzip` utility under the `gnu` directory on the product software disc so that you can uncompress the Storage Foundation patches.

Upgrading using SMIT is not available for Storage Foundation HA.

To uncompress the packages:

- 1 Log in as superuser.
- 2 Create an installation directory on your system large enough for all the Storage Foundation patches. Refer to the disk space requirements in the system requirements section.

```
# mkdir /tmp/install
```
- 3 Place the Veritas software disc into a DVD drive connected to your system.
See [“Mounting a software disc”](#) on page 44.
- 4 Mount the disk by determining the device access name of the DVD drive. The format for the device access name is `cdN` where *N* is the device number. After inserting the disk into the DVD drive, enter:

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

- 5 Change to the directory containing the Storage Foundation patches:

```
# cd /mnt/cdrom/storage_foundation/patches
```
- 6 Copy the compressed patch files and the table of contents (`.toc`) file from the software disc to the temporary directory.

```
# cp -r * /tmp/install  
# cp .toc /tmp/install/
```

The `.toc` specifies the order in which the Storage Foundation components must be installed, and is used by the `installp` command. In general `VRTSvcki`, `VRTSvxvm`, and `VRTSvxfs` must be installed first in the specified order.

- 7 Change to the temporary directory and unzip the compressed package files:

```
# cd /tmp/install  
# gunzip VRTS*.gz
```

- 8 Invoke SMIT from the command line to upgrade the system. First, upgrade the already installed components of Storage Foundation (formerly known as Foundation Suite):

```
# cd /tmp/install  
# smit update_all
```

- 9 Once the existing components have been upgraded, add the new components added to the 5.1 release with this command:

```
# smit install
```

- 10 After successful upgrade, you must reboot the system. Reboot using the command:

```
# shutdown -r
```

- 11 To take advantage of new features, upgrade the VxVM disk group version (90) to the latest (140).

See the `vxdg` manual pages for more details.

Upgrading the AIX operating system

Use this procedure to upgrade the AIX operating system if Storage Foundation 5.1 is installed. You must upgrade to a version that Storage Foundation 5.1 supports.

Before you upgrade AIX (after installing or upgrading Storage Foundation), you must temporarily disable Storage Foundation to prevent it from starting, until the AIX upgrade is complete. It is necessary to avoid Storage Foundation operation until AIX is upgraded to the required maintenance level. After the AIX upgrade is complete, you can then enable Storage Foundation operation.

To upgrade the AIX operating system

- 1 Create the install-db file.

```
# touch /etc/vx/reconfig.d/state.d/install-db
```

- 2 Stop activity to all file systems and raw volumes, for example by unmounting any file systems that have been created on volumes.

```
# umount mnt_point
```

- 3 Stop all the volumes by entering the following command for each disk group:

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

- 4 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
```

- 5 Stop the VEA backend service by entering the following command:

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

- 6 Upgrade the AIX operating system. See the operating system documentation for more information.

- 7 Apply the necessary APARs.

For information about APARs required for Storage Foundation 5.1, refer to the *Storage Foundation Release Notes*.

- 8 Enable Storage Foundation to start after you reboot.

```
# rm /etc/vx/reconfig.d/state.d/install-db
```

- 9 Reboot the system.

```
# shutdown -Fr
```

Upgrading Veritas Volume Replicator

If a previous version of Veritas Volume Replicator (VVR) is configured, the product installer upgrades VVR automatically when you upgrade the Storage Foundation products.

When upgrading from 4.0 or later, you have the option to upgrade without disrupting replication.

See [“Upgrading VVR without disrupting replication”](#) on page 186.

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 26.

When both the Primary and the Secondary have the previous version of VVR installed, the upgrade can be performed either on the Primary or on the Secondary. We recommend that the Secondary hosts be upgraded before the Primary host in the RDS. This section includes separate sets of steps, for the Primary upgrade and for the Secondary upgrade.

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

Upgrading VVR on the Secondary

Follow these instructions to upgrade the Secondary hosts.

To upgrade the Secondary

- 1 Stop replication to the Secondary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvgname
```

- 2 Upgrade from VVR 4.0 to VVR 5.1 on the Secondary.

- 3 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname sec_hostname
```

Upgrading VVR on the Primary

After you upgrade the Secondary, use the Veritas product installer to upgrade the Primary.

Note: Reduce application downtime while upgrading by planning your upgrade.

See [“Planning an upgrade from the previous VVR version”](#) on page 26.

Post-upgrade tasks

The tasks in the following sections must be performed after upgrade, to restore the previous configurations and set up Storage Foundation 5.1 correctly. Perform the tasks required for the products and features that are relevant to your installation.

Optional configuration steps

After the upgrade is complete, additional tasks may need to be performed.

You can perform the following optional configuration steps:

- 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:

```
# vradmin changeip newpri=v6 newsec=v6
```

where *v6* is the IPv6 address.

- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Veritas Volume Manager Administrator's Guide*.
- 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. After you upgrade, perform one of the following steps:
 - Obtain a valid license key and run the `vxlicinst` command to add it to your system.
See [“Installing Veritas product license keys”](#) on page 130.
 - Use the `vxkeyless` command to update the license keys to the keyless license model.
See [“Setting or changing the product level for keyless licensing”](#) on page 129.

- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
See [“Upgrading disk layout versions”](#) on page 191.
See [“Upgrading VxVM disk group versions”](#) on page 198.

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 were frozen in step 6 of the section [Preparing for the upgrade when VCS agents are configured](#) by typing 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 176.

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 that you noted in step [\[Unresolved xref\]](#).

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.

```
# hagrps -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 Configure an Installed Product. Or use the installation script with the `-configure` option.

- 6 Bring online the RVGLogowner group on the master:

```
# hagr -online RVGLogownerGrp -sys masterhost
```

- 7 Start and bring online the failover service groups on the remaining host:

```
# hagr -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 and Version 7. No prior versions can be created or mounted.

Use the `vxfsconvert` or `vxupgrade` utilities to upgrade older disk layout versions to disk layout Version 7.

The `vxfsconvert` command converts an unmounted file system from disk layout Version 5 or prior to disk layout Version 7.

The `vxupgrade` command upgrades a mounted file system from disk layout Version 6 to Version 7. The `vxupgrade` cannot upgrade any previous versions because those versions cannot be mounted.

See the `vxfsconvert` or `vxupgrade` man pages.

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

Post upgrade tasks for migrating the SFDB repository database

If you plan to continue using the Storage Foundation for Databases (SFDB) tools which you are using with Storage Foundation for Oracle 5.0 or 4.x, you must perform one of the following procedures:

- Migrating a 5.0 SFDB repository database
- Migrating a 4.x SFDB repository database
- Upgrading without migrating existing Storage Checkpoints and DBDST parameters

If you plan to continue using the Database Storage Checkpoints and the Database Dynamic Storage Tiering parameters which you created in 5.0 or 4.x, you must migrate the repository database to 5.1 after upgrading Storage Foundation to 5.1.

Migrating from a 5.0 repository database to 5.1

For clustered environments, perform the following on one node only.

To migrate from a 5.0 repository database to 5.1

- 1 As root, set the Oracle group permission for various directories used by Oracle.

```
# /opt/VRTSdbed/common/bin/sfua_db_config
```

- 2 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
```

- 3 On the same node that you ran `sfua_rept_migrate` run the following command as Oracle user. For each Oracle instance, migrate the old repository data to the SQLite repository.

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
```

- 4 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
```

- 5 Remove the VRTSdbms3 package.

Warning: After you remove the VRTSdbms3 package, you will not be able to migrate the repository of any Oracle instance that you have not yet migrated.

If one of the following conditions applies:

- If you have not migrated the repositories for all the Oracle instances, and you do not need to migrate them.
- If you do not have a repository that you need to migrate from 5.0.

Perform the following step for the node that you have run `sfua_rept_migrate` to remove the VRTSdbms3 package manually.

```
# /usr/sbin/installp -u VRTSdbms3
```

If you have a clustered environment, perform the following steps for the rest of the nodes.

- As root, set the Oracle group permission for various directories used by Oracle.

```
# /opt/VRTSdbed/common/bin/sfua_db_config
```

- Remove the VRTSdbms3 package manually on the rest of the nodes before running the SFDB tools:

```
# /usr/sbin/installp -u VRTSdbms3
```

After VRTSdbms3 package is removed, you can destroy the repository disk group that was created at 5.0 version if you do not need it anymore.

- 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
```

Remove the VRTSdbms3 package manually on the secondary host.

```
# /usr/sbin/installp -u VRTSdbms3
```

- 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

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

- 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*

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, you need to upgrade existing disk groups.

After upgrading from Storage Foundation 4.x to 5.1, you must upgrade any existing disk groups which 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, 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 a version earlier than 5.1.

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:

```
# vxdctl defaultdg diskgroup
```

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

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 239.

Upgrading Storage Foundation using an alternate disk

This chapter includes the following topics:

- [About upgrading Storage Foundation using an alternate disk](#)
- [Supported upgrade scenarios](#)
- [Supported upgrade paths](#)
- [Preparing to upgrade Storage Foundation on an alternate disk](#)
- [Upgrading Storage Foundation on an alternate disk](#)
- [Upgrading Storage Foundation and AIX on an alternate rootvg that is enabled for DMP on an LVM root disk](#)
- [Configuring fencing for an ADI upgrade](#)
- [Verifying the upgrade](#)

About upgrading Storage Foundation using an alternate disk

Use the alternate disk installation process to upgrade the operating system and Storage Foundation on a production server while the server runs. Perform the upgrade on an alternate or inactive boot environment. After the upgrade, you restart the system to use the updated environment. The instructions in this section assume a working knowledge of the alternate disk installation process. See the

operating system documentation for detailed information on alternate disk installations.

Note: Only Technology Level (TL) and Service Pack (SP) releases of the operating system can be upgraded using this procedure.

Upgrading Storage Foundation on an alternate disk has the following advantages:

- The server remains active during the time the new boot environment is created and upgraded on the alternate boot device.
- The actual downtime for the upgrade is reduced to the period of time required for a single reboot.
- The original boot environment is still available for use if the updated environment fails to become active.

Upgrading Storage Foundation on an alternate disk involves the following steps:

Preparing to upgrade Storage Foundation on an alternate disk See [“Preparing to upgrade Storage Foundation on an alternate disk”](#) on page 203.

Upgrading Storage Foundation on an alternate disk See [“Upgrading Storage Foundation on an alternate disk”](#) on page 204.

Verifying the upgrade See [“Verifying the upgrade”](#) on page 217.

Supported upgrade scenarios

The following upgrade scenarios are supported on an alternate disk:

- Upgrading only Storage Foundation
- Upgrading only the operating system (Technology Level (TL) and Service Pack (SP) releases)
- Upgrading the operating system (Technology Level (TL) and Service Pack (SP) releases) and Storage Foundation

Supported upgrade paths

You can upgrade the operating system and Storage Foundation using an alternate disk from the following versions:

AIX version Technology Level and Service Pack releases of AIX 5.3 and later

Storage Foundation 4.0 MP4 and later
version

Preparing to upgrade Storage Foundation on an alternate disk

Complete the preparatory steps in the following procedure before you upgrade Storage Foundation on an alternate disk.

To prepare to upgrade Storage Foundation on an alternate disk

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

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

```
# lspv
```

```
Output similar to the following displays:
```

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

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

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

- 4 Make sure that the following filesets are installed on the primary disk:

```
bos.alt_disk_install.boot_images, bos.alt_disk.install.rte
```

- 5 Mount the Storage Foundation installation media.

Determine the filesets you want to install on the alternate disk by running the following command:

```
# ./installsf -install_option
```

where `install_option` is one of the following:

-minpkgs: For installing the minimum set of packages

-recpkgs: For installing the recommended packages

-allpkgs: For installing all packages

Copy the required filesets to a directory on the primary boot disk, for example `/usr`

If you are upgrading the operating system along with Storage Foundation, copy the necessary operating system filesets and the Storage Foundation filesets to a directory on the primary disk, for example `/usr`.

Upgrading Storage Foundation on an alternate disk

Use one of the following options to upgrade Veritas Storage Foundation on an alternate disk when not in an HA environment:

SMIT interface See [“To upgrade Storage Foundation on an alternate disk using the SMIT interface”](#) on page 205.

Manual See [“To upgrade Storage Foundation on an alternate disk manually”](#) on page 206.

Use one of the following options to upgrade Veritas Storage Foundation High Availability on an alternate disk:

SMIT interface See [“To upgrade Storage Foundation HA on an alternate disk using the SMIT interface”](#) on page 207.

Manual See [“To upgrade Storage Foundation HA on an alternate disk manually”](#) on page 210.

The procedure provides instructions to clone the primary boot environment to the alternate disk, upgrade Storage Foundation on the alternate disk, and reboot the system to start from the alternate disk.

Note: The alternate disk upgrade is a manual procedure that does not require the Veritas installation program. As a result, the installation and uninstallation scripts are not available in the `/opt/VRTS/install` directory. You can access these scripts from the installation media.

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

To upgrade Storage Foundation on an alternate disk using the SMIT interface

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

Start the SMIT menu :

```
# smit alt_clone
```

Provide information for the following fields in the SMIT menu.

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

```
hdisk1
```

Fileset(s) to install Enter the following:

```
all
```

Directory or Device with images Enter the full path of the directory that contains the filesets to be upgraded.

```
/usr/
```

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

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

- 2 Press **Enter** to start the upgrade on the alternate disk.

The upgrade process takes some time.

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

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

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

```
# shutdown -r
```

- 5 Verify the upgrade.
See “[Verifying the upgrade](#)” on page 217.

- 6 Complete the post-upgrade tasks.

To upgrade Storage Foundation on an alternate disk manually

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

```
# /usr/sbin/alt_disk_copy -I "acNgXY" -P "all" -l "/usr" -w \
"all" -d "hdisk1"
```

Where:

- `-d` indicates the name of the target disk on which you clone the primary disk.
- `-l` indicates the full path of the directory that contains the filesets to be upgraded
- `-w` indicates the list of Storage Foundation filesets that you want to upgrade on the alternate boot disk. The option `all` indicates that all the filesets contained in the directory you specified (using option `-l`) must be installed to the alternate boot disk.

- 2 Mount the alternate disk at `/altroot.5.10`.

```
# mount /dev/hdisk1 /altroot.5.10
```

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

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

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

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

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

```
# shutdown -r
```

- Verify the upgrade.
See “[Verifying the upgrade](#)” on page 217.
- Complete the post-upgrade tasks.

To upgrade Storage Foundation HA on an alternate disk using the SMIT interface

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

Start the SMIT menu :

```
# smit alt_clone
```

Provide information for the following fields in the SMIT menu.

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

```
hdisk1
```

Fileset(s) to install Enter the following:

```
all
```

Directory or Device with images Enter the full path of the directory that contains the filesets to be upgraded.

```
/usr/
```

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

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

- Press **Enter** to start the upgrade on the alternate disk.
The upgrade process takes some time.
- Use the following command to wake up the volume group on the alternate boot disk (`hdisk1`) that you cloned.

```
# /usr/sbin/alt_rootvg_op -W -d hdisk1
```

4 Verify that the alternate disk is created:

```
# lspv
Output similar to the following displays:
hdisk0          0009710fa9c79877    rootvg
hdisk1          0009710f0b90db93    altinst_rootvg
```

5 Run the following command to export the root path installation environment variable.

```
# export INSTALL_ROOT_PATH=/alt_inst
```

6 Run the following command on the alternate root path of any one node in the cluster to configure a Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure \
-use_llthost
```

The `-use_llthost` option indicates that the `/etc/llthost` file is used to determine the names of the nodes in the cluster. Alternatively, you can specify the node names instead of the file name.

7 Confirm that the you have created the Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -display \
-use_llthost
```

The output should resemble:

```
Finding existing UUID information ...

nodeA .... exist.

nodeB .... exist.

Done

Valid uuid exist on nodeA nodeB

{ef228450-1dd1-11b2-a7bb-5938100f2199} : nodeA nodeB
```

8 Change directory to `/alt_inst/etc/VRTSvcs/conf/config`.

```
# cd /alt_inst/etc/VRTSvcs/conf/config
```

- 9 Open the main.cf file and delete the include vvrtypes.cf line, which is deprecated. The VVR agents are now included in the updated 5.1 types.cf file. The main.cf file is in the /alt_inst/etc/VRTSvcs/conf/config directory.
- 10 Back up a copy of the old types.cf file and copy the new one for Storage Foundation to use.

```
# mv types.cf types.cf.ORIG
```

```
# cp ../types.cf .
```

- 11 If you have a secure cluster, perform the instructions in the following section:
See [“Upgrading a cluster that is in secure mode”](#) on page 214.
- 12 If you did not configure fencing in the 5.0 MP3 cluster, but want to use it in your 5.1 cluster, perform the instructions in the following section:
See [“Configuring fencing for an ADI upgrade”](#) on page 215.
- 13 Move to root and run the `alt_rootvg_op -S` command to put the alternate root to sleep.

```
# cd /
```

```
# alt_rootvg_op -S
```

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

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

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

```
# shutdown -r
```

16 Verify the upgrade.

See “[Verifying the upgrade](#)” on page 217.

17 Remove the obsolete filesets from the alternate disk to complete the upgrade.

You can determine the obsolete packages by comparing the list of installed filesets with the filesets that are displayed on running the following command:

```
# ./installsf -install_option
```

where `install_option` is one of the following:

-minpkgs: Displays the minimum set of packages

-recpkgs: Displays the recommended packages

-allpkgs: Displays all packages

If you find obsolete packages, remove them by running the following command:

```
# installp -u pkg_name
```

Note: Make sure that you do not remove the filesets of other Veritas products, such as Veritas Cluster Server Management Console, that may be installed on your system. If you need assistance, contact Symantec Technical Support.

To upgrade Storage Foundation HA on an alternate disk manually

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

```
# /usr/sbin/alt_disk_copy -I "acNgXY" -P "all" -l "/usr" -w \
"all" -d "hdisk1"
```

Where:

- `-d` indicates the name of the target disk on which you clone the primary disk.
- `-l` indicates the full path of the directory that contains the filesets to be upgraded
- `-w` indicates the list of Storage Foundation filesets that you want to upgrade on the alternate boot disk. The option `all` indicates that all the

filesets contained in the directory you specified (using option `-l`) must be installed to the alternate boot disk.

- 2 Use the following command to wake up the volume group on the alternate boot disk (`hdisk1`) that you cloned.

```
# alt_rootvg_op -W -d hdisk1
```

- 3 Verify that the alternate disk is created:

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

- 4 Run the following command to export the root path installation environment variable.

```
# export INSTALL_ROOT_PATH=/alt_inst
```

- 5 Run the following command on the alternate root path of any one node in the cluster to configure a Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -configure \
-use_llthost
```

The `-use_llthost` option indicates that the `/etc/llthost` file is used to determine the names of the nodes in the cluster. Alternatively, you can specify the node names instead of the file name.

- 6 Confirm that the you have created the Universal Unique Identifier:

```
# /alt_inst/opt/VRTSvcs/bin/uuidconfig.pl -clus -display \  
-use_llthost
```

The output should resemble:

```
Finding existing UUID information ...
```

```
nodeA .... exist.
```

```
nodeB .... exist.
```

```
Done
```

```
Valid uuid exist on nodeA nodeB
```

```
{ef228450-1dd1-11b2-a7bb-5938100f2199} : nodeA nodeB
```

- 7 Change directory to `cd /alt_inst/etc/VRTSvcs/conf/config`.

```
# cd /alt_inst/etc/VRTSvcs/conf/config
```

- 8 Open the `main.cf` file and delete the `include vvrtypes.cf` line, which is deprecated. The VVR agents are now included in the updated 5.1 `types.cf` file. The `main.cf` file is in the `/alt_inst/etc/VRTSvcs/conf/config` directory.
- 9 Back up a copy of the old `types.cf` file and copy the new one for Storage Foundation to use.

```
# mv types.cf types.cf.ORIG
```

```
# cp ../types.cf .
```

- 10 If you have a secure cluster, perform the instructions in the following section:
See [“Upgrading a cluster that is in secure mode”](#) on page 214.
- 11 If you did not configure fencing in the 5.0 MP3 cluster, but want to use it in your 5.1 cluster, perform the instructions in the following section:
See [“Configuring fencing for an ADI upgrade”](#) on page 215.

- 12 Move to root and run the `alt_rootvg_op -S` command to put the alternate root to sleep.

```
# cd /  
  
# alt_rootvg_op -S
```

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

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

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

```
# shutdown -r
```

- 15 Verify the upgrade.

See [“Verifying the upgrade”](#) on page 217.

- 16 Remove the obsolete filesets from the alternate disk to complete the upgrade. You can determine the obsolete packages by comparing the list of installed filesets with the filesets that are displayed on running the following command:

```
# ./installsf -install_option
```

where `install_option` is one of the following:

-minpkgs: Displays the minimum set of packages

-recpkgs: Displays the recommended packages

-allpkgs: Displays all packages

If you find obsolete packages, remove them by running the following command:

```
# installp -u pkg_name
```

Note: Make sure that you do not remove the filesets of other Veritas products, such as Veritas Cluster Server Management Console, that may be installed on your system. If you need assistance, contact Symantec Technical Support.

Upgrading a cluster that is in secure mode

If you had a secure cluster in your 5.0 MP3 cluster and want to keep it for 5.1, perform the following procedure.

To enable security for the upgraded secure cluster

1 Change directory to VRTSat.

```
# cd /alt_inst/var/VRTSat
```

2 Edit the /alt_inst/var/VRTSat/ABAuthSource file. Delete all HA_SERVICES-related entries in it. Remove text similar to the following:

```
[HA_SERVICES@symantecexample]
"PD_state"=dword:00000001
"PD_expiryinterval"=dword:00000000
[HA_SERVICES@symantecexample\admin]
"PD_password"=hex:8d,ab,d2,a3,fe, . . . c4,17,5d,6f,35,3c,12,40
"IsDomainAdmin"=dword:00000001
[HA_SERVICES@symantecexample\_HA_VCS_symantecexample]
"PD_principaltype"=dword:00000002
"PD_password"=hex:7f,31,af,c0,b2, . . . 6c,48,33,fe,13,2d,4e,56
"IsBrokerAdmin"=dword:00000000
"IsDomainAdmin"=dword:00000000
"CanAcceptProxyFlag"=dword:00000000
"CanProxyFlag"=dword:00000000
[HA_SERVICES@symantecexample\_CMDSERVER_VCS_symantecexample]
"PD_principaltype"=dword:00000002
"PD_password"=hex:da,79,b1,9d,fe, . . . 24,54,e1,90,fb,fb,fb,82
"IsDomainAdmin"=dword:00000000
"IsBrokerAdmin"=dword:00000000
"CanProxyFlag"=dword:00000000
"CanAcceptProxyFlag"=dword:00000000
[HA_SERVICES@symantecexample\webserver_VCS_symantecexample.com]
"PD_principaltype"=dword:00000002
"PD_password"=hex:38,29,ba,6d,57, . . . d1,c1,1d,ca,34,0c,82,9f
"IsDomainAdmin"=dword:00000000
"IsBrokerAdmin"=dword:00000000
"CanProxyFlag"=dword:00000001
"CanAcceptProxyFlag"=dword:00000000
"PD_expiryinterval"=dword:00000000
```

- 3 Touch /alt_inst/var/VRTSat/LocalAuthSource.

```
# touch /alt_inst/var/VRTSat/LocalAuthSource
```

- 4 Ensure that DNS has the IPv6 address for localhost. Add the following lines into /etc/hosts.

```
127.0.0.1 localhost
::1 localhost
```

Upgrading Storage Foundation and AIX on an alternate rootvg that is enabled for DMP on an LVM root disk

To upgrade Storage Foundation and AIX on an alternate rootvg that is enabled for DMP on an LVM root disk, perform the following procedure.

To upgrade Storage Foundation on altinst_rootvg on a DMP-enabled rootvg

- 1 Clone the currently running system to the alternate disk. For example:

```
# alt_disk_copy -d <hdisk##>
```

- 2 Boot from the alternate disk.
- 3 Upgrade AIX and Storage Foundation according to the instructions for alternate root upgrade.

See “[Upgrading Storage Foundation on an alternate disk](#)” on page 204.

- 4 Enable DMP on the root disk.

See “[Configuring DMP support for booting over a SAN](#)” on page 127.

Configuring fencing for an ADI upgrade

If you did not configure fencing in 5.0 MP3, the configuration files that fencing requires are missing on the primary boot disk. Since these configuration files are missing, ADI cannot install them when it installs the 5.1 packages on the alternate disk. The absence of these files means that fencing cannot start after the upgrade.

Configuring fencing in disabled mode for an ADI upgrade

Perform the following procedure to configure I/O fencing in disabled mode for the ADI upgrade.

To configure I/O fencing in disabled mode for an ADI upgrade

- ◆ On all the nodes in the cluster type:

```
# cp /alt_inst/etc/vxfen.d/vxfenmode_disabled \  
/alt_inst/etc/vxfenmode
```

Configuring fencing in SCSI-3 mode for an ADI upgrade

Perform the following procedures to configure I/O fencing in SCSI-3 mode for the ADI upgrade.

To update the I/O fencing files and start I/O fencing for an ADI upgrade using SCSI-3

- 1 Perform the tasks in the following section.

See [“Setting up coordinator disk groups”](#) on page 140.

When complete return to this procedure.

- 2 Create the `/alt_inst/etc/vxfendg` file, which includes the name of the coordinator disk group. On each node, type:

```
# echo "vxfencoordg" > /alt_inst/etc/vxfendg
```

Where *vxfencoordg* is an example name and changes based on the name that you give to the coordinator disk group. Do not use spaces between the quotes in the "vxfencoordg" text.

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

- For a DMP configuration:

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

- For a raw device configuration:

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

- 4 Check the updated `/etc/vxfenmode` configuration. Enter the following command on one of the nodes. For example:

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

Modifying VCS configuration to use I/O fencing

After you add coordinator disks and configure I/O fencing, add the `UseFence = SCSI3` cluster attribute to the VCS configuration file `/alt_inst/etc/VRTSvcs/conf/config/main.cf`. If you reset this attribute to `UseFence = None`, VCS will 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 on the alternate disk

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

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

- 2 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
)
```

- 3 Save and close the file.
- 4 Verify the syntax of the file `/alt_inst/etc/VRTSvcs/conf/config/main.cf`:

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

Verifying the upgrade

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

To verify the upgrade

- ◆ Verify that the alternate boot environment is active:

```
# lspv
Output similar to the following displays:
hdisk0          0009710fa9c79877    old_rootvg
hdisk1          0009710f0b90db93    rootvg
```


Performing a phased upgrade of Storage Foundation High Availability

This chapter includes the following topics:

- [About phased upgrade](#)
- [Performing a phased upgrade from Storage Foundation 5.0 MP3](#)

About phased upgrade

Perform a phased upgrade to minimize the downtime for the cluster. Depending on the situation, you can calculate the approximate downtime as follows:

You can fail over all your service groups to the nodes that are up.

Downtime equals the time that is taken to offline and online the service groups.

You have a service group that you cannot fail over to a node that runs during upgrade.

Downtime for that service group equals the time that is taken to perform an upgrade and restart the node.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade.

Planning for a phased upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group.

Some rough guidelines follow:

- Split the cluster in half. If the cluster has an odd number of nodes, calculate $(n+1)/2$, and start the upgrade with the even number of nodes.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.

Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules.
- When you start the installer, only select Storage Foundation.
- While you perform the upgrades, do not add or remove service groups to any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.

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:

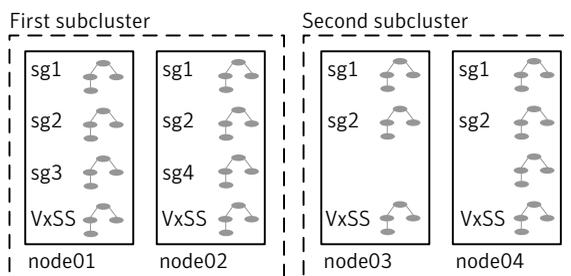
- sg1 and sg2 are parallel service groups and run on all the nodes.

- 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 13-1 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.

See [“Performing a phased upgrade from Storage Foundation 5.0 MP3”](#) on page 222.

Performing a phased upgrade from Storage Foundation 5.0 MP3

This section explains how to perform a phased upgrade of Storage Foundation on four nodes with four service groups. Note that in this scenario, the service groups cannot stay online during the upgrade of the second subcluster. Do not add, remove, or change resources or service groups on any nodes during the upgrade. These changes are likely to get lost after the upgrade. The following example illustrates the steps to perform a phased upgrade. The phased upgrade is from Storage Foundation 5.0 MP3 in a secure cluster to Storage Foundation 5.1 in a secure cluster.

See [“About phased upgrade”](#) on page 219.

Moving the service groups to the second subcluster

Perform the following steps to establish the service group's status and to switch the service groups.

To move service groups to the second subcluster

- 1 On the first subcluster, determine where the service groups are online.

```
# hagr -state
```

The output resembles:

```
#Group  Attribute System Value
sg1     State     node01 | ONLINE |
sg1     State     node02 | ONLINE |
sg1     State     node03 | ONLINE |
sg1     State     node04 | ONLINE |
sg2     State     node01 | ONLINE |
sg2     State     node02 | ONLINE |
sg2     State     node03 | ONLINE |
sg2     State     node04 | ONLINE |
sg3     State     node01 | ONLINE |
sg3     State     node02 | OFFLINE |
sg3     State     node03 | OFFLINE |
sg3     State     node04 | OFFLINE |
sg4     State     node01 | OFFLINE |
sg4     State     node02 | ONLINE |
sg4     State     node03 | OFFLINE |
sg4     State     node04 | OFFLINE |
VxSS    State     node01 | ONLINE |
VxSS    State     node02 | ONLINE |
VxSS    State     node03 | ONLINE |
VxSS    State     node04 | ONLINE |
```

- 2 Offline the parallel service groups (sg1 and sg2) and the VXSS group from the first subcluster. Switch the failover service groups (sg3 and sg4) from the first subcluster (node01 and node02) to the nodes on the second subcluster (node03 and node04).

```
# hagr -offline sg1 -sys node01
# hagr -offline sg2 -sys node01
# hagr -offline sg1 -sys node02
# hagr -offline sg2 -sys node02
# hagr -offline VxSS -sys node01
# hagr -offline VxSS -sys node02
# hagr -switch sg3 -to node03
# hagr -switch sg4 -to node04
```

- 3 Unmount all the VxFS file systems that Storage Foundation does not manage, for example:

```
# df -k
```

Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/hd4	20971520	8570080	60%	35736	2%	/
/dev/hd2	5242880	2284528	57%	55673	9%	/usr
/dev/hd9var	4194304	3562332	16%	5877	1%	/var
/dev/hd3	6291456	6283832	1%	146	1%	/tmp
/dev/hd1	262144	261408	1%	62	1%	/home
/dev/hd11admin	262144	184408	30%	6	1%	/admin
/proc	-	-	-	-	-	/proc
/dev/hd10opt	20971520	5799208	73%	65760	5%	/opt
/dev/vx/dsk/dg2/dg2vol1		10240	7600	26%	4	1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2		10240	7600	26%	4	1% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3		10240	7600	26%	4	1% /mnt/dg2/dg2vol3

```
# umount /mnt/dg2/dg2vol1  
# umount /mnt/dg2/dg2vol2  
# umount /mnt/dg2/dg2vol3
```

- 4 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

- 5 Freeze the nodes.

```
# hasys -freeze -persistent node01  
# hasys -freeze -persistent node02
```

- 6 Dump the configuration and make it read-only.

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

- 7 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|
```

- 8 Perform this step on node01 and node02 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
```

- 9 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 installsf.

```
# cd /storage_foundation
```

- 3 Make sure that Storage Foundation is running. Start the installsf program, specify the nodes in the first subcluster (node1 and node2).

```
# ./installsf node1 node2
```

The program starts with a copyright message and specifies the directory where it creates the logs.

4 Review the available installation options.

- 1 Selects Storage Foundation (SF)
- 2 Selects Storage Foundation and High Availability (SFHA).

For this example, select **2** to select SFHA.

```
Select a product to install? [1-2,q,?] (1) 2
```

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 EULA.pdf file present on media? [y,n,q,?]  
(y) y
```

6 Review the available installation options.

- 1 Installs only the minimal required Storage Foundation filesets that provides basic functionality of the product.
- 2 Installs the recommended Storage Foundation filesets that provides complete functionality of the product.
Note that this option is the default.
- 3 Installs all the Storage Foundation filesets.
You must choose this option to configure any optional Storage Foundation feature.
- 4 Displays the Storage Foundation filesets for each option.

For this example, select **3** for all filesets.

```
Select the filesets 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.

8 When you are prompted, reply **y** to continue with the upgrade.

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

9 When you are prompted, reply **y** to continue with the upgrade.

```
Do you want to upgrade SFHA? If you answer n then only SF
will be upgraded. [y,n,q] (y) y
```

10 When you are prompted, reply **y** to stop appropriate processes.

```
Do you want to stop Storage Foundation processes? [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 Storage Foundation does not manage, for example:

```
# df -k
```

Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on
Filesystem	1024-blocks	Free	%Used	Iused	%Iused	Mounted on
/dev/hd4	20971520	8570080	60%	35736	2%	/
/dev/hd2	5242880	2284528	57%	55673	9%	/usr
/dev/hd9var	4194304	3562332	16%	5877	1%	/var
/dev/hd3	6291456	6283832	1%	146	1%	/tmp
/dev/hd1	262144	261408	1%	62	1%	/home
/dev/hd11admin	262144	184408	30%	6	1%	/admin
/proc	-	-	-	-	-	/proc
/dev/hd10opt	20971520	5799208	73%	65760	5%	/opt
/dev/vx/dsk/dg2/dg2vol1		10240	7600	26%	4	1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2		10240	7600	26%	4	1% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3		10240	7600	26%	4	1% /mnt/dg2/dg2vol3

```
# umount /mnt/dg2/dg2vol1  
# umount /mnt/dg2/dg2vol2  
# umount /mnt/dg2/dg2vol3
```

- 3 Make the configuration writable on the second subcluster.

```
# haconf -makerw
```

- 4 Unfreeze the service groups.

```
# hagrps -unfreeze sg1 -persistent  
# hagrps -unfreeze sg2 -persistent  
# hagrps -unfreeze sg3 -persistent  
# hagrps -unfreeze sg4 -persistent  
# hagrps -unfreeze VxSS -persistent
```

- 5 Dump the configuration and make it read-only.

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

6 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
```

7 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

8 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
```

9 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

```
# /opt/VRTSvcs/bin/hastop -local
# /etc/init.d/vxfen.rc stop
# /etc/init.d/gab.rc stop
# /etc/init.d/llt.rc stop
```

10 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 not loaded.

```
# /sbin/vxfenconfig -l
VXFEN vxfenconfig ERROR V-11-2-1087 There are 0 active coordination points

# /sbin/gabconfig -l
GAB Driver Configuration
Driver state      : Unconfigured
Partition arbitration: Disabled
Control port seed : Disabled
Halt on process death: Disabled
Missed heartbeat halt: Disabled
Halt on rejoin   : Disabled
Keep on killing  : Disabled
Quorum flag      : Disabled
Restart          : Disabled
Node count       : 0
Disk HB interval (ms): 1000
Disk HB miss count : 4
IOFENCE timeout (ms) : 15000
Stable timeout (ms) : 5000

# /usr/sbin/strload -q -d /usr/lib/drivers/pse/llt
/usr/lib/drivers/pse/llt: no
```

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 -r
```

- 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.

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 `installsf`.

```
# cd /storage_foundation
```

- 3 Confirm that Storage Foundation is stopped on node03 and node04. Start the `installsf` program, specify the nodes in the second subcluster (node3 and node4).

```
# ./installsf node3 node4
```

The program starts with a copyright message and specifies the directory where it creates the logs.

- 4 Review the available installation options.

- 1 Selects Storage Foundation (SF)
- 2 Selects Storage Foundation and High Availability (SFHA).

For this example, select **2** to select SFHA.

```
Select a product to install? [1-2,q,?] (1) 2
```

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

6 Review the available installation options.

- 1 Installs only the minimal required Storage Foundation filesets that provides basic functionality of the product.
- 2 Installs the recommended Storage Foundation filesets that provides complete functionality of the product.
 Note that this option is the default.
- 3 Installs all the Storage Foundation filesets.
 You must choose this option to configure any optional Storage Foundation feature.
- 4 Displays the Storage Foundation filesets for each option.

For this example, select **3** for all filesets.

```
Select the filesets 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.

8 When you are prompted, reply **y to continue with the upgrade.**

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

9 When you are prompted, reply **y to continue with the upgrade.**

```
Do you want to upgrade SFHA? If you answer n then only SF
will be upgraded. [y,n,q] (y) y
```

10 When you are prompted, reply **y to stop appropriate processes.**

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

11 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 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
```

- 2 Reboot the node03 and node04 in the second subcluster.

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

The nodes in the second subcluster join the nodes in the first subcluster.

3 Check to see if Storage Foundation 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
```

4 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
```

In this example, you have performed a phased upgrade of Storage Foundation. The service groups were down when you took them offline on node03 and node04, to the time Storage Foundation brought them online on node01 or node02.

Verifying the Storage Foundation installation

This chapter includes the following topics:

- [Verifying that the products were installed](#)
- [Installation log files](#)
- [Starting and stopping processes for the Veritas products](#)
- [Checking Volume Manager processes](#)
- [Checking Veritas File System installation](#)
- [About enabling LDAP authentication for clusters that run in secure mode](#)
- [About the LLT and GAB configuration files](#)
- [Verifying the LLT, GAB, and VCS configuration files](#)
- [Verifying LLT, GAB, and cluster operation](#)

Verifying that the products were installed

Verify that the Veritas Storage Foundation products are installed.

Use the `ls1pp` command to check which packages have been installed.

```
# ls1pp -L | grep VRTS
```

The packages should be in the COMMITTED state.

Use the following sections to further verify the product installation.

Installation log files

After every product installation, the installer creates three text files:

- Installation log file
- Response file
- Summary file

The name and location of each file is displayed at the end of a product installation, and are always located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep the files for auditing, debugging, and future use.

Using the installation log file

The installation log file contains all commands executed during the procedure, their output, and errors generated by the commands. This file is for debugging installation problems and can be used for analysis by Veritas Support.

Using the summary file

The summary file contains the results of the installation by the common product installer or product installation scripts. The summary includes the list of the packages, and the status (success or failure) of each package. The summary also indicates which processes were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

To stop the processes

- ◆ Use the `-stop` option to the product installation script.

For example, to stop VCS processes, enter the following command:

```
# ./installvcs -stop
```

To start the processes

- ◆ Use the `-start` option to the product installation script.

For example: To start VCS processes, enter the following command:

```
# ./installvcs -start
```

Checking Volume Manager processes

Use the following procedure to verify that Volume Manager processes are running.

To confirm that key Volume Manager processes are running

- ◆ Type the following command:

```
# ps -ef | grep vx
```

Entries for the `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxcached`, and `vxconfigbackupd` processes should appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

Checking Veritas File System installation

After the Storage Foundation software has been successfully installed, you can confirm successful Veritas File System installation.

To confirm the File System installation

- ◆ Use the `lsvfs` command as follows:

```
# lsvfs vxfs
```

Entries for these processes appear in output similar to the following:

```
vxfs 32 /sbin/helpers/vxfs /sbin/helpers/vxfs
```

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.

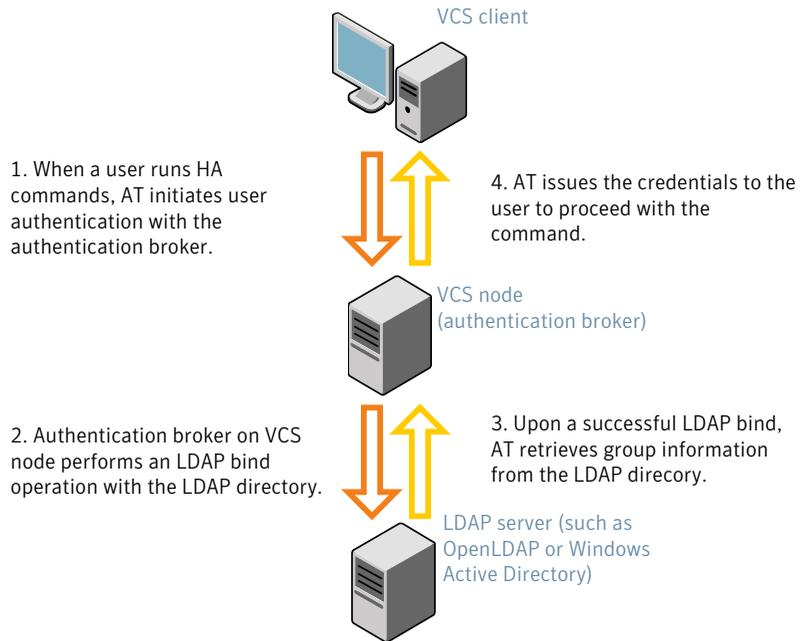
See “[Enabling LDAP authentication for clusters that run in secure mode](#)” on page 243.

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 14-1](#) depicts the Storage Foundation cluster communication with the LDAP servers when clusters run in secure mode.

Figure 14-1 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 Storage Foundation 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 Storage Foundation 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 Storage Foundation 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.

About the LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires `/etc/llthosts` and `/etc/llttab` files. GAB requires `/etc/gabtab` file.

LLT and GAB also require the initialization configuration files:

- `/etc/default/llt`
- `/etc/default/gab`

The information that these LLT and GAB configuration files contain is as follows:

- The `/etc/default/llt` file

This file stores the start and stop environment variables for LLT:

- `LLT_START`—Defines the startup behavior for the LLT module after a system reboot. Valid values include:
 - 1—Indicates that LLT is enabled to start up.
 - 0—Indicates that LLT is disabled to start up.
- `LLT_STOP`—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include:
 - 1—Indicates that LLT is enabled to shut down.
 - 0—Indicates that LLT is disabled to shut down.

The installer sets the value of these variables to 1 at the end of Storage Foundation HA configuration.

- The `/etc/llthosts` file

The file `llthosts` is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file

must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.

For example, the file `/etc/llthosts` contains the entries that resemble:

```
0      galaxy
1      nebula
```

■ The `/etc/llttab` file

The file `llttab` contains the information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the LLT network links that correspond to the specific system.

For example, the file `/etc/llttab` contains the entries that resemble:

```
set-node galaxy
set-cluster 2
link en1 /dev/dlpi/en:1 - ether - -
link en2 /dev/dlpi/en:2 - ether - -

set-node galaxy
set-cluster 2
link en1 /dev/en:1 - ether - -
link en2 /dev/en:2 - ether - -
```

The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the `link` command. These lines identify the two network cards that the LLT protocol uses.

If you configured a low priority link under LLT, the file also includes a "link-lopri" line.

Refer to the `llttab(4)` manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the `llttab` file.

■ The `/etc/default/gab` file

This file stores the start and stop environment variables for GAB:

- `GAB_START`—Defines the startup behavior for the GAB module after a system reboot. Valid values include:

- 1—Indicates that GAB is enabled to start up.
- 0—Indicates that GAB is disabled to start up.

- `GAB_STOP`—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include:

- 1—Indicates that GAB is enabled to shut down.
- 0—Indicates that GAB is disabled to shut down.

The installer sets the value of these variables to 1 at the end of Storage Foundation HA configuration.

- The `/etc/gabtab` file

After you install Storage Foundation, the file `/etc/gabtab` contains a `gabconfig(1)` command that configures the GAB driver for use.

The file `/etc/gabtab` contains a line that resembles:

```
/sbin/gabconfig -c -nN
```

The `-c` option configures the driver for use. The `-nN` specifies that the cluster is not formed until at least `N` nodes are ready to form the cluster. Symantec recommends that you set `N` to be the total number of nodes in the cluster.

Note: Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. Using `-c -x` can lead to a split-brain condition.

Verifying the LLT, GAB, and VCS configuration files

Make sure that the LLT, GAB, and VCS configuration files contain the information you provided during VCS installation and configuration.

To verify the LLT, GAB, and VCS configuration files

- 1 Navigate to the location of the configuration files:

- LLT
`/etc/llthosts`
`/etc/llttab`
- GAB
`/etc/gabtab`
- VCS
`/etc/VRTSvcs/conf/config/main.cf`

- 2 Verify the content of the configuration files.

See [“About the LLT and GAB configuration files”](#) on page 249.

Verifying LLT, GAB, and cluster operation

Verify the operation of LLT, GAB, and the cluster using the VCS commands.

To verify LLT, GAB, and cluster operation

- 1 Log in to any node in the cluster as superuser.
- 2 Make sure that the PATH environment variable is set to run the VCS commands.
- 3 Verify LLT operation.
See “[Verifying LLT](#)” on page 252.
- 4 Verify GAB operation.
- 5 Verify the cluster operation.
See “[Verifying the cluster](#)” on page 255.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

- 1 Log in as superuser on the node galaxy.
- 2 Run the `lltstat` command on the node galaxy to view the status of LLT.

```
lltstat -n
```

The output on galaxy resembles:

```
LLT node information:
Node           State      Links
*0 galaxy      OPEN      2
 1 nebula      OPEN      2
```

Each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
Node           State      Links
* 0 galaxy      OPEN      2
 1 nebula      OPEN      2
 2 saturn       OPEN      1
```

- 3 Log in as superuser on the node nebula.
- 4 Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
Node           State      Links
 0 galaxy      OPEN      2
*1 nebula      OPEN      2
```

- 5 To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv | more
```

The output on galaxy resembles:

Node	State	Link	Status	Address
*0 galaxy	OPEN			
		en1	UP	08:00:20:93:0E:34
		en2	UP	08:00:20:93:0E:34
1 nebula	OPEN			
		en1	UP	08:00:20:8F:D1:F2
		en2	DOWN	
2	CONNWAIT			
		en1	DOWN	
		en2	DOWN	
3	CONNWAIT			
		en1	DOWN	
		en2	DOWN	
.				
.				
.				
31	CONNWAIT			
		en1	DOWN	
		/dev/en:2	DOWN	

Note that the output lists 32 nodes. The command reports the status on the two nodes in the cluster, galaxy and nebula, along with the details for the non-existent nodes.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- A MAC address for each link

However, the output in the example shows different details for the node nebula. The private network connection is possibly broken or the information in the `/etc/llttab` file may be incorrect.

- 6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node galaxy in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
  Port  Usage      Cookie
  ---  ---
  0     gab        0x0
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects: 0 1
  7     gab        0x7
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects: 0 1
  31    gab        0x1F
        opens:    0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects: 0 1
```

Verifying the cluster

Verify the status of the cluster using the `hastatus` command. This command returns the system state and the group state.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for a description of system states and the transitions between them.

To verify the cluster

- 1 To verify the status of the cluster, type the following command:

```
hastatus -summary
```

The output resembles:

```
-- SYSTEM STATE
-- System                State                Frozen

A galaxy                 RUNNING             0
A nebula                 RUNNING             0

-- GROUP STATE
-- Group                System                Probed  AutoDisabled  State
```

- 2 Review the command output for the following information:

- The system state

If the value of the system state is RUNNING, the cluster is successfully started.

Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

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

Refer to the *Veritas Cluster Server Administrator's Guide* for information about the system attributes for VCS.

To verify the cluster nodes

- ◆ On one of the nodes, type the `hasys -display` command:

```
hasys -display
```

The example shows the output when the command is run on the node galaxy. The list continues with similar information for nebula (not shown) and any other nodes in the cluster.

```
#System  Attribute                Value
galaxy   AgentsStopped           0
```

#System	Attribute	Value
galaxy	AvailableCapacity	100
galaxy	CPUUsage	0
galaxy	CPUUsageMonitoring	Enabled 0 ActionThreshold 0 ActionTimeLimit 0 Action NONE NotifyThreshold 0 NotifyTimeLimit 0
galaxy	Capacity	100
galaxy	ConfigBlockCount	117
galaxy	ConfigChecksum	10844
galaxy	ConfigDiskState	CURRENT
galaxy	ConfigFile	/etc/VRTSvcs/conf/config
galaxy	ConfigInfoCnt	0
galaxy	ConfigModDate	Wed 14 Oct 2009 17:22:48
galaxy	ConnectorState	Down
galaxy	CurrentLimits	
galaxy	DiskHbStatus	
galaxy	DynamicLoad	0
galaxy	EngineRestarted	0
galaxy	EngineVersion	5.1.00.0
galaxy	Frozen	0
galaxy	GUIIPAddr	
galaxy	HostUtilization	CPU 0 Swap 0
galaxy	LLTNodeId	0
galaxy	LicenseType	DEMO
galaxy	Limits	
galaxy	LinkHbStatus	link2 UP link3 UP
galaxy	LoadTimeCounter	0

#System	Attribute	Value
galaxy	LoadTimeThreshold	600
galaxy	LoadWarningLevel	80
galaxy	NoAutoDisable	0
galaxy	NodeId	0
galaxy	OnGrpCnt	1
galaxy	ShutdownTimeout	600
galaxy	SourceFile	./main.cf
galaxy	SysInfo	Aix:galaxy,5,2,00023BDA4C00
galaxy	SysName	galaxy
galaxy	SysState	RUNNING
galaxy	SystemLocation	
galaxy	SystemOwner	
galaxy	TFrozen	0
galaxy	TRSE	0
galaxy	UpDownState	Up
galaxy	UserInt	0
galaxy	UserStr	
galaxy	VCSFeatures	DR
galaxy	VCSMode	

Adding and removing nodes in Storage Foundation and High Availability clusters

This chapter includes the following topics:

- [About adding and removing nodes](#)
- [Adding nodes using the Storage Foundation installer](#)
- [Manually adding a node to a cluster](#)
- [Removing a node from a cluster](#)
- [Adding a node to a single-node cluster](#)

About adding and removing nodes

After you install Storage Foundation and create a cluster, you can add and remove nodes from the cluster. You can create a cluster of up to 32 nodes.

The Veritas product installer supports adding a node. You can also add a node manually. The Veritas product installer does not support removing a node. You must remove a node manually.

Adding nodes using the Storage Foundation installer

The Storage Foundation 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/clusuuid
 - /etc/default/llt
 - /etc/default/gab
 - /etc/default/vxfen
- Configures security on the new node if the existing cluster is a secure 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 Storage Foundation 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 Storage Foundation cluster using the Storage Foundation installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the Storage Foundation installer with the `-addnode` option.

```
# cd /opt/VRTS/install  
  
# ./installsf -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 Storage Foundation cluster. The installer uses the node information to identify the existing cluster.

```
Enter a node name in the Storage Foundation cluster to which  
you want to add a node: galaxy
```

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

```
Enter the system names separated by spaces  
to add to the cluster: saturn
```

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

- 6 Enter the name of the network interface that you want to configure as the first private heartbeat link.

Note: The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

```
Enter the NIC for the first private heartbeat  
link on saturn: [b,q,?] en1
```

- 7 Enter y to configure a second private heartbeat link.

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

```
Would you like to configure a second private  
heartbeat link? [y,n,q,b,?] (y)
```

- 8 Enter the name of the network interface that you want to configure as the second private heartbeat link.

```
Enter the NIC for the second private heartbeat link  
on saturn: [b,q,?] en2
```

- 9 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

- 10 Review and confirm the information.

- 11 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: en3
```

- 12 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 Storage Foundation 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)
```

- 13 The installer starts the SF HA cluster processes on the new node. The new node is now part of the cluster.

Manually adding a node to a cluster

The system you add to the cluster must meet the hardware and software requirements.

Table 15-1 specifies the tasks that are involved in adding a cluster. The example demonstrates how to add a node saturn to already existing nodes, galaxy and nebula.

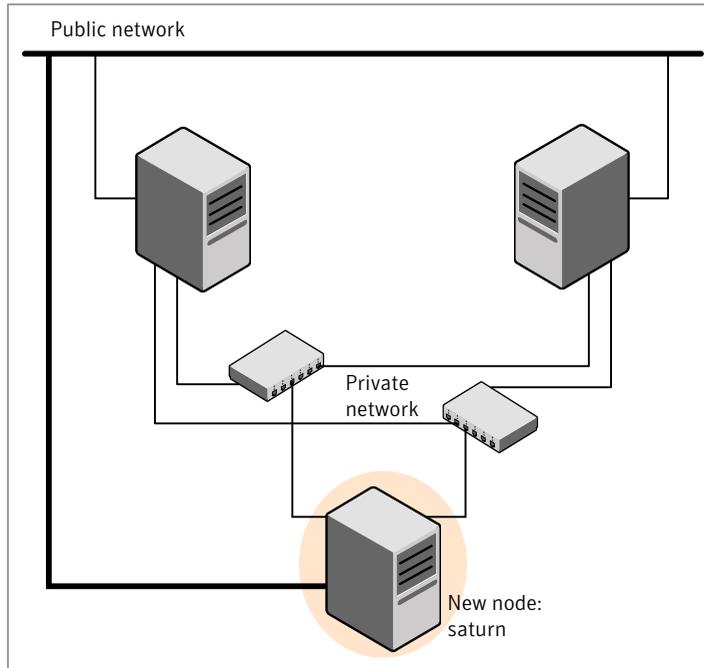
Table 15-1 Tasks that are involved in adding a node to a cluster

Task	Reference
Set up the hardware.	See “Setting up the hardware” on page 263.
Install the software manually.	
Add a license key.	See “Setting or changing the product level for keyless licensing” on page 129.
If the existing cluster runs in secure mode, set up the new node to run in secure mode.	See “Setting up the node to run in secure mode” on page 265.
Configure LLT and GAB.	See “Configuring LLT and GAB” on page 267.
If the existing cluster is configured for I/O fencing, configure I/O fencing on the new node.	See “Configuring I/O fencing on the new node” on page 269.
Add the node to the existing cluster.	See “Adding the node to the existing cluster” on page 274.
Start Storage Foundation and verify the cluster.	See “Starting Storage Foundation and verifying the cluster” on page 275.

Setting up the hardware

Figure 15-1 shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

Figure 15-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the Storage Foundation private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a two-node cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a two-node cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

[Figure 15-1](#) illustrates a new node being added to an existing two-node cluster using two independent hubs.

- 2 Connect the system to the shared storage, if required.

Installing the Storage Foundation software manually when adding a node

Install the Veritas Storage Foundation 5.1 filesets manually and add a license key.

For more information, see the following:

- See [“Setting or changing the product level for keyless licensing”](#) on page 129.

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.

See [“Configuring LLT and GAB”](#) on page 267.

[Table 15-2](#) uses the following information for the following command examples.

Table 15-2 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.

See [“Configuring the authentication broker on node saturn”](#) on page 266.

- 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.

See “[Setting up Storage Foundation related security configuration](#)” on page 267.

- 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 Storage Foundation related security configuration

Perform the following steps to configure Storage Foundation related security settings.

Setting up Storage Foundation related security configuration

- 1 Start /opt/VRTSat/bin/vxatd process.
- 2 Create HA_SERVICES domain for Storage Foundation.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```
- 3 Add Storage Foundation and webserver principal to AB on node saturn.

```
# vssat addprpl --pdrtype 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
```

Configuring LLT and GAB

Create the LLT and GAB configuration files on the new node and update the files on the existing nodes.

To configure LLT

- 1 Create the file /etc/llthosts on the new node. You must also update it on each of the current nodes in the cluster.

For example, suppose you add saturn to a cluster consisting of galaxy and nebula:

- If the file on one of the existing nodes resembles:

```
0 galaxy
1 nebula
```

- Update the file for all nodes, including the new one, resembling:

```
0 galaxy
1 nebula
2 saturn
```

- 2 Create the file `/etc/llttab` on the new node, making sure that line beginning `"set-node"` specifies the new node.

The file `/etc/llttab` on an existing node can serve as a guide.

The following example describes a system where node saturn is the new node on cluster ID number 2:

```
set-node saturn
set-cluster 2
link en1 /dev/dlpi/en:1 - ether - -
link en2 /dev/dlpi/en:2 - ether - -
```

- 3 Copy the following file from one of the nodes in the existing cluster to the new node:

```
/etc/default/llt
```

- 4 On the new system, run the command:

```
# /sbin/lltconfig -c
```

To configure GAB

- 1 Create the file `/etc/gabtab` on the new system.

- If the `/etc/gabtab` file on the existing nodes resembles:

```
/sbin/gabconfig -c
```

The file on the new node should be the same. Symantec recommends that you use the `-c -nN` option, where *N* is the total number of cluster nodes.

- If the `/etc/gabtab` file on the existing nodes resembles:

```
/sbin/gabconfig -c -n2
```

The file on all nodes, including the new node, should change to reflect the change in the number of cluster nodes. For example, the new file on each node should resemble:

```
/sbin/gabconfig -c -n3
```

The `-n` flag indicates to Storage Foundation the number of nodes that must be ready to form a cluster before Storage Foundation starts.

- 2 Copy the following file from one of the nodes in the existing cluster to the new node:

```
/etc/default/gab
```

- 3 On the new node, to configure GAB run the command:

```
# /sbin/gabconfig -c
```

To verify GAB

- 1 On the new node, run the command:

```
# /sbin/gabconfig -a
```

The output should indicate that port a membership shows all nodes including the new node. The output should resemble:

```
GAB Port Memberships
=====
Port a gen a3640003 membership 012
```

- 2 Run the same command on the other nodes (galaxy and nebula) to verify that the port a membership includes the new node:

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen a3640003 membership 012
Port h gen fd570002 membership 01
Port h gen fd570002 visible ; 2
```

Configuring I/O fencing on the new node

- Prepare to configure I/O fencing on the new node.
 See [“Preparing to configure I/O fencing on the new node”](#) on page 270.

- If the existing cluster runs server-based fencing, configure server-based fencing on the new node.
See [“Configuring server-based fencing on the new node”](#) on page 270.
- Copy the I/O fencing files from an existing node to the new node and start I/O fencing on the new node.
See [“Starting I/O fencing on the new node”](#) on page 273.

Preparing to configure I/O fencing on the new node

Perform the following tasks before you configure and start I/O fencing on the new node.

To prepare to configure I/O fencing on the new node

- ◆ Determine whether the existing cluster runs disk-based or server-based fencing mechanism. On one of the nodes in the existing cluster, run the following command:

```
# vxfsadm -d
```

If the fencing mode in the output is SCSI3, then the cluster uses disk-based fencing.

If the fencing mode in the output is CUSTOMIZED, then the cluster uses server-based fencing.

Configuring server-based fencing on the new node

This section describes the procedures to configure server-based fencing on a new node. Depending on whether server-based fencing is configured in secure or non-secure mode on the existing cluster, perform the tasks in one of the following procedures:

- Server-based fencing in non-secure mode:
[To configure server-based fencing in non-secure mode on the new node](#)
- Server-based fencing in secure mode:
[To configure server-based fencing with security on the new node](#)

To configure server-based fencing in non-secure mode on the new node

- 1 Log in to each CP server as the root user.
- 2 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 3 Verify that the new node is added to the CP server configuration:

```
# cpsadm -s mycps1.symantecexample.com \  
-a list_nodes
```

The new node must be listed in the command output.

- 4 Add the VCS user cpsclient@saturn to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e cpsclient@saturn \  
-f cps_operator -g vx
```

```
User cpsclient@saturn successfully added
```

Perform the following procedure for a secure configuration.

To configure server-based fencing with security on the new node

- 1 As the root user, create the VCS user and the domain on the new node:
 - Create a dummy configuration file /etc/VRTSvcs/conf/config/main.cf that resembles the following example:

```
# cat main.cf  
  
include "types.cf"  
cluster clus1 {  
    SecureClus = 1  
}  
  
system saturn {  
}
```

- Verify the dummy configuration file:

```
# cd /etc/VRTSvcs/conf/config  
# /opt/VRTSvcs/bin/hacf -verify .
```

- Start VCS in one node mode on the new node:

```
# /opt/VRTSvcs/bin/hastart -onenode
```

- 2 Verify that the VCS user and the domain are created on the new node:

```
# /opt/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 -local
```

- 4 If the root broker for the CP server and the new node are different, run the following command to establish trust between the authentication broker of the CP Server and the new node:

```
# /usr/bin/echo y | /opt/VRTSvcs/bin/cpsat setuptrust \  
-b mycps1.symantecexample.com -s high
```

- 5 Log in to each CP server as the root user.

- 6 Update each CP server configuration with the new node information:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_node -c clus1 -h saturn -n2
```

```
Node 2 (saturn) successfully added
```

- 7 Verify that the new node is added to the CP server configuration:

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

The new node must be listed in the output.

- 8 Add the VCS user `_HA_VCS_saturn@HA_SERVICES@saturn.symantec.com` to each CP server:

```
# cpsadm -s mycps1.symantecexample.com \  
-a add_user -e _HA_VCS_saturn@HA_SERVICES@saturn.symantec.com \  
-f cps_operator -g vx
```

User `_HA_VCS_saturn@HA_SERVICES@saturn.symantec.com` successfully added

Adding the new node to the vxfen service group

Perform the steps in the following procedure to add the new node to the vxfen service group.

To add the new node to the vxfen group using the CLI

- 1 On one of the nodes in the existing SF HA cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the node saturn to the existing vxfen group.

```
# hagrps -modify vxfen SystemList -add saturn 2
```

- 3 Save the configuration by running the following command from any node in the SF HA cluster:

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

Starting I/O fencing on the new node

Copy the I/O fencing files from an existing node to the new node and start I/O fencing on the new node. This task starts I/O fencing based on the fencing mechanism that is configured in the existing cluster.

To start I/O fencing on the new node

- 1 Copy the following I/O fencing configuration files from one of the nodes in the existing cluster to the new node:

- `/etc/vxfenmode`

- /etc/vxfendg—This file is required only for disk-based fencing.
 - /etc/default/vxfen
- 2 Start I/O fencing on the new node.

```
# /etc/init.d/vxfen.rc start
```
 - 3 Run the GAB configuration command on the new node to verify that the port b membership is formed.

```
# gabconfig -a
```

Adding the node to the existing cluster

Perform the tasks on one of the existing nodes in the cluster.

To add the new node to the existing cluster

- 1 Copy the cluster UUID from the one of the nodes in the existing cluster to the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -copy -from_sys \  
node_name_in_running_cluster -to_sys new_sys1 ... new_sysn
```

Where you are copying the cluster UUID from a node in the cluster (*node_name_in_running_cluster*) to systems from *new_sys1* through *new_sysn* that you want to join the cluster.

- 2 Copy the following file from one of the nodes in the existing cluster to the new node:

```
/etc/default/vcs
```

- 3 Enter the command:

```
# haconf -makerw
```

- 4 Add the new system to the cluster:

```
# hasys -add saturn
```

- 5 Copy the main.cf file from an existing node to your new node:

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

- 6 Check the VCS configuration file. No error message and a return value of zero indicates that the syntax is legal.

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

- 7 If necessary, modify any new system attributes.
- 8 Enter the command:

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

Starting Storage Foundation and verifying the cluster

Start Storage Foundation after adding the new node to the cluster and verify the cluster.

To start Storage Foundation and verify the cluster

- 1 Start Storage Foundation on the newly added system:

```
# hastart
```

- 2 Run the GAB configuration command on each node to verify that port a and port h include the new node in the membership:

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
Port a gen a3640003 membership 012
Port h gen fd570002 membership 012
```

If the cluster uses I/O fencing, then the GAB output also shows port b membership.

Removing a node from a cluster

[Table 15-3](#) 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 15-3 Tasks that are involved in removing a node

Task	Reference
<ul style="list-style-type: none"> ■ Back up the configuration file. ■ Check the status of the nodes and the service groups. 	<p>See “Verifying the status of nodes and service groups” on page 276.</p>
<ul style="list-style-type: none"> ■ Switch or remove any Storage Foundation service groups on the node departing the cluster. ■ Delete the node from Storage Foundation configuration. 	<p>See “Deleting the departing node from Storage Foundation configuration” on page 277.</p>
<p>Modify the llhosts and gabtab files to reflect the change.</p>	<p>See “Modifying configuration files on each remaining node” on page 280.</p>
<p>If the existing cluster is configured to use server-based I/O fencing, remove the node configuration from the CP server.</p>	<p>See “Removing the node configuration from the CP server” on page 280.</p>
<p>For a cluster that is running in a secure mode, remove the security credentials from the leaving node.</p>	<p>See “Removing security credentials from the leaving node ” on page 281.</p>
<p>On the node departing the cluster:</p> <ul style="list-style-type: none"> ■ Modify startup scripts for LLT, GAB, and Storage Foundation to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. ■ Remove the Storage Foundation filesets. 	<p>See “Unloading LLT and GAB and removing VCS on the departing node” on page 281.</p>

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 Storage Foundation 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 Storage Foundation on the departing node:

```
# hastop -sys saturn
```

- 5 Check the status again. The state of the departing node should be EXITED. Make sure that any service group that you want to fail over is online on other nodes.

```
# hastatus -summary
```

```
-- SYSTEM STATE  
-- System      State          Frozen  
A galaxy      RUNNING       0  
A nebula      RUNNING       0  
A saturn      EXITED        0  
  
-- GROUP STATE  
-- Group      System        Probed  AutoDisabled  State  
B grp1       galaxy        Y       N              ONLINE  
B grp1       nebula        Y       N              OFFLINE  
B grp2       galaxy        Y       N              ONLINE  
B grp3       nebula        Y       N              ONLINE  
B grp3       saturn        Y       Y              OFFLINE  
B grp4       saturn        Y       N              OFFLINE
```

- 6** Delete the departing node from the SystemList of service groups grp3 and grp4.

```
# hagrps -modify grp3 SystemList -delete saturn
# hagrps -modify grp4 SystemList -delete saturn
```

- 7** For the service groups that run only on the departing node, delete the resources from the group before you delete the group.

```
# hagrps -resources grp4
    processx_grp4
    processy_grp4
# hares -delete processx_grp4
# hares -delete processy_grp4
```

- 8** Delete the service group that is configured to run on the departing node.

```
# hagrps -delete grp4
```

- 9** Check the status.

```
# hastatus -summary
-- SYSTEM STATE
-- System      State          Frozen
A galaxy       RUNNING        0
A nebula       RUNNING        0
A saturn       EXITED         0

-- GROUP STATE
-- Group      System      Probed  AutoDisabled  State
B grp1       galaxy     Y       N              ONLINE
B grp1       nebula     Y       N              OFFLINE
B grp2       galaxy     Y       N              ONLINE
B grp3       nebula     Y       N              ONLINE
```

- 10** Delete the node from the cluster.

```
# hasys -delete saturn
```

- 11** Save the configuration, making it read only.

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

Modifying configuration files on each remaining node

Perform the following tasks on each of the remaining nodes of the cluster.

To modify the configuration files on a remaining node

- 1 If necessary, modify the `/etc/gabtab` file.

No change is required to this file if the `/sbin/gabconfig` command has only the argument `-c`. Symantec recommends using the `-nN` option, where *N* is the number of cluster systems.

If the command has the form `/sbin/gabconfig -c -nN`, where *N* is the number of cluster systems, make sure that *N* is not greater than the actual number of nodes in the cluster. When *N* is greater than the number of nodes, GAB does not automatically seed.

Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`.

- 2 Modify `/etc/llhosts` file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 galaxy
1 nebula
2 saturn
```

To:

```
0 galaxy
1 nebula
```

Removing the node configuration from the CP server

After removing a node from a SF HA cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

To remove the node configuration from the CP server

- 1 Log into the CP server as the root user.
- 2 View the list of VCS users on the CP server, using the following command:

```
# cpsadm -s cp_server -a list_users
```

Where *cp_server* is the virtual IP/ virtual hostname of the CP server.

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e _HA_VCS_saturn@HA_SERVICES@saturn.nodes.example.com \  
-f cps_operator -g vx
```

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@saturn -f cps_operator -g vx
```

- 4 Remove the node entry from the CP server:

```
cpsadm -s cp_server -a rm_node -h saturn -c clus1 -n 2
```

- 5 View the list of nodes on the CP server to ensure that the node entry was removed:

```
cpsadm -s cp_server -a list_nodes
```

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node saturn. Perform the following steps.

To remove the security credentials

- 1 Kill /opt/VRTSat/bin/vxatd process.
- 2 Remove the root credentials on node saturn.

```
# vssat deletecred --domain type:domainname --prplname prplname
```

Unloading LLT and GAB and removing VCS on the departing node

On the node departing the cluster, unconfigure and unload the LLT and GAB utilities, and remove the VCS filesets.

If you have configured Storage Foundation HA as part of the Storage Foundation and High Availability products, you may have to delete other dependent filesets before you can delete all of the following ones.

Adding a node to a single-node cluster

All nodes in the new cluster must run the same version of Storage Foundation. The example procedure refers to the existing single-node Storage Foundation node as Node A. The node that is to join Node A to form a multiple-node cluster is Node B.

[Table 15-4](#) specifies the activities that you need to perform to add nodes to a single-node cluster.

Table 15-4 Tasks to add a node to a single-node cluster

Task	Reference
Set up Node B to be compatible with Node A.	See “Setting up a node to join the single-node cluster” on page 283.
<ul style="list-style-type: none"> ■ Add Ethernet cards for private heartbeat network for Node B. ■ If necessary, add Ethernet cards for private heartbeat network for Node A. ■ Make the Ethernet cable connections between the two nodes. 	See “Installing and configuring Ethernet cards for private network” on page 283.
Connect both nodes to shared storage.	See “Configuring the shared storage” on page 284.
<ul style="list-style-type: none"> ■ Bring up Storage Foundation on Node A. ■ Edit the configuration file. 	See “Bringing up the existing node” on page 284.
<p>If necessary, install Storage Foundation on Node B and add a license key.</p> <p>Make sure Node B is running the same version of Storage Foundation as the version on Node A.</p>	See “Installing the Storage Foundation software manually when adding a node to a single node cluster” on page 285.
Edit the configuration files on Node B.	See “Creating configuration files” on page 285.
Start LLT and GAB on Node B.	See “Starting LLT and GAB” on page 286.
<ul style="list-style-type: none"> ■ Start LLT and GAB on Node A. ■ Restart VCS on Node A. ■ Modify service groups for two nodes. 	See “Reconfiguring Storage Foundation on the existing node” on page 286.

Table 15-4 Tasks to add a node to a single-node cluster (*continued*)

Task	Reference
<ul style="list-style-type: none"> ■ Start VCS on Node B. ■ Verify the two-node cluster. 	See “Verifying configuration on both nodes” on page 287.

Setting up a node to join the single-node cluster

The new node to join the existing single node that runs Storage Foundation must run the same operating system.

To set up a node to join the single-node cluster

- 1 Do one of the following tasks:
 - If Storage Foundation is not currently running on Node B, proceed to step 2.
 - If the node you plan to add as Node B is currently part of an existing cluster, remove the node from the cluster. After you remove the node from the cluster, remove the Storage Foundation filesets and configuration files.
See [“Removing a node from a cluster”](#) on page 275.
 - If the node you plan to add as Node B is also currently a single Storage Foundation node, uninstall Storage Foundation.
 - If you renamed the LLT and GAB startup files, remove them.
- 2 If necessary, install VxVM and VxFS.
See [“Installing VxVM or VxFS if necessary”](#) on page 283.

Installing VxVM or VxFS if necessary

If you have either VxVM or VxFS with the cluster option installed on the existing node, install the same version on the new node.

Refer to the appropriate documentation for VxVM and VxFS to verify the versions of the installed products. Make sure the same version runs on all nodes where you want to use shared storage.

Installing and configuring Ethernet cards for private network

Both nodes require Ethernet cards (NICs) that enable the private network. If both Node A and Node B have Ethernet cards installed, you can ignore this step.

For high availability, use two separate NICs on each node. The two NICs provide redundancy for heartbeating.

To install and configure Ethernet cards for private network

- 1 Shut down Storage Foundation on Node A.

```
# hastop -local
```

- 2 Shut down the nodes.

```
# shutdown -F
```

- 3 Install the Ethernet card on Node A.

If you want to use aggregated interface to set up private network, configure aggregated interface.

- 4 Install the Ethernet card on Node B.

If you want to use aggregated interface to set up private network, configure aggregated interface.

- 5 Configure the Ethernet card on both nodes.

- 6 Make the two Ethernet cable connections from Node A to Node B for the private networks.

- 7 Restart the nodes.

Configuring the shared storage

Make the connection to shared storage from Node B. Configure VxVM on Node B and reboot the node when you are prompted.

See [“Setting up shared storage”](#) on page 41.

Bringing up the existing node

Bring up the node.

To bring up the node

- 1 Log in as superuser.

- 2 Make the Storage Foundation configuration writable.

```
# haconf -makerw
```

- 3 Display the service groups currently configured.

```
# hagrps -list
```

- 4 Freeze the service groups.

```
# hagrps -freeze group -persistent
```

Repeat this command for each service group in step 3.

- 5 Make the configuration read-only.

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

- 6 Stop Storage Foundation on Node A.

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

- 7 Enable the GAB and LLT startup files so they can be used.

```
# mv /etc/rc.d/rc2.d/x92gab /etc/rc.d/rc2.d/S92gab
```

```
# mv /etc/rc.d/rc2.d/x7011t /etc/rc.d/rc2.d/S7011t
```

Installing the Storage Foundation software manually when adding a node to a single node cluster

Install the Veritas Storage Foundation 5.1 filesets manually and install the license key.

Refer to the following sections:

-
-
- See [“Setting or changing the product level for keyless licensing”](#) on page 129.

Creating configuration files

Create the configuration files for your cluster.

To create the configuration files

- 1 Create the file `/etc/llttab` that lists both the nodes.
- 2 Create the file `/etc/llthosts`. Set up `/etc/llthosts` for a two-node cluster.
- 3 Create the file `/etc/gabtab`.

Starting LLT and GAB

On the new node, start LLT and GAB.

To start LLT and GAB

- 1 Start LLT on Node B.

```
# /etc/init.d/llt.rc start
```
- 2 Start GAB on Node B.

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

Reconfiguring Storage Foundation on the existing node

Reconfigure Storage Foundation on the existing nodes.

To reconfigure Storage Foundation on existing nodes

- 1 On Node A, create the files `/etc/llttab`, `/etc/llthosts`, and `/etc/gabtab`. Use the files that are created on Node B as a guide, customizing the `/etc/llttab` for Node A.
- 2 Start LLT on Node A.

```
# /etc/init.d/llt.rc start
```
- 3 Start GAB on Node A.

```
# /etc/init.d/gab.rc start
```
- 4 Check the membership of the cluster.

```
# gabconfig -a
```
- 5 Copy the cluster UUID from the existing node to the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -copy -from_sys \  
node_name_in_running_cluster -to_sys new_sys1 ... new_sysn
```

Where you are copying the cluster UUID from a node in the cluster (`node_name_in_running_cluster`) to systems from `new_sys1` through `new_sysn` that you want to join the cluster.
- 6 Start Storage Foundation on Node A.

```
# hastart
```

7 Make the Storage Foundation configuration writable.

```
# haconf -makerw
```

8 Add Node B to the cluster.

```
# hasys -add sysB
```

9 Add Node B to the system list of each service group.

■ List the service groups.

```
# hagr -list
```

■ For each service group that is listed, add the node.

```
# hagr -modify group SystemList -add sysB 1
```

Verifying configuration on both nodes

Verify the configuration for the nodes.

To verify the nodes' configuration

1 On Node B, check the cluster membership.

```
# gabconfig -a
```

2 Start the Storage Foundation on Node B.

```
# hstart
```

3 Verify that Storage Foundation is up on both nodes.

```
# hastatus
```

4 List the service groups.

```
# hagr -list
```

5 Unfreeze the service groups.

```
# hagrps -unfreeze group -persistent
```

6 Implement the new two-node configuration.

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

Uninstalling Storage Foundation

This chapter includes the following topics:

- [About removing Veritas Storage Foundation](#)
- [Preparing to uninstall a Storage Foundation product](#)
- [Disabling the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling Storage Foundation with the Veritas Web-based installer](#)
- [Uninstalling Storage Foundation filesets using the script-based installer](#)
- [Removing Storage Foundation products using SMIT](#)
- [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 a Storage Foundation product

Complete the following preparations to uninstall a Storage Foundation product.

Warning: Failure to follow the preparations that are outlined in this chapter can result in loss of data.

To remove Veritas Storage Foundation, complete the following preparations before the uninstallation:

- Back up all VxFS file systems in full and move the files in all VxFS file systems to native file systems backed with LVM logical volumes. Raw application data stored in VxVM logical volumes must be moved to LVM logical volumes.
- Remove all but one copy of file systems and databases.
- Remove all but one plex from volumes that contain multiple plexes (mirrors). To display a list of all volumes, use the command:

```
# vxprint -Ath
```

To remove a plex, use the command:

```
# vxplex -g diskgroup -o rm dis plex
```

- If a remaining plex contains multiple subdisks, consolidate the subdisks into a single subdisk using the commands:

```
# vxassist -g diskgroup mirror volume layout=contig  
# vxplex -g diskgroup -o rm dis plex
```

Sufficient space on another disk is required for this operation to complete.

- Modify `/etc/filesystems` to remove or change entries for VxFS file systems that were moved to native file systems.
- Move all data from volumes created from multiple regions of storage, including striped or spanned volumes, onto a single disk or appropriate LVM logical volume. This can be done using one of the following three methods:
 - Back up the system to tape or other media and recover the system from this.
 - Move volumes incrementally (evacuate) onto logical volumes. Evacuation moves subdisks from the source disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to LVM volumes. See [“Moving volumes to physical disks”](#) on page 291.

Moving volumes to physical disks

You can use the following steps to move data off of VxVM volumes.

To move data off of VxVM volumes

- 1 Evacuate as many disks as possible by using one of the following methods:
 - the "Remove a disk" option in `vxdiskadm`
 - the Veritas Enterprise Administrator
 - the `vxevac` script from the command line.
- 2 Remove the evacuated disks from Veritas Volume Manager control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name  
# /usr/lib/vxvm/bin/vxdiskunsetup -C disk_access_name  
# vxdisk rm disk_access_name
```

For example:

```
# vxdg -g mydg rmdisk mydg01  
# /usr/lib/vxvm/bin/vxdiskunsetup -C hdisk1  
# vxdisk rm hdisk01
```

- 3 Decide which volume to move first. If the volume to be moved is mounted, unmount it. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume has been synchronized.
- 4 On the free disk space, create an LVM logical volume that is the same size as the VxVM volume. If there is not enough free space for the logical volume, add a new disk to the system for the first volume to be removed. For subsequent volumes, you can use the free space generated by the removal of the first volume.
- 5 Copy the data on the volume onto the newly created LVM logical volume using the following command:

```
# dd if=/dev/vx/dsk/diskgroup/volume of=/dev/vgvol
```

where *diskgroup* is the name of a VxVM disk group, *volume* is the old volume in that disk group, and *vgvol* is a newly created LVM volume.

If the volume contains a VxFS file system, the user data managed by VxFS in the volume must be backed up or copied to a native AIX file system in an LVM logical volume.

6 The entries in `/etc/filesystems` for volumes holding VxFS file systems, that were copied to native file systems in step 5, must be modified according to the change in step 5.

7 Mount the disk if the corresponding volume was previously mounted.

8 Remove the volume from VxVM using the following command:

```
# vxedit -g diskgroup -rf rm volume
```

9 Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

```
# vxprint -g diskgroup -F "%snum" disk_media_name
```

10 If the return code is not 0, there are still some subdisks on this disk that must be subsequently removed. If the return code is 0, remove the disk from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name  
# vxdisk rm disk_access_name
```

11 Copy the data in the next volume to be removed to the newly created free space.

12 Reboot the system after all volumes have been converted successfully. Verify that no open volumes remain after the system reboot using the following command:

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

13 If any volumes remain open, repeat the steps listed above.

Disabling the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

To disable the agents

- 1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
# hagrps -state service_group -sys system_name
```

If none of the service groups is online, skip to 3.

- 2 If the service group is online, take it offline.

To take the service group offline without bringing it online on any other system in the cluster, enter:

```
# hagrps -offline service_group -sys system_name
```

- 3 Stop the agent on the system by entering:

```
# haagent -stop agent_name -sys system_name
```

When you get the message `Please look for messages in the log file, check the file /var/VRTSvcs/log/engine_A.log for a message confirming that each agent has stopped.`

You can also use the `ps` command to confirm that the agent is stopped.

- 4 Remove the system from the `SystemList` of the service group. If you disable the agent on all the systems in the `SystemList`, you can also remove the service groups and resource types from the VCS configuration.

Read information on administering VCS from the command line.

Refer to the *Veritas Cluster Server User's Guide*.

Removing the Replicated Data Set

This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

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 Storage Foundation with the Veritas Web-based installer

This section describes uninstalling Storage Foundation or Storage Foundation High Availability with the Veritas Web-based installer.

To uninstall Storage Foundation

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 In an HA configuration, stop VCS processes on either the local system or all systems.

To stop VCS processes on the local system:

```
# hastop -local
```

To stop VCS processes on all systems:

```
# hastop -all
```

- 3 Start the Web-based installer.
See [“Starting the Veritas Web-based installer”](#) on page 57.
- 4 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 5 Select Storage Foundation or Storage Foundation High Availability from the Product drop-down list, and click **Next**.
- 6 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Validate**.
- 7 After the validation completes successfully, click **Uninstall** to uninstall Storage Foundation on the selected system.
- 8 If there are any processes running on the target system, the installer stops the processes. Click **Next**.
- 9 After the installer stops the processes, the installer removes the products from the specified system.
Click **Next**.
- 10 After the uninstall completes, the installer displays the location of the log and summary files. If required, view the files to confirm the status of the removal.
- 11 Click **Finish**. The webinstaller prompts you for another task.

Uninstalling Storage Foundation filesets using the script-based installer

Use the following procedure to remove Storage Foundation products.

Not all filesets may be installed on your system depending on the choices that you made when you installed the software.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 35.

To shut down and remove the installed Storage Foundation filesets

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/vfstab`. Failing to remove these entries could result in system boot problems later.

- 2 Unmount all mount points for VxFS file systems.

```
# umount /mount_point
```

- 3 If the VxVM package (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

- 4 Stop the VEA Service.

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

```
# ./uninstallsf -ha
```

- 8 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall Storage Foundation, for example, `host1`:

```
Enter the system names separated by spaces from which to  
uninstall Storage Foundation: host1
```

- 9 The uninstall script prompts you to select Storage Foundation or Storage Foundation High Availability.
- 10 The uninstall script prompts you to confirm the uninstall. If you respond yes, the processes are stopped and the filesets are uninstalled.

The uninstall script creates log files and displays the location of the log files.
- 11 Most filesets have kernel components. In order to ensure complete removal, a system reboot is recommended after all filesets have been removed.

Removing Storage Foundation products using SMIT

Use the following procedure to remove Storage Foundation products using SMIT.

To remove the packages using SMIT:

- 1 Enter this command to invoke SMIT:

```
# smit
```

- 2 In SMIT, select **Software Installation and Maintenance> Software Maintenance and Utilities> Remove Installed Software**.
- 3 Under the "SOFTWARE name" menu, press F4 or Esc-4 to list all software installed on the system.
- 4 Enter "/" for Find, type "VRTS" to find all Veritas packages, and select the packages that you want to remove.

- 5 Reboot the system after removing all Storage Foundation packages.
- 6 Depending on the choices that were made when Storage Foundation was originally installed, you may find that not all of the listed Storage Foundation packages are installed on the system. You may also choose to remove the `VRTSvlic` licensing package unless this is required by other Veritas software.

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 VRTSeps package is used to remove the CP server configuration. When using the configuration utility, a configuration removal script is run and the following tasks are performed:

- All CP server configuration files are removed
- The VCS configuration for CP server is removed

After running the utility and script, you can then uninstall VCS from the node or cluster.

Note: The configuration script has to run only once per CP server (which can be on a single node or SFHA cluster), when removing the CP server configuration.

The configuration utility performs the following steps to remove the CP server configuration:

- Offlines the CP server service group (CPSSG), if it is online
- Removes the CPSSG service group from the VCS configuration

The following procedure describes how to remove the CP server configuration.

To remove the CP server configuration

- 1** To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@mycps1.symantecexample.com # /opt/VRTScps/bin/configure_cps.pl
```

- 2** The Veritas Coordination Point Server Configuration utility appears with an option menu.

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

Select one of the following:

- [1] Configure Coordination Point Server on single node VCS system
- [2] Configure Coordination Point Server on SFHA cluster
- [3] Unconfigure Coordination Point Server

- 3** Select option 3 to unconfigure the Coordination Point Server.
- 4** A warning appears and prompts you to confirm the action to unconfigure the Coordination Point Server.

Enter "y" to proceed.

```
Unconfiguring Coordination Point Server stops the vxcperv process.  
VCS clusters using this server for coordination purpose  
will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)[Default:n] :y
```

- 5 After entering "y" to proceed, messages appear informing you of the progress in removing the CP server configuration.

When the CP server configuration has been unconfigured, a success message appears.

For an example of the messages from a single node VCS cluster:

```
A single node VCS cluster is currently configured.  
Stopping the CP server ...
```

```
Removing the CP Server from VCS configuration..
```

```
Removing resource dependencies...  
Deleting the resources configured under CPSSG service group...  
Deleting the CPSSG service group...
```

```
Successfully unconfigured the Veritas Coordination Point Server.
```

For an example of the messages from a CP server on an SFHA cluster:

```
A multinode CP Server cluster is currently configured.  
Stopping the CP server ...
```

```
Removing the CP Server from VCS configuration..
```

```
Removing resource dependencies...  
Deleting the resources configured under CPSSG service group...  
Deleting the CPSSG service group...
```

```
Successfully unconfigured the Veritas Coordination Point Server.
```

- 6 You are then prompted to delete the CP server database. Enter "y" to delete the database.

For example:

```
Do you want to delete the CP Server database? (y/n) (Default:n) :
```

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

- 7 You are then prompted to delete the CP server configuration file and log files. Enter "y" to delete these files.

For example:

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files (in /var/VRTScps)? (y/n)
(Default:n) : y
```

- 8 Run the following `hagrp -state` command to ensure that the CPSSG resource has been removed from the node.

For example:

```
root@mycps1.symantecexample.com # hagrp -state CPSSG

VCS WARNING V-16-1-40131 Group CPSSG does not exist
in the local cluster
```

Removing the Storage Foundation for Databases (SFDB) repository after removing the product

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file will disable the SFDB tools.

To remove the SFDB repository

- 1 Change directories to the location of the local lookup information for the Oracle SID.

For example:

```
# cd /var/vx/vxdba/$ORACLE_SID
```

- 2 Identify the SFDB repository file and any associated links:

For example:

```
# ls -al
```

```
lrwxrwxrwx  1 oracle  oinstall      26 Jul 21 13:58 .sfdb_rept -> \  
/ora_data1/TEST/.sfdb_rept
```

```
# cd /ora_data1/TEST
```

Follow the symlink of .sfdb_rept.

- 3 Remove the repository directory containing the repository file and all backups.

For example:

```
# rm -rf .sfdb_rept
```

- 4 Remove the local lookup directory for the Oracle SID:

```
# cd /var/vx/vxdba
```

```
# rm -rf $ORACLE_SID
```

This completes the removal of the SFDB repository.

Installation scripts

This appendix includes the following topics:

- [About installation scripts](#)
- [Installation script options](#)

About installation scripts

Veritas Storage Foundation and High Availability Solutions 5.1 provides several installation scripts.

To install a fresh installation on a system, or to upgrade from Veritas Storage Foundation and High Availability Solutions version prior to 5.1, the recommended installation method is to use the common product installer. To use the common product installer, run the `installer` command.

See “[About the common product installer](#)” on page 46.

An alternative to the `installer` script is to use a product-specific installation script. If you obtained a Veritas product from an electronic download site, which does not include the common product installer, use the appropriate product installation script.

The following product installation scripts are available:

Veritas Cluster Server (VCS)	<code>installvcs</code>
Veritas Storage Foundation (SF)	<code>installsf</code>
Veritas Storage Foundation Cluster File System (SFCFS)	<code>installsfcfs</code>
Veritas Storage Foundation for Oracle RAC (SFRAC)	<code>installsfrac</code>

Symantec Product Authentication Service (AT) `installat`

Veritas Volume Manager `installvm`

To use the installation script, enter the script name at the prompt. For example, to install Veritas Storage Foundation, type `./installsf` at the prompt.

Installation script options

Table A-1 shows command line options for the product installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Veritas Storage Foundation product scripts, except where otherwise noted.

See “[About installation scripts](#)” on page 303.

Table A-1 Available command line options

Command Line Option	Function
<i>system1 system2...</i>	Specifies the systems on which to run the installation options. A system name is required for all options. If not specified, the command prompts for a system name.
<code>-addnode</code>	Adds a node to a high availability cluster.
<code>-allpkgs</code>	Displays all packages and patches required for the specified product. The packages and patches are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
<code>-configure</code>	Configures the product after installation.
<code>-fencing</code>	Configures I/O fencing in a running cluster.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-ha	<p>Specifies that the Storage Foundation High Availability software is installed or displayed. Otherwise, the list of Storage Foundation packages excludes the Veritas Cluster Server packages.</p> <p>This option only applies to the <code>installsf</code> script when one of the following options is specified:</p> <ul style="list-style-type: none"> ■ <code>-allpkgs</code> ■ <code>-recpkgs</code> ■ <code>-minpkgs</code>
-hostfile <i>full_path_to_file</i>	Specifies the location of a file that contains a list of hostnames on which to install.
-installallpkgs	Specifies that all packages are installed.
-installminpkgs	Specifies that the minimum package set is installed.
-installrecpkgs	Specifies that the required package set is installed.
-keyfile <i>ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i ssh_key_file</code> to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
-logpath <i>log_path</i>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.
-makeresponsefile	Generates a response file without performing an actual installation, configuration, or uninstallation.
-minpkgs	Displays the minimal packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-pkginfo	Displays a list of filesets and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS filesets.
-pkgpath <i>package_path</i>	Designates the path of a directory that contains all packages to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
-pkgset	Discovers and lists the 5.1 packages installed on the systems that you specify.
-pkgtable	Displays the Storage Foundation 5.1 packages in the correct installation order.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
-recpkgs	Displays the recommended packages and patches required for the specified product. The packages and patches are listed in correct installation order. Optional packages are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See <code>allpkgs</code> option.
-redirect	Displays progress details without showing the progress bar.
-responsefile <i>response_file</i>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.

Table A-1 Available command line options (*continued*)

Command Line Option	Function
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See “ Configuring secure shell (ssh) or remote shell before installing products ” on page 35.
-security	Enable or disable Symantec Product Authentication Service in a VCS cluster that is running. Install and configure Root Broker for Symantec Product Authentication Service. You can specify this option with the <code>installvcs</code> , <code>installsf</code> or <code>installsfefs</code> scripts. For more information about Symantec Product Authentication Service in a VCS cluster, see the <i>Veritas Cluster Server Installation Guide</i> .
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-tmppath <i>tmp_path</i>	Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.

Response files

This appendix includes the following topics:

- [About response files](#)
- [About the installation simulator](#)
- [Installing Storage Foundation using response files](#)
- [Configuring Storage Foundation using response files](#)
- [Upgrading Storage Foundation using response files](#)
- [Uninstalling Storage Foundation using response files](#)
- [Syntax in the response file](#)
- [Response file variable definitions](#)
- [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 response file option.

About the installation simulator

The product installer includes an option to simulate installing, configuring, or uninstalling the selected Veritas product. The simulation option steps through the installation script, including all of the preinstallation checks on the systems. However, the simulation does not actually install the packages, uninstall previously installed packages, or start or stop any processes.

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

Use the installation simulator in the following situations:

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

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

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

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

- To create a response file for your system.

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

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

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

Installing Storage Foundation using response files

Typically, you can use the response file that the installer generates after you perform Storage Foundation installation on one cluster to install Storage Foundation on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To install Storage Foundation using response files

- 1 Make sure the systems where you want to install Storage Foundation 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 Storage Foundation.
- 4 Edit the values of the response file variables as necessary.
- 5 Mount the product disc and navigate to the folder that contains the installation program.
- 6 Start the installation from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file  
# ./installsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Configuring Storage Foundation using response files

Typically, you can use the response file that the installer generates after you perform Storage Foundation configuration on one cluster to configure Storage Foundation on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure Storage Foundation using response files

- 1 Make sure the Storage Foundation filesets are installed on the systems where you want to configure Storage Foundation.
- 2 Copy the response file to one of the cluster systems where you want to configure Storage Foundation.
- 3 Edit the values of the response file variables as necessary.
To configure optional features, you must define appropriate values for all the response file variables that are related to the optional feature.
- 4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsf -responsefile /tmp/response_file
```

Where */tmp/response_file* is the response file's full path name.

Upgrading Storage Foundation using response files

Typically, you can use the response file that the installer generates after you perform Storage Foundation upgrade on one cluster to upgrade Storage Foundation on other clusters. You can also create a response file using the `-make_responsefile` option of the installer.

To perform automated Storage Foundation upgrade

- 1 Make sure the systems where you want to upgrade Storage Foundation 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 Storage Foundation.
- 4 Edit the values of the response file variables as necessary.

- 5 Mount the product disc and navigate to the folder that contains the installation program.
- 6 Start the upgrade from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file  
# ./installsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Uninstalling Storage Foundation using response files

Typically, you can use the response file that the installer generates after you perform Storage Foundation uninstallation on one cluster to uninstall Storage Foundation on other clusters.

To perform automated Storage Foundation uninstallation

- 1 Make sure that you meet the pre-requisites to uninstall Storage Foundation.
- 2 Copy the response file to one of the cluster systems where you want to uninstall Storage Foundation.
- 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/installsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Syntax in the response file

The syntax of the Perl statements that are included in the response file variables varies. It can depend on whether the variables require scalar or list values.

For example, in the case of a string value:

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value", "value", "value"];
```

Response file variable definitions

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (CSGNIC, CSGVIP, and CSGNETMASK) must be defined if any are defined. The same is true for the SMTP notification (SMTPSERVER, SMTPRECP, and SMTPRSEV), the SNMP trap notification (SNMPPORT, SNMPCONS, and SNMPCSEV), and the Global Cluster Option (CGONIC, GCOVIP, and GCONETMASK).

[Table B-1](#) lists the variables that are used in the response file and their definitions.

Table B-1 Response file variables

Variable	Description
CFG{opt}{install}	Installs Storage Foundation filesets. Configuration can be performed at a later time using the <code>-configure</code> option. List or scalar: scalar Optional or required: optional
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required
\$CFG{opt}{vxkeyless}	Installs the product with keyless license. List or scalar: scalar Optional or required: optional
CFG{systems}	List of systems on which the product is to be installed, uninstalled, or configured. List or scalar: list Optional or required: required
CFG{systemscfs}	List of systems for configuration if secure environment prevents the installer to install Storage Foundation on all systems at once. List or scalar: list Optional or required: required

Table B-1 Response file variables (*continued*)

Variable	Description
CFG{product}	<p>Defines the product to be installed, uninstalled, or configured.</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 filesets. The location must be accessible from all target systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{tmppath}	<p>Defines the location where a working directory is created to store temporary files and the depots that are needed during the install. The default location is <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>

Table B-1 Response file variables (*continued*)

Variable	Description
CFG{donotinstall} {fileset}	Instructs the installation to not install the optional filesets in the list. List or scalar: list Optional or required: optional
CFG{donotremove} {fileset}	Instructs the uninstallation to not remove the optional filesets in the list. List or scalar: list Optional or required: optional
\$CFG{sfcs_fencingenabled}	When SFCFS is configured, defines if fencing is enabled. Scalar Required 0 or 1
CFG{vcs_clustername}	Defines the name of the cluster. List or scalar: scalar Optional or required: required
CFG{vcs_clusterid}	An integer between 0 and 65535 that uniquely identifies the cluster. List or scalar: scalar Optional or required: required
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
CFG{opt}{configure}	Performs the configuration after the filesets are installed using the <code>-install</code> option. List or scalar: scalar Optional or required: optional

Table B-1 Response file variables (*continued*)

Variable	Description
CFG{vcs_lltlink#} {system}	<p>Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (LLTLINK1 and LLTLINK2). Up to four LLT links can be configured.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>
CFG{vcs_lltlinklowpri} {system}	<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>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{vcs_csgnic}	<p>Defines the NIC for Cluster Management Console to use on a system. 'ALL' can be entered as a system value if the same NIC is used on all systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{csgvip}	<p>Defines the virtual IP address that the Cluster Management Console uses.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{vcs_csgnetmask}	<p>Defines the Netmask of the virtual IP address that the Cluster Management Console uses.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{vcs_smtpserver}	<p>Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>

Table B-1 Response file variables (*continued*)

Variable	Description
CFG{vcs_smtprecip}	<p>List of full email addresses (example: user@symantecexample.com) of SMTP recipients.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{vcs_smtprsev}	<p>Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{vcs_snmpport}	<p>Defines the SNMP trap daemon port (default=162).</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{vcs_snmpcons}	<p>List of SNMP console system names</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{vcs_snmpcsev}	<p>Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names.</p> <p>List or scalar: list</p> <p>Optional or required: optional</p>
CFG{vcs_gconic} {system}	<p>Defines the NIC for the Virtual IP that the Global Cluster Option uses. 'ALL' can be entered as a system value if the same NIC is used on all systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>

Table B-1 Response file variables (*continued*)

Variable	Description
CFG{vcs_gcovip}	Defines the virtual IP address to that the Global Cluster Option uses. List or scalar: scalar Optional or required: optional
CFG{vcs_gconetmask}	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. List or scalar: scalar Optional or required: optional
CFG{vcs_userenpw}	List of encoded passwords for users List or scalar: list Optional or required: optional
CFG{vcs_username}	List of names of users List or scalar: list Optional or required: optional
\$CFG{vcs_securitymenuopt}	Specifies the menu option to choose to configure the cluster in secure mode. List or scalar: scalar <ul style="list-style-type: none"> ■ 1–Automatic ■ 2–Semi-automatic ■ 3–Manual Optional or required: optional
\$CFS{vcs_clustername}	Defines the name of the cluster. Optional or required: optional
CFG{vcs_userpriv}	List of privileges for users List or scalar: list Optional or required: optional
\$CFG{opt}{prodmode}	List of modes for product List or scalar: list Optional or required: optional

Table B-1 Response file variables (*continued*)

Variable	Description
CFG{opt}{upgrade}	Upgrades all filesets installed, without configuration. List or scalar: list Optional or required: optional
CFG{opt}{uninstall}	Uninstalls Storage Foundation filesets. List or scalar: scalar Optional or required: optional

Sample response file for SFHA configuration

The following example shows a response file for configuring Storage Foundation High Availability.

Sample response file for SFHA install

The following example shows a response file for installing Storage Foundation High Availability.

```
#####
#Auto generated sfha responsefile #
#####

our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
$CFG{opt}{prodmode}="SF Enterprise HA";
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{opt}{ha}=1;
$CFG{upi}="SF";
$CFG{prod}="SF51";
$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}{ha}=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.

Sample response file for SFHA upgrade

Configuring I/O fencing using a response file

This appendix includes the following topics:

- [Response file variables to configure disk-based I/O fencing](#)
- [Sample response file for configuring disk-based I/O fencing](#)
- [Configuring I/O fencing using response files](#)
- [Response file variables to configure server-based I/O fencing](#)

Response file variables to configure disk-based I/O fencing

[Table C-1](#) lists the response file variables that specify the required information to configure disk-based I/O fencing for Storage Foundation.

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. (Required)

Table C-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG{vxfen_config_fencing_option}	Scalar	Specifies the I/O fencing configuration mode. <ul style="list-style-type: none"> ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode (Required)
CFG {vxfen_config_fencing_mechanism}	Scalar	Specifies the I/O fencing mechanism. (Optional)
CFG{vxfen_config_fencing_dg}	Scalar	Specifies the disk group for I/O fencing. (Optional) Note: You must define either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.
CFG{vxfen_config_fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) Note: You must define either the vxfen_config_fencing_dg variable or the vxfen_config_fencing_newdg_disks variable.

Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions. See [“Response file variables to configure disk-based I/O fencing”](#) on page 323.

```
#
# Configuration Values:
#
```

```
our %CFG;

$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;

$CFG{prod}="SF51";

$CFG{systems}=[ qw(galaxy nebula) ];
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{vxfen_config_fencing_dg}="fendg";
$CFG{vxfen_config_fencing_mechanism}="dmp";
$CFG{vxfen_config_fencing_newdg_disks}=
  [ qw(rhdisk75 rhdisk76 rhdisk77) ];
$CFG{vxfen_config_fencing_option}=2;
```

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for Storage Foundation HA. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure I/O fencing using response files

- 1 Make sure that Storage Foundation HA 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 87.
- 3 Copy the response file to one of the cluster systems where you want to configure I/O fencing.
See [“Sample response file for configuring disk-based I/O fencing”](#) on page 324.
See [“Sample response file for configuring server-based I/O fencing”](#) on page 328.

- 4 Edit the values of the response file variables as necessary.
See “Response file variables to configure disk-based I/O fencing” on page 323.
See “Response file variables to configure server-based I/O fencing” on page 326.
- 5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsf -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file’s full path name.

Response file variables to configure server-based I/O fencing

You can use a CP server response file to configure server-based customized I/O fencing. The installer uses the CP server response file for the following types of I/O fencing configurations:

- Client cluster fencing (server-based I/O fencing configuration itself)
The installer configures server-based customized I/O fencing on the 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>Note: 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 <code>"\$CFG{fencing_cpc_reusedg}=0"</code> or <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. 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 Entering a "1" indicates that security is enabled. Entering a "0" indicates that security has not been enabled.

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing :

Storage Foundation and High Availability components

This appendix includes the following topics:

- [Veritas Storage Foundation installation filesets](#)
- [Veritas Cluster Server installation filesets](#)
- [Veritas Storage Foundation obsolete and reorganized installation filesets](#)

Veritas Storage Foundation installation filesets

[Table D-1](#) shows the fileset name and contents for each English language fileset for Veritas Storage Foundation. The table also gives you guidelines for which filesets to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and Veritas Cluster Server (VCS) filesets, the combined functionality is called Storage Foundation and High Availability.

See [“Veritas Cluster Server installation filesets”](#) on page 331.

Table D-1 Veritas Storage Foundation filesets

fileset	Contents	Configuration
VRTSaslapm	Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries Required for the support and compatibility of various storage arrays.	Minimum
VRTSat	Symantec Product Authentication Service Installs the Symantec Product Authentication Service, which provides authentication services to other Symantec products. This package contains a server and client component. The server provides services for a root broker, authentication broker, or both. The client allows Symantec products to communicate with the brokers. Required to use Symantec Product Authentication Service.	All
VRTSperl	Perl 5.10.0 for Veritas	Minimum
VRTSveki	Veritas Kernel Interface Contains a common set of modules that other Veritas drivers use.	Minimum
VRTSvlic	Veritas License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Minimum
VRTSvxfs	Veritas File System binaries Required for VxFS file system support.	Minimum
VRTSvxvm	Veritas Volume Manager binaries	Minimum
VRTSdbed	Veritas Storage Foundation for Oracle	Recommended
VRTSob	Veritas Enterprise Administrator	Recommended

Table D-1 Veritas Storage Foundation filesets (*continued*)

fileset	Contents	Configuration
VRTSodm	ODM Driver for VxFS 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 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: http://www.symantec.com/business/storage-foundation-manager	Recommended
VRTSspt	Veritas Software Support Tools	Recommended
VRTSfssdk	Veritas File System Software Developer Kit For VxFS APIs, the package contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.	All

Veritas Cluster Server installation filesets

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

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

See “[Veritas Storage Foundation installation filesets](#)” on page 329.

Table D-2 VCS installation filesets

fileset	Contents	Configuration
VRTSgab	Veritas Cluster Server group membership and atomic broadcast services	Minimum
VRTSllt	Veritas Cluster Server low-latency transport	Minimum
VRTSvc	Veritas Cluster Server	Minimum
VRTSvcsg	Veritas Cluster Server Bundled Agents	Minimum
VRTSvxfen	Veritas I/O Fencing	Minimum
VRTScutil	Veritas Cluster Server Utilities	Recommended
VRTSvcsea	Consolidated database and enterprise agent packages	Recommended
VRTScps	Veritas Coordination Point Server The Coordination Point Server is an alternate mechanism for I/O fencing. It implements I/O fencing through a client/server architecture and can provide I/O fencing for multiple VCS clusters.	Advanced

Veritas Storage Foundation obsolete and reorganized installation filesets

[Table D-3](#) lists the filesets that are obsolete or reorganized for Storage Foundation and Storage Foundation High Availability.

Table D-3 Veritas Storage Foundation obsolete and reorganized filesets

fileset	Description
Infrastructure	
SYMClma	Obsolete
VRTSaa	Included in VRTSsmfh
VRTSccg	Included in VRTSsmfh

Table D-3 Veritas Storage Foundation obsolete and reorganized filesets
(continued)

fileset	Description
VRTSdbms3	Obsolete
VRTSicsco	Obsolete
VRTSjre	Obsolete
VRTSjre15	Obsolete
VRTSmh	Included in VRTSsmfh
VRTSobc33	Obsolete
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product packages	
VRTSacclib	Obsolete
VRTSalloc	Obsolete
VRTScmccc	Obsolete
VRTScmcs	Obsolete
VRTScscm	Included in VRTScutil
VRTScscw	Included in VRTScutil
VRTScsocw	Included in VRTScutil
VRTScssim	Included in VRTScutil
VRTSd2gui	Included in VRTSdbed
VRTSdb2ed	Included in VRTSdbed
VRTSdbcom	Included in VRTSdbed
VRTSdbed	Included in VRTSdbed
VRTSdcli	Obsolete

Table D-3 Veritas Storage Foundation obsolete and reorganized filesets
(continued)

fileset	Description
VRTSddlpr	Obsolete
VRTSdsa	Obsolete
VRTSfsman	Included in mainpkg
VRTSfsmnd	Included in mainpkg
VRTSfspro	Included in VRTSsmfh
VRTSgapms	Obsolete
VRTSmapro	Included in VRTSsmfh
VRTSorgui	Obsolete
VRTSvail	Obsolete
VRTSvcscdb	Included in VRTSvcsea
VRTSvcscor	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvcsc
VRTSvdid	Obsolete
VRTSvmman	Included in mainpkg
VRTSvmpro	Included in VRTSsmfh
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 an installation on AIX](#)
- [Incorrect permissions for root on remote system](#)
- [Resource temporarily unavailable](#)
- [Inaccessible system](#)
- [Storage Foundation Cluster File System problems](#)

Restarting the installer after a failed connection

If an installation is killed because of a failed connection, you can restart the installer to resume the installation. The installer detects the existing installation. The installer prompts you whether you want to resume the installation. If you resume the installation, the installation proceeds from the point where the installation failed.

What to do if you see a licensing reminder

In this release, you can install without a license key. In order to comply with the End User License Agreement, you must either install a license key or make the host managed by a Management Server. If you do not comply with these terms within 60 days, the following warning messages result:

```
WARNING V-365-1-1 This host is not entitled to run Veritas Storage
Foundation/Veritas Cluster Server.As set forth in the End User
License Agreement (EULA) you must complete one of the two options
set forth below. To comply with this condition of the EULA and
stop logging of this message, you have <nn> days to either:
- make this host managed by a Management Server (see
  http://go.symantec.com/sfhakeyless for details and free download),
  or
- add a valid license key matching the functionality in use on this host
  using the command 'vxlicinst' and validate using the command
  'vxkeyless set NONE'
```

To comply with the terms of the EULA, and remove these messages, you must do one of the following within 60 days:

- Install a valid license key corresponding to the functionality in use on the host. See “[Installing Veritas product license keys](#)” on page 130.

After you install the license key, you must validate the license key using the following command:

```
# vxkeyless set NONE
```

- Continue with keyless licensing by managing the server or cluster with a management server.

For more information about keyless licensing, see the following URL: <http://go.symantec.com/sfhakeyless>

Troubleshooting an installation on AIX

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

See “[Preparing to uninstall a Storage Foundation product](#)” on page 290.

Then reinstall.

See “[About the common product installer](#)” on page 46.

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Checking system communication ..... Done
System verification did not complete successfully
The following errors were discovered on the systems:
The ssh permission denied on host1 rsh permission denied on host1
either ssh or rsh is needed to be setup between the local node and host1
for communication. Enter the AIX system names separated by
spaces: q,? (host1)
```

Suggested solution: You need to set up the systems to allow remote access using ssh or rsh.

See [“Configuring secure shell \(ssh\) or remote shell before installing products”](#) on page 35.

Note: Remove remote shell permissions after completing the Storage Foundation installation and configuration.

Resource temporarily unavailable

If the installation fails with the following error message on the console:

```
fork() failed: Resource temporarily unavailable
```

The value of maximum number of processes allowed per user may not be large enough. This kernel attribute is a tunable and can be changed on any node of the cluster.

To determine the current value of "Maximum number of PROCESSES allowed per user", enter:

```
# lsattr -H -E -l sys0 -a maxuproc
```

To see the default value of this tunable and its valid range of values, enter:

```
# odmget -q "attribute=maxuproc" PdAt
```

If necessary, you can change the value of the tunable using the smitty interface:

```
# smitty chgsys
```

You can also directly change the CuAt class using the following command:

```
# chdev -l sys0 -a maxuproc=600
```

Increasing the value of the parameter takes effect immediately; otherwise the change takes effect after a reboot.

See the `smitty` and `chdev` manual pages.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
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 AIX 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.

Storage Foundation Cluster File System problems

If there is a device failure or controller failure to a device, the file system may become disabled cluster-wide. To address the problem, unmount file system on all the nodes, then run a full `fsck`. When the file system check completes, mount all nodes again.

High availability issues

This section describes high availability issues.

Network partition/jeopardy

Network partition (or split brain) is a condition where a network failure can be misinterpreted as a failure of one or more nodes in a cluster. If one system in the cluster incorrectly assumes that another system failed, it may restart applications already running on the other system, thereby corrupting data. CFS tries to prevent this by having redundant heartbeat links.

At least one link must be active to maintain the integrity of the cluster. If all the links go down, after the last network link is broken, the node can no longer communicate with other nodes in the cluster. Thus the cluster is in one of two possible states. Either the last network link is broken (called a network partition condition), or the last network link is okay, but the node crashed, in which case

it is not a network partition problem. It is not possible to identify whether it is the first or second state, so a kernel message is issued to indicate that a network partition may exist and there is a possibility of data corruption.

Jeopardy is a condition where a node in the cluster has a problem connecting to other nodes. In this situation, the link or disk heartbeat may be down, so a jeopardy warning may be displayed. Specifically, this message appears when a node has only one remaining link to the cluster and that link is a network link. This is considered a critical event because the node may lose its only remaining connection to the network.

Warning: Do not remove the communication links while shared storage is still connected.

Low memory

Under heavy loads, software that manages heartbeat communication links may not be able to allocate kernel memory. If this occurs, a node halts to avoid any chance of network partitioning. Reduce the load on the node if this happens frequently.

A similar situation may occur if the values in the `/etc/llttab` files on all cluster nodes are not correct or identical.

Troubleshooting cluster installation

This appendix includes the following topics:

- [Unmount failures](#)
- [Command failures](#)
- [Installer cannot create UUID for the cluster](#)
- [The vxfsenthdw utility fails when SCSI TEST UNIT READY command fails](#)
- [Troubleshooting on the CP server](#)
- [Troubleshooting server-based I/O fencing on the SF HA cluster](#)
- [Troubleshooting server-based I/O fencing in mixed mode](#)

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 43.

- 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 Storage Foundation configuration, upgrade, or when you add a node to the cluster using the installer.

Workaround: To start Storage Foundation, you must run the `uuidconfig.pl` script manually to configure the UUID on each cluster node.

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

The `vxfcntlshdw` utility fails when `SCSI TEST UNIT READY` command fails

While running the `vxfcntlshdw` utility, you may see a message that resembles as follows:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node  
FAILED.  
Contact the storage provider to have the hardware configuration  
fixed.
```

The disk array does not support returning success for a `SCSI TEST UNIT READY` command when another host has the disk reserved using SCSI-3 persistent reservations. This happens with the Hitachi Data Systems 99XX arrays if bit 186 of the system mode option is not enabled.

Troubleshooting on the CP server

All the CP server operations and messages are logged in the `/var/VRTScps/log` directory in a detailed and easy to read format. The entries are sorted by date and time. The logs can be used for troubleshooting purposes or to review for any possible security issue on the single node VCS or SFHA cluster hosting the CP server.

The following files contain logs and text files that may be useful in understanding and troubleshooting a CP server:

- `/var/VRTScps/log/cpsrvr_[ABC].log`
- `/var/VRTSat/vrtsat_broker.txt` (Security related)

If the `vxcperv` process fails on the CP server, then review the following diagnostic files:

- `/var/VRTScps/diag/FFDC_CPS_<pid>_vxcperv.log`
- `/var/VRTScps/diag/stack_<pid>_vxcperv.txt`

Note: If the `vxcperv` process fails on the CP server, these files are present in addition to a core file. VCS restarts `vxcperv` process automatically in such situations.

CP server service group issues

If you cannot bring up the CPSSG service group after the CP server configuration, verify that the CPSSG service group and its resources are valid and properly configured in the VCS configuration.

Check the VCS engine log to see if any of the CPSSG service group resources are FAULTED. The engine log is located in the following directory:

`/var/VRTSvcs/log/engine_[ABC].log`

The resources that are configured under the CPSSG service groups are displayed in the following figures:

- CPSSG group and dependency figure for CP server hosted on a single node VCS cluster:
- CPSSG group and dependency figure for CP server hosted on an SFHA cluster:

Note: For information about general VCS troubleshooting procedures, refer to the Veritas™ Cluster Server User's Guide, Version 5.1.

Testing the connectivity of the CP server

The connectivity of the CP server can be tested using the `cpsadm` command. The following `cpsadm` command tests whether a CP server is up and running at a process level:

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

where `cp_server` is the virtual IP address or virtual hostname on which the CP server is listening.

Issuing the command on the SF HA cluster nodes requires the environment variables `CPS_USERNAME` and `CPS_DOMAINTYPE` to be set.

Troubleshooting server-based I/O fencing on the SF HA cluster

The file `/var/VRTSvc/log/vxfen/vxfend_[ABC].log` contains logs and text files that may be useful in understanding and/or troubleshooting fencing-related issues on a SF HA cluster node.

Issues during server-based fencing start up on SF HA cluster node

The following issues may occur during fencing start up on the SF HA cluster node:

- `cpsadm` command on the SF HA cluster gives connection error
- Authentication failure
- Authorization failure
- Preexisting split-brain

cpsadm command on the SF HA cluster node gives connection error

If you receive a connection error message after issuing the `cpsadm` command on the SF HA cluster, perform the following actions:

- Ensure that the CP server is reachable from all the SF HA cluster nodes.
- Check that the correct CP server virtual IP/virtual hostname and port number are being used by the SF HA cluster nodes.
Check the `/etc/vxfenmode` file.
- Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.

Authorization failure

Authorization failure occurs when the CP server's SF HA cluster nodes or users are not added in the CP server configuration. Therefore, fencing on the SF HA cluster node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points. To resolve this issue, add the SF HA cluster node and user in the CP server configuration and restart fencing. Refer to the following section:

See “[Preparing the CP servers manually for use by the SF HA cluster](#)” on page 155.

Preexisting split-brain

To illustrate preexisting split-brain, assume there are three CP servers acting as coordination points. One of the three CP servers then becomes inaccessible. While in this state, also one client node leaves the cluster. When the inaccessible CP server restarts, it has a stale registration from the node which left the 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. The situation is similar to that of preexisting split-brain, with coordinator disks, where the problem is solved by the administrator running the `vxfcntlclearpre` command. A similar solution is required using the `cpsadm` command.

The following `cpsadm` command can be used to clear a registration on a CP server:

```
# cpsadm -s cp_server -a unreg_node -c cluster_name -n nodeid
```

where `cp_server` is the virtual IP address or virtual hostname on which the CP server is listening, `cluster_name` is the VCS name for the SF HA cluster, and `nodeid` specifies the node id of SF HA cluster node.

After removing all stale registrations, the joiner node will be able to join the cluster.

Issues during online migration of coordination points

During online migration of coordination points using the `vxfwswap` 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).
- The cluster or 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

The following procedure can be use to troubleshoot a mixed I/O fencing configuration (configuration using both coordinator disks and CP server for I/O fencing). This procedure involves using the following commands to obtain I/O fencing information:

- To obtain I/O fencing cluster information on the coordinator disks, run the following command on one of the cluster nodes:

```
# vxfenadm -s diskname
```

Any keys other than the valid keys used by the cluster nodes that appear in the command output are spurious keys.

- To obtain I/O fencing cluster information on the CP server, run the following command on one of the cluster nodes:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp_server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster_name* is the VCS name for the 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.

Troubleshooting mixed I/O fencing configuration (coordinator disks and CP server)

- 1 Review the current I/O fencing configuration by accessing and viewing the information in the `vxfenmode` file.

Enter the following command on one of the 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.

Enter the `vxfenadm -s` command on each of the SF HA cluster nodes.

```
# vxfenadm -s /dev/vx/rdmp/3pardata0_190
```

```
Device Name: /dev/vx/rdmp/3pardata0_190
Total Number Of Keys: 2
key[0]:
    [Numeric Format]: 86,70,66,69,65,68,48,48
    [Character Format]: VFBEAD00
    [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
    [Numeric Format]: 86,70,66,69,65,68,48,49
    [Character Format]: VFBEAD01
*    [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

```
# vxfenadm -s /dev/vx/rdmp/3pardata0_191
```

```
Device Name: /dev/vx/rdmp/3pardata0_191
Total Number Of Keys: 2
key[0]:
    [Numeric Format]: 86,70,66,69,65,68,48,48
    [Character Format]: VFBEAD00
    [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
    [Numeric Format]: 86,70,66,69,65,68,48,49
    [Character Format]: VFBEAD01
*    [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

4 Review the CP server information about the cluster nodes.

On the CPS server, run the `cpsadm list nodes` command to review a list of nodes in the cluster.

The command syntax is as follows:

```
# cpsadm -s cp_server -a list_nodes
```

where *cp server* is the virtual IP address or virtual hostname on which the CP server is listening.

For example:

```
# /opt/VRTS/bin/cpsadm -s 10.140.94.101 -a list_nodes
```

ClusName	UUID	Hostname(Node ID)	Registered
gl-rh2	{25aeb8c6-1dd2-11b2-95b5-a82227078d73}	node_101(0)	0
gl-rh2	{25aeb8c6-1dd2-11b2-95b5-a82227078d73}	node_102(1)	0
cpstest	{a0cf10e8-1dd1-11b2-87dc-080020c8fa36}	node_220(0)	0
cpstest	{a0cf10e8-1dd1-11b2-87dc-080020c8fa36}	node_240(1)	0
ictwo	{f766448a-1dd1-11b2-be46-5d1da09d0bb6}	node_330(0)	0
ictwo	{f766448a-1dd1-11b2-be46-5d1da09d0bb6}	sassette(1)	0
fencing	{e5288862-1dd1-11b2-bc59-0021281194de}	CDC-SFLAB-CD-01(0)	0
fencing	{e5288862-1dd1-11b2-bc59-0021281194de}	CDC-SFLAB-CD-02(1)	0
gl-su2	{8f0a63f4-1dd2-11b2-8258-d1bcc1356043}	gl-win03(0)	0
gl-su2	{8f0a63f4-1dd2-11b2-8258-d1bcc1356043}	gl-win04(1)	0
gl-su1	{2d2d172e-1dd2-11b2-bc31-045b4f6a9562}	gl-win01(0)	0
gl-su1	{2d2d172e-1dd2-11b2-bc31-045b4f6a9562}	gl-win02(1)	0
gl-ax4	{c17cf9fa-1dd1-11b2-a6f5-6dbd1c4b5676}	gl-ax06(0)	0
gl-ax4	{c17cf9fa-1dd1-11b2-a6f5-6dbd1c4b5676}	gl-ax07(1)	0
gl-ss2	{da2be862-1dd1-11b2-9fb9-0003bac43ced}	galaxy(0)	1
gl-ss2	{da2be862-1dd1-11b2-9fb9-0003bac43ced}	nebula(1)	1

5 Review the CP server list membership.

On the CP server, run the following command to review the list membership. The command syntax is as follows:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp_server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster_name* is the VCS name for the 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.

Sample SF HA cluster setup diagrams for CP server-based I/O fencing

This appendix includes the following topics:

- [Configuration diagrams for setting up server-based I/O fencing](#)

Configuration diagrams for setting up server-based I/O fencing

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:

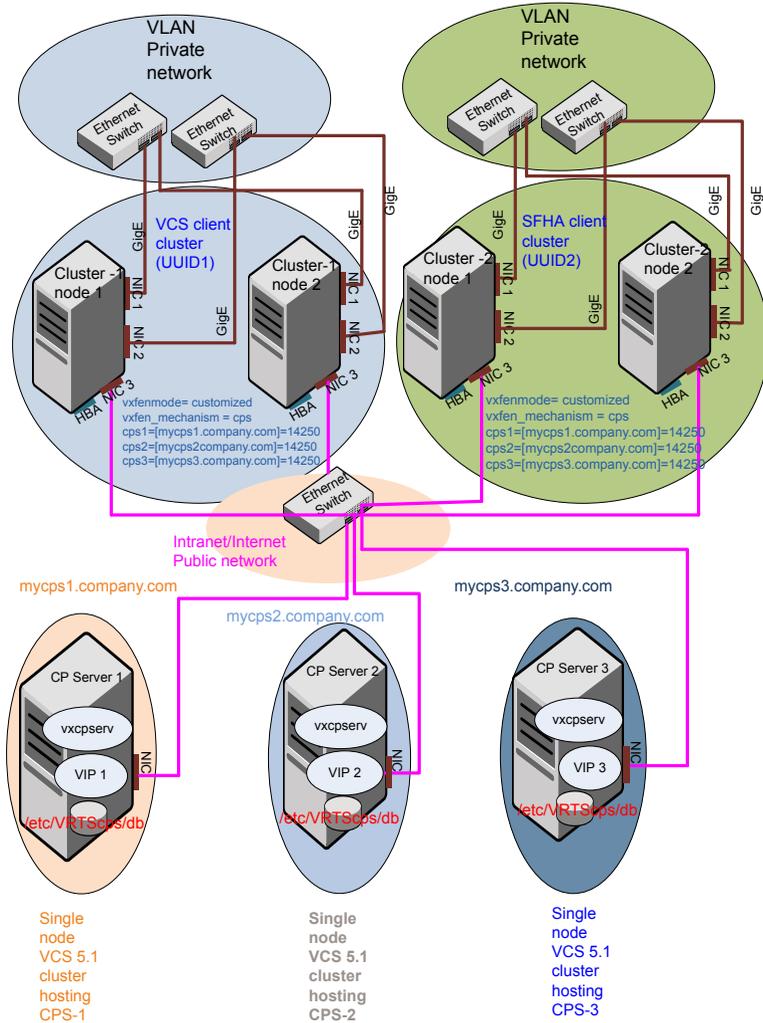
- Two unique client clusters that are served by 3 CP servers:
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
See [Figure G-2](#) on page 355.
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
See [Figure G-3](#) on page 357.
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:

Two unique client clusters served by 3 CP servers

[\[Unresolved xref\]](#) displays a configuration where two unique client clusters are being served by 3 CP servers (coordination points). Each client cluster has its own unique user ID (UUID1 and UUID2).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

Figure G-1 Two unique client clusters served by 3 CP servers



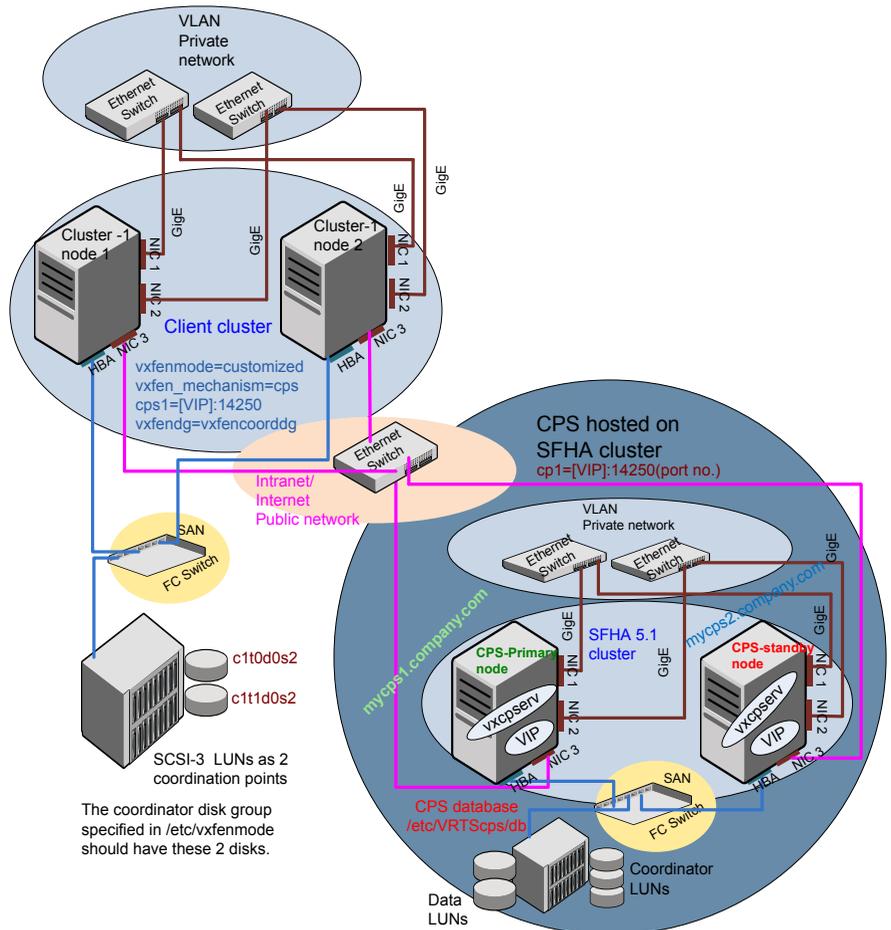
Client cluster served by highly available CPS and 2 SCSI-3 disks

Figure G-2 displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to customized with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure G-2 Client cluster served by highly available CP server and 2 SCSI-3 disks



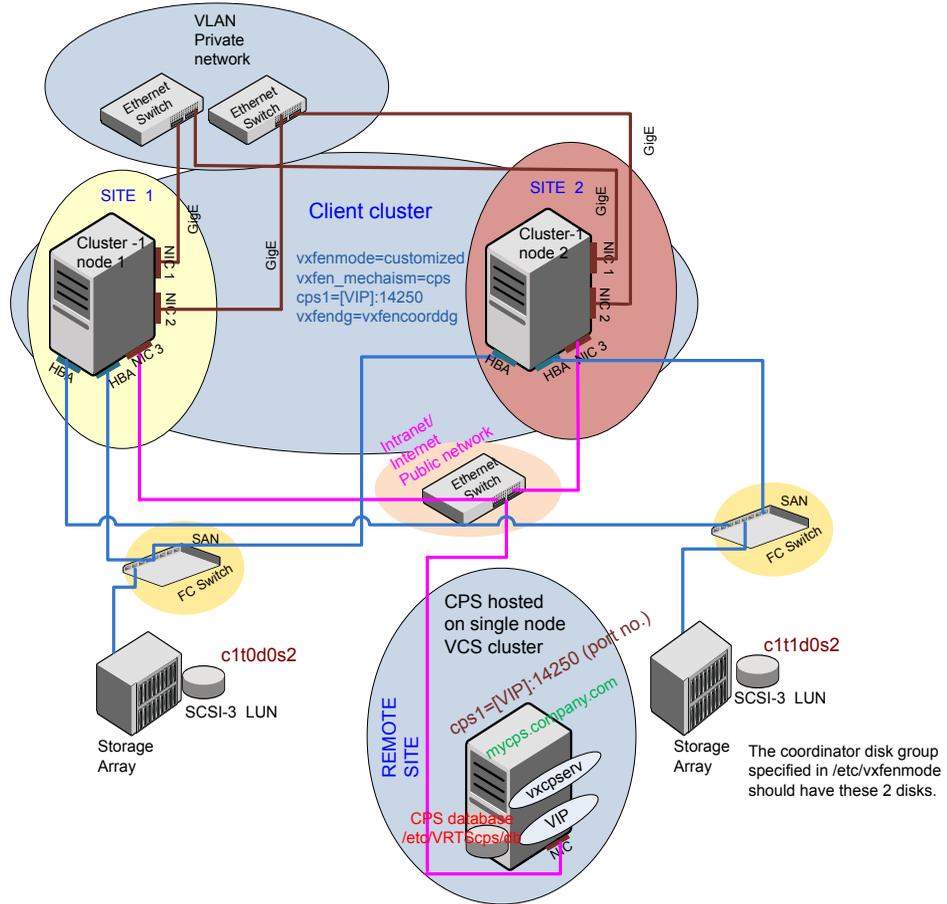
Two node campus cluster served by remote CP server and 2 SCSI-3 disks

[Figure G-3](#) displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoorddg`. The third coordination point is a CP server on a single node VCS cluster.

Figure G-3 Two node campus cluster served by remote CP server and 2 SCSI-3



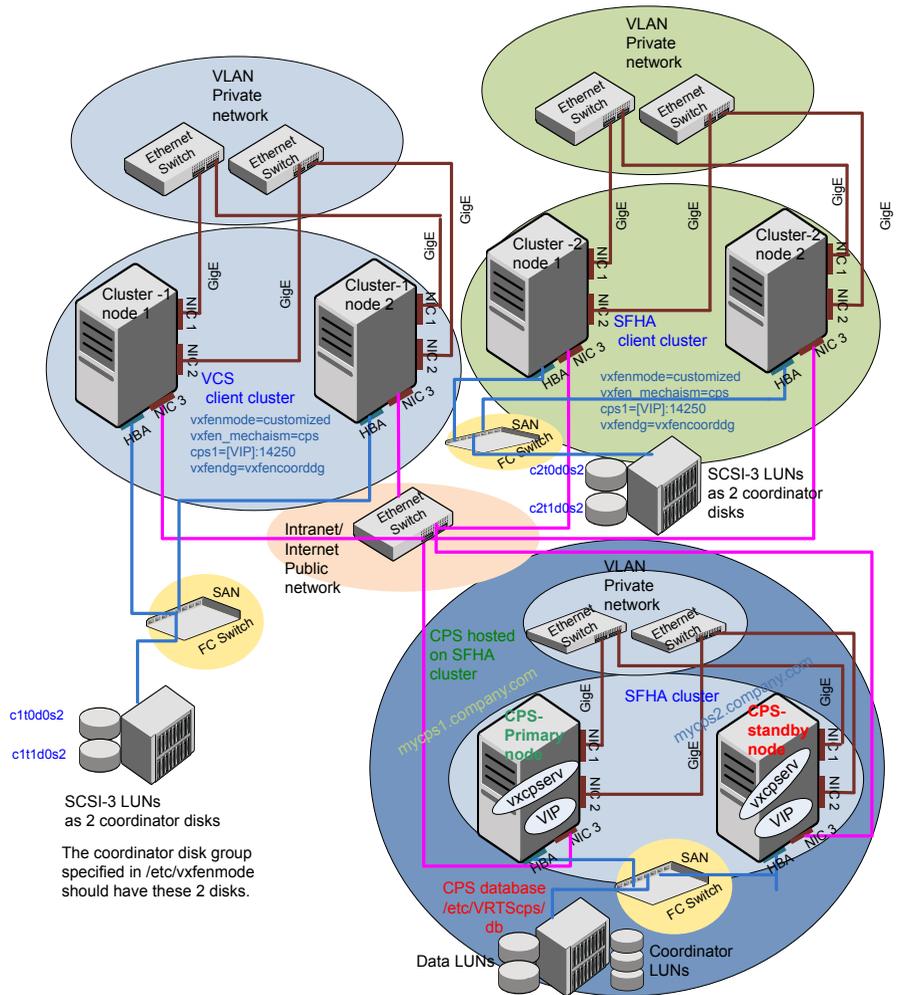
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

[Unresolved xref] displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNS (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure G-4 Multiple client clusters served by highly available CP server and 2 SCSI-3 disks



Changing NFS server major numbers for VxVM volumes

This appendix includes the following topics:

- [Changing NFS server major numbers for VxVM volumes](#)

Changing NFS server major numbers for VxVM volumes

Use the `haremajor` command to determine and reassign the major number that a system uses for shared VxVM volume block devices. For Veritas Volume Manager, the major number is set to the `vxio` driver number. To be highly available, each NFS server in a VCS cluster must have the same `vxio` driver number, or major number.

To list the major number currently in use on a system

- ◆ Use the command:

```
# haremajor -v
55
```

Run this command on each cluster node. If major numbers are not the same on each node, you must change them on the nodes so that they are identical.

To list the available major numbers for a system

- ◆ Use the command:

```
# haremajor -a
54, 56..58, 60, 62..
```

The output shows the numbers that are not in use on the system where the command is issued.

To reset the major number on a system

- ◆ You can reset the major number to an available number on a system. For example, to set the major number to 75 type:

```
# haremajor -s 75
```

Configuring LLT over UDP using IPv4

This appendix includes the following topics:

- [Using the UDP layer for LLT](#)
- [Configuring LLT over UDP](#)

Using the UDP layer for LLT

Veritas Storage Foundation 5.1 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

LLT over UDP is slower than LLT over Ethernet. Use LLT over UDP only when the hardware configuration makes it necessary.

Configuring LLT over UDP

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on different physical networks.

If the LLT private links are not on different physical networks, then make sure that the links are on separate subnets. Set the broadcast address in `/etc/llttab` explicitly depending on the subnet for each link.

See [“Broadcast address in the `/etc/llttab` file”](#) on page 362.

- 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 364.
- Set the broadcast address correctly for direct-attached (non-routed) links.
See [“Sample configuration: direct-attached links”](#) on page 366.
- 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 367.

Broadcast address in the `/etc/llttab` file

The broadcast address is set explicitly for each link in the following example.

- Display the content of the `/etc/llttab` file on the first node galaxy:

```
galaxy # cat /etc/llttab

set-node galaxy
set-cluster 1
link link1 /dev/xti/udp - udp 50000 - 192.168.9.1 192.168.9.255
link link2 /dev/xti/udp - udp 50001 - 192.168.10.1 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

- Display the content of the `/etc/llttab` file on the second node nebula:

```
nebula # cat /etc/llttab

set-node nebula
set-cluster 1
link link1 /dev/xti/udp - udp 50000 - 192.168.9.2 192.168.9.255
link link2 /dev/xti/udp - udp 50001 - 192.168.10.2 192.168.10.255
```

Verify the subnet mask using the `ifconfig` command to ensure that the two links are on separate subnets.

The link command in the /etc/llttab file

Review the link command information in this section for the /etc/llttab file. See the following information for sample configurations:

- See “[Sample configuration: direct-attached links](#)” on page 366.
- See “[Sample configuration: links crossing IP routers](#)” on page 367.

[Table I-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 I-1 Field description for link command in /etc/llttab

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example /dev/xti/udp.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “ Selecting UDP ports ” on page 364.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IP address</i>	IP address of the link on the local node.
<i>bcast-address</i>	<ul style="list-style-type: none"> ■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the /etc/llttab file

The `set-addr` command in the /etc/llttab file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 367.

[Table I-2](#) describes the fields of the set-addr command.

Table I-2 Field description for set-addr command in /etc/llttab

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.
<i>link tag-name</i>	The string that LLT uses to identify the link; for example link1, link2,....
<i>address</i>	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -a | more
UDP
  Local Address          Remote Address         State
  -----
  *.*                   *.*                   Unbound
  *.32771                *.*                   Idle
  *.32776                *.*                   Idle
  *.32777                *.*                   Idle
  *.name                 *.*                   Idle
  *.biff                 *.*                   Idle
  *.talk                 *.*                   Idle
  *.32779                *.*                   Idle
  .
  .
  .
  *.55098                *.*                   Idle
  *.syslog                *.*                   Idle
  *.58702                *.*                   Idle
  *.*                    *.*                   Unbound
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# set_parms ip_address
```

For example:

- For the first network interface on the node galaxy:

```
IP address=192.168.9.1, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

For the first network interface on the node nebula:

```
IP address=192.168.9.2, Broadcast address=192.168.9.255,  
Netmask=255.255.255.0
```

- For the second network interface on the node galaxy:

```
IP address=192.168.10.1, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

For the second network interface on the node nebula:

```
IP address=192.168.10.2, Broadcast address=192.168.10.255,  
Netmask=255.255.255.0
```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in `/etc/llttab` depending on the subnet that the links are on.

An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

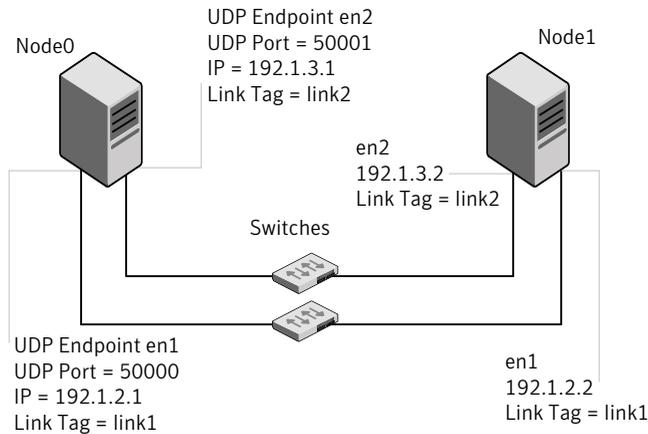
```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100
```

```
link link1 /dev/xti/udp - udp 50000 - 192.168.30.1  
192.168.30.255  
link link2 /dev/xti/udp - udp 50001 - 192.168.31.1  
192.168.31.255
```

Sample configuration: direct-attached links

Figure I-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure I-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0  
set-cluster 1  
#configure Links  
#link tag-name device node-range link-type udp port MTU \  
IP-address boast-address  
link link1 /dev/xti/udp - udp 50000 - 192.1.2.1 192.1.2.255  
link link2 /dev/xti/udp - udp 50001 - 192.1.3.1 192.1.3.255
```

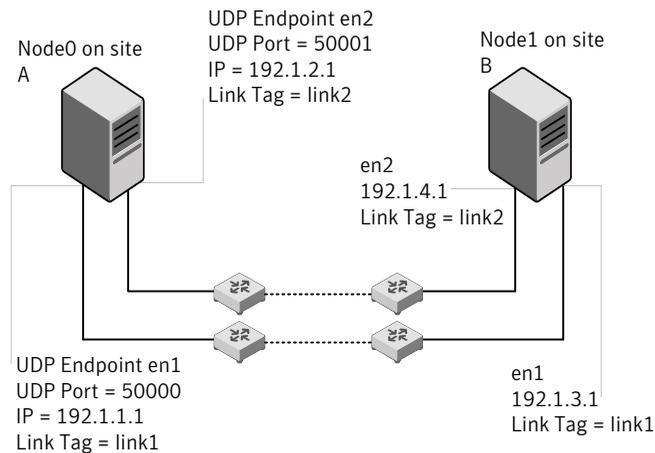
The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
# configure Links
# link tag-name device node-range link-type udp port MTU \
IP-address broadcast-address
link link1 /dev/xti/udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 /dev/xti/udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

[Figure I-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 I-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1
link link1 /dev/xti/udp - udp 50000 - 192.1.3.1 -
link link2 /dev/xti/udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
```

```
set-addr      0 link1 192.1.1.1
set-addr      0 link2 192.1.2.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3
```

```
#disable LLT broadcasts
set-bcasthb   0
set-arp       0
```

The /etc/llttab file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 /dev/xti/udp - udp 50000 - 192.1.1.1 -
link link2 /dev/xti/udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr      1 link1 192.1.3.1
set-addr      1 link2 192.1.4.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb   0
set-arp       0
```

Configuring LLT over UDP using IPv6

This appendix includes the following topics:

- [Using the UDP layer of IPv6 for LLT](#)
- [Configuring LLT over UDP using IPv6](#)
- [Configuring LLT over UDP using IPv6](#)

Using the UDP layer of IPv6 for LLT

Veritas Storage Foundation 5.1 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

Configuring LLT over UDP using IPv6

The following checklist is to configure LLT over UDP:

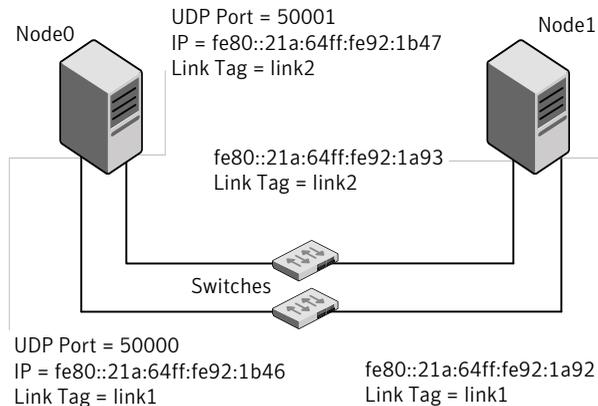
- For UDP6, the multicast address is set to "-".

- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.
- Make sure the IPv6 addresses in the `/etc/llttab` files are consistent with the IPv6 addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
See “[Selecting UDP ports](#)” on page 374.
- 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 371.

Sample configuration: direct-attached links

Figure J-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure J-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```

set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

```

The file for Node 1 resembles:

```

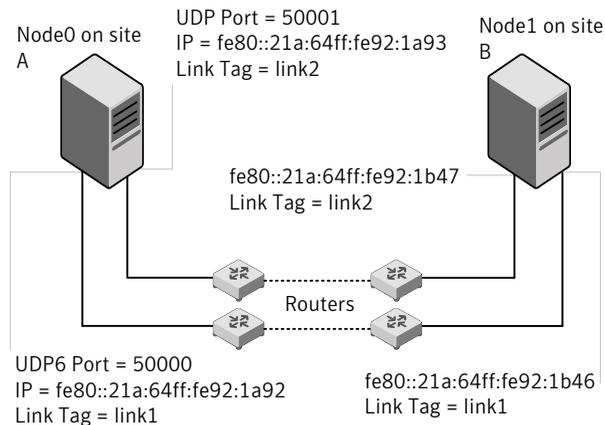
set-node Node1
set-cluster 1
# configure Links
# link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

```

Sample configuration: links crossing IP routers

Figure J-2 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure J-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 0 link1 fe80::21a:64ff:fe92:1b46
set-addr 0 link2 fe80::21a:64ff:fe92:1b47
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

The `/etc/llttab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 1 link1 fe80::21a:64ff:fe92:1a92
set-addr 1 link2 fe80::21a:64ff:fe92:1a93
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

Configuring LLT over UDP using IPv6

The following checklist is to configure LLT over UDP:

- For UDP6, the multicast address is set to "-".
- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.
- Make sure the IPv6 addresses in the `/etc/llttab` files are consistent with the IPv6 addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
See “[Selecting UDP ports](#)” on page 374.
- 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 371.

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 370.
- See “[Sample configuration: links crossing IP routers](#)” on page 371.

Note that some of the fields in [Table J-1](#) differ from the command for standard LLT links.

[Table J-1](#) describes the fields of the link command that are shown in the `/etc/llttab` file examples.

Table J-1 Field description for link command in `/etc/llttab`

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example <code>/dev/xti/udp6</code> .
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp6" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “ Selecting UDP ports ” on page 374.
<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.

Table J-1 Field description for link command in `/etc/llttab` (*continued*)

Field	Description
<i>IPv6 address</i>	IPv6 address of the link on the local node.
<i>mcast-address</i>	"-" is the default for clusters spanning routers.

The set-addr command in the `/etc/llttab` file

The `set-addr` command in the `/etc/llttab` file is required when the multicast feature of LLT is disabled, such as when LLT must cross IP routers.

See [“Sample configuration: links crossing IP routers”](#) on page 371.

[Table J-2](#) describes the fields of the `set-addr` command.

Table J-2 Field description for set-addr command in `/etc/llttab`

Field	Description
<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>	IPv6 address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file `/etc/services`. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -a | more
Active Internet connections (including servers)
Proto Recv-Q Send-Q Local Address Foreign Address (state)
tcp          0      0 *.32778      *.*          LISTEN
```

```
tcp      0      0 *.32781          *.*             LISTEN
udp4     0      0 *.daytime        *.*
udp4     0      0 *.time           *.*
udp4     0      0 *.sunrpc         *.*
udp      0      0 *.snmp           *.*
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the `/etc/services` file, its associated name is displayed rather than the port number in the output.

Index

A

- adding
 - users 117
- adding node
 - to a one-node cluster 282
- agents
 - disabling 292
- applications, stopping 176
- attributes
 - UseFence 142

C

- cables
 - cross-over Ethernet 264
- cabling shared devices 42
- CFS
 - troubleshooting 338
- cluster
 - removing a node from 275
 - verifying operation 255
- command failures 342
- commands
 - hastart 275
 - hastatus 255
 - hasys 256
 - lltconfig 249
 - lltstat 252
 - vxdisksetup (initializing disks) 131
- configuration
 - restoring the original 189
- configuring VCS
 - adding users 117
 - event notification 117, 119
 - global clusters 121
 - secure mode 113
 - starting 107
- coordinator disks
 - DMP devices 83
 - for I/O fencing 83
 - setting up 140
- CP server I/O fencing configuration procedures 91

D

- data disks
 - for I/O fencing 83
- deinstalling the Volume Manager 290
- disabling the agents 292
- disks
 - adding and initializing 131
 - coordinator 140
 - testing with vxfcntlshdw 133
 - verifying node access 134

E

- Ethernet controllers 264

F

- freezing service groups 176

G

- gabtab file
 - verifying after installation 249
- global clusters
 - configuration 121

H

- hastart 275
- hastatus -summary command 255
- hasys -display command 256
- high availability issues 339
 - low memory 339
 - network partition 339
- hubs
 - independent 264

I

- I/O fencing
 - checking disks 133
 - setting up 139
 - shared storage 133

installing

- post 122
- Root Broker 77

J

jeopardy 338–339

L

links

- private network 249

LLT

- verifying 252

lltconfig command 249

llthosts file

- verifying after installation 249

lltstat command 252

llttab file

- verifying after installation 249

log files 343

M

manual pages

- potential problems 341
- troubleshooting 341

N

network partition 338

O

original configuration

- restoring the 189

P

PATH variable

- VCS commands 251

persistent reservations

- SCSI-3 41

phased 219

phased upgrade 219

- example 220

planning to upgrade VVR 25

preinstallation 25

preparing to upgrade VVR 176

problems

- accessing manual pages 341
- executing file system commands 342

R

removing

- the Replicated Data Set 293

removing a system from a cluster 275

remsh 108

Replicated Data Set

- removing the 293

restoring the original configuration 189

Root Broker

- installing 77

rsh 108

S

SCSI

- changing initiator IDs 41

SCSI ID

- changing 42
- verifying 42

SCSI-3

- persistent reservations 41

SCSI-3 persistent reservations

- verifying 139

service groups

- freezing 176
- unfreezing 188

setup

- cabling shared devices 42
- SCSI Initiator ID 41

Shared storage

- Fibre Channel 41

shared storage

- setting SCSI initiator ID 41

single-node cluster

- adding a node to 282

SMTP email notification 117

SNMP trap notification 119

split brain 338

ssh 108

starting configuration

- installvcs program 108
- Veritas product installer 108

stopping

- applications 176

Storage Foundation

- coordinator disks 140

Storage Foundation installation

- verifying

- cluster operations 251
- GAB operations 251

Storage Foundation installation *(continued)*
 verifying *(continued)*
 LLT operations 251
Symantec Product Authentication Service 77, 113
system state attribute value 255

T

troubleshooting
 accessing manual pages 341
 executing file system commands 342

U

unfreezing service groups 188
upgrade
 phased 219
 planning 168
upgrading
 phased 219
upgrading VVR
 from 4.0 26
 planning 25
 preparing 176

V

VCS
 command directory path variable 251
vradmin
 delpri 294
 stoprep 294
VVR 4.0
 planning an upgrade from 26
vvr_upgrade_finish script 190
vxdisksetup command 131