Symantec™ Storage Foundation and High Availability 6.2 Installation Guide - Linux

January 2016

VERITAS™

Symantec™ Storage Foundation and High Availability Installation Guide

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

Product version: 6.2

Document version: 6.2 Rev 4

Legal Notice

Copyright © 2016 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

http://www.symantec.com

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

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

Documentation

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

https://sort.symantec.com/documents

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

doc_feedback@symantec.com

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

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

About Symantec Connect

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

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

Contents

Technical Support 4

Section 1 Installation overview and planning 22

Chapter 1 Introducing Storage Foundation and High Availability 23

About Storage Foundation High Availability 23
About Symantec Replicator Option 24
About Veritas Operations Manager 25
About Storage Foundation and High Availability features 25
About LLT and GAB 26
About I/O fencing 26
About global clusters 27
About Symantec Operations Readiness Tools 28
About configuring SFHA clusters for data integrity 29
About I/O fencing for SFHA in virtual machines that do not support SCSI-3 PR 30
About I/O fencing components 30

Chapter 2 System requirements 34

Release notes 34
Important preinstallation information for SFHA 35
Supported operating systems 35
Storage Foundation memory requirements 35
Disk space requirements 35
Checking installed product versions and downloading maintenance releases and patches 36
Obtaining installer patches 37
Disabling external network connection attempts 38
Database requirements 38
VxVM licenses 38
Cross-Platform Data Sharing licensing 39
I/O fencing requirements 39
Coordinator disk requirements for I/O fencing 40

CP server requirements 40
Non-SCSI-3 I/O fencing requirements 44
Number of nodes supported 44

Chapter 3 Planning to install SFHA 45

About installation and configuration methods 45
About response files 47
Downloading the Storage Foundation and High Availability software 48
About the VRTSspt RPM troubleshooting tools 49

Chapter 4 Licensing SFHA 50

About Symantec product licensing 50
Setting or changing the product level for keyless licensing 51
Installing Symantec product license keys 53

Section 2 Preinstallation tasks 55

Chapter 5 Preparing to install Storage Foundation High Availability 56

Installation preparation overview 56
About using ssh or rsh with the installer 57
Setting up the private network 58
Setting up shared storage 60
Setting up shared storage: SCSI 60
Setting up shared storage: Fibre Channel 62
Setting environment variables 63
Setting the kernel.hung_task_panic tunable 64
Optimizing LLT media speed settings on private NICs 65
Guidelines for setting the media speed of the LLT interconnects 65
Mounting the product disc 65
Assessing the system for installation readiness 66
Prechecking your systems using the installer 66

Section 3 Installation using the script-based installer 68

Chapter 6 Installing SFHA 69

About the script-based installer 69
Installing Storage Foundation and High Availability using the script-based installer 71

Chapter 7 Preparing to configure SFHA clusters for data integrity 75

About planning to configure I/O fencing 75
Typical Storage Foundation High Availability configuration with server-based I/O fencing 79
Recommended CP server configurations 80
Setting up the CP server 83
Planning your CP server setup 83
Installing the CP server using the installer 85
Configuring the CP server cluster in secure mode 85
Setting up shared storage for the CP server database 86
Configuring the CP server using the installer program 87
Configuring the CP server manually 99
Configuring CP server using response files 105
Verifying the CP server configuration 110
Configuring the CP server using the web-based installer 111

Chapter 8 Configuring SFHA 113

Configuring Storage Foundation High Availability using the installer 113
Overview of tasks to configure SFHA using the script-based installer 113
Required information for configuring Storage Foundation and High Availability Solutions 114
Starting the software configuration 115
Specifying systems for configuration 116
Configuring the cluster name 117
Configuring private heartbeat links 117
Configuring the virtual IP of the cluster 123
Configuring Storage Foundation and High Availability in secure mode 124
Configuring a secure cluster node by node 125

Adding VCS users 130
Configuring SMTP email notification 131
Configuring SNMP trap notification 132
Configuring global clusters 134
Completing the SFHA configuration 135
Verifying and updating licenses on the system 136
Configuring SFDB 138

Chapter 9 Manually configuring SFHA clusters for data integrity 139

Setting up disk-based I/O fencing using installsfha 139
Initializing disks as VxVM disks 139
Checking shared disks for I/O fencing 140
Configuring disk-based I/O fencing using installsfha 144
Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installsfha 146
Setting up server-based I/O fencing using installsfha 148
Refreshing keys or registrations on the existing coordination points for server-based fencing using the installsfha 156
Setting the order of existing coordination points for server-based fencing using the installsfha 158
Setting up non-SCSI-3 I/O fencing in virtual environments using installsfha 161
Setting up majority-based I/O fencing using installsfha 163
Enabling or disabling the preferred fencing policy 165

Section 4 Installation using the web-based installer 168

Chapter 10 Installing SFHA 169

About the web-based installer 169
Before using the web-based installer 170
Starting the web-based installer 170
Obtaining a security exception on Mozilla Firefox 171
Performing a preinstallation check with the web-based installer 172
Setting installer options with the web-based installer 172
Installing SFHA with the web-based installer 173

Chapter 11 Configuring SFHA 175

Configuring SFHA using the web-based installer 175
Configuring SFHA for data integrity using the web-based installer 181

Section 5 Automated installation using response files 195

Chapter 12 Performing an automated SFHA installation 196

Installing SFHA using response files 196
Response file variables to install Storage Foundation and High Availability 197
Sample response file for SFHA install 199

Chapter 13 Performing an automated SFHA configuration 201

Configuring SFHA using response files 201
Response file variables to configure Storage Foundation and High Availability 202
Sample response file for SFHA configuration 212

Chapter 14 Performing an automated I/O fencing configuration using response files 213

Configuring I/O fencing using response files 213
Response file variables to configure disk-based I/O fencing 214
Sample response file for configuring disk-based I/O fencing 217
Response file variables to configure server-based I/O fencing 217
Sample response file for configuring server-based I/O fencing 219
Sample response file for configuring non-SCSI-3 I/O fencing 220
Response file variables to configure non-SCSI-3 I/O fencing 220
Response file variables to configure majority-based I/O fencing 222
Sample response file for configuring majority-based I/O fencing 222

Section 6 Installation using operating system-specific methods 224

Chapter 15 Installing SFHA using operating system-specific methods 225

About installing SFHA using operating system-specific methods 225
Installing SFHA using Kickstart 226
Sample Kickstart configuration file 228
Installing Storage Foundation and High Availability using yum 230
Installing SFHA using the Red Hat Satellite server 235
Using Red Hat Satellite server to install SFHA products 236

Chapter 16 Configuring SFHA clusters for data integrity 238

Setting up disk-based I/O fencing manually 238
Removing permissions for communication 239
Identifying disks to use as coordinator disks 239
Setting up coordinator disk groups 239
Creating I/O fencing configuration files 240
Modifying VCS configuration to use I/O fencing 241
Verifying I/O fencing configuration 243
Setting up server-based I/O fencing manually 244
Preparing the CP servers manually for use by the Storage Foundation High Availability 244
Generating the client key and certificates manually on the client nodes 247
Configuring server-based fencing on the Storage Foundation High Availability manually 249
Configuring CoordPoint agent to monitor coordination points 256
Verifying server-based I/O fencing configuration 257
Setting up non-SCSI-3 fencing in virtual environments manually 258
Sample /etc/vxfenmode file for non-SCSI-3 fencing 260
Setting up majority-based I/O fencing manually 264
Creating I/O fencing configuration files 264
Modifying VCS configuration to use I/O fencing 264
Verifying I/O fencing configuration 266

Section 7 Managing your Symantec deployments 268

Chapter 17 Performing centralized installations using the Deployment Server 269

About the Deployment Server 270
Deployment Server overview 271
Installing the Deployment Server 272
Setting up a Deployment Server 273
Setting deployment preferences 276
Specifying a non-default repository location 278
Downloading the most recent release information 278
Loading release information and patches on to your Deployment Server 279
Viewing or downloading available release images 280
Viewing or removing repository images stored in your repository 285
Deploying Symantec product updates to your environment 287
Finding out which releases you have installed, and which upgrades or updates you may need 288
Defining Install Bundles 290
Creating Install Templates 295
Deploying Symantec releases 297
Connecting the Deployment Server to SORT using a proxy server 300

Section 8 Upgrade of SFHA 301

Chapter 18 Planning to upgrade SFHA 302

Upgrade methods for SFHA 302
Supported upgrade paths for SFHA 6.2 303
Considerations for upgrading SFHA to 6.2 on systems configured with an Oracle resource 306
About using the installer to upgrade when the root disk is encapsulated 307
Preparing to upgrade SFHA 308
Getting ready for the upgrade 308
Creating backups 309
Determining if the root disk is encapsulated 310
Pre-upgrade tasks for migrating the SFDB repository database 311

Pre-upgrade planning for Volume Replicator 311
Preparing to upgrade VVR when VCS agents are configured 314
Upgrading the array support 317
Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches 318

Chapter 19 Upgrading Storage Foundation and High Availability 321

Upgrading Storage Foundation and High Availability from previous versions to 6.2 321
Upgrading Storage Foundation and High Availability using the script-based installer 322
Upgrading SFHA using the web-based installer 326
Upgrading Volume Replicator 328
Upgrading VVR without disrupting replication 328
Upgrading SFDB 329

Chapter 20 Performing a rolling upgrade of SFHA 331

About rolling upgrades 331
Supported rolling upgrade paths 334
Performing a rolling upgrade using the script-based installer 334
Performing a rolling upgrade of SFHA using the web-based installer 338

Chapter 21 Performing a phased upgrade of SFHA 341

About phased upgrade 341
Prerequisites for a phased upgrade 341
Planning for a phased upgrade 342
Phased upgrade limitations 342
Phased upgrade example 342
Phased upgrade example overview 343
Performing a phased upgrade using the script-based installer 344
Moving the service groups to the second subcluster 344
Upgrading the operating system on the first subcluster 347
Upgrading the first subcluster 348
Preparing the second subcluster 349
Activating the first subcluster 353
Upgrading the operating system on the second subcluster 354
Upgrading the second subcluster 354
Finishing the phased upgrade 355

Chapter 22 Performing an automated SFHA upgrade using response files 359

Upgrading SFHA using response files 359
Response file variables to upgrade Storage Foundation and High Availability 360
Sample response file for SFHA upgrade 363
Performing rolling upgrade of SFHA using response files 364
Response file variables to upgrade SFHA using rolling upgrade 364
Sample response file for SFHA using rolling upgrade 366

Chapter 23 Performing post-upgrade tasks 368

Optional configuration steps 368
Re-joining the backup boot disk group into the current disk group 369
Reverting to the backup boot disk group after an unsuccessful upgrade 369
Post upgrade tasks for migrating the SFDB repository database 370
Migrating from a 5.0 repository database to 6.2 370
Migrating from a 5.1 or higher repository database to 6.2 373
Migrating SFDB from 5.0x to 6.2 376
Recovering VVR if automatic upgrade fails 376
Post-upgrade tasks when VCS agents for VVR are configured 377
Unfreezing the service groups 377
Restoring the original configuration when VCS agents are configured 378
Upgrading disk layout versions 380
Upgrading VxVM disk group versions 381
Updating variables 382
Setting the default disk group 382
About enabling LDAP authentication for clusters that run in secure mode 382
Enabling LDAP authentication for clusters that run in secure mode 384
Verifying the Storage Foundation and High Availability upgrade 388

Section 9 Post-installation tasks 389

Chapter 24 Performing post-installation tasks 390

Switching on Quotas 390
About configuring authentication for SFDB tools 390
Configuring vxdbd for SFDB tools authentication 391

Chapter 25 Verifying the SFHA installation 392

Upgrading the disk group version 392
Performing a postcheck on a node 393
Verifying that the products were installed 394
Installation log files 394
Using the installation log file 394
Using the summary file 394
Starting and stopping processes for the Symantec products 395
Checking Veritas Volume Manager processes 395
Verifying the LLT, GAB, and VCS configuration files 396
Verifying LLT, GAB, and cluster operation 396
Verifying LLT 397
Verifying the cluster 399
Verifying the cluster nodes 400

Section 10 Uninstallation of SFHA 403

Chapter 26 Uninstalling Storage Foundation and High Availability 404

Removing VxFS file systems 404
Removing rootability 405
Moving volumes to disk partitions 406
Moving volumes onto disk partitions using VxVM 406
Disabling VCS agents for VVR the agents on a system 408
Removing the Replicated Data Set 409
Uninstalling SFHA RPMs using the script-based installer 410
Uninstalling SFHA with the web-based installer 412
Removing license files (Optional) 413
Removing the CP server configuration using the installer program 413
Removing the Storage Foundation for Databases (SFDB) repository 415

Chapter 27 Uninstalling SFHA using response files 417

Uninstalling SFHA using response files 417
Response file variables to uninstall Storage Foundation and High Availability 418
Sample response file for SFHA uninstallation 419

Section 11 Adding and removing nodes 420

Chapter 28 Adding a node to SFHA clusters 421

About adding a node to a cluster 421
Before adding a node to a cluster 422
Adding a node to a cluster using the SFHA installer 424
Adding a node using the web-based installer 427
Adding the node to a cluster manually 428
Starting Veritas Volume Manager (VxVM) on the new node 429
Configuring cluster processes on the new node 429
Setting up the node to run in secure mode 431
Starting fencing on the new node 432
Configuring the ClusterService group for the new node 432
Adding a node using response files 433
Response file variables to add a node to a SFHA cluster 433
Sample response file for adding a node to a SFHA cluster 434
Configuring server-based fencing on the new node 434
Adding the new node to the vxfen service group 435
After adding the new node 436
Adding nodes to a cluster that is using authentication for SFDB tools 436
Updating the Storage Foundation for Databases (SFDB) repository after adding a node 437

Chapter 29 Removing a node from SFHA clusters 439

Removing a node from a SFHA cluster 439
Verifying the status of nodes and service groups 440
Deleting the departing node from SFHA configuration 441
Modifying configuration files on each remaining node 444
Removing the node configuration from the CP server 444
Removing security credentials from the leaving node 445
Unloading LLT and GAB and removing VCS on the departing node 446
Updating the Storage Foundation for Databases (SFDB) repository after removing a node 447

Section 12 Installation reference 448

Appendix A SFHA services and ports 449

About SFHA services and ports 449

Appendix B Installation scripts 451

Installation script options 451
About using the postcheck option 457

Appendix C Tunable files for installation 460

About setting tunable parameters using the installer or a response file 460
Setting tunables for an installation, configuration, or upgrade 461
Setting tunables with no other installer-related operations 462
Setting tunables with an un-integrated response file 463
Preparing the tunables file 464
Setting parameters for the tunables file 464
Tunables value parameter definitions 465

Appendix D Configuration files 473

About the LLT and GAB configuration files 473
About the AMF configuration files 476
About the VCS configuration files 477
Sample main.cf file for VCS clusters 478
Sample main.cf file for global clusters 480
About I/O fencing configuration files 482
Sample configuration files for CP server 484
Sample main.cf file for CP server hosted on a single node that runs VCS 485
Sample main.cf file for CP server hosted on a two-node SFHA cluster 487
Sample CP server configuration (/etc/vxcps.conf) file output 490

Appendix E Configuring the secure shell or the remote shell for communications 491

About configuring secure shell or remote shell communication modes before installing products 491
Manually configuring passwordless ssh 492
Setting up ssh and rsh connection using the installer -comsetup command 495
Setting up ssh and rsh connection using the pwdutil.pl utility 497
Restarting the ssh session 500
Enabling rsh for Linux 500

Appendix F Storage Foundation and High Availability components 503

Storage Foundation and High Availability installation RPMs 503
Symantec Cluster Server installation RPMs 506
Symantec Storage Foundation obsolete and reorganized installation RPMs 507

Appendix G Troubleshooting installation issues 511

Restarting the installer after a failed connection 511
What to do if you see a licensing reminder 511
Incorrect permissions for root on remote system 512
Inaccessible system 513
Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.2 (2184482) 513
Troubleshooting the webinstaller 514

Appendix H Troubleshooting cluster installation 515

Unmount failures 515
Command failures 515
Installer cannot create UUID for the cluster 516
The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails 516
Troubleshooting CP server 517
Troubleshooting issues related to the CP server service group 518
Checking the connectivity of CP server 518
Troubleshooting server-based fencing on the SFHA cluster nodes 518
Issues during fencing startup on Storage Foundation High Availability nodes set up for server-based fencing 519
Issues during online migration of coordination points 520
Vxfen service group activity after issuing the vxfenswap command 520

Appendix I Sample Storage Foundation High Availability setup diagrams for CP server-based I/O fencing 522

Configuration diagrams for setting up server-based I/O fencing 522
Two unique client clusters served by 3 CP servers 523

Client cluster served by highly available CPS and 2 SCSI-3 disks 523
Two node campus cluster served by remote CP server and 2 SCSI-3 disks 524
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks 526

Appendix J Configuring LLT over UDP 527

Using the UDP layer for LLT 527
When to use LLT over UDP 527
Manually configuring LLT over UDP using IPv4 527
Broadcast address in the /etc/lltab file 528
The link command in the /etc/lltab file 529
The set-addr command in the /etc/lltab file 529
Selecting UDP ports 530
Configuring the netmask for LLT 530
Configuring the broadcast address for LLT 531
Sample configuration: direct-attached links 531
Sample configuration: links crossing IP routers 533
Using the UDP layer of IPv6 for LLT 534
When to use LLT over UDP 534
Manually configuring LLT over UDP using IPv6 535
Sample configuration: direct-attached links 535
Sample configuration: links crossing IP routers 536

Appendix K Using LLT over RDMA 539

Using LLT over RDMA 539
About RDMA over RoCE or InfiniBand networks in a clustering environment 539
How LLT supports RDMA capability for faster interconnects between applications 540
Using LLT over RDMA: supported use cases 541
Configuring LLT over RDMA 541
Choosing supported hardware for LLT over RDMA 542
Installing RDMA, InfiniBand or Ethernet drivers and utilities 543
Configuring RDMA over an Ethernet network 544
Configuring RDMA over an InfiniBand network 546
Tuning system performance 550
Manually configuring LLT over RDMA 552
LLT over RDMA sample /etc/lltab 556
Verifying LLT configuration 556
Troubleshooting LLT over RDMA 557

IP addresses associated to the RDMA NICs do not automatically plumb on node restart 557
Ping test fails for the IP addresses configured over InfiniBand interfaces. 558
After a node restart, by default the Mellanox card with Virtual Protocol Interconnect (VPI) gets configured in InfiniBand mode 558
The LLT module fails to start 558

Appendix L Compatibility issues when installing Storage Foundation High Availability with other products 560

Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present 561
Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present 561
Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present 561
Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present 562

Index 563

Section 1

Installation overview and planning

- Chapter 1. Introducing Storage Foundation and High Availability
- Chapter 2. System requirements
- Chapter 3. Planning to install SFHA
- Chapter 4. Licensing SFHA

Chapter 1

Introducing Storage Foundation and High Availability

This chapter includes the following topics:

- About Storage Foundation High Availability
- About Veritas Operations Manager
- About Storage Foundation and High Availability features
- About Symantec Operations Readiness Tools
- About configuring SFHA clusters for data integrity

About Storage Foundation High Availability

Symantec Storage Foundation High Availability by Symantec (SFHA) includes the following:

Symantec Storage Foundation	Symantec Storage Foundation includes the following: ■ Veritas File System by Symantec (VxFS). Veritas File System is a high-performance journaling file system that provides easy management and quick-recovery for applications. Veritas File System delivers scalable performance, continuous availability, increased I/O throughput, and structural integrity. ■ Veritas Volume Manager by Symantec (VxVM). Veritas Volume Manager removes the physical limitations of disk storage. You can configure, share, manage, and optimize storage I/O performance online without interrupting data availability. Veritas Volume Manager also provides easy-to-use, online storage management tools to reduce downtime. VxFS and VxVM are a part of all Symantec Storage Foundation products. Do not install or update VxFS or VxVM as individual components.
Symantec Cluster Server (VCS)	Symantec Cluster Server by Symantec is a clustering solution that provides the following benefits: ■ Reduces application downtime ■ Facilitates the consolidation and the failover of servers ■ Manages a range of applications in heterogeneous environments
Veritas agents	Veritas agents provide high availability for specific resources and applications. Each agent manages resources of a particular type. For example, the Oracle agent manages Oracle databases. Agents typically start, stop, and monitor resources and report state changes.

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

About Symantec Replicator Option

Symantec Replicator Option is an optional, separately-licensable feature.

File Replicator enables replication at the file level over IP networks. File Replicator leverages data duplication, provided by Veritas File System, to reduce the impact of replication on network resources.

Volume Replicator replicates data to remote locations over any standard IP network to provide continuous data availability and disaster recovery.

Volume Replicator is available with Storage Foundation, Storage Foundation High Availability, Storage Foundation Cluster File System, Storage Foundation for Oracle RAC, and Storage Foundation for SybaseCE.

Before installing this option, read the Release Notes for the product.

To install the option, follow the instructions in the Installation Guide for the product.

About Veritas Operations Manager

Veritas Operations Manager provides a centralized management console for Symantec Storage Foundation and High Availability products. You can use Veritas Operations Manager to monitor, visualize, and manage storage resources and generate reports.

Symantec recommends using Veritas Operations Manager (VOM) to manage Storage Foundation and Cluster Server environments.

You can download Veritas Operations Manager from http://go.symantec.com/vom.

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

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

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

About Storage Foundation and High Availability features

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

About LLT and GAB

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

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

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

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.

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 installer installs the I/O fencing driver, part of VRTSvxfen RPM, when you install SFHA. To protect data on shared disks, you must configure I/O fencing after you install and configure SFHA.

I/O fencing modes - disk-based and server-based I/O fencing - use coordination points for arbitration in the event of a network partition. Whereas, majority-based I/O fencing mode does not use coordination points for arbitration. With majority-based I/O fencing you may experience loss of high availability in some cases. You can configure disk-based, server-based, or majority-based I/O fencing:

Disk-based I/O fencing	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.

Server-based I/O fencing	I/O fencing that uses at least one CP server system is referred to as server-based I/O fencing. Server-based fencing can include only CP servers, or a mix of CP servers and coordinator disks. Server-based I/O fencing ensures data integrity in clusters. In virtualized environments that do not support SCSI-3 PR, SFHA supports non-SCSI-3 I/O fencing. See "About I/O fencing for SFHA in virtual machines that do not support SCSI-3 PR" on page 30.
Majority-based I/O fencing	Majority-based I/O fencing mode does not need coordination points to provide protection against data corruption and data consistency in a clustered environment. Symantec designed majority-based I/O fencing mode to be used in stand-alone appliances. You can configure I/O fencing in majority-based mode, but as a best practice that where possible, utilize Coordination Point servers and or shared SCSI-3 disks to be used as coordination points.

See " About planning to configure I/O fencing" on page 75.

Note: Symantec recommends that you use I/O fencing to protect your cluster against split-brain situations.

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

About global clusters

Global clusters provide the ability to fail over applications between geographically distributed clusters when disaster occurs. You require a separate license to configure global clusters. You must add this license during the installation. The installer only asks about configuring global clusters if you have used the global cluster license.

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

About Symantec Operations Readiness Tools

Symantec Operations Readiness Tools (SORT) is a website that automates and simplifies some of the most time-consuming administrative tasks. It helps you identify risks in your datacenters and improve operational efficiency, enabling you to manage the complexity that is associated with datacenter architectures and scale.

Table 1-1 lists three major datacenter tasks and the SORT tools that can help you accomplish them.

Table 1-1 Datacenter tasks and the SORT tools

Task	SORT tools
Prepare for installations and upgrades	■ Installation and Upgrade checklists Display system requirements including memory, disk space, and architecture. ■ Installation and Upgrade custom reports Create reports that determine if you're ready to install or upgrade a Symantec enterprise product. ■ Array-specific Module Finder List the latest Array Support Libraries (ASLs) and Array Policy Modules (APMs) for UNIX servers, and Device Driver Installers (DDIs) and Device Discovery Layers (DDLs) for Windows servers. ■ High Availability Agents table Find and download the agents for applications, databases, replication, and Symantec partners.
Identify risks and get server-specific recommendations	■ Patch notifications Receive automatic email notifications about patch updates. (Sign in required.) ■ Risk Assessment check lists Display configuration recommendations based on your Symantec product and platform. ■ Risk Assessment custom reports Create reports that analyze your system and give you recommendations about system availability, storage use, performance, and best practices. ■ Error code descriptions and solutions Display detailed information on thousands of Symantec error codes.

Table 1-1 Datacenter tasks and the SORT tools *(continued)*

Task	SORT tools
Improve efficiency	■ Patch Finder List and download patches for your Symantec enterprise products. ■ License/Deployment custom reports Create custom reports that list your installed Symantec products and license keys. Display licenses by product, platform, server tier, and system. ■ Symantec Performance Value Unit (SPVU) Calculator Use the calculator to assist you with the pricing meter transition. ■ Documentation List and download Symantec product documentation, including manual pages, product guides, and support articles. ■ Related links Display links to Symantec product support, forums, customer care, and vendor information on a single page.

SORT is available at no additional charge.

To access SORT, go to:

https://sort.symantec.com

About configuring SFHA clusters for data integrity

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

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

- Broken set of private networks
  If a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects. The remaining node then takes corrective action. The failure of the private interconnects, instead of the actual nodes, presents identical symptoms and causes each node to determine its peer has departed. This situation typically results in data corruption because both nodes try to take control of data storage in an uncoordinated manner.
- System that appears to have a system-hang

If a system is so busy that it appears to stop responding, the other nodes could declare it as dead. This declaration may also occur for the nodes that use the hardware that supports a "break" and "resume" function. When a node drops to PROM level with a break and subsequently resumes operations, the other nodes may declare the system dead. They can declare it dead even if the system later returns and begins write operations.

I/O fencing is a feature that prevents data corruption in the event of a communication breakdown in a cluster. SFHA uses I/O fencing to remove the risk that is associated with split-brain. I/O fencing allows write access for members of the active cluster. It blocks access to storage from non-members so that even a node that is alive is unable to cause damage.

After you install and configure SFHA, you must configure I/O fencing in SFHA to ensure data integrity.

See " About planning to configure I/O fencing" on page 75.

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

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

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

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

See "Setting up non-SCSI-3 I/O fencing in virtual environments using installsfha" on page 161.

See "Setting up non-SCSI-3 fencing in virtual environments manually" on page 258.

About I/O fencing components

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

- Data disks—Store shared data
  See "About data disks" on page 31.
- Coordination points—Act as a global lock during membership changes
  See "About coordination points" on page 31.

About data disks

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

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

Note: Disk based fencing is possible only if VxVM is also present long with VCS.

About coordination points

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

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

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

- Coordinator disks
  Disks that act as coordination points are called coordinator disks. Coordinator disks are three standard disks or LUNs set aside for I/O fencing during cluster reconfiguration. Coordinator disks do not serve any other storage purpose in the SFHA configuration.
  You can configure coordinator disks to use Veritas Volume Manager's Dynamic Multi-pathing (DMP) feature. Dynamic Multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. So, you can configure I/O fencing to use DMP devices. I/O

fencing uses SCSI-3 disk policy that is dmp-based on the disk device that you use.

Note: The dmp disk policy for I/O fencing supports both single and multiple hardware paths from a node to the coordinator disks. If few coordinator disks have multiple hardware paths and few have a single hardware path, then we support only the dmp disk policy. For new installations, Symantec only supports dmp disk policy for IO fencing even for a single hardware path.

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

- Coordination point servers

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

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

  In short, the CP server functions as another arbitration mechanism that integrates within the existing I/O fencing module.

  Note: With the CP server, the fencing arbitration logic still remains on the SFHA cluster.

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

About preferred fencing

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

You can configure the preferred fencing policy using the cluster-level attribute PreferredFencingPolicy for the following:

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

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

See "Enabling or disabling the preferred fencing policy" on page 165.

Chapter 2

System requirements

This chapter includes the following topics:

- Release notes
- Important preinstallation information for SFHA
- Supported operating systems
- Storage Foundation memory requirements
- Disk space requirements
- Checking installed product versions and downloading maintenance releases and patches
- Obtaining installer patches
- Disabling external network connection attempts
- Database requirements
- VxVM licenses
- Cross-Platform Data Sharing licensing
- I/O fencing requirements
- Number of nodes supported

Release notes

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

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

https://sort.symantec.com/documents

Important preinstallation information for SFHA

Before you install SFHA, make sure that you have reviewed the following information:

- Preinstallation checklist for your configuration. Go to the SORT installation checklist tool. From the drop-down lists, select the information for the Symantec product you want to install, and click **Generate Checklist**.
- Hardware compatibility list for information about supported hardware: http://www.symantec.com/docs/TECH211575
- For important updates regarding this release, review the Late-Breaking News Technote on the Symantec Technical Support website: http://www.symantec.com/docs/TECH225259

Supported operating systems

For information on supported operating systems for various components of SFHA, see the *Storage Foundation and High Availability Release Notes*.

Storage Foundation memory requirements

A minimum of 1 GB of memory is strongly recommended.

Disk space requirements

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

Use the **Perform a Preinstallation Check (P)** menu for the web-based installer to determine whether there is sufficient space.

Or, go to the installation directory and run the installer with the `-precheck` option.

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

See "About the script-based installer" on page 69.

Checking installed product versions and downloading maintenance releases and patches

Symantec provides a means to check the Symantec RPMs you have installed, and download any needed maintenance releases and patches.

Use the `installer` command with the `-version` option to determine what is installed on your system, and download any needed maintenance releases or patches. After you have installed the current version of the product, you can use the `showversion` script in the `/opt/VRTS/install` directory to find product information.

The `version` option or the `showversion` script checks the specified systems and discovers the following:

- SFHA product versions that are installed on the system
- All the required RPMs and the optional Symantec RPMs installed on the system
- Any required or optional RPMs (if applicable) that are not present
- Installed patches
- Available base releases (major or minor)
- Available maintenance releases
- Available patch releases

To check your systems and download maintenance releases and patches

1. Mount the media, or navigate to the installation directory.
2. Start the installer with the `-version` option.

   ```
   # ./installer -version sys1 sys2
   ```

   For each system, the installer lists all of the installed base releases, maintenance releases, and patches, followed by the lists of available downloads.
3. If you have Internet access, follow the prompts to download the available maintenance releases and patches to the local system.
4. If you do not have Internet access, you can download any needed maintenance releases and patches from the Symantec Operations Readiness Tools (SORT) Patch Finder page at:

   https://sort.symantec.com/patch/finder

   You can obtain installer patches automatically or manually.

   See “Obtaining installer patches” on page 37.

Downloading maintenance releases and patches requires the installer to make outbound networking calls. You can also disable external network connection attempts.

See "Disabling external network connection attempts" on page 38.

Obtaining installer patches

Symantec occasionally finds issues with the Storage Foundation and High Availability installer, and posts public installer patches on the Symantec Operations Readiness Tools (SORT) website's Patch Finder page at:

https://sort.symantec.com/patch/finder

You can access installer patches automatically or manually.

To download installer patches automatically

- Starting with Storage Foundation and High Availability version 6.1, installer patches are downloaded automatically. No action is needed on your part.

  If you are running Storage Foundation and High Availability version 6.1 or later, and your system has Internet access, the installer automatically imports any needed installer patch, and begins using it.

Automatically downloading installer patches requires the installer to make outbound networking calls. You can also disable external network connection attempts.

See "Disabling external network connection attempts" on page 38.

If your system does not have Internet access, you can download installer patches manually.

To download installer patches manually

1. Go to the Symantec Operations Readiness Tools (SORT) website's Patch Finder page, and save the most current Symantec patch on your local system.
2. Navigate to the directory where you want to unzip the file you downloaded in step 1.
3. Unzip the patch tar file. For example, run the following command:

   ```
   # gunzip cpi-6.2P2-patches.tar.gz
   ```

4. Untar the file. For example, enter the following:

   ```
   # tar -xvf cpi-6.2P2-patches.tar
   patches/
   patches/CPI62P2.pl
   README
   ```

5 Navigate to the installation media or to the installation directory.

6 To start using the patch, run the `installer` command with the `-require` option. For example, enter the following:

```
# ./installer -require /target_directory/patches/CPI62P2.pl
```

Disabling external network connection attempts

When you execute the `installer` command, the installer attempts to make an outbound networking call to get information about release updates and installer patches. If you know your systems are behind a firewall, or do not want the installer to make outbound networking calls, you can disable external network connection attempts by the installer.

To disable external network connection attempts

- Disable inter-process communication (IPC).

  To disable IPC, run the installer with the `-noipc` option.

  For example, to disable IPC for system1 (sys1) and system2 (sys2) enter the following:

  ```
  # ./installer -noipc sys1 sys2
  ```

Database requirements

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

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

Note: SFHA supports running Oracle and Sybase on VxFS and VxVM.

SFHA does not support running SFDB tools with Sybase.

VxVM licenses

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

Table 2-1 describes the levels of licensing in Veritas Volume Manager and supported features.

Table 2-1 Levels of licensing in Veritas Volume Manager and supported features

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

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

To see the license features that are enabled in VxVM

- Enter the following command:

```
# vxdctl license
```

Cross-Platform Data Sharing licensing

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

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

I/O fencing requirements

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

- Coordinator disks
  See "Coordinator disk requirements for I/O fencing" on page 40.
- CP servers

See "CP server requirements" on page 40.

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

See "Non-SCSI-3 I/O fencing requirements" on page 44.

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 at least three coordinator disks or there must be odd number of coordinator disks.
- The coordinator disks must be DMP devices.
- Each of the coordinator disks must use a physically separate disk or LUN. Symantec recommends using the smallest possible LUNs for coordinator disks.
- Each of the coordinator disks should exist on a different disk array, if possible.
- The coordinator disks must support SCSI-3 persistent reservations.
- Coordinator devices can be attached over iSCSI protocol but they must be DMP devices and 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.
- The coordinator disk size must be at least 128 MB.

CP server requirements

SFHA 6.2 clusters (application clusters) support coordination point servers (CP servers) that are hosted on the following VCS and SFHA versions:

- VCS 6.1 or later single-node cluster
- SFHA 6.1 or later cluster

Upgrade considerations for CP servers

- Upgrade VCS or SFHA on CP servers to version 6.2 if the current release version is prior to version 6.1.
- You do not need to upgrade CP servers to version 6.2 if the release version is 6.1.

- CP servers on version 6.2 support HTTPS-based communication with application clusters on version 6.1 or later.
- CP servers on version 6.2 support IPM-based communication with application clusters on versions before 6.1.
- You need to configure VIPs for HTTPS-based communication if release version of application clusters is 6.1 or later.
- You need to configure VIPs for IPM-based communication if release version of application clusters is before 6.1.

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

See the *Symantec Cluster Server Installation Guide*.

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

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

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

Table 2-2 lists additional requirements for hosting the CP server.

Table 2-2 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: ■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB) ■ 300 MB in /usr ■ 20 MB in /var ■ 10 MB in /etc (for the CP server database) See "Disk space requirements" on page 35.

Table 2-2 CP server hardware requirements *(continued)*

Hardware required	Description
Storage	When CP server is hosted on an SFHA cluster, there must be shared storage between the nodes of this SFHA cluster.
RAM	Each CP server requires at least 512 MB.
Network	Network hardware capable of providing TCP/IP connection between CP servers and SFHA clusters (application clusters).

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

Table 2-3 CP server supported operating systems and versions

CP server	Operating system and version

CP server hosted on a VCS single-node cluster or on an SFHA cluster	CP server supports any of the following operating systems:
	■ AIX 6.1 and 7.1
	■ Linux:
	  ■ RHEL 6
	  ■ RHEL 7
	  ■ SLES 11
	■ Oracle Solaris 10
	■ Oracle Solaris 11
	Review other details such as supported operating system levels and architecture for the supported operating systems.
	See the *Symantec Cluster Server Release Notes* or the *Symantec Storage Foundation High Availability Release Notes* for that platform.

Following are the CP server networking requirements and recommendations:

- Symantec recommends that network access from the application clusters to the CP servers should be made highly-available and redundant. The network connections require either a secure LAN or VPN.
- The CP server uses the TCP/IP protocol to connect to and communicate with the application clusters by these network paths. The CP server listens for messages from the application clusters using TCP port 443 if the communication happens over the HTTPS protocol. TCP port 443 is the default port that can be changed while you configure the CP server. The CP server listens for messages from the application clusters over the IPM-based protocol using the TCP port

14250. Unlike HTTPS protocol, which is a standard protocol, IPM (Inter Process Messaging) is a VCS-specific communication protocol.

Symantec recommends that you configure multiple network paths to access a CP server. If a network path fails, CP server does not require a restart and continues to listen on all the other available virtual IP addresses.

- The CP server supports either Internet Protocol version 4 (IPv4 addresses) or IPv6 addresses when communicating with the application clusters over the IPM-based protocol. The CP server only supports Internet Protocol version 4 (IPv4) when communicating with the application clusters over the HTTPS protocol.
- When placing the CP servers within a specific network configuration, you must take into consideration the number of hops from the different application cluster nodes to the CP servers. As a best practice, Symantec recommends that the number of hops and network latency from the different application cluster nodes to the CP servers should be equal. This ensures that if an event occurs that results in an I/O fencing scenario, there is no bias in the race due to difference in number of hops or network latency between the CPS and various nodes.

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

Table 2-4 Supported communication modes between SFHA cluster (application cluster) and CP server

Communication mode	CP server (HTTPS-based communication)	CP server (IPM-based secure communication)	CP server (IPM-based non-secure communication)
SFHA cluster (release version 6.1 or later)	Yes	No	No
SFHA cluster (release version prior to 6.1)	No	Yes	Yes

For secure communications between the SFHA and CP server over the IPM-based protocol, consider the following requirements and suggestions:

- In a secure communication environment, all CP servers that are used by the application cluster must be configured with security enabled. A configuration where the application cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.
- For non-secure communication between CP server and application clusters, there is no need to configure Symantec Product Authentication Service. In non-secure mode, authorization is still provided by CP server for the application

cluster users. The authorization that is performed only ensures that authorized users can perform appropriate actions as per their user privileges on the CP server.

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

Non-SCSI-3 I/O fencing requirements

Supported virtual environment for non-SCSI-3 fencing:

- VMware Server ESX 4.0, 5.0, 5.1, and 5.5 on AMD Opteron or Intel Xeon EM64T (x86_64)
  Guest operating system: See the *Storage Foundation and High Availability Release Notes* for the list of supported Linux operating systems.

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

- SFHA must be configured with Cluster attribute UseFence set to SCSI3
- For server-based I/O fencing, all coordination points must be CP servers

Number of nodes supported

SFHA supports cluster configurations with up to 64 nodes.

SmartIO writeback caching only supports cluster configurations with up to 2 nodes.

Chapter 3

Planning to install SFHA

This chapter includes the following topics:

- About installation and configuration methods
- Downloading the Storage Foundation and High Availability software
- About the VRTSspt RPM troubleshooting tools

About installation and configuration methods

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

Table 3-1 shows the installation and configuration methods that SFHA supports.

Table 3-1 Installation and configuration methods

Method	Description

The script-based installer	Using the script-based installer, you can install Symantec products from a driver system running a supported platform to target computers running any supported platform.
	
	To install your Symantec product using the installer, choose one of the following:
	
	■ The general product installer:
	`installer`
	The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
	
	■ Product-specific installation scripts:
	`installsfha<version>`
	The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installsfha` script is identical to running the general product installer and specifying SFHA from the list of products to install.
	Use the product-specific installation scripts to install or configure individual products you download electronically.
	
	See "About the script-based installer" on page 69.

The web-based installer	Using the web-based installer, you can install Symantec products from a driver system running a supported platform to target computers running any supported platform The web-based installer provides an interface to manage the installation and configuration from a remote site using a standard web browser. `webinstaller`
Deployment Server	Using the Deployment Server, you can store multiple release images in one central location and deploy them to systems of any supported platform. See "About the Deployment Server" on page 270.

Table 3-1 Installation and configuration methods *(continued)*

Method	Description
Silent installation using response files	Response files automate installation and configuration by using the information that is stored in a specified file instead of prompting you for information. 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 option to install silently on one or more systems.
Install Bundles	Beginning with version 6.1, you can easily install or upgrade your systems directly to a base, maintenance, or patch level in one step using Install Bundles. The installer installs both releases as if they were combined in the same release image. The various scripts, RPMs, and patch components are merged, and multiple releases are installed together as if they are one combined release. See "Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches" on page 318.
Kickstart (For RedHat Linux systems)	Kickstart lets you automatically install systems based on predefined customized configurations.
yum (For RedHat Linux systems)	You can install SFHARPMs into the yum repository. yum is supported for Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7.

About response files

The installer script or product installation script generates a response file during any installation, configuration, upgrade, or uninstall procedure. The response file contains the configuration information that you entered during the procedure. When the procedure completes, the installation script displays the location of the response files.

You can use the response file for future installation procedures by invoking an installation script with the `-responsefile` option. The response file passes arguments to the script to automate the installation of that product. You can edit the file to automate installation and configuration of additional systems.

You can generate a response file using the `-makeresponsefile` option.

See "Installation script options" on page 451.

Syntax in the response file

The syntax of the Perl statements that is 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 1 ", "value 2 ", "value 3 "];
```

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 website.

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

To download the trialware version of the software

1. Open the following link in your browser:

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

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

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

See "About the script-based installer" on page 69.

To download the software

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

   The estimated space for download, gunzip, and tar extract is 1 GB.

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

   See "Disk space requirements" on page 35.

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

   ```
   # df -k filesystem
   ```

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

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

About the VRTSspt RPM troubleshooting tools

The VRTSspt RPM provides a group of tools for troubleshooting a system and collecting information on its configuration. If you install and use the VRTSspt RPM, it will be easier for Symantec Support to diagnose any issues you may have.

The tools can gather Veritas File System and Veritas Volume Manager metadata information and establish various benchmarks to measure file system and volume manager performance. Although the tools are not required for the operation of any Symantec product, Symantec recommends installing them should a support case be needed to be opened with Symantec Support. Use caution when you use the VRTSspt RPM, and always use it in concert with Symantec Support.

Chapter 4

Licensing SFHA

This chapter includes the following topics:

- About Symantec product licensing
- Setting or changing the product level for keyless licensing
- Installing Symantec product license keys

About Symantec product licensing

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

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

http://www.symantec.com/products-solutions/licensing/activating-software/detail.jsp?detail_id=licensing_portal

The 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 systems with a management server. If you do not comply with the above terms, continuing to use the Symantec product is a violation of your End User License Agreement, and results in warning messages

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

http://go.symantec.com/sfhakeyless

If you upgrade to this release from a previous release of the Symantec software, the installer asks whether you want to upgrade the key to the new version. 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 51.
  See the `vxkeyless(1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.
  See "Installing Symantec product license keys" on page 53.
  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: To change from one product group to another, you may need to perform additional steps.

Setting or changing the product level for keyless licensing

The keyless licensing method uses product levels to determine the Symantec products and functionality that are licensed.

For more information to use keyless licensing 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.

When you upgrade from a previous release, the product installer prompts you to update the `vxkeyless` license product level to the current release level. If you update the `vxkeyless` license product level during the upgrade process, no further action is required. If you do not update the `vxkeyless` license product level, the output you see when you run the `vxkeyless display` command includes the previous release's `vxkeyless` license product level. Each `vxkeyless` license product level name includes the suffix *_previous_release_version*. For example, DMP_6.0, or SFENT_VR_5.1SP1, or VCS_GCO_5.1. If there is no suffix, it is the current release version.

You would see the suffix *_previous_release_version* if you did not update the `vxkeyless` product level when prompted by the product installer. Symantec highly recommends that you always use the current release version of the product levels. To do so, use the `vxkeyless set` command with the desired product levels. If you see SFENT_60, VCS_60, use the `vxkeyless set SFENT,VCS` command to update the product levels to the current release.

After you install or upgrade, 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. Change your current working directory:

   ```
   # export PATH=$PATH:/opt/VRTSvlic/bin
   ```

2. View the current setting for the product level.

   ```
   # vxkeyless -v display
   ```

3. View the possible settings for the product level.

   ```
   # vxkeyless displayall
   ```

4. Set the desired product level.

   ```
   # vxkeyless set prod_levels
   ```

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

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

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

See "Installing Symantec product license keys" on page 53.

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 Symantec product license keys

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

`vxlicinst`	Installs a license key for a Symantec product
`vxlicrep`	Displays the currently installed licenses
`vxlictest`	Retrieves the features and their descriptions that are 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 or change a license

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

   ```
   # cd /opt/VRTS/bin

   # ./vxlicinst -k license key
   ```

2. Run the following Veritas Volume Manager (VxVM) command to recognize the new license:

   ```
   # vxdctl license init
   ```

   See the `vxdctl(1M)` manual page.

If you have `vxkeyless` licensing, you can view or update the keyless product licensing levels.

See "Setting or changing the product level for keyless licensing" on page 51.

You can install SFHA if you install a pair of valid VCS and SF keys. Even if your VCS keys and SF keys do not show when you run the `vxkeyless display` command, you can still install and configure SFHA.

Section 2

Preinstallation tasks

- Chapter 5. Preparing to install Storage Foundation High Availability

Chapter 5

Preparing to install Storage Foundation High Availability

This chapter includes the following topics:

- Installation preparation overview
- About using ssh or rsh with the installer
- Setting up the private network
- Setting up shared storage
- Setting environment variables
- Setting the kernel.hung_task_panic tunable
- Optimizing LLT media speed settings on private NICs
- Guidelines for setting the media speed of the LLT interconnects
- Mounting the product disc
- Assessing the system for installation readiness

Installation preparation overview

Table 5-1 provides an overview of an installation using the product installer.

Table 5-1 Installation overview

Installation task	Section
Obtain product licenses.	See "About Symantec product licensing" on page 50.
Download the software, or insert the product DVD.	See "Downloading the Storage Foundation and High Availability software" on page 48. See "Mounting the product disc" on page 65.
Set environment variables.	See "Setting environment variables" on page 63.
Configure the Secure Shell (ssh) or Remote Shell (rsh) on all nodes.	See "About configuring secure shell or remote shell communication modes before installing products" on page 491.
Verify that hardware, software, and operating system requirements are met.	See "Release notes" on page 34.
Check that sufficient disk space is available.	See "Disk space requirements" on page 35.
Use the installer to install the products.	See "About the script-based installer" on page 69.

About using ssh or rsh with the installer

The installer uses passwordless Secure Shell (ssh) or Remote Shell (rsh) communications among systems. The installer uses the ssh daemon or rsh daemon that comes bundled with the operating system. During an installation, you choose the communication method that you want to use. Or, you can run the `installer -comsetup` command to set up ssh or rsh explicitly. You then provide the installer with the superuser passwords for the systems where you plan to install. When the installation process completes, the installer asks you if you want to remove the password-less connection. If installation terminated abruptly, use the installation script's `-comcleanup` option to remove the ssh configuration or rsh configuration from the systems.

See "Installation script options" on page 451.

In most installation, configuration, upgrade (where necessary), and uninstallation scenarios, the installer can configure ssh or rsh on the target systems. In the following scenarios, you need to set up ssh or rsh manually, or use the `installer -comsetup` option to set up an ssh or rsh configuration from the systems.

- When you perform installer sessions using a response file.

See “About configuring secure shell or remote shell communication modes before installing products” on page 491.

Setting up the private network

VCS requires you to set up a private network between the systems that form a cluster. You can use either NICs or aggregated interfaces to set up private network.

You can use network switches instead of hubs.

Refer to the *Symantec Cluster Server Administrator's Guide* to review VCS performance considerations.

Figure 5-1 shows two private networks for use with VCS.

Figure 5-1 Private network setups: two-node and four-node clusters

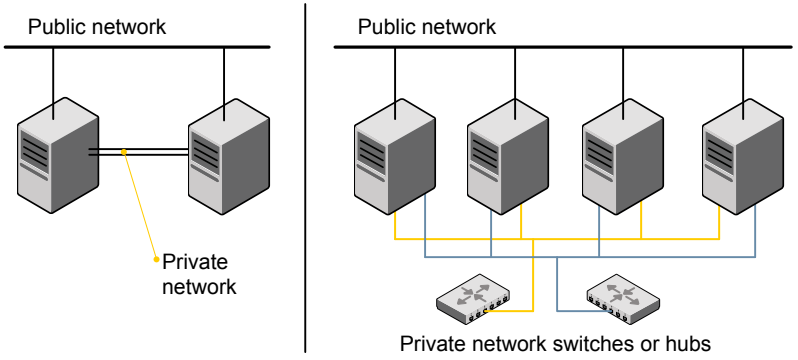

You need to configure at least two independent networks between the cluster nodes with a network switch for each network. You can also interconnect multiple layer 2 switches for advanced failure protection. Such connections for LLT are called cross-links.

Figure 5-2 shows a private network configuration with crossed links between the network switches.

Figure 5-2 Private network setup with crossed links

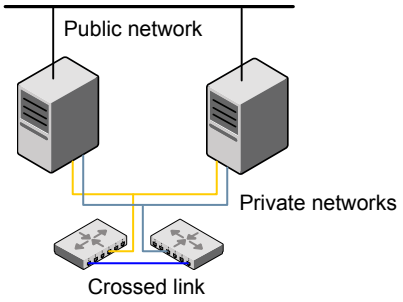

Symantec recommends one of the following two configurations:

- Use at least two private interconnect links and one public link. The public link can be a low priority link for LLT. The private interconnect link is used to share cluster status across all the systems, which is important for membership arbitration and high availability. The public low priority link is used only for heartbeat communication between the systems.
- If your hardware environment allows use of only two links, use one private interconnect link and one public low priority link. If you decide to set up only two links (one private and one low priority link), then the cluster must be configured to use I/O fencing, either disk-based or server-based fencing configuration. With only two links, if one system goes down, I/O fencing ensures that other system can take over the service groups and shared file systems from the failed node.

To set up the private network

1. Install the required network interface cards (NICs).

   Create aggregated interfaces if you want to use these to set up private network.

2. Connect the SFHA private NICs on each system.
3. Use crossover Ethernet cables, switches, or independent hubs for each SFHA communication network. Note that the crossover Ethernet cables are supported only on two systems.

   Ensure that you meet the following requirements:

   - The power to the switches or hubs must come from separate sources.
   - On each system, you must use two independent network cards to provide redundancy.
   - If a network interface is part of an aggregated interface, you must not configure the network interface under LLT. However, you can configure the aggregated interface under LLT.

- When you configure Ethernet switches for LLT private interconnect, disable the spanning tree algorithm on the ports used for the interconnect.

During the process of setting up heartbeat connections, consider a case where a failure removes all communications between the systems.

Note that a chance for data corruption exists under the following conditions:

- The systems still run, and
- The systems can access the shared storage.

4 Test the network connections. Temporarily assign network addresses and use `telnet` or `ping` to verify communications.

LLT uses its own protocol, and does not use TCP/IP. So, you must ensure that the private network connections are used only for LLT communication and not for TCP/IP traffic. To verify this requirement, unplumb and unconfigure any temporary IP addresses that are configured on the network interfaces.

The installer configures the private network in the cluster during configuration.

You can also manually configure LLT.

Setting up shared storage

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

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

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

See “ About planning to configure I/O fencing” on page 75.

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

Setting up shared storage: SCSI

Perform the following steps to set up shared storage.

To set up shared storage

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

   Set the following:

   - Set Host adapter SCSI ID = 7, or to an appropriate value for your configuration.
   - Set Host Adapter BIOS in Advanced Configuration Options to Disabled.
6. Format the shared disk and create required partitions on it.

   Perform the following:

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

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

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

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

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

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

     For example, enter the following command:

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

     Where the name of the disk group is dg, the name of the volume is vol01, and the file system type is vxfs.
7. Power off the disk.
8. Remove the terminator from the disk and connect the disk to the other cluster system.
9. Power on the disk.

10 Boot the second system. The system can now detect the disk.

11 Press Ctrl+A to bring up the SCSI BIOS settings for the disk.

Set the following:

- Set Host adapter SCSI ID = 6, or to an appropriate value for your configuration. Note that the SCSI ID should be different from the one configured on the first cluster system.
- Set Host Adapter BIOS in Advanced Configuration Options to Disabled.

12 Verify that you can view the shared disk using the `fdisk` command.

Setting up shared storage: Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up shared storage for Fibre Channel

1 Connect the Fibre Channel disk to a cluster system.

2 Boot the system and change the settings of the Fibre Channel. Perform the following tasks for all QLogic adapters in the system:

- Press Alt+Q to bring up the QLogic adapter settings menu.
- Choose **Configuration Settings**.
- Click Enter.
- Choose **Advanced Adapter Settings**.
- Click Enter.
- Set the Enable Target Reset option to **Yes** (the default value).
- Save the configuration.
- Reboot the system.

3 Verify that the system detects the Fibre Channel disks properly.

4 Create volumes. Format the shared disk and create required partitions on it and perform the following:

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

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

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

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

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

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

  For example, enter the following command:

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

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

5 Repeat step 2 and step 3 for all nodes in the clusters that require connections with Fibre Channel.

6 Power off this cluster system.

7 Connect the same disks to the next cluster system.

8 Turn on the power for the second system.

9 Verify that the second system can see the disk names correctly—the disk names should be the same.

See "Verifying that the nodes have access to the same disk" on page 141.

Setting environment variables

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

After installation, SFHA commands are in `/opt/VRTS/bin`. SFHA manual pages are stored in `/opt/VRTS/man`.

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

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

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

- If you want to use a shell such as `sh` or `bash`, enter the following:

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

- If you want to use a shell such as `csh` or `tcsh`, enter the following:

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

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

- If you want to use a shell such as `sh` or `bash`, enter the following:

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

- If you want to use a shell such as `csh` or `tcsh`, enter the following:

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

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

Setting the kernel.hung_task_panic tunable

By default, in the Linux kernel the `kernel.hung_task_panic` tunable is enabled and the `kernel.hung_task_timeout_secs` tunable is set to a default non-zero value.

To ensure that the node does not panic, the `kernel.hung_task_panic` tunable must be disabled. If `kernel.hung_task_panic` is enabled, then it causes the kernel to panic when any of the following kernel threads waits for more than the `kernel.hung_task_timeout_secs` value:

- The vxfenconfig thread in the vxfen configuration path waits for GAB to seed.
- The vxfenswap thread in the online coordinator disks replacement path waits for the snapshot of peer nodes of the new coordinator disks.
- The vxfs_thread waits for a reconfiguration event to happen in the cluster.
- The vxglm_thread waits for GAB to seed.

To disable the `kernel.hung_task_panic` tunable:

- Set the `kernel.hung_task_panic` tunable to zero (0) in the `/etc/sysctl.conf` file. This step ensures that the change is persistent across node restarts.
- Run the command on each node.

  # `sysctl -w kernel.hung_task_panic=0`

To verify the `kernel.hung_task_panic tunable` value, run the following command:

- # `sysctl -a | grep hung_task_panic`

Optimizing LLT media speed settings on private NICs

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

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

Guidelines for setting the media speed of the LLT interconnects

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

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

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

Mounting the product disc

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

To mount the product disc

1. Log in as superuser on a system where you want to install SFHA.

   The system from which you install SFHA does not need to be part of the cluster. The systems must be in the same subnet.

2. Insert the product disc with the SFHA software into a drive that is connected to the system.

   The disc is automatically mounted.

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

   ```
   # mount -t iso9660 /dev/sr0 /media
   ```

4. Navigate to the location of the RPMs.

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

   Where *dist* is rhel6, rhel7 or sles11, and *arch* is x86_64 for RHEL and SLES.

Assessing the system for installation readiness

Symantec provides the following tools for assessing your system, to ensure that the system meets the requirements for installing Symantec Storage Foundation 6.2.

Symantec Operations Readiness Tools	Symantec Operations Readiness Tools (SORT) is a web-based application that is designed to support Symantec enterprise products. See "About Symantec Operations Readiness Tools" on page 28.
Prechecking your systems using the installer	Performs a preinstallation check on the specified systems. The product installer reports whether the specified systems meet the minimum requirements for installing Symantec Storage Foundation 6.2. See "Prechecking your systems using the installer" on page 66.

Prechecking your systems using the installer

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

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

To use the precheck option

1. Start the script-based or web-based installer.

   See "Installing Storage Foundation and High Availability using the script-based installer" on page 71.

   See "Installing SFHA with the web-based installer" on page 173.

2. Select the precheck option:

   - From the web-based installer, select the **Perform a Pre-Installation Check** from the Task pull-down menu.
   - In the script-based installer, from root on the system where you want to perform the check, start the installer.

     ```
     # ./installer
     ```

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

3. Enter the system name or the IP address of the system that you want to check.
4. Review the output and make the changes that the installer recommends.

Section 3

Installation using the script-based installer

- Chapter 6. Installing SFHA
- Chapter 7. Preparing to configure SFHA clusters for data integrity
- Chapter 8. Configuring SFHA
- Chapter 9. Manually configuring SFHA clusters for data integrity

Chapter 6

Installing SFHA

This chapter includes the following topics:

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

About the script-based installer

You can use the script-based installer to install Symantec products (version 6.1 and later) from a driver system that runs any supported platform to a target system that runs different supported platforms.

To install your Symantec product, use one of the following methods:

- The general product installer (`installer`). The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
  See "Installing Storage Foundation and High Availability using the script-based installer" on page 71.
- Product-specific installation scripts (`installsfha`). The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installsfha` script is identical to running the general product installer and specifying SFHA from the list of products to install. Use the product-specific installation scripts to install or configure individual products you download electronically. You can find these scripts at the root of the product media. These scripts are also installed with the product.

Table 6-1 lists all the SFHA Solutions product installation scripts. The list of product-specific installation scripts that you find on your system depends on the product that you install on your system.

Table 6-1 Product installation scripts

Symantec product name	Script name in the media	Script name after an installation
For all SFHA Solutions products	`installer`	N/A
Symantec ApplicationHA	`installapplicationha`	`installapplicationha<version>`
Symantec Cluster Server (VCS)	`installvcs`	`installvcs<version>`
Symantec Storage Foundation (SF)	`installsf`	`installsf<version>`
Symantec Storage Foundation and High Availability (SFHA)	`installsfha`	`installsfha<version>`
Symantec Storage Foundation Cluster File System High Availability (SFCFSHA)	`installsfcfsha`	`installsfcfsha<version>`
Symantec Storage Foundation for Oracle RAC (SF Oracle RAC)	`installsfrac`	`installsfrac<version>`
Symantec Dynamic Multi-pathing (DMP)	`installdmp`	`installdmp<version>`

When you install from the installation media, the script name does not include a product version.

When you configure the product after an installation, the installation scripts include the product version in the script name.

For example, for the 6.2 version:

```
# /opt/VRTS/install<productname>62 -configure
```

Note: The general product installer (`installer`) script does not include the product version.

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

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

See "Installation script options" on page 451.

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

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

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

To install Storage Foundation and High Availability

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

   See "About configuring secure shell or remote shell communication modes before installing products" on page 491.

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

   See "Mounting the product disc" on page 65.

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

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

   Where *dist* is rhel6, rhel7 or sles11, and *arch* is x86_64 for RHEL and SLES.

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

   ```
   # ./installer
   ```

5 Press I to install and press Enter.

6 When the list of available products is displayed, select Storage Foundation and High Availability, enter the corresponding number, and press Enter.

7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).

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

8 Select from one of the following installation options:

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

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

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

```
Enter the system names separated by spaces:
[q,?] sys1 sys2
```

10 After the system checks complete, the installer displays a list of the RPMs to be installed. Press Enter to continue with the installation.

11 If the coomunication fails during the precheck, the installer can configure remote shell or secure shell communications for you among systems, however each system needs to have rsh or ssh servers installed. You also need to provide the superuser passwords for the systems. Note that for security reasons, the installation program neither stores nor caches these passwords.

12 You need to synchronize the system clocks of your application servers or have them point to an NTP server. After the system check, if the nodes have time difference, the installer prompts:

```
Do you want to synchronize system clock with NTP server(s)?
[y,n,q] (y)
Enter the NTP server names separated by spaces: [b] megami.veritas.com

 Synchronizing system clock on sys1 ..... Done
 Synchronizing system clock on sys2..... Done

System clock synchronized on systems
```

13 The installer may prompt to restore previous Veritas Volume Manager configurations.

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

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

---

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

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

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

16 If you selects product licensing mode as 2 (SF Enterprise), the installer prompts you to decide to enable replication or not:

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

Enter your option.

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

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

17 If Veritas Volume Manager (VxVM) is started and the installer detects the presence of a Solid State Drive (SSD) device, the installer displays the following message:

```
SSD devices have been detected on systemname.
It is strongly recommended that you use the SmartIO feature to
accelerate I/O performance.  See the Storage Foundation and
High Availability Solutions documentation for more information
on using the SmartIO feature.
```

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

```
Installation procedures and diagnostic information were saved in the
log files under directory /var/tmp/installer-<platform>-<uuid>.
Analyzing this information helps Symantec discover and
fix failed operations performed by the installer.
Would you like to send the information about this installation to
Symantec to help improve installation in the future? [y,n,q,?] (y) y
```

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

Chapter 7

Preparing to configure SFHA clusters for data integrity

This chapter includes the following topics:

- About planning to configure I/O fencing
- Setting up the CP server

About planning to configure I/O fencing

After you configure SFHA with the installer, you must configure I/O fencing in the cluster for data integrity. Application clusters on release version 6.2 (HTTPS-based communication) only support CP servers on release version 6.1 and later.

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

The coordination points in server-based fencing can include only CP servers or a mix of CP servers and coordinator disks.

Symantec also supports server-based fencing with a single coordination point which is a single highly available CP server that is hosted on an SFHA cluster.

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

You use majority fencing mechanism if you do not want to use coordination points to protect your cluster. Symantec recommends that you configure I/O fencing in majority mode if you have a smaller cluster environment and you do not want to invest additional disks or servers for the purposes of configuring fencing.

Note: Majority-based I/O fencing is not as robust as server-based or disk-based I/O fencing in terms of high availability. With majority-based fencing mode, in rare cases, the cluster might become unavailable.

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

See Figure 7-2 on page 78.

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

Figure 7-1 Workflow to configure I/O fencing

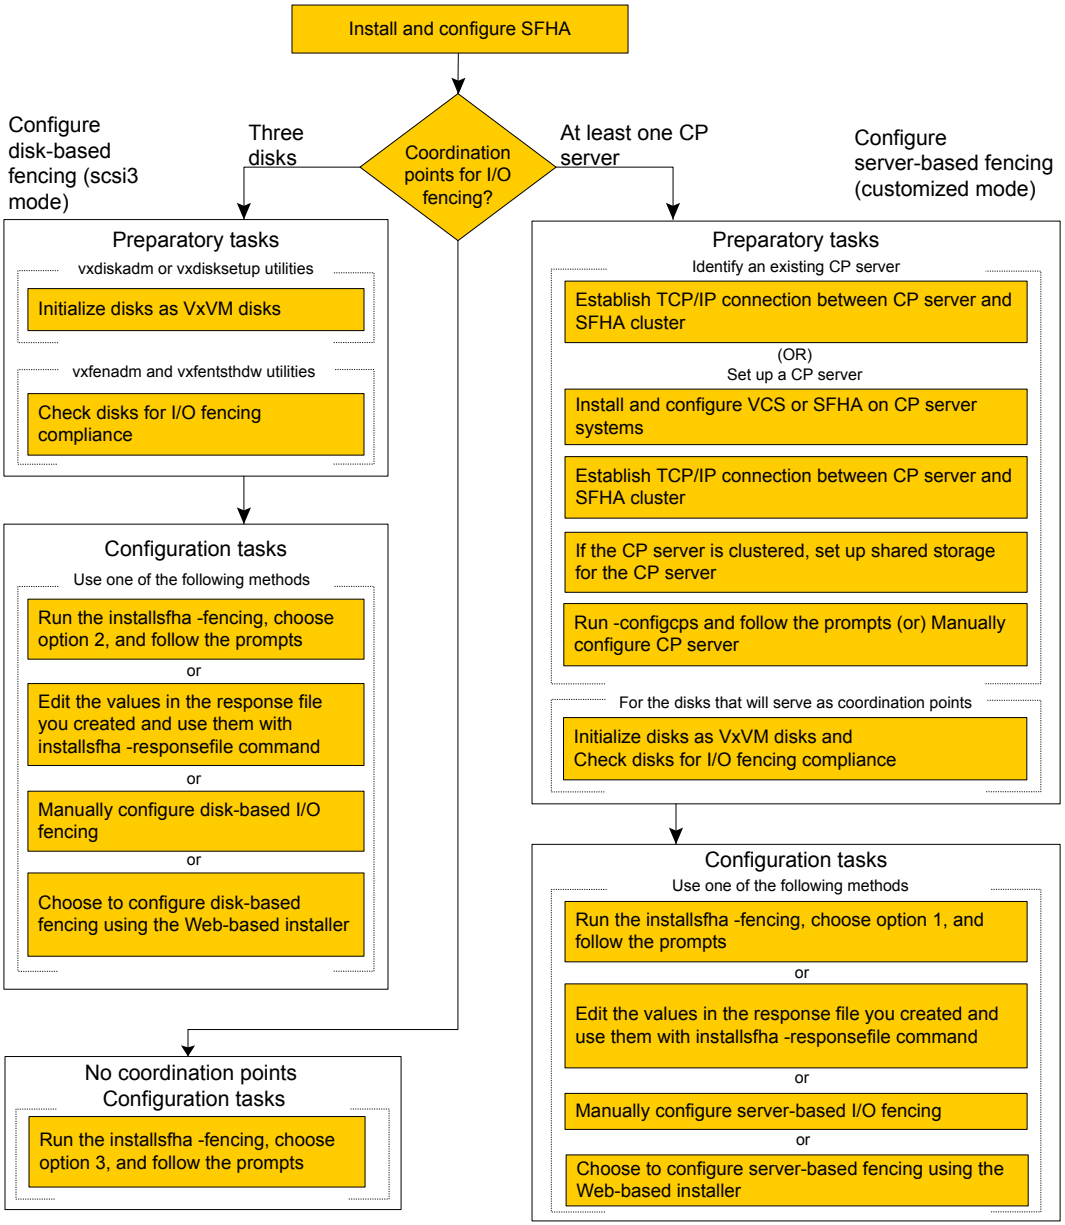

Figure 7-2 illustrates a high-level flowchart to configure non-SCSI-3 I/O fencing for the SFHA cluster in virtual environments that do not support SCSI-3 PR.

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

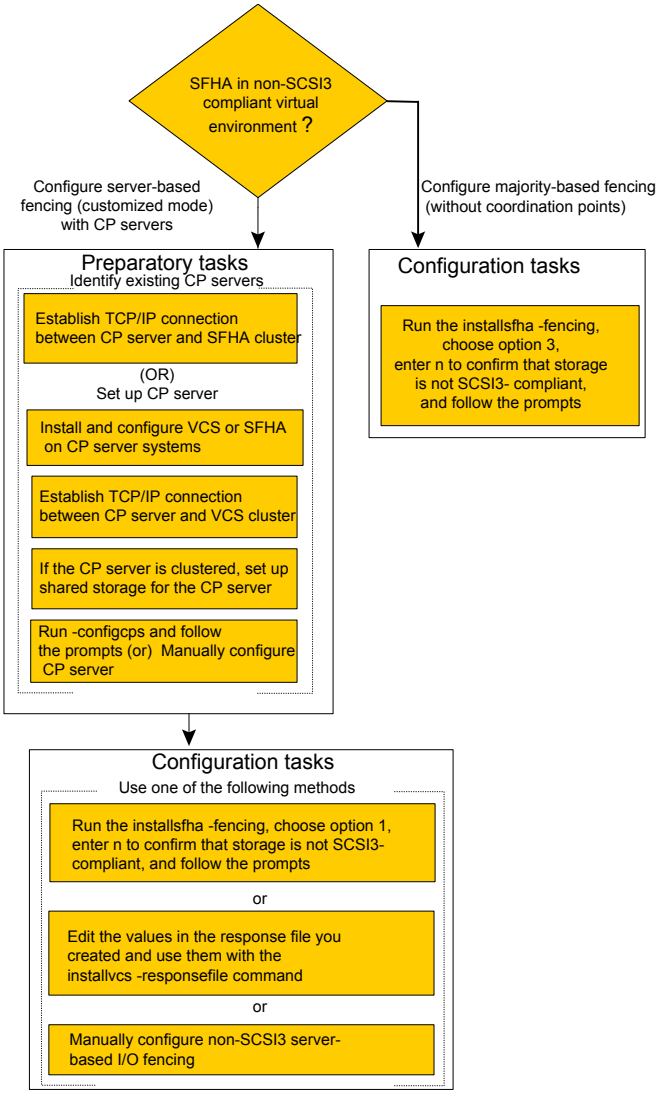

After you perform the preparatory tasks, you can use any of the following methods to configure I/O fencing:

Using the installsfha	See "Setting up disk-based I/O fencing using installsfha" on page 139. See "Setting up server-based I/O fencing using installsfha" on page 148. See "Setting up non-SCSI-3 I/O fencing in virtual environments using installsfha" on page 161. See "Setting up majority-based I/O fencing using installsfha" on page 163.
Using the web-based installer	See "Configuring SFHA for data integrity using the web-based installer" on page 181.
Using response files	See "Response file variables to configure disk-based I/O fencing" on page 214. See "Response file variables to configure server-based I/O fencing" on page 217. See "Response file variables to configure non-SCSI-3 I/O fencing" on page 220. See "Response file variables to configure majority-based I/O fencing" on page 222. See "Configuring I/O fencing using response files" on page 213.
Manually editing configuration files	See "Setting up disk-based I/O fencing manually" on page 238. See "Setting up server-based I/O fencing manually" on page 244. See "Setting up non-SCSI-3 fencing in virtual environments manually" on page 258. See "Setting up majority-based I/O fencing manually " on page 264.

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

See the *Symantec Storage foundation High Availability Administrator's Guide* for more details.

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

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

Figure 7-3 CP server, Storage Foundation High Availability, and coordinator disks

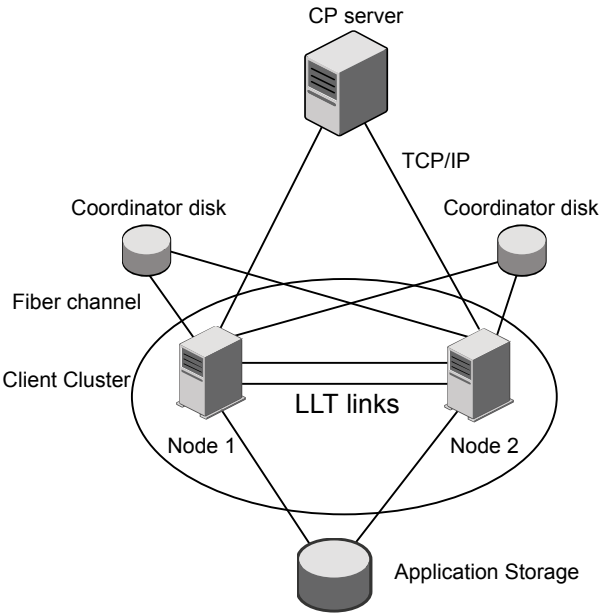

Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points
  See Figure 7-4 on page 81.
- Multiple application clusters use a single CP server and single or multiple pairs of coordinator disks (two) as their coordination points
  See Figure 7-5 on page 82.
- Multiple application clusters use a single CP server as their coordination point
  This single coordination point fencing configuration must use a highly available CP server that is configured on an SFHA cluster as its coordination point.
  See Figure 7-6 on page 82.

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

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

Figure 7-4 displays a configuration using three CP servers that are connected to multiple application clusters.

Figure 7-4 Three CP servers connecting to multiple application clusters

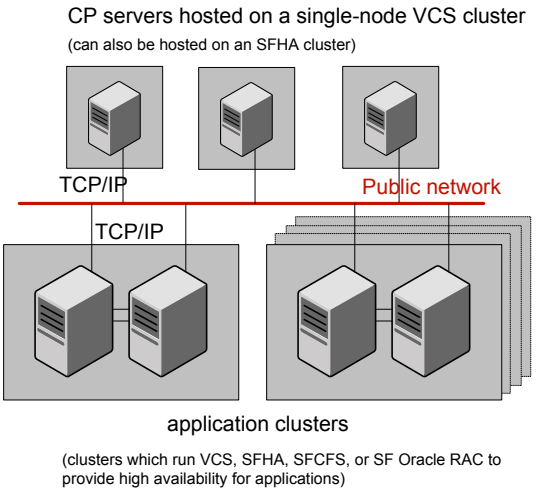

Figure 7-5 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.

Figure 7-5 Single CP server with two coordinator disks for each application cluster

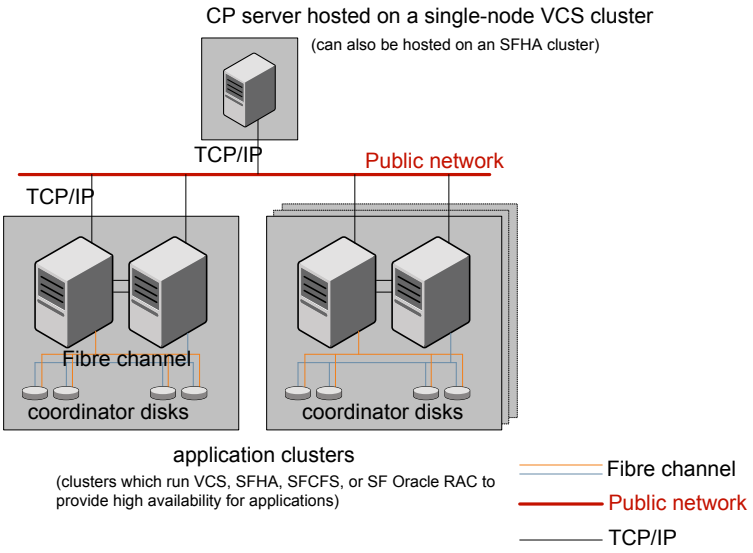

Figure 7-6 displays a configuration using a single CP server that is connected to multiple application clusters.

Figure 7-6 Single CP server connecting to multiple application clusters

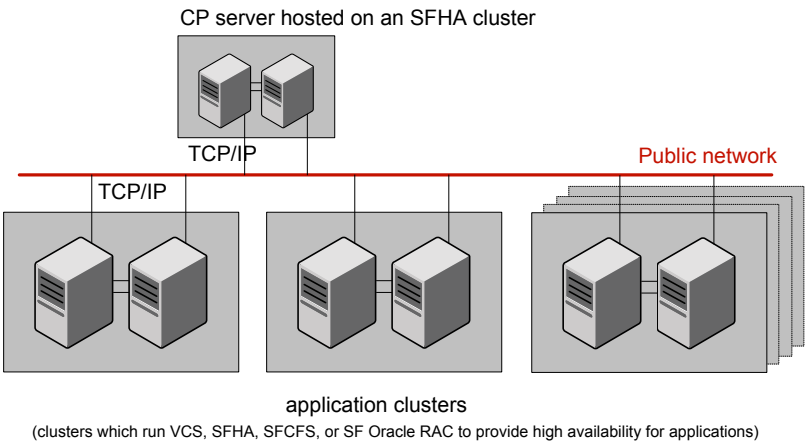

See "Configuration diagrams for setting up server-based I/O fencing" on page 522.

Setting up the CP server

Table 7-1 lists the tasks to set up the CP server for server-based I/O fencing.

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

Task	Reference
Plan your CP server setup	See "Planning your CP server setup" on page 83.
Install the CP server	See "Installing the CP server using the installer" on page 85.
Configure the CP server cluster in secure mode	See "Configuring the CP server cluster in secure mode" on page 85.
Set up shared storage for the CP server database	See "Setting up shared storage for the CP server database" on page 86.
Configure the CP server	See " Configuring the CP server using the installer program" on page 87. See "Configuring the CP server using the web-based installer" on page 111. See "Configuring the CP server manually" on page 99. See "Configuring CP server using response files" on page 105.
Verify the CP server configuration	See "Verifying the CP server configuration" on page 110.

Planning your CP server setup

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

To plan your CP server setup

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

   Symantec recommends hosting the CP server on an SFHA cluster to make the CP server highly available.

2. If you host the CP server on an SFHA cluster, review the following information. Make sure you make the decisions and meet these prerequisites when you set up the CP server:

- You must set up shared storage for the CP server database during your CP server setup.
- Decide whether you want to configure server-based fencing for the Storage Foundation High Availability (application cluster) with a single CP server as coordination point or with at least three coordination points.
  Symantec recommends using at least three coordination points.

3 Decide whether you want to configure the CP server cluster for IPM-based communication or HTTPS communication or both.

  For IPM-based communication, the CP server on release 6.1 and later supports clients prior to 6.1 release. When you configure the CP server, you are required to provide VIPs for IPM-based clients.

  For HTTPS-based communication, the CP server on release 6.1 and later only supports clients on release 6.1 and later.

4 Decide whether you want to configure the CP server cluster in secure mode for IPM-based communication.

  Symantec recommends configuring the CP server cluster in secure mode for IPM-based secure communication between the CP server and its clients (SFHA clusters). Note that you use IPM-based communication if you want the CP server to support clients that are installed with a release version prior to 6.1 release.

5 Set up the hardware and network for your CP server.

  See “CP server requirements” on page 40.

6 Have the following information handy for CP server configuration:

- Name for the CP server
  The CP server name should not contain any special characters. CP server name can include alphanumeric characters, underscore, and hyphen.
- Port number for the CP server
  Allocate a TCP/IP port for use by the CP server.
  Valid port range is between 49152 and 65535. The default port number for HTTPS-based communication is 443 and for IPM-based secure communication is 14250.
- Virtual IP address, network interface, netmask, and networkhosts for the CP server
  You can configure multiple virtual IP addresses for the CP server.

Installing the CP server using the installer

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

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

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

CP server setup uses a single system	Install and configure VCS to create a single-node VCS cluster. During installation of VCS 6.2, VRTScps will come under recommended set of RPMs. See the *Symantec Cluster Server Installation Guide* for instructions on installing and configuring VCS. Proceed to configure the CP server. See " Configuring the CP server using the installer program" on page 87. See "Configuring the CP server manually" on page 99.
CP server setup uses multiple systems	Install and configure SFHA to create an SFHA cluster. This makes the CP server highly available. Meet the following requirements for CP server: ■ During installation of SFHA 6.2, VRTScps will come under recommended set of RPMs. Proceed to set up shared storage for the CP server database.

Configuring the CP server cluster in secure mode

You must configure security on the CP server only if you want IPM-based (Symantec Product Authentication Service) secure communication between the CP server and the SFHA cluster (CP server clients). However, IPM-based communication enables the CP server to support application clusters prior to release 6.1.

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

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

To configure the CP server cluster in secure mode

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

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

  ```
  # /opt/VRTS/install/installvcs<version>  -security
  ```

  Where *<version>* is the specific release version.

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

  ```
  # /opt/VRTS/install/installsfha<version>  -security
  ```

  Where *<version>* is the specific release version.

  See "About the script-based installer" on page 69.

Setting up shared storage for the CP server database

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

The installer can set up shared storage for the CP server database when you configure CP server for the SFHA cluster.

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

To set up shared storage for the CP server database

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

   For example:

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

2. Create a mirrored volume over the disk group.

   For example:

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

3. Create a file system over the volume.

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

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

AIX	`# mkfs -V vxfs /dev/vx/rdsk/cps_dg/cps_volume`
Linux	`# mkfs -t vxfs /dev/vx/rdsk/cps_dg/cps_volume`
Solaris	`# mkfs -F vxfs /dev/vx/rdsk/cps_dg/cps_volume`

Configuring the CP server using the installer program

Use the configcps option available in the installer program to configure the CP server.

Perform one of the following procedures:

For CP servers on single-node VCS cluster:	See "To configure the CP server on a single-node VCS cluster" on page 88.
For CP servers on an SFHA cluster:	See "To configure the CP server on an SFHA cluster" on page 93.

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

1. Verify that the `VRTScps` RPM is installed on the node.
2. Run the installvcs<*version*> program with the configcps option.

   ```
   # /opt/VRTS/install/installvcs<version> -configcps
   ```

   Where `<version>` is the specific release version.

   See "About the script-based installer" on page 69.
3. Installer checks the cluster information and prompts if you want to configure CP Server on the cluster.

   Enter **y** to confirm.
4. Select an option based on how you want to configure Coordination Point server.

   ```
   1) Configure Coordination Point Server on single node VCS system
   2) Configure Coordination Point Server on SFHA cluster
   3) Unconfigure Coordination Point Server
   ```
5. Enter the option: [1-3,q] **1**.

   The installer then runs the following preconfiguration checks:

   - Checks to see if a single-node VCS cluster is running with the supported platform.
     The CP server requires VCS to be installed and configured before its configuration.

   The installer automatically installs a license that is identified as a CP server-specific license. It is installed even if a VCS license exists on the node. CP server-specific key ensures that you do not need to use a VCS license on the single-node. It also ensures that Veritas Operations Manager (VOM) identifies the license on a single-node coordination point server as a CP server-specific license and not as a VCS license.
6. Restart the VCS engine if the single-node only has a CP server-specific license.

   ```
   A single node coordination point server will be configured and
   VCS will be started in one node mode, do you want to
   continue?  [y,n,q]  (y)
   ```

7 Communication between the CP server and application clusters is secured by HTTPS from release 6.1.0 onwards. However, clusters on earlier release versions (prior to 6.1.0) that are using IPM-based communication are still supported.

Enter the name of the CP Server.

```
Enter the name of the CP Server: [b]   cps1
```

8 Enter valid virtual IP addresses for the CP Server with HTTPS-based secure communication. A CP Server can be configured with more than one virtual IP address. For HTTPS-based communication, only IPv4 addresses are supported. For IPM-based communication, both IPv4 and IPv6 addresses are supported.

```
Enter Virtual IP(s) for the CP server for HTTPS,
separated by a space:  [b]  10.200.58.231 10.200.58.232
10.200.58.233
```

Note: Ensure that the virtual IP address of the CP server and the IP address of the NIC interface on the CP server belongs to the same subnet of the IP network. This is required for communication to happen between client nodes and CP server.

---

9 Enter the corresponding CP server port number for each virtual IP address or press **Enter** to accept the default value (443).

```
Enter the default port '443' to be used for all the
virtual IP addresses for HTTPS communication or assign the
corresponding port number in the range [49152, 65535] for
each virtual IP address. Ensure that each port number is
separated by a single
space: [b]  (443) 54442 54443 54447
```

10 Decide if you want to support clusters that are on releases prior to 6.1.0. These clusters use the Symantec Product Authentication Services (AT) (secure IPM-based protocol) to securely communicate with the CP servers.

```
Do you want to support older (prior to 6.1.0)
clusters? [y,n,q,b]  (y)
```

11 Enter virtual IPs for the CP Server for IPM-based secure communication.

```
Enter Virtual IP(s) for the CP server for IPM,
separated by a space [b]  10.182.36.8 10.182.36.9
```

Note that both IPv4 and IPv6 addresses are supported.

12 Enter corresponding port number for each Virtual IP address or accept the default port.

```
Enter the default port '14250' to be used for all the
virtual IP addresses for IPM-based communication, or assign
the corresponding port number in the range [49152, 65535]
for each virtual IP address.
Ensure that each port number is separated by a single space:
[b] (14250) 54448 54449
```

13 Decide if you want to enable secure communication between the CP server and application clusters.

```
Symantec 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,q,b] (y) n
```

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

```
Enter absolute path of the database: [b] (/etc/VRTScps/db)
```

15 Verify and confirm the CP server configuration information.

```
CP Server configuration verification:
----------------------------------------------
CP Server Name:  cps1
CP Server Virtual IP(s) for HTTPS: 10.200.58.231, 10.200.58.232,
10.200.58.233
CP Server Virtual IP(s) for IPM: 10.182.36.8, 10.182.36.9
CP Server Port(s) for HTTPS: 54442, 54443, 54447
CP Server Port(s) for IPM: 54448, 54449
CP Server Security for IPM: 0
CP Server Database Dir: /etc/VRTScps/db

----------------------------------------------

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

16 The installer proceeds with the configuration process, and creates a vxcps.conf configuration file.

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

17 Configure the CP Server Service Group (CPSSG) for this cluster.

```
Enter how many NIC resources you want to configure (1 to 2): 2
```

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

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

```
Enter a valid network interface on sys1 for NIC resource - 1: eth0
Enter a valid network interface on sys1 for NIC resource - 2:  eth1
```

19 Enter the NIC resource you want to associate with the virtual IP addresses.

```
Enter the NIC resource you want to associate with the virtual IP 10.200.58.231 (1 to 2): 1
Enter the NIC resource you want to associate with the virtual IP 10.200.58.232 (1 to 2): 2
```

20 Enter the networkhosts information for each NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure NIC resource
to be always online

Do you want to add NetworkHosts attribute for the NIC device eth0
on system sys1? [y,n,q] y
Enter a valid IP address to configure NetworkHosts for NIC eth0
on system sys1: 10.200.56.22

Do you want to add another Network Host? [y,n,q] n
```

21 Enter the netmask for virtual IP addresses. If you entered an IPv6 address, enter the prefix details at the prompt.

Note that if you are using HTTPS-based communication, only IPv4 addresses are supported.

```
Enter the netmask for virtual IP for
HTTPS 192.169.0.220: (255.255.252.0)
Enter the netmask for virtual IP for
IPM 192.169.0.221: (255.255.252.0)
```

22 Installer displays the status of the Coordination Point Server configuration. After the configuration process has completed, a success message appears.

```
For example:
Updating main.cf with CPSSG service group.. Done
Successfully added the CPSSG service group to VCS configuration.
Trying to bring CPSSG service group
ONLINE and will wait for upto 120 seconds

The Symantec coordination point server is ONLINE

The Symantec coordination point server has
been configured on your system.
```

23 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.... |ONLINE|
```

It also generates the configuration file for CP server (/etc/vxcps.conf). The vxcpserv process and other resources are added to the VCS configuration in the CP server service group (CPSSG).

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

To configure the CP server on an SFHA cluster

1 Verify that the `VRTScps` RPM is installed on each node.

2 Ensure that you have configured passwordless ssh or rsh on the CP server cluster nodes.

3 Run the installsfha<*version*> program with the configcps option.

```
# ./installsfha<version> -configcps
```

Where `<version>` is the specific release version.

See "About the script-based installer" on page 69.

4 Installer checks the cluster information and prompts if you want to configure CP Server on the cluster.

Enter **y** to confirm.

5 Select an option based on how you want to configure Coordination Point server.

```
1)  Configure Coordination Point Server on single node VCS system
2)  Configure Coordination Point Server on SFHA cluster
3)  Unconfigure Coordination Point Server
```

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

The installer then runs the following preconfiguration checks:

- Checks to see if an SFHA cluster is running with the supported platform. The CP server requires SFHA to be installed and configured before its configuration.

7 Communication between the CP server and application clusters is secured by HTTPS from Release 6.1.0 onwards. However, clusters on earlier release versions (prior to 6.1.0) that are using IPM-based communication are still supported.

Enter the name of the CP server.

```
Enter the name of the CP Server: [b]  cps1
```

8 Enter valid virtual IP addresses for the CP Server. A CP Server can be configured with more than one virtual IP address. For HTTPS-based communication, only IPv4 addresses are supported. For IPM-based communication, both IPv4 and IPv6 addresses are supported

```
Enter Virtual IP(s) for the CP server for HTTPS,
 separated by a space: [b] 10.200.58.231 10.200.58.232 10.200.58.233
```

9 Enter the corresponding CP server port number for each virtual IP address or press Enter to accept the default value (443).

```
Enter the default port '443' to be used for all the virtual IP addresses
for HTTPS communication or assign the corresponding port number in the range [49152,
65535] for each virtual IP address. Ensure that each port number is separated by
a single space: [b] (443) 65535 65534 65537
```

10 Decide if you want to support clusters that are on releases prior to 6.1.0. These clusters use the Symantec Product Authentication Services (AT) (secure IPM-based protocol) to securely communicate with the CP servers.

```
Do you want to support older (prior to 6.1.0) clusters? [y,n,q,b] (y)
```

11 Enter Virtual IPs for the CP Server for IPM-based secure communication. Both IPv4 and IPv6 addresses are supported.

```
Enter Virtual IP(s) for the CP server for IPM, separated by a space:
[b] 10.182.36.8 10.182.36.9
```

12 Enter corresponding port number for each Virtual IP address or accept the default port.

```
Enter the default port '14250' to be used for all
the virtual IP addresses for IPM-based communication, or
assign the corresponding port number in the range [49152, 65535]
for each virtual IP address.
Ensure that each port number is separated by a single space:
[b]  (14250) 54448 54449
```

13 Decide if you want to enable secure communication between the CP server and application clusters.

```
Symantec 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,q,b] (y)
```

14 Enter absolute path of the database.

```
CP Server uses an internal database to store the client information.
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: [b] /cpsdb
```

15 Verify and confirm the CP server configuration information.

```
CP Server configuration verification:

CP Server Name: cps1
CP Server Virtual IP(s) for HTTPS: 10.200.58.231, 10.200.58.232,
10.200.58.233
CP Server Virtual IP(s) for IPM: 10.182.36.8, 10.182.36.9
CP Server Port(s) for HTTPS: 65535, 65534, 65537
CP Server Port(s) for IPM: 54448, 54449
CP Server Security for IPM: 1
CP Server Database Dir: /cpsdb

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

16 The installer proceeds with the configuration process, and creates a vxcps.conf configuration file.

```
Successfully generated the /etc/vxcps.conf configuration file
Copying configuration file /etc/vxcps.conf to sys0....Done
Creating mount point /cps_mount_data on sys0. ... Done
Copying configuration file /etc/vxcps.conf to sys0. ... Done
Press Enter to continue.
```

17 Configure CP Server Service Group (CPSSG) for this cluster.

```
Enter how many NIC resources you want to configure (1 to 2): 2

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

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

```
Enter a valid network interface on sys1 for NIC resource - 1: eth0
Enter a valid network interface on sys1 for NIC resource - 2: eth1
```

19 Enter the NIC resource you want to associate with the virtual IP addresses.

```
Enter the NIC resource you want to associate with the virtual IP 10.200.58.231 (1 to 2): 1
Enter the NIC resource you want to associate with the virtual IP 10.200.58.232 (1 to 2): 2
```

20 Enter the networkhosts information for each NIC resource.

```
Symantec recommends configuring NetworkHosts attribute to ensure NIC resource
to be always online


Do you want to add NetworkHosts attribute for the NIC device eth0
on system sys1? [y,n,q] y
Enter a valid IP address to configure NetworkHosts for NIC eth0
on system sys1: 10.200.56.22


Do you want to add another Network Host? [y,n,q] n
Do you want to apply the same NetworkHosts for all systems? [y,n,q] (y)
```

21 Enter the netmask for virtual IP addresses. If you entered an IPv6 address, enter the prefix details at the prompt.

Note that if you are using HTTPS-based communication, only IPv4 addresses are supported.

```
Enter the netmask for virtual IP for
HTTPS 192.168.0.111: (255.255.252.0)
Enter the netmask for virtual IP for
IPM 192.168.0.112: (255.255.252.0)
```

22 Configure a disk group for CP server database. You can choose an existing disk group or create a new disk group.

```
Symantec recommends to use the disk group that has at least
two disks on which mirrored volume can be created.
Select one of the options below for CP Server database disk group:

1)  Create a new disk group
2)  Using an existing disk group

Enter the choice for a disk group: [1-2,q]  2
```

23 Select one disk group as the CP Server database disk group.

```
Select one disk group as CP Server database disk group: [1-3,q] 3
1)  mycpsdg
2)  cpsdg1
3)  newcpsdg
```

24 Select the CP Server database volume.

You can choose to use an existing volume or create new volume for CP Server database. If you chose newly created disk group, you can only choose to create new volume for CP Server database.

```
Select one of the options below for CP Server database volume:
 1)  Create a new volume on disk group newcpsdg
 2)  Using an existing volume on disk group newcpsdg
```

25 Enter the choice for a volume: [1-2,q] **2**.

26 Select one volume as CP Server database volume [1-1,q] **1**

```
1) newcpsvol
```

27 After the VCS configuration files are updated, a success message appears.

```
For example:
Updating main.cf with CPSSG service group .... Done
Successfully added the CPSSG service group to VCS configuration.
```

28 If the cluster is secure, installer creates the softlink `/var/VRTSvcs/vcsauth/data/CPSERVER` to `/cpsdb/CPSERVER` and check if credentials are already present at `/cpsdb/CPSERVER`. If not, installer creates credentials in the directory, otherwise, installer asks if you want to reuse exsting credentials.

```
Do you want to reuse these credentials? [y,n,q] (y)
```

29 After the configuration process has completed, a success message appears.

```
For example:
Trying to bring CPSSG service group ONLINE and will wait for upto 120 seconds
The Symantec Coordination Point Server is ONLINE
The Symantec Coordination Point Server has been configured on your system.
```

30 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 cps1 |ONLINE|
CPSSG State cps2 |OFFLINE|
```

It also generates the configuration file for CP server (`/etc/vxcps.conf`). The vxcpserv process and other resources are added to the VCS configuration in the CP server service group (CPSSG).

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

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

The CP server supports both IPM-based secure communication and HTTPS-based secure communication. CP servers that are configured for IPM-based secure communication support client nodes that are running prior to 6.1 versions of the product. However, CP servers that are configured for HTTP-based communication only support client nodes that are running the 6.1 or later version of the product. Client nodes with product versions prior to 6.1 are not supported for HTTPS-based communication.

You need to manually generate certificates for the CP server and its client nodes to configure the CP server for HTTPS-based communication.

Table 7-2 Tasks to configure the CP server manually

Task	Reference
Configure CP server manually for IPM-based secure communication	See "Configuring the CP server manually for IPM-based secure communication" on page 100.
Configure CP server manually for HTTPS-communication	See "Configuring the CP server manually for HTTPS-based communication" on page 101. See "Generating the key and certificates manually for the CP server" on page 102. See "Completing the CP server configuration" on page 105.

Configuring the CP server manually for IPM-based secure communication

Perform the following steps to manually configure the CP server in the Symantec Product Authentication Services (AT) (IPM-based) secure mode.

To manually configure the CP server

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

   ```
   # hastop -local
   ```

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

   See "Sample configuration files for CP server" on page 484.

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

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

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

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

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

   Based on whether you configured the CP server using the Symantec Product Authentication Services (AT) protocol (IPM-based) in secure mode or not, do one of the following:

   - For a CP server cluster which is configured in secure mode, edit the `/etc/vxcps.conf` file to set `security`=1.

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

5 Start VCS on all the cluster nodes.

```
# hastart
```

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

```
# hagrp -state CPSSG
```

Output similar to the following appears:

```
# Group Attribute  System                      Value
  CPSSG State      cps1.symantecexample.com  |ONLINE|
```

Configuring the CP server manually for HTTPS-based communication

Perform the following steps to manually configure the CP server in the Symantec Product Authentication Services (AT) (IPM-based) secure mode.

To manually configure the CP server

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

```
# hastop -local
```

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

See "Sample configuration files for CP server" on page 484.

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

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

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

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

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

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

If you configured the CP server in HTTPS mode, do the following:

- Edit the `/etc/vxcps.conf` file to set `vip_https` with the virtual IP addresses required for HTTPS communication.
- Edit the `/etc/vxcps.conf` file to set `port_https` with the ports used for HTTPS communication.

5 Manually generate keys and certificates for the CP server.

See "Generating the key and certificates manually for the CP server" on page 102.

Generating the key and certificates manually for the CP server

CP server uses the HTTPS protocol to establish secure communication with client nodes. HTTPS is a secure means of communication, which happens over a secure communication channel that is established using the SSL/TLS protocol.

HTTPS uses x509 standard certificates and the constructs from a Public Key Infrastructure (PKI) to establish secure communication between the CP server and client. Similar to a PKI, the CP server, and its clients have their own set of certificates signed by a Certification Authority (CA). The server and its clients trust the certificate.

Every CP server acts as a certification authority for itself and for all its client nodes. The CP server has its own CA key and CA certificate and a server certificate generated, which is generated from a server private key. The server certificate is issued to the Universally Unique Identifier (UUID) of the CP server. All the IP addresses or domain names that the CP server listens on are mentioned in the Subject Alternative Name section of the CP server's server certificate

The OpenSSL library must be installed on the CP server to create the keys or certificates.. If OpenSSL is not installed, then you cannot create keys or certificates. The vxcps.conf file points to the configuration file that determines which keys or certificates are used by the CP server when SSL is initialized. The configuration value is stored in the `ssl_conf_file` and the default value is `/etc/vxcps_ssl.properties`.

To manually generate keys and certificates for the CP server:

1 Create directories for the security files on the CP server.

```
# mkdir -p /var/VRTScps/security/keys /var/VRTScps/security/certs
```

2 Generate an OpenSSL config file, which includes the VIPs.

The CP server listens to requests from client nodes on these VIPs. The server certificate includes VIPs, FQDNs, and host name of the CP server. Clients can reach the CP server by using any of these values. However, Symantec

recommends that client nodes use the IP address to communicate to the CP server.

The sample configuration uses the following values:

- Config file name: *https_ssl_cert.conf*
- VIP: *192.168.1.201*
- FQDN: *cpsone.company.com*
- Host name: *cpsone*

Note the IP address, VIP, and FQDN values used in the [alt_names] section of the configuration file are sample values. Replace the sample values with your configuration values. Do not change the rest of the values in the configuration file.

```
[req]
distinguished_name = req_distinguished_name
req_extensions = v3_req

[req_distinguished_name]
countryName = Country Name (2 letter code)
countryName_default = US
localityName = Locality Name (eg, city)
organizationalUnitName = Organizational Unit Name (eg, section)
commonName = Common Name (eg, YOUR name)
commonName_max = 64
emailAddress = Email Address
emailAddress_max = 40

[v3_req]
keyUsage = keyEncipherment, dataEncipherment
extendedKeyUsage = serverAuth
subjectAltName = @alt_names

[alt_names]
DNS.1 = cpsone.company.com
DNS.2 = cpsone
DNS.3 = 192.168.1.201
```

3 Generate a 4096-bit CA key that is used to create the CA certificate.

The key must be stored at `/var/VRTScps/security/keys/ca.key`. Ensure that only root users can access the CA key, as the key can be misused to create fake certificates and compromise security.

```
# /usr/bin/openssl genrsa -out /var/VRTScps/security/keys/ca.key
4096
```

4 Generate a self-signed CA certificate.

```
# /usr/bin/openssl req -new -x509 -days days -key
/var/VRTScps/security/keys/ca.key -subj \
'/C=countryname/L=localityname/OU=COMPANY/CN=CACERT' -out \
/var/VRTScps/security/certs/ca.crt
```

Where, *days* is the days you want the certificate to remain valid, *countryname* is the name of the country, localityname is the city, *CACERT* is the certificate name.

5 Generate a 2048-bit private key for CP server.

The key must be stored at `/var/VRTScps/security/keys/server_private key`.

```
# /usr/bin/openssl genrsa -out \
/var/VRTScps/security/keys/server_private.key 2048
```

6 Generate a Certificate Signing Request (CSR) for the server certificate.

The Certified Name (CN) in the certificate is the UUID of the CP server.

```
# /usr/bin/openssl genrsa -out
/var/VRTScps/security/keys/server_private.key 2048
```

7 Generate a Certificate Signing Request (CSR) for the server certificate.

The Certified Name (CN) in the certificate is the UUID of the CP server.

```
# /usr/bin/openssl req -new -key
/var/VRTScps/security/keys/server_private.key \
-config https_ssl_cert.conf -subj \
'/C=CountryName/L=LocalityName/OU=COMPANY/CN=UUID' \
-out /var/VRTScps/security/certs/server.csr
```

Where, *countryname* is the name of the country, *localityname* is the city, *UUID* is the certificate name.

8 Generate the server certificate by using the key certificate of the CA.

```
# /usr/bin/openssl x509 -req -days days -in
/var/VRTScps/security/certs/server.csr \
-CA /var/VRTScps/security/certs/ca.crt -CAkey \
/var/VRTScps/security/keys/ca.key \
-set_serial 01 -extensions v3_req -extfile https_ssl_cert.conf \
-out /var/VRTScps/security/certs/server.crt
```

Where, *days* is the days you want the certificate to remain valid, *https_ssl_cert.conf* is the configuration file name.

You successfully created the key and certificate required for the CP server.

9 Ensure that no other user except the root user can read the keys and certificates.

10 Complete the CP server configuration.

See "Completing the CP server configuration" on page 105.

Completing the CP server configuration

To verify the service groups and start VCS perform the following steps:

1 Start VCS on all the cluster nodes.

```
# hastart
```

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

```
# hagrp -state CPSSG
```

Output similar to the following appears:

```
# Group Attribute  System                    Value
  CPSSG State      cps1.symantecexample.com  |ONLINE|
```

Configuring CP server using response files

You can configure a CP server using a generated responsefile.

On a single node VCS cluster:

- Run the `installvcs<version>` command with the responsefile option to configure the CP server on a single node VCS cluster.

  ```
  # /opt/VRTS/install/installvcs<version> -responsefile
  '/tmp/sample1.res'
  ```

  Where *<version>* is the specific release version.

  See "About the script-based installer" on page 69.

On a SFHA cluster:

- Run the `installsfha<version>` command with the responsefile option to configure the CP server on a SFHA cluster.

  ```
  # /opt/VRTS/install/installsfha<version> -responsefile
  '/tmp/sample1.res'
  ```

  Where *<version>* is the specific release version.

  See "About the script-based installer" on page 69.

Response file variables to configure CP server

Table 7-3 describes the response file variables to configure CP server.

Table 7-3 describes response file variables to configure CP server

Variable	List or Scalar	Description
CFG{opt}{configcps}	Scalar	This variable performs CP server configuration task
CFG{cps_singlenode_config}	Scalar	This variable describes if the CP server will be configured on a singlenode VCS cluster
CFG{cps_sfha_config}	Scalar	This variable describes if the CP server will be configured on a SFHA cluster
CFG{cps_unconfig}	Scalar	This variable describes if the CP server will be unconfigured
CFG{cpsname}	Scalar	This variable describes the name of the CP server
CFG{cps_db_dir}	Scalar	This variable describes the absolute path of CP server database

Table 7-3 describes response file variables to configure CP server *(continued)*

Variable	List or Scalar	Description
CFG{cps_security}	Scalar	This variable describes if security is configured for the CP server
CFG{cps_reuse_cred}	Scalar	This variable describes if reusing the existing credentials for the CP server
CFG{cps_https_vips}	List	This variable describes the virtual IP addresses for the CP server configured for HTTPS-based communication
CFG{cps_ipm_vips}	List	This variable describes the virtual IP addresses for the CP server configured for IPM-based communication
CFG{cps_https_ports}	List	This variable describes the port number for the virtual IP addresses for the CP server configured for HTTPS-based communication
CFG{cps_ipm_ports}	List	This variable describes the port number for the virtual IP addresses for the CP server configured for IPM-based communication
CFG{cps_nic_list}{cpsvip<n>}	List	This variable describes the NICs of the systems for the virtual IP address
CFG{cps_netmasks}	List	This variable describes the netmasks for the virtual IP addresses
CFG{cps_prefix_length}	List	This variable describes the prefix length for the virtual IP addresses
CFG{cps_network_hosts}{cpsnic<n>}	List	This variable describes the network hosts for the NIC resource
CFG{cps_vip2nicres_map}{<vip>}	Scalar	This variable describes the NIC resource to associate with the virtual IP address
CFG{cps_diskgroup}	Scalar	This variable describes the disk group for the CP server database
CFG{cps_volume}	Scalar	This variable describes the volume for the CP server database

Table 7-3 describes response file variables to configure CP server *(continued)*

Variable	List or Scalar	Description
CFG{cps_newdg_disks}	List	This variable describes the disks to be used to create a new disk group for the CP server database
CFG{cps_newvol_volsize}	Scalar	This variable describes the volume size to create a new volume for the CP server database
CFG{cps_delete_database}	Scalar	This variable describes if deleting the database of the CP server during the unconfiguration
CFG{cps_delete_config_log}	Scalar	This variable describes if deleting the config files and log files of the CP server during the unconfiguration
CFG{cps_reconfig}	Scalar	This variable defines if the CP server will be reconfigured

Sample response file for configuring the CP server on single node VCS cluster

Review the response file variables and their definitions.

See Table 7-3 on page 106.

```
#
# Configuration Values:
#
our %CFG;

$CFG{cps_db_dir}="/etc/VRTScps/db";
$CFG{cps_https_ports}=[ qw(443) ];
$CFG{cps_https_vips}=[ qw(192.169.0.220) ];
$CFG{cps_ipm_ports}=[ qw(14250) ];
$CFG{cps_ipm_vips}=[ qw(192.169.0.221) ];
$CFG{cps_netmasks}=[ qw(255.255.252.0 255.255.252.0) ];
$CFG{cps_nic_list}{cpsvip1}=[ qw(eth0) ];
$CFG{cps_nic_list}{cpsvip2}=[ qw(eth0) ];
$CFG{cps_security}="0";
```

```
$CFG{cps_singlenode_config}=1;
$CFG{cps_vip2nicres_map}{"192.169.0.220"}=1;
$CFG{cps_vip2nicres_map}{"192.169.0.221"}=1;
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(cps1) ];
$CFG{vcs_clusterid}=64505;
$CFG{vcs_clustername}="single";

1;
```

Sample response file for configuring the CP server on SFHA cluster

Review the response file variables and their definitions.

See Table 7-3 on page 106.

```
#
# Configuration Values:
#
our %CFG;

$CFG{cps_db_dir}="/cpsdb";
$CFG{cps_diskgroup}="cps_dg1";
$CFG{cps_https_ports}=[ qw(50006 50007) ];
$CFG{cps_https_vips}=[ qw(10.198.90.6 10.198.90.7) ];
$CFG{cps_ipm_ports}=[ qw(14250) ];
$CFG{cps_ipm_vips}=[ qw(10.198.90.8) ];
$CFG{cps_netmasks}=[ qw(255.255.248.0 255.255.248.0 255.255.248.0) ];
$CFG{cps_network_hosts}{cpsnic1}=[ qw(10.198.88.18) ];
$CFG{cps_network_hosts}{cpsnic2}=[ qw(10.198.88.18) ];
$CFG{cps_newdg_disks}=[ qw(emc_clariion0_249) ];
$CFG{cps_newvol_volsize}=10;
$CFG{cps_nic_list}{cpsvip1}=[ qw(eth0 eth0) ];
$CFG{cps_nic_list}{cpsvip2}=[ qw(eth0 eth0) ];
$CFG{cps_nic_list}{cpsvip3}=[ qw(eth0 eth0) ];
$CFG{cps_security}="0";
$CFG{cps_sfha_config}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.6"}=1;
```

```
$CFG{cps_vip2nicres_map}{"10.198.90.7"}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.8"}=1;
$CFG{cps_volume}="volcps";
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{noipc}=1;
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw(cps1 cps2) ];
$CFG{vcs_clusterid}=49604;
$CFG{vcs_clustername}="sfha2233";

1;
```

Verifying the CP server configuration

Perform the following steps to verify the CP server configuration.

To verify the CP server configuration

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

   If the application cluster is configured for HTTPS-based communication, no need to provide the port number assigned for HTTP communication.

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

   For IPM-based communication, you need to specify 14250 as the port number.

   ```
   # cpsadm -s cp_server -p 14250  -a ping_cps
   ```

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

Configuring the CP server using the web-based installer

Perform the following steps to configure the CP server using the web-based installer.

To configure SFHA on a cluster

1. Start the web-based installer.

   See "Starting the web-based installer" on page 170.

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

Task	Configure CP server
Product	Storage Foundation and High Availability

   Click **Next**.

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

4. In the Confirmation dialog box, verify cluster information is correct and choose whether or not to configure CP server.

   - To configure CP server, click **Yes**.
   - To configure CP server later, click **No**.

5. On the Select Option page, select Configure CP Server on a single-node VCS system or SFHA cluster and click **Next**.

6. On the Configure CP Server page, provide CP server information, such as, name, virtual IPs, port numbers, and absolute path of the database to store the configuration details.

   Click **Next**.

7. Configure the CP Server Service Group (CPSSG), select the number of NIC resources, and associate NIC resources to virtual IPs that are going to be used to configure the CP Server.

   Click **Next**.

8. Configure network hosts for the CP server.

   Click **Next**.

9. Configure disk group for the CP server.

   Click **Next**.

   Note: This step is not applicable for a single node cluster.

10 Configure volume for the disk group associated to the CP server.

Click **Next**.

Note: This step is not applicable for a single node cluster.

---

11 Click **Finish** to complete configuring the CP server.

Chapter 8

Configuring SFHA

This chapter includes the following topics:

- Configuring Storage Foundation High Availability using the installer
- Configuring SFDB

Configuring Storage Foundation High Availability using the installer

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

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

Table 8-1 lists the tasks that are involved in configuring SFHA using the script-based installer.

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

Task	Reference
Start the software configuration	See "Starting the software configuration" on page 115.
Specify the systems where you want to configure SFHA	See "Specifying systems for configuration" on page 116.
Configure the basic cluster	See "Configuring the cluster name" on page 117. See "Configuring private heartbeat links" on page 117.

Table 8-1 Tasks to configure SFHA using the script-based installer *(continued)*

Task	Reference
Configure virtual IP address of the cluster (optional)	See "Configuring the virtual IP of the cluster" on page 123.
Configure the cluster in secure mode (optional)	See "Configuring Storage Foundation and High Availability in secure mode" on page 124.
Add VCS users (required if you did not configure the cluster in secure mode)	See "Adding VCS users" on page 130.
Configure SMTP email notification (optional)	See "Configuring SMTP email notification" on page 131.
Configure SNMP email notification (optional)	See "Configuring SNMP trap notification" on page 132.
Configure global clusters (optional) **Note:** You must have enabled global clustering when you installed SFHA.	See "Configuring global clusters" on page 134.
Complete the software configuration	See "Completing the SFHA configuration" on page 135.

Required information for configuring Storage Foundation and High Availability Solutions

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

See also the *Symantec Cluster Server Installation Guide*.

- A unique Cluster name
- A unique Cluster ID number between 0-65535
- Two or more NIC cards per system used for heartbeat links
  One or more heartbeat links are configured as private links and one heartbeat link may be configured as a low priority link.

You can configure Storage Foundation High Availability in secure mode.

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

SFHA usernames and passwords are no longer used when a cluster is running in Secure Mode.

The following information is required to configure SMTP notification:

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

The following information is required to configure SNMP notification:

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

Starting the software configuration

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

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

To configure SFHA using the product installer

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

   ```
   # ./installer
   ```

   The installer starts the product installation program with a copyright message and specifies the directory where the logs are created.
3. From the opening Selection Menu, choose: `C` for "Configure an Installed Product."
4. From the displayed list of products to configure, choose the corresponding number for your product:

   Storage Foundation and High Availability

To configure SFHA using the installsfha program

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

   ```
   # /opt/VRTS/install/installsfha<version> -configure
   ```

   Where `<version>` is the specific release version.

   See "About the script-based installer" on page 69.

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

Specifying systems for configuration

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

To specify system names for configuration

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

   ```
   Enter the operating_system system names separated
   by spaces:  [q,?] (sys1) sys1 sys2
   ```

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. If ssh binaries cannot communicate with remote nodes, the installer tries rsh binaries. And if both ssh and rsh binaries fail, the installer prompts to help the user to setup ssh or rsh binaries.
   - Makes sure that the systems are running with the supported operating system
   - Checks whether SFHA is installed

- Exits if SFHA 6.2 is not installed

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

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

See " About planning to configure I/O fencing" on page 75.

Configuring the cluster name

Enter the cluster information when the installer prompts you.

To configure the cluster

1 Review the configuration instructions that the installer presents.

2 Enter a unique cluster name.

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

Configuring private heartbeat links

You now configure the private heartbeat links that LLT uses.

See "Setting up the private network" on page 58.

VCS provides the option to use LLT over Ethernet or LLT over UDP (User Datagram Protocol) or LLT over RDMA. Symantec recommends that you configure heartbeat links that use LLT over Ethernet or LLT over RDMA for high performance, unless hardware requirements force you to use LLT over UDP. If you want to configure LLT over UDP, make sure you meet the prerequisites.

You must not configure LLT heartbeat using the links that are part of aggregated links. For example, link1, link2 can be aggregated to create an aggregated link, aggr1. You can use aggr1 as a heartbeat link, but you must not use either link1 or link2 as heartbeat links.

See "Using the UDP layer for LLT" on page 527.

See "Using LLT over RDMA: supported use cases " on page 541.

The following procedure helps you configure LLT heartbeat links.

To configure private heartbeat links

1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or LLT over UDP or LLT over RDMA.

- Option 1: Configure the heartbeat links using LLT over Ethernet (answer installer questions)

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

Skip to step 2.

- Option 2: Configure the heartbeat links using LLT over UDP (answer installer questions)

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

  Skip to step 3.

- Option 3: Configure the heartbeat links using LLT over RDMA (answer installer questions)

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

  Skip to step 4.

- Option 4: Automatically detect configuration for LLT over Ethernet

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

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

  Skip to step 7.

Note: Option 4 is not available when the configuration is a single node configuration.

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

The installer discovers and lists the network interface cards.

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

```
Enter the NIC for the first private heartbeat link on sys1:
[b,q,?] eth1
eth1 has an IP address configured on it. It could be a
public NIC on sys1.
Are you sure you want to use eth1 for the first private
heartbeat link? [y,n,q,b,?] (n) y
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on sys1:
[b,q,?] eth2
eth2 has an IP address configured on it. It could be a
public NIC on sys1.
Are you sure you want to use eth2 for the second private
heartbeat link? [y,n,q,b,?] (n) y
Would you like to configure a third private heartbeat link?
[y,n,q,b,?](n)

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

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

```
Enter the NIC for the first private heartbeat
link on sys1: [b,q,?] private_NIC1
Do you want to use address 192.168.0.1 for the
first private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the first private heartbeat
link on sys1: [b,q,?] (50000)
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat
link on sys1: [b,q,?] private_NIC2
Do you want to use address 192.168.1.1 for the
second private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the second private heartbeat
link on sys1: [b,q,?] (50001)
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n) y
Enter the NIC for the low priority heartbeat
link on sys1: [b,q,?] (private_NIC0)
Do you want to use address 192.168.3.1 for
the low priority heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the UDP port for the low priority heartbeat
link on sys1: [b,q,?] (50004)
```

4 If you chose option 3, choose the interconnect type to configure RDMA.

```
1)  Converged Ethernet (RoCE)
2)  InfiniBand
b)  Back to previous menu

Choose the RDMA interconnect type [1-2,b,q,?] (1) 2
```

The system displays the details such as the required OS files, drivers required for RDMA , and the IP addresses for the NICs.

A sample output of the IP addresses assigned to the RDMA enabled NICs using InfiniBand network. Note that with RoCE, the RDMA NIC values are represented as eth0, eth1, and so on.

```
System        RDMA NIC          IP Address
================================================================
sys1           ib0               192.168.0.1
sys1           ib1               192.168.3.1
sys2           ib0               192.168.0.2
sys2           ib1               192.168.3.2
```

5 If you chose option 3, enter the NIC details for the private heartbeat links. This step uses RDMA over an InfiniBand network. With RoCE as the interconnect type, RDMA NIC is represented as Ethernet (eth).

```
Enter the NIC for the first private heartbeat
link (RDMA) on sys1: [b,q,?] <ib0>

Do you want to use address 192.168.0.1 for the
first private heartbeat link on sys1: [y,n,q,b,?] (y)

Enter the port for the first private heartbeat
link (RDMA) on sys1: [b,q,?] (50000) ?

Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link (RDMA) on sys1:
[b,q,?] (ib1)

Do you want to use the address 192.168.3.1 for the second
private heartbeat link on sys1: [y,n,q,b,?] (y)
Enter the port for the second private heartbeat link (RDMA) on sys1:
[b,q,?] (50001)

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

6 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 sys1, 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.

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

7 If you chose option 4, the installer detects NICs on each system and network links, and sets link priority.

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

See step 2 for option 1, or step 3 for option 2 or step 4 for option 3.

8 Enter a unique cluster ID:

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

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

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

9 Verify and confirm the information that the installer summarizes.

Configuring the virtual IP of the cluster

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

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

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

To configure the virtual IP of the cluster

1 Review the required information to configure the virtual IP of the cluster.

2 When the system prompts whether you want to configure the virtual IP, enter `y`.

3 Confirm whether you want to use the discovered public NIC on the first system.

Do one of the following:

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

```
Active NIC devices discovered on sys1: eth0
Enter the NIC for Virtual IP of the Cluster to use on sys1:
[b,q,?](eth0)
```

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 eth0 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

If you want to set up trust relationships for your secure cluster, refer to the following topics:

See "Configuring a secure cluster node by node" on page 125.

Configuring Storage Foundation and High Availability in secure mode

Configuring SFHA in secure mode ensures that all the communication between the systems is encrypted and users are verified against security credentials. SFHA user names and passwords are not used when a cluster is running in secure mode.

To configure SFHA in secure mode

1 To install SFHA in secure mode, run the command:

```
# installsfha<version> -security
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 69.

2 The installer displays the following question before the install stops the product processes:

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.

- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

3 To verify the cluster is in secure mode after configuration, run the command:

```
# haclus -value SecureClus
```

The command returns 1 if cluster is in secure mode, else returns 0.

Configuring a secure cluster node by node

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

Table 8-2 lists the tasks that you must perform to configure a secure cluster.

Table 8-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See "Configuring the first node" on page 125.
Configure security on the remaining nodes	See "Configuring the remaining nodes" on page 126.
Complete the manual configuration steps	See "Completing the secure cluster configuration" on page 127.

Configuring the first node

Perform the following steps on one node in your cluster.

To configure security on the first node

1 Ensure that you are logged in as superuser.

2 Enter the following command:

```
# /opt/VRTS/install/installsfha<version> -securityonenode
```

Where *<version>* is the specific release version.

See "About the script-based installer" on page 69.

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

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

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

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

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

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

3 The installer completes the secure configuration on the node. It specifies the location of the security configuration files and prompts you to copy these files to the other nodes in the cluster. The installer also specifies the location of log files, summary file, and response file.

4 Copy the security configuration files from the location specified by the installer to temporary directories on the other nodes in the cluster.

Configuring the remaining nodes

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

To configure security on each remaining node

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

   ```
   # /opt/VRTS/install/installsfha<version> -securityonenode
   ```

   Where *<version>* is the specific release version.

   See "About the script-based installer" on page 69.

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

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

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

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

   Select the option you would like to perform [1-2,q.?]  2
   Enter the security conf file directory: [b]
   ```

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

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

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

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

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

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

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

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

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

3. On all nodes, stop the CmdServer.

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

4 To grant access to all users, add or modify `SecureClus=1` and `DefaultGuestAccess=1` in the cluster definition.

For example:

To grant read access to everyone:

```
Cluster clus1 (
SecureClus=1
DefaultGuestAccess=1
)
```

Or

To grant access to only root:

```
Cluster clus1 (
SecureClus=1
)
```

Or

To grant read access to specific user groups, add or modify SecureClus=1 and GuestGroups={} to the cluster definition.

For example:

```
cluster clus1 (
SecureClus=1
GuestGroups={staff, guest}
```

5 Modify `/etc/VRTSvcs/conf/config/main.cf` file on the first node, and add `-secure` to the WAC application definition if GCO is configured.

For example:

```
Application wac (
                StartProgram = "/opt/VRTSvcs/bin/wacstart -secure"
                StopProgram = "/opt/VRTSvcs/bin/wacstop"
                MonitorProcesses = {"/opt/VRTSvcs/bin/wac -secure"}
                RestartLimit = 3
                )
```

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

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

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

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

8 On all nodes, start CmdServer.

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

9 On the first node, unfreeze the service groups.

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

# /opt/VRTSvcs/bin/hagrp -list Frozen=1

# /opt/VRTSvcs/bin/hagrp -unfreeze groupname -persistent

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

Adding VCS users

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

To add VCS users

1 Review the required information to add VCS users.

2 Reset the password for the Admin user, if necessary.

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

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

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

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

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

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

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

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

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

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation.

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

To configure SMTP email notification

1 Review the required information to configure the SMTP email notification.

2 Specify whether you want to configure the SMTP notification.

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

See "Configuring SNMP trap notification" on page 132.

3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the SMTP server's host name.

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

- Enter the email address of each recipient.

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

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

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

4 Add more SMTP recipients, if necessary.

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

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

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

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

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

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

5 Verify and confirm the SMTP notification information.

```
SMTP Address: smtp.example.com
Recipient: ozzie@example.com receives email for Warning or
higher events
Recipient: harriet@example.com receives email for Error or
higher events

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

Configuring SNMP trap notification

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

Note that you can also configure the notification after installation.

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

To configure the SNMP trap notification

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

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

   See "Configuring global clusters" on page 134.

3. Provide information to configure SNMP trap notification.

   Provide the following information:

   - 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,?] sys5
     ```

   - 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 sys5 [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,?] sys4
     Enter the minimum severity of events for which SNMP traps
     should be sent to sys4 [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.

```
SNMP Port: 162
Console: sys5 receives SNMP traps for Error or
higher events
Console: sys4 receives SNMP traps for SevereError or
higher events

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

Configuring global clusters

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

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

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

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

To configure the global cluster option

1 Review the required information to configure the global cluster option.

2 Specify whether you want to configure the global cluster option.

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.

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

Completing the SFHA configuration

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

To complete the SFHA configuration

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

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

2. Review the output as the installer stops various processes and performs the configuration. The installer then restarts SFHA and its related processes.

3. Enter y at the prompt to send the installation information to Symantec.

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

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

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

summary file	Describes the cluster and its configured resources.
log file	Details the entire configuration.
response file	Contains the configuration information that can be used to perform secure or unattended installations on other systems. See "Configuring SFHA using response files" on page 201.

Verifying the NIC configuration

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

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

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

Set the PERSISTENT_NAME for all the NICs.

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

Verifying and updating licenses on the system

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

See "Checking licensing information on the system" on page 136.

See "Updating product licenses" on page 137.

Checking licensing information on the system

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

To check licensing information

1. Navigate to the `/sbin` folder containing the `vxlicrep` program and enter:

   ```
   # vxlicrep
   ```

2. Review the following output to determine the following information:
   - The license key
   - The type of license
   - The product for which it applies
   - Its expiration date, if any. Demo keys have expiration dates. Permanent keys and site keys do not have expiration dates.

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

Features :=
Platform                   = Linux
Version                    = 6.0
```

```
Tier                    = 0
Reserved                = 0
Mode                    = VCS
```

Updating product licenses

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

See "Replacing a SFHA demo license with a permanent license" on page 137.

To update product licenses using the installer command

1 On each node, enter the license key using the command:

```
# ./installer -license
```

2 At the prompt, enter your license number.

To update product licenses using the vxlicinst command

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

```
# vxlicinst -k license key
```

Replacing a SFHA demo license with a permanent license

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

To replace a demo key

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

2 Shut down SFHA on all nodes in the cluster:

```
# hastop -all -force
```

This command does not shut down any running applications.

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

```
# vxlicinst -k license key
```

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

```
# vxlicrep
```

5 Start SFHA on each node:

```
# hastart
```

Configuring SFDB

By default, SFDB tools are disabled that is the vxdbd daemon is not configured. You can check whether SFDB tools are enabled or disabled using the`/opt/VRTS/bin/sfae_config status` command.

To enable SFDB tools

1 Log in as root.

2 Run the following command to configure and start the vxdbd daemon. After you perform this step, entries are made in the system startup so that the daemon starts on a system restart.

```
#/opt/VRTS/bin/sfae_config enable
```

To disable SFDB tools

1 Log in as root.

2 Run the following command:

```
#/opt/VRTS/bin/sfae_config disable
```

Chapter 9

Manually configuring SFHA clusters for data integrity

This chapter includes the following topics:

- Setting up disk-based I/O fencing using installsfha
- Setting up server-based I/O fencing using installsfha
- Setting up non-SCSI-3 I/O fencing in virtual environments using installsfha
- Setting up majority-based I/O fencing using installsfha
- Enabling or disabling the preferred fencing policy

Setting up disk-based I/O fencing using installsfha

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

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

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

   ```
   # fdisk -l
   ```

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

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

  ```
  # vxdisksetup -i device_name
  ```

  The example specifies the CDS format:

  ```
  # vxdisksetup -i sdr format=cdsdisk
  ```

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

Checking shared disks for I/O fencing

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

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

The vxfentsthdw 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 *Symantec Cluster Server Administrator's Guide*.

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

- Verifying the Array Support Library (ASL)
  See "Verifying Array Support Library (ASL)" on page 140.
- Verifying that nodes have access to the same disk
  See "Verifying that the nodes have access to the same disk" on page 141.
- Testing the shared disks for SCSI-3
  See "Testing the disks using vxfentsthdw utility" on page 142.

Verifying Array Support Library (ASL)

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

To verify Array Support Library (ASL)

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

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

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

   The following output is a sample:

   ```
   # vxddladm listsupport all


   LIBNAME             VID                 PID
   ============================================================
   libvxhitachi.so     HITACHI             DF350, DF400, DF400F,
                                           DF500, DF500F
   libvxxp1281024.so   HP                  All
   libvxxp12k.so       HP                  All
   libvxddns2a.so      DDN                 S2A 9550, S2A 9900,
                                           S2A 9700
   libvxpurple.so      SUN                 T300
   libvxxiotechE5k.so  XIOTECH             ISE1400
   libvxcopan.so       COPANSYS            8814, 8818
   libvxibmds8k.so     IBM                 2107
   ```

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

   ```
   # vxdisk scandisks
   ```

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

Verifying that the nodes have access to the same disk

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

To verify that the nodes have access to the same disk

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

   ```
   # vxfenadm -i diskpath
   ```

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

   For example, an EMC disk is accessible by the /dev/sdx path on node A and the /dev/sdy path on node B.

   From node A, enter:

   ```
   # vxfenadm -i /dev/sdx

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

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

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

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

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

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

Testing the disks using vxfentsthdw utility

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

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

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

```
The disk /dev/sdx is ready to be configured for I/O Fencing on
node sys1
```

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

To test the disks using vxfentsthdw utility

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

   See "About configuring secure shell or remote shell communication modes before installing products" on page 491.

2. From one node, start the utility.

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: sys1
   Enter the second node of the cluster: sys2
   ```

4. Review the output as the utility performs the checks and reports its activities.

5. If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node sys1.

   ```
   The disk is now ready to be configured for I/O Fencing on node
   sys1

   ALL tests on the disk /dev/sdx have PASSED
   The disk is now ready to be configured for I/O fencing on node
   sys1
   ```

6. Run the vxfentsthdw utility for each disk you intend to verify.

Note: Only dmp disk devices can be used as coordinator disks.

Configuring disk-based I/O fencing using installsfha

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

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

1 Start the installsfha with `-fencing` option.

```
# /opt/VRTS/install/installsfha<version> -fencing
```

Where *<version>* is the specific release version.

See "About the script-based installer" on page 69.

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

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

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

The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 6.2 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 [1-7,b,q] 2
```

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

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

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

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

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.
  The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
- To create a new disk group, perform the following steps:
  - Enter the number corresponding to the **Create a new disk group** option.
    The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks.
    Symantec recommends that you use three disks as coordination points for disk-based I/O fencing.
  - If the available VxVM CDS disks are less than the required, installer asks whether you want to initialize more disks as VxVM disks. Choose the disks you want to initialize as VxVM disks and then use them to create new disk group.
  - Enter the numbers corresponding to the disks that you want to use as coordinator disks.
  - Enter the disk group name.

6 Verify that the coordinator disks you chose meet the I/O fencing requirements.

You must verify that the disks are SCSI-3 PR compatible using the vxfentsthdw utility and then return to this configuration program.

See "Checking shared disks for I/O fencing" on page 140.

7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.

8 Verify and confirm the I/O fencing configuration information that the installer summarizes.

9 Review the output as the configuration program does the following:

- Stops VCS and I/O fencing on each node.
- Configures disk-based I/O fencing and starts the I/O fencing process.
- Updates the VCS configuration file main.cf if necessary.
- Copies the /etc/vxfenmode file to a date and time suffixed file /etc/vxfenmode-*date-time*. This backup file is useful if any future fencing configuration fails.
- Updates the I/O fencing configuration file /etc/vxfenmode.

- Starts VCS on each node to make sure that the SFHA is cleanly configured to use the I/O fencing feature.

10 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.

11 Configure the Coordination Point Agent.

```
Do you want to configure Coordination Point Agent on
the client cluster? [y,n,q] (y)
```

12 Enter a name for the service group for the Coordination Point Agent.

```
Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen) vxfen
```

13 Set the level two monitor frequency.

```
Do you want to set LevelTwoMonitorFreq? [y,n,q] (y)
```

14 Decide the value of the level two monitor frequency.

```
Enter the value of the LevelTwoMonitorFreq attribute: [b,q,?] (5)
```

Installer adds Coordination Point Agent and updates the main configuration file.

See "Configuring CoordPoint agent to monitor coordination points" on page 256.

Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installsfha

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for disk-based I/O fencing using the installsfha

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

```
# /opt/VRTS/install/installsfha<version> -fencing
```

where, *<version>* is the specific release version.

See "About the script-based installer" on page 69.

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

Note down the location of log files that you can access if there is a problem with the configuration process.

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

The program checks that the local node running the script can communicate with the remote nodes and checks whether SFHA 6.2 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type the number corresponding to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,q]
```

4 Ensure that the disk group constitution that is used by the fencing module contains the same disks that are currently used as coordination disks.

5 Verify the coordination points.

```
For example,
Disk Group: fendg
Fencing disk policy: dmp
Fencing disks:
 emc_clariion0_62
 emc_clariion0_65
 emc_clariion0_66
```

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

```
Successfully completed the vxfenswap operation
```

The keys on the coordination disks are refreshed.

6 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

7 Do you want to view the summary file? [y,n,q] **(n)**.

Setting up server-based I/O fencing using installsfha

You can configure server-based I/O fencing for the SFHA cluster using the installsfha.

With server-based fencing, you can have the coordination points in your configuration as follows:

- Combination of CP servers and SCSI-3 compliant coordinator disks
- CP servers only
  Symantec also supports server-based fencing with a single highly available CP server that acts as a single coordination point.

See " About planning to configure I/O fencing" on page 75.

See "Recommended CP server configurations" on page 80.

This section covers the following example procedures:

Mix of CP servers and coordinator disks	See "To configure server-based fencing for the SFHA cluster (one CP server and two coordinator disks)" on page 148.
Single CP server	See "To configure server-based fencing for the SFHA cluster (single CP server)" on page 153.

To configure server-based fencing for the SFHA cluster (one CP server and two coordinator disks)

1 Depending on the server-based configuration model in your setup, make sure of the following:

- CP servers are configured and are reachable from the SFHA cluster. The SFHA cluster is also referred to as the application cluster or the client cluster.
  See "Setting up the CP server" on page 83.
- The coordination disks are verified for SCSI3-PR compliance.

See "Checking shared disks for I/O fencing" on page 140.

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

```
# /opt/VRTS/install/installsfha<version> -fencing
```

Where <version> is the specific release version. The installsfha starts with a copyright message and verifies the cluster information.

See "About the script-based installer" on page 69.

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

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

The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 6.2 is configured properly.

4 Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,b,q] 1
```

5 Make sure that the storage supports SCSI3-PR, and answer y at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

6 Provide the following details about the coordination points at the installer prompt:

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

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

- Enter the total number of coordinator disks among the coordination points.

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

7 Provide the following CP server details at the installer prompt:

- Enter the total number of virtual IP addresses or the total number of fully qualified host names for each of the CP servers.

```
How many IP addresses would you like to use to communicate
 to Coordination Point Server #1?: [b,q,?] (1) 1
```

- Enter the virtual IP addresses or the fully qualified host name for each of the CP servers. The installer assumes these values to be identical as viewed from all the application cluster nodes.

  ```
  Enter the Virtual IP address or fully qualified host name #1
  for the HTTPS Coordination Point Server #1:
  [b] 10.209.80.197
  ```

  The installer prompts for this information for the number of virtual IP addresses you want to configure for each CP server.

- Enter the port that the CP server would be listening on.

  ```
  Enter the port that the coordination point server 10.198.90.178
  would be listening on or accept the default port
  suggested: [b] (443)
  ```

8 Provide the following coordinator disks-related details at the installer prompt:

- Choose the coordinator disks from the list of available disks that the installer displays. Ensure that the disk you choose is available from all the SFHA (application cluster) nodes.
  The number of times that the installer asks you to choose the disks depends on the information that you provided in step 6. For example, if you had chosen to configure two coordinator disks, the installer asks you to choose the first disk and then the second disk:

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

  1) sdx
  2) sdy
  3) sdz

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

- If you have not already checked the disks for SCSI-3 PR compliance in step 1, check the disks now.
  The installer displays a message that recommends you to verify the disks in another window and then return to this configuration procedure.
  Press Enter to continue, and confirm your disk selection at the installer prompt.

- Enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):
[b] (vxfencoorddg)
```

9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 3
Coordination Point Server ([VIP or FQHN]:Port):
    1. 10.209.80.197 ([10.209.80.197]:443)
SCSI-3 disks:
    1. sdx
    2. sdy
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk policy used for customized fencing: dmp
```

The installer initializes the disks and the disk group and deports the disk group on the SFHA (application cluster) node.

10 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

11 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

```
Updating client cluster information on Coordination Point Server 10.209.80.197

Adding the client cluster to the Coordination Point Server 10.209.80.197 .......... Done

Registering client node sys1 with Coordination Point Server 10.209.80.197...... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.209.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 ..Done

Updating /etc/vxfenmode file on sys1 ................................. Done
Updating /etc/vxfenmode file on sys2 ......... ....................... Done
```

See “About I/O fencing configuration files” on page 482.

12 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

13 Configure the CP agent on the SFHA (application cluster). The Coordination Point Agent monitors the registrations on the coordination points.

```
Do you want to configure Coordination Point Agent on
the client cluster? [y,n,q] (y)

Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen)
```

14 Additionally the coordination point agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidently deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles.

Note that for the LevelTwoMonitorFreq attribute to be applicable there must be disks as part of the Coordinator Disk Group.

```
Enter the value of the LevelTwoMonitorFreq attribute: (5)

Adding Coordination Point Agent via sys1 .... Done
```

15 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

16 Verify the fencing configuration using:

```
# vxfenadm -d
```

17 Verify the list of coordination points.

```
# vxfenconfig -l
```

To configure server-based fencing for the SFHA cluster (single CP server)

1 Make sure that the CP server is configured and is reachable from the SFHA cluster. The SFHA cluster is also referred to as the application cluster or the client cluster.

2 See “Setting up the CP server” on page 83.

3 Start the installsfha with `-fencing` option.

```
# /opt/VRTS/install/installsfha<version>  -fencing
```

Where <version> is the specific release version. The installsfha starts with a copyright message and verifies the cluster information.

See “About the script-based installer” on page 69.

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

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

The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 6.2 is configured properly.

5 Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,q] 1
```

6 Make sure that the storage supports SCSI3-PR, and answer y at the following prompt.

```
Does your storage environment support SCSI3 PR? [y,n,q] (y)
```

7 Enter the total number of coordination points as **1**.

```
Enter the total number of co-ordination points including both
Coordination Point servers and disks: [b] (3) 1
```

Read the installer warning carefully before you proceed with the configuration.

8 Provide the following CP server details at the installer prompt:

- Enter the total number of virtual IP addresses or the total number of fully qualified host names for each of the CP servers.

  ```
  How many IP addresses would you like to use to communicate
  to Coordination Point Server #1? [b,q,?] (1) 1
  ```

- Enter the virtual IP address or the fully qualified host name for the CP server. The installer assumes these values to be identical as viewed from all the application cluster nodes.

  ```
  Enter the Virtual IP address or fully qualified host name
  #1 for the Coordination Point Server #1:
  [b] 10.209.80.197
  ```

  The installer prompts for this information for the number of virtual IP addresses you want to configure for each CP server.

- Enter the port that the CP server would be listening on.

  ```
  Enter the port in the range [49152, 65535] which the
  Coordination Point Server 10.209.80.197
  would be listening on or simply accept the default
  port suggested: [b] (443)
  ```

9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 1
Coordination Point Server ([VIP or FQHN]:Port):
    1. 10.209.80.197 ([10.209.80.197]:443)
```

10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the SFHA (application cluster).

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

11 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

12 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

The installer also populates the `/etc/vxfenmode` file with the entry `single_cp=1` for such single CP server fencing configuration.

```
Updating client cluster information on Coordination Point Server 10.209.80.197

Adding the client cluster to the Coordination Point Server 10.209.80.197 .......... Done

Registering client node sys1 with Coordination Point Server 10.209.80.197...... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.209.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Updating /etc/vxfenmode file on sys1 ................................. Done
Updating /etc/vxfenmode file on sys2 ......... ....................... Done
```

See “About I/O fencing configuration files” on page 482.

13 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

14 Configure the CP agent on the SFHA (application cluster).

```
Do you want to configure Coordination Point Agent on the
client cluster? [y,n,q] (y)

Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen)

Adding Coordination Point Agent via sys1 ... Done
```

15 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

Refreshing keys or registrations on the existing coordination points for server-based fencing using the installsfha

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss might occur because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose registrations of the cluster nodes, the cluster might panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for server-based I/O fencing using the installsfha

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

```
# /opt/VRTS/install/installsfha<version> -fencing
```

where `<version>` is the specific release version.

See "About the script-based installer" on page 69.

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

Note the location of log files that you can access if there is a problem with the configuration process.

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

The program checks that the local node running the script can communicate with the remote nodes and checks whether SFHA 6.2 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type the number corresponding to the option that suggests to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q] 5
```

4 Ensure that the `/etc/vxfentab` file contains the same coordination point servers that are currently used by the fencing module.

Also, ensure that the disk group mentioned in the `/etc/vxfendg` file contains the same disks that are currently used by the fencing module as coordination disks.

5 Verify the coordination points.

```
For example,
Total number of coordination points being used: 3
Coordination Point Server ([VIP or FQHN]:Port):
      1. 10.198.94.146 ([10.198.94.146]:443)
      2. 10.198.94.144 ([10.198.94.144]:443)
SCSI-3 disks:
      1. emc_clariion0_61
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk policy used for customized fencing: dmp
```

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

```
Updating client cluster information on Coordination Point Server
 IPaddress

Successfully completed the vxfenswap operation
```

The keys on the coordination disks are refreshed.

7 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

8 Do you want to view the summary file? [y,n,q] **(n)**.

Setting the order of existing coordination points for server-based fencing using the installsfha

This section describes the reasons, benefits, considerations, and the procedure to set the order of the existing coordination points for server-based fencing.

About deciding the order of existing coordination points

You can decide the order in which coordination points can participate in a race during a network partition. In a network partition scenario, I/O fencing attempts to contact coordination points for membership arbitration based on the order that is set in the `vxfentab` file.

When I/O fencing is not able to connect to the first coordination point in the sequence it goes to the second coordination point and so on. To avoid a cluster panic, the surviving subcluster must win majority of the coordination points. So, the order must begin with the coordination point that has the best chance to win the race and must end with the coordination point that has the least chance to win the race.

For fencing configurations that use a mix of coordination point servers and coordination disks, you can specify either coordination point servers before coordination disks or disks before servers.

Note: Disk-based fencing does not support setting the order of existing coordination points.

Considerations to decide the order of coordination points

- Choose the coordination points based on their chances to gain membership on the cluster during the race and hence gain control over a network partition. In effect, you have the ability to save a partition.

- First in the order must be the coordination point that has the best chance to win the race. The next coordination point you list in the order must have relatively lesser chance to win the race. Complete the order such that the last coordination point has the least chance to win the race.

Setting the order of existing coordination points using the installsfha

To set the order of existing coordination points

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

   ```
   # /opt/VRTS/install/installsfha<version> -fencing
   ```

   where *<version>* is the specific release version.

   See "About the script-based installer" on page 69.

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

   Note the location of log files that you can access if there is a problem with the configuration process.

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

   The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 6.2 is configured properly.

3. Review the I/O fencing configuration options that the program presents. Type the number corresponding to the option that suggests to set the order of existing coordination points.

   For example:

   ```
   Select the fencing mechanism to be configured in this
   Application Cluster [1-6,q] 6


   Installer will ask the new order of existing coordination points.
   Then it will call vxfenswap utility to commit the
   coordination points change.
   ```

   Warning: The cluster might panic if a node leaves membership before the coordination points change is complete.

4 Review the current order of coordination points.

```
Current coordination points order:
(Coordination disks/Coordination Point Server)
Example,
1)  /dev/vx/rdmp/emc_clariion0_65,/dev/vx/rdmp/emc_clariion0_66,
/dev/vx/rdmp/emc_clariion0_62
2)  [10.198.94.144]:443
3)  [10.198.94.146]:443
b)  Back to previous menu
```

5 Enter the new order of the coordination points by the numbers and separate the order by space [1-3,b,q] **3 1 2**.

```
New coordination points order:
(Coordination disks/Coordination Point Server)
Example,
1) [10.198.94.146]:443
2) /dev/vx/rdmp/emc_clariion0_65,/dev/vx/rdmp/emc_clariion0_66,
/dev/vx/rdmp/emc_clariion0_62
3) [10.198.94.144]:443
```

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

```
Preparing vxfenmode.test file on all systems...
Running vxfenswap...
Successfully completed the vxfenswap operation
```

7 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

8 Do you want to view the summary file? [y,n,q] **(n)**.

9 Verify that the value of `vxfen_honor_cp_order` specified in the `/etc/vxfenmode` file is set to **1**.

```
For example,
vxfen_mode=customized
vxfen_mechanism=cps
port=443
scsi3_disk_policy=dmp
cps1=[10.198.94.146]
vxfendg=vxfencoorddg
cps2=[10.198.94.144]
vxfen_honor_cp_order=1
```

10 Verify that the coordination point order is updated in the output of the `vxfenconfig -l` command.

```
For example,
I/O Fencing Configuration Information:
======================================

single_cp=0
[10.198.94.146]:443 {e7823b24-1dd1-11b2-8814-2299557f1dc0}
/dev/vx/rdmp/emc_clariion0_65  60060160A38B1600386FD87CA8FDDD11
/dev/vx/rdmp/emc_clariion0_66  60060160A38B1600396FD87CA8FDDD11
/dev/vx/rdmp/emc_clariion0_62  60060160A38B16005AA00372A8FDDD11
[10.198.94.144]:443 {01f18460-1dd2-11b2-b818-659cbc6eb360}
```

Setting up non-SCSI-3 I/O fencing in virtual environments using installsfha

If you have installed VCS in virtual environments that do not support SCSI-3 PR-compliant storage, you can configure non-SCSI-3 fencing.

To configure I/O fencing using the installsfha in a non-SCSI-3 PR-compliant setup

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

   ```
   # /opt/VRTS/install/installsfha<version> -fencing
   ```

   Where *<version>* is the specific release version.

   See "About the script-based installer" on page 69.

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

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

   The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA 6.2 is configured properly.

3. For server-based fencing, review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

   ```
   Select the fencing mechanism to be configured in this
   Application Cluster
   [1-7,q] 1
   ```

4. Enter **n** to confirm that your storage environment does not support SCSI-3 PR.

   ```
   Does your storage environment support SCSI3 PR?
   [y,n,q] (y) n
   ```

5. Confirm that you want to proceed with the non-SCSI-3 I/O fencing configuration at the prompt.

6. For server-based fencing, enter the number of CP server coordination points you want to use in your setup.

7. For server-based fencing, enter the following details for each CP server:

   - Enter the virtual IP address or the fully qualified host name.
   - Enter the port address on which the CP server listens for connections. The default value is 443. You can enter a different port address. Valid values are between 49152 and 65535.

   The installer assumes that these values are identical from the view of the Storage Foundation High Availability nodes that host the applications for high availability.

8. For server-based fencing, verify and confirm the CP server information that you provided.
9. Verify and confirm the Storage Foundation High Availability configuration information.

   Review the output as the installer performs the following tasks:

   - Updates the CP server configuration files on each CP server with the following details for only server-based fencing, :
     - Registers each node of the Storage Foundation High Availability with the CP server.
     - Adds CP server user to the CP server.
     - Adds Storage Foundation High Availability to the CP server user.
   - Updates the following configuration files on each node of the Storage Foundation High Availability
     - `/etc/vxfenmode` file
     - `/etc/vxenviron` file
     - `/etc/sysconfig/vxfen` file
     - `/etc/llttab` file
     - `/etc/vxfentab` (only for server-based fencing)
10. Review the output as the installer stops SFHA on each node, starts I/O fencing on each node, updates the VCS configuration file main.cf, and restarts SFHA with non-SCSI-3 fencing.

    For server-based fencing, confirm to configure the CP agent on the Storage Foundation High Availability.
11. Confirm whether you want to send the installation information to Symantec.
12. After the installer configures I/O fencing successfully, note the location of summary, log, and response files that installer creates.

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

Setting up majority-based I/O fencing using installsfha

You can configure majority-based fencing for the cluster using the installsfha .

Perform the following steps to confgure majority-based I/O fencing

1 Start the installsfha with the -fencing option.

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

Where *version* is the specific release version. The installsfha starts with a copyright message and verifies the cluster information.

See "About the script-based installer" on page 69.

Note: Make a note of the log file location which you can access in the event of any issues with the configuration process.

2 Confirm that you want to proceed with the I/O fencing configuration at the prompt. The program checks that the local node running the script can communicate with remote nodes and checks whether SFHA is configured properly.

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

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,b,q] 3
```

Note: The installer will ask the following question. Does your storage environment support SCSI3 PR? [y,n,q,?] Input 'y' if your storage environment supports SCSI3 PR. Other alternative will result in installer configuring non-SCSI3 fencing(NSF).

4 The installer then populates the /etc/vxfenmode file with the appropriate details in each of the application cluster nodes.

```
Updating /etc/vxfenmode file on sys1 ..... Done
Updating /etc/vxfenmode file on sys2 ..... Done
```

5 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

6 Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

7 Verify the fencing configuration.

```
# vxfenadm -d
```

Enabling or disabling the preferred fencing policy

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

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

Preferred fencing is not applicable to majority-based I/O fencing.

See "About preferred fencing" on page 32.

To enable preferred fencing for the I/O fencing configuration

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

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

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

- Make the VCS configuration writable.

```
# haconf -makerw
```

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

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

- Set the value of the system-level attribute FencingWeight for each node in the cluster.
  For example, in a two-node cluster, where you want to assign sys1 five times more weight compared to sys2, run the following commands:

```
# hasys -modify sys1 FencingWeight 50
# hasys -modify sys2 FencingWeight 10
```

- Save the VCS configuration.

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

- Verify fencing node weights using:

  ```
  # vxfenconfig -a
  ```

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

- Make the VCS configuration writable.

  ```
  # haconf -makerw
  ```

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

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

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

  ```
  # hagrp -modify service_group Priority 1
  ```

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

- Save the VCS configuration.

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

5 To enable site-based race policy, perform the following steps:

- Make the VCS configuration writable.

  ```
  # haconf -makerw
  ```

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

  ```
  # haclus -modify PreferredFencingPolicy Site
  ```

- Set the value of the site-level attribute Preference for each site.

  ```
  For example,
  # hasite -modify Pune Preference 2
  ```

- Save the VCS configuration.

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

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

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

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

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

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

```
# haconf -makerw
# haclus -modify PreferredFencingPolicy Disabled
# haconf -dump -makero
```

Section 4

Installation using the web-based installer

- Chapter 10. Installing SFHA
- Chapter 11. Configuring SFHA

Chapter 10

Installing SFHA

This chapter includes the following topics:

- About the web-based installer
- Before using the web-based installer
- Starting the web-based installer
- Obtaining a security exception on Mozilla Firefox
- Performing a preinstallation check with the web-based installer
- Setting installer options with the web-based installer
- Installing SFHA with the web-based installer

About the web-based installer

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

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

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

The web installer creates log files whenever the web installer operates. While the installation processes operate, 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. Symantec recommends that you keep these files for auditing, debugging, and future use.

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

See "Before using the web-based installer" on page 170.

See "Starting the web-based installer" on page 170.

Before using the web-based installer

The web-based installer requires the following configuration.

Table 10-1 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Symantec products.	Must be a supported platform for Symantec Storage Foundation 6.2.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must be at one of the supported operating system update levels.
Administrative system	The system where you run the web browser to perform the installation.	Must have a web browser. Supported browsers: ■ Internet Explorer 6, 7, and 8 ■ Firefox 3.x and later

Starting the web-based installer

This section describes starting the web-based installer.

To start the web-based installer

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

```
# ./webinstaller start
```

The `webinstaller` script displays a URL. Note this URL.

Note: If you do not see the URL, please check your firewall and iptables settings. If you have configured a firewall, ensure that the firewall settings allow access to the port 14172. You can alternatively use the `-port` option to use a free port instead.

You can use the following command to display the details about ports used by webinstaller and its status:

```
# ./webinstaller status
```

2 On the administrative server, start the web browser.

3 Navigate to the URL that the script displayed.

4 Certain browsers may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

5 Log in as superuser.

Obtaining a security exception on Mozilla Firefox

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

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

To obtain a security exception

1 Click **Or you can add an exception** link.

2 Click **I Understand the Risks**, or **You can add an exception**.

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 preinstallation check with the web-based installer

This section describes performing a preinstallation check with the web-based installer.

To perform a preinstallation check

1 Start the web-based installer.

See "Starting the web-based installer" on page 170.

2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.

3 Indicate the systems on which to perform the precheck. Enter one or more system names, separated by spaces. Click **Next**.

4 The installer performs the precheck and displays the results.

5 If the validation completes successfully, click **Next**. The installer prompts you to begin the installation. Click **Yes** to install on the selected system. Click **No** to install later.

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

Setting installer options with the web-based installer

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

The supported options follow:

- `-serial`
- `-require` *`path_to_patch_file`*
- `-mediapath` *`directory_path_to_install_media`*
- `-logpath` *`directory_path_to_save_logs`*
- `-tmppath` *`directory_path_to_save_temp_files`*

See "Installation script options" on page 451.

To use installer options

1. On the web-installer's entry page, click the **Advanced Options** link.
2. In the Command line options field, enter the option that you want to use.

   For example, if you want to use the serial option and the logpath option, enter:

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

   Where */opt/VRTS/install/advlogs* is the path that you want to use. Separate the command with a space.
3. Click the **OK** button and proceed.

Installing SFHA with the web-based installer

This section describes installing SFHA with the Symantec web-based installer.

To install SFHA using the web-based installer

1. Perform preliminary steps.

   See "Performing a preinstallation check with the web-based installer" on page 172.
2. Start the web-based installer.

   See "Starting the web-based installer" on page 170.
3. On the License agreement page, read the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
4. Indicate the systems where you want to install. Separate multiple system names with spaces. Click **Next**.
5. If you have not yet configured a communication mode among systems, you have the option to let the installer configure ssh or rsh. If you choose to allow this configuration, select the communication mode and provide the superuser passwords for the systems.
6. After the validation completes successfully, click **Next** to install SFHA on the selected system.
7. After the installation completes, you must choose your licensing method.

   On the license page, select one of the following radio buttons:

   - Enable keyless licensing and complete system licensing later

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

For more information, go to the following website:

http://go.symantec.com/sfhakeyless

Click **Next**

- Enter a valid license key
  If you have a valid license key, input the license key and click **Next**.

8 For Storage Foundation and High Availability, click **Next**.If the installer prompts you to restart the system, then restart the system and invoke the web-based installer again for configuration. If the installer does not require a restart, go to step 9.

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

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

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

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

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

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

10 If you are prompted, enter the option to specify whether you want to send your installation information to Symantec.

```
Installation procedures and diagnostic information were saved in
the log files under directory
/var/tmp/installer-<platform>-<uuid>. Analyzing this information
helps Symantec discover and fix failed operations performed by
the installer. Would you like to send the information about this
installation to Symantec to help improve installation in the
future? [y,n,q,?]
```

Click **Finish**.

Chapter 11

Configuring SFHA

This chapter includes the following topics:

- Configuring SFHA using the web-based installer

Configuring SFHA using the web-based installer

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

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

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

To configure SFHA on a cluster

1. Start the web-based installer.

   See "Starting the web-based installer" on page 170.

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

Task	Configure a Product
Product	Symantec Storage Foundation and High Availability

   Click **Next**.

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

Example: **sys1 sys2**

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

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

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

Would you like to configure I/O fencing on the cluster?, click **Yes**.

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

See "Configuring SFHA for data integrity using the web-based installer" on page 181.

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

Where `<version>` is the specific release version.

See "About the script-based installer" on page 69.

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

Cluster Name	Enter a unique cluster name.
Cluster ID	Enter a unique cluster ID. Note that you can have the installer check to see if the cluster ID is unique. Symantec recommends that you use the installer to check for duplicate cluster IDs in multi-cluster environments.
Check duplicate cluster ID	Select the check box if you want the installer to verify if the given cluster ID is unique in your private network. The verification is performed after you specify the heartbeat details in the following pages. The verification takes some time to complete.
LLT Type	Select an LLT type from the list. You can choose to configure LLT over UDP or LLT over Ethernet or LLT over RDMA.
Number of Heartbeats	Choose the number of heartbeat links you want to configure. See "Setting up the private network" on page 58.
Additional Low Priority Heartbeat NIC	Select the check box if you want to configure a low priority link. The installer configures one heartbeat link as low priority link. See "Setting up the private network" on page 58.
Unique Heartbeat NICs per system	For LLT over Ethernet, select the check box if you do not want to use the same NIC details to configure private heartbeat links on other systems. For LLT over UDP, this check box is selected by default. For LLT over RDMA, this check box is selected by default.

Click **Next**.

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

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

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

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

Click **Next**.

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

Security	To configure a secure SFHA cluster, select the **Configure secure cluster** check box. If you want to perform this task later, do not select the **Configure secure cluster** check box. You can use the `-security` option of the `installsfha<version>`.
Virtual IP	■ Select the **Configure Virtual IP** check box. ■ If each system uses a separate NIC, select the **Configure NICs for every system separately** check box. ■ Select the interface on which you want to configure the virtual IP. ■ Enter a virtual IP address and value for the netmask. You can use an IPv4 or an IPv6 address.
VCS Users	■ Reset the password for the Admin user, if necessary. ■ Select the **Configure VCS users** option. ■ Click **Add** to add a new user. Specify the user name, password, and user privileges for this user.

SMTP

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

SNMP

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

GCO

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

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

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

Click **Next**.

8 The installer displays the following question before the install stops the product processes:

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

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

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

If you did not choose to configure I/O fencing in step 4, then skip to step 12. Go to step 11 to configure fencing.

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

Configure Coordination Point client based fencing	Choose this option to configure server-based I/O fencing.
Configure disk based fencing	Choose this option to configure disk-based I/O fencing.
Configure majority based fencing	Choose this option to configure majority based I/O fencing.

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

See “Configuring SFHA for data integrity using the web-based installer” on page 181.

12 Click **Next** to complete the process of configuring SFHA.

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

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

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

Configuring SFHA for data integrity using the web-based installer

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

See “Configuring SFHA using the web-based installer” on page 175.

See “ About planning to configure I/O fencing” on page 75.

Ways to configure I/O fencing using the web-based installer:

- See “Configuring disk-based fencing for data integrity using the web-based installer” on page 182.
- See “Configuring server-based fencing for data integrity using the web-based installer” on page 184.
- See “Configuring fencing in disabled mode using the web-based installer” on page 186.
- See “Configuring fencing in majority mode using the web-based installer” on page 187.
- See “Replacing, adding, or removing coordination points using the web-based installer” on page 188.

- See “Refreshing keys or registrations on the existing coordination points using web-based installer” on page 190.
- See “Setting the order of existing coordination points using the web-based installer” on page 192.

Configuring disk-based fencing for data integrity using the web-based installer

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

See “Configuring SFHA using the web-based installer” on page 175.

See “ About planning to configure I/O fencing” on page 75.

To configure SFHA for data integrity

1. Start the web-based installer.

   See “Starting the web-based installer” on page 170.

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

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

   Click **Next**.

3. Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4. On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5. On the Select Fencing Type page, select the `Configure disk-based fencing` option.

6. In the Confirmation dialog box that appears, confirm whether your storage environment supports SCSI-3 PR.

   You can configure non-SCSI-3 fencing in a virtual environment that is not SCSI-3 PR compliant.

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

Click **Next**.

8 On the Configure Fencing page, specify the following information:

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

- If you selected one of the disk groups that is listed, choose the fencing disk policy for the disk group.
- If you selected the **Create a new disk group** option, make sure you have SCSI-3 PR enabled disks, and click **Yes** in the confirmation dialog box.

  Click **Next**.

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

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

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

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

10 Verify and confirm the I/O fencing configuration information.

The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

11 If you want to configure the Coordination Point agent on the client cluster, do the following:

- At the prompt for configuring the Coordination Point agent on the client cluster, click **Yes** and enter the Coordination Point agent service group name.
- If you want to set the LevelTwoMonitorFreq attribute, click Yes at the prompt and enter a value (0 to 65535).
- Follow the rest of the prompts to complete the Coordination Point agent configuration.

12 Click **Next** to complete the process of configuring I/O fencing.

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

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

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

Configuring server-based fencing for data integrity using the web-based installer

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

See "Configuring SFHA using the web-based installer" on page 175.

See " About planning to configure I/O fencing" on page 75.

To configure SFHA for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 170.

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

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

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 On the Select Fencing Type page, select the `Configure Coordination Point client based fencing` option.

6 In the Confirmation dialog box that appears, confirm whether your storage environment supports SCSI-3 PR.

You can configure non-SCSI-3 fencing in a virtual environment that is not SCSI-3 PR compliant.

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

Click **Next**.

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

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

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

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

10 In the Confirmation dialog box that appears, confirm whether the coordination points information you provided is correct, and click **Yes**.

11 Verify and confirm the I/O fencing configuration information.

The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

12 If you want to configure the Coordination Point agent on the client cluster, do the following:

- At the prompt for configuring the Coordination Point agent on the client cluster, click **Yes** and enter the Coordination Point agent service group name.
- Follow the rest of the prompts to complete the Coordination Point agent configuration.

13 Click **Next** to complete the process of configuring I/O fencing.

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

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

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

Configuring fencing in disabled mode using the web-based installer

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

See "Configuring SFHA using the web-based installer" on page 175.

See " About planning to configure I/O fencing" on page 75.

To configure SFHA for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 170.

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

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

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 Fencing may be enabled, installer may prompt whether you want to reconfigure it.

Click **Yes**.

6 On the Select Fencing Type page, select the `Configure fencing in disabled mode` option.

7 Installer stops VCS before applying the selected fencing mode to the cluster.

Note: Unfreeze any frozen service group and unmount any file system that is mounted in the cluster.

Click **Yes**.

8 Installer restarts VCS on all systems of the cluster. I/O fencing is disabled.

9 Verify and confirm the I/O fencing configuration information.

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

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

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

Configuring fencing in majority mode using the web-based installer

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

See "Configuring SFHA using the web-based installer" on page 175.

See " About planning to configure I/O fencing" on page 75.

To configure SFHA for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 170.

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

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

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 Fencing may be enabled, installer may prompt whether you want to reconfigure it.

Click **Yes**.

6 On the Select Fencing Type page, select the `Configure fencing in majority mode` option.

7 Installer stops VCS before applying the selected fencing mode to the cluster.

Note: Unfreeze any frozen service group and unmount any file system that is mounted in the cluster.

Click **Yes**.

8 Installer restarts VCS on all systems of the cluster. I/O fencing is in majority mode.

9 Verify and confirm the I/O fencing configuration information.

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

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

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

Replacing, adding, or removing coordination points using the web-based installer

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

This procedure does not apply to majority-based I/O fencing.

See “Configuring SFHA using the web-based installer” on page 175.

See “ About planning to configure I/O fencing” on page 75.

To configure SFHA for data integrity

1. Start the web-based installer.

   See "Starting the web-based installer" on page 170.

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

Task	I/O Fencing configuration
Product	Symantec Storage Foundation and High Availability

   Click **Next**.

3. Verify the cluster information that the installer presents and confirm whether you want to configure I/O Fencing on the cluster.

4. On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5. Fencing may be enabled, installer may prompt whether you want to reconfigure it.

   Click **Yes**.

6. On the Select Fencing Type page, select the `Replace/Add/Remove coordination points` option.

7. The installer prompts to select the coordination points you want to remove from the currently configured coordination points.

   Click **Next**.

8. Provide the number of Coordination point server and disk coordination points to be added to the configuration.

   Click **Next**.

9. Provide the number of virtual IP addresses or Fully Qualified Host Name (FQHN) used for each coordination point server.

   Click **Next**.

10. Provide the IP or FQHN and port number for each coordination point server.

    Click **Next**.

11 Installer prompts to confirm the online migration coordination point servers.

Click **Yes**.

12 Installer proceeds with migration of the new coordination point servers. VCS is restarted during configuration.

Click **Next**.

13 You can add a Coordination Point agent to the client cluster and also provide name to the agent.

14 Click **Next**.

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

16 Select the check box to specify whether you want to send your installation information to Symantec.

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

Refreshing keys or registrations on the existing coordination points using web-based installer

This procedure does not apply to majority-based I/O fencing.

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points using web-based installer

1. Start the web-based installer.

   See "Starting the web-based installer" on page 170.

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

Task	I/O Fencing configuration
Product	Symantec Storage Foundation and High Availability

   Click **Next**.

3. Verify the cluster information that the installer presents and click **Yes** to confirm whether you want to configure I/O fencing on the cluster.

4. On the **Select Cluster** page, enter the system name and click **Yes** to confirm cluster information.

5. On the **Select Cluster** page, click **Next** when the installer completes the cluster verification successfully.

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

6. The installer may prompt you to reconfigure fencing if it is already enabled. Click **Yes** to reconfigure fencing.

7. On the **Select Fencing Type** page, select the `Refresh keys/registrations on the existing coordination points` option.

8. Ensure that the `/etc/vxfenmode` file contains the same coordination point servers that are currently used by the fencing module.

9. Ensure that the disk group mentioned in the `/etc/vxfenmode` file contains the same disks that are currently used by the fencing module as coordination disks.

10. Installer lists the reasons for the loss of registrations.

    Click **OK**.

11. Verify the coordination points.

    Click **Yes** if the information is correct.

12 Installer updates the client cluster information on the coordination point servers.

Click **Next**.

Installer prepares the `vxfenmode` file on all nodes and runs the `vxfenswap` utility to refresh registrations on the coordination points.

13 On the **Completion** page, view the `summary` file, `log` file, or `response` file to confirm the configuration.

14 Select the check box to specify whether you want to send your installation information to Symantec.

Click **Finish**.

Setting the order of existing coordination points using the web-based installer

This section describes the reasons, benefits, considerations, and the procedure to set the order of the existing coordination points using the web-based installer.

It does not apply to majority-based I/O fencing.

About deciding the order of existing coordination points

You can decide the order in which coordination points can participate in a race during a network partition. In a network partition scenario, I/O fencing attempts to contact coordination points for membership arbitration based on the order that is set in the `vxfenmode` file.

When I/O fencing is not able to connect to the first coordination point in the sequence it goes to the second coordination point and so on. To avoid a cluster panic, the surviving subcluster must win majority of the coordination points. So, the order must begin with the coordination point that has the best chance to win the race and must end with the coordination point that has the least chance to win the race

For fencing configurations that use a mix of coordination point servers and coordination disks, you can either specify coordination point servers before coordination point disks or disks before servers.

Note: Disk-based fencing does not support setting the order of existing coordination points.

Considerations to decide the order of coordination points

- Choose coordination points based on their chances gain membership on the cluster during the race and hence gain control over a network partition. In effect, you have the ability to save a partition.

- First in the order must be the coordination point that has the best chance to win the race. The next coordination point you list in the order must have relatively lesser chance to win the race. Complete the order such that the last coordination point has the least chance to win the race.

Setting the order of existing coordination points using the web-based installer

To set the order of existing coordination points for server-based fencing using the web-based installer

1. Start the web-based installer.

   See “Starting the web-based installer” on page 170.

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

Task	I/O Fencing configuration
Product	Symantec Storage Foundation and High Availability

   Click **Next**.

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

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

6. The installer may prompt you to reconfigure fencing if it is already enabled. Click **Yes** to reconfigure fencing.

   Click **Yes**.

7. On the **Select Fencing Type** page, select the `Set the order of existing coordination points` option.
8. Confirm **OK** at the installer message about the procedure.
9. Decide the new order by moving the existing coordination points to the box on the window in the order you want. If you want to change the current order of coordination points, click **Reset** and start again.
10. Click **Next** if the information is correct.

11 On the **Confirmation** window, click **Yes**.

Installer prepares the `vxfenmode` file on all nodes and runs the `vxfenswap` utility to update the new order of coordination points.

12 On the **Completion** page, view the summary file, log file, or response file to confirm the configuration.

13 Select the check box to specify whether you want to send your installation information to Symantec.

Click **Finish**.

Section 5

Automated installation using response files

- Chapter 12. Performing an automated SFHA installation
- Chapter 13. Performing an automated SFHA configuration
- Chapter 14. Performing an automated I/O fencing configuration using response files

Chapter 12

Performing an automated SFHA installation

This chapter includes the following topics:

- Installing SFHA using response files
- Response file variables to install Storage Foundation and High Availability
- Sample response file for SFHA install

Installing SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA installation on one cluster to install SFHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

```
# ./installer -makeresponsefile
```

See "About the script-based installer" on page 69.

To install SFHA using response files

1. Make sure the systems where you want to install SFHA meet the installation requirements.
2. Make sure that the preinstallation tasks are completed.
3. Copy the response file to one of the cluster systems where you want to install SFHA.
4. Edit the values of the response file variables as necessary.

5 Mount the product disc and navigate to the directory that contains the installation program.

6 Start the installation from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file

# ./installsfha -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

7 Complete the SFHA post-installation tasks.

For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.

Response file variables to install Storage Foundation and High Availability

Table 12-1 lists the response file variables that you can define to install SFHA.

Table 12-1 Response file variables for installing SFHA

Variable	Description
CFG{opt}{install}	Installs SFHA RPMs. Configuration can be performed at a later time using the `-configure` option. List or scalar: scalar Optional or required: optional
CFG{opt}{installallpkgs} or CFG{opt}{installrecpkgs} or CFG{opt}{installminpkgs}	Instructs the installer to install SFHA RPMs based on the variable that has the value set to 1: ■ installallpkgs: Installs all RPMs ■ installrecpkgs: Installs recommended RPMs ■ installminpkgs: Installs minimum RPMs **Note:** Set only one of these variable values to 1. In addition to setting the value of one of these variables, you must set the variable `CFG{opt}{install}` to 1. List or scalar: scalar Optional or required: required

Table 12-1 Response file variables for installing SFHA *(continued)*

Variable	Description
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{opt}{license}	Installs the product with permanent license. List or scalar: scalar Optional or required: optional
CFG{keys}{hostname}	List of keys to be registered on the system if the variable `CFG{opt}{vxkeyless}` is set to 0 or if the variable `CFG{opt}{licence}` is set to 1. List or scalar: scalar Optional or required: optional
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional

Table 12-1 Response file variables for installing SFHA *(continued)*

Variable	Description
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{updatekeys}	Updates the keyless license to the current version. List or scalar: scalar Optional or required: optional
CFG{opt}{rsh}	Defines that *rsh* must be used instead of ssh as the communication method between systems. List or scalar: scalar Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
CFG{opt}{prodmode}	List of modes for product List or scalar: list Optional or required: optional

Sample response file for SFHA install

The following example shows a response file for installing Storage Foundation High Availability.

```
##############################################
#Auto generated sfha responsefile #
##############################################


our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
```

```
$CFG{opt}{prodmode}="SF Enterprise HA";
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw( sys1 sys2 ) ];
$CFG{keys}{sys1}=["XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX"];
$CFG{keys}{sys2}=["XXXX-XXXX-XXXX-XXXX-XXXX-XXXX-XXXX"];
$CFG{opt}{logpath}="/opt/VRTS/install/logs/HxRT-601-xxxx";

1;
```

Chapter 13

Performing an automated SFHA configuration

This chapter includes the following topics:

- Configuring SFHA using response files
- Response file variables to configure Storage Foundation and High Availability
- Sample response file for SFHA configuration

Configuring SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA configuration on one cluster to configure SFHA on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

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

```
# ./installsfha -makeresponsefile -configure
```

To configure SFHA using response files

1. Make sure the SFHA RPMs are installed on the systems where you want to configure SFHA.
2. Copy the response file to one of the cluster systems where you want to configure SFHA.

3 Edit the values of the response file variables as necessary.

To configure optional features, you must define appropriate values for all the response file variables that are related to the optional feature.

See "Response file variables to configure Storage Foundation and High Availability" on page 202.

4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsfha<version>
-responsefile /tmp/response_file
```

Where *<version>* is the specific release version, and */tmp/response_file* is the response file's full path name.

See "About the script-based installer" on page 69.

Response file variables to configure Storage Foundation and High Availability

Table 13-1 lists the response file variables that you can define to configure SFHA.

Table 13-1 Response file variables specific to configuring Storage Foundation and High Availability

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the RPMs are already installed. (Required) Set the value to 1 to configure SFHA.
CFG{accepteula}	Scalar	Specifies whether you agree with EULA.pdf on the media. (Required)
CFG{systems}	List	List of systems on which the product is to be configured. (Required)

Table 13-1 Response file variables specific to configuring Storage Foundation and High Availability *(continued)*

Variable	List or Scalar	Description
CFG{prod}	Scalar	Defines the product to be configured. The value is SFHA62 for SFHA (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{secusrgrps}	List	Defines the user groups which get read access to the cluster. (Optional)
CFG {rootsecusrgrps}	Scalar	Defines the read access to the cluster only for root and other users or usergroups which are granted explicit privileges on VCS objects. (Optional)
CFG{opt}{rsh}	Scalar	Defines that *rsh* must be used instead of ssh as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. **Note:** The installer copies the response files and summary files also to the specified *logpath* location. (Optional)

Table 13-1 Response file variables specific to configuring Storage Foundation and High Availability *(continued)*

Variable	List or Scalar	Description
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec website. The value 0 indicates that the installation logs are not uploaded to the Symantec website. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtprecp, and smtprsev), the SNMP trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Table 13-2 lists the response file variables that specify the required information to configure a basic SFHA cluster.

Table 13-2 Response file variables specific to configuring a basic SFHA cluster

Variable	List or Scalar	Description
CFG{vcs_clusterid}	Scalar	An integer between 0 and 65535 that uniquely identifies the cluster. (Required)
CFG{vcs_clustername}	Scalar	Defines the name of the cluster. (Required)
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)

Table 13-2 Response file variables specific to configuring a basic SFHA cluster *(continued)*

Variable	List or Scalar	Description
CFG{fencingenabled}	Scalar	In a SFHA configuration, defines if fencing is enabled. Valid values are 0 or 1. (Required)

Table 13-3 lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 13-3 Response file variables specific to configuring private LLT over Ethernet

Variable	List or Scalar	Description
CFG{vcs_lltlink#} {"system"}	Scalar	Defines the NIC to be used for a private heartbeat link on each system. Atleast two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. See "Setting up the private network" on page 58. You must enclose the system name within double quotes. (Required)
CFG{vcs_lltlinklowpri#} {"system"}	Scalar	Defines a low priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication. If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on. You must enclose the system name within double quotes. (Optional)

Table 13-4 lists the response file variables that specify the required information to configure LLT over UDP.

Table 13-4 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltoverudp}=1	Scalar	Indicates whether to configure heartbeat link using LLT over UDP. (Required)
CFG{vcs_udplink<n>_address} {<sys1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG {vcs_udplinklowpri<n>_address} {<sys1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<sys1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)

Table 13-4 Response file variables specific to configuring LLT over UDP *(continued)*

Variable	List or Scalar	Description
CFG{vcs_udplinklowpri<n>_port} {<sys1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_netmask} {<sys1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG {vcs_udplinklowpri<n>_netmask} {<sys1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 13-5 lists the response file variables that specify the required information to configure virtual IP for SFHA cluster.

Table 13-5 Response file variables specific to configuring virtual IP for SFHA cluster

Variable	List or Scalar	Description
CFG{vcs_csgnic} {system}	Scalar	Defines the NIC device to use on a system. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_csgvip}	Scalar	Defines the virtual IP address for the cluster. (Optional)
CFG{vcs_csgnetmask}	Scalar	Defines the Netmask of the virtual IP address for the cluster. (Optional)

Table 13-6 lists the response file variables that specify the required information to configure the SFHA cluster in secure mode.

Table 13-6 Response file variables specific to configuring SFHA cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonenode}	Scalar	Specifies that the securityonenode option is being used.
CFG{securityonenode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. ■ 1—Configure the first node ■ 2—Configure the other node
CFG{secusrgrps}	List	Defines the user groups which get read access to the cluster. List or scalar: list Optional or required: optional

Table 13-6 Response file variables specific to configuring SFHA cluster in secure mode *(continued)*

Variable	List or Scalar	Description
CFG{rootsecusrgrps}	Scalar	Defines the read access to the cluster only for root and other users or user groups which are granted explicit privileges in VCS objects. (Optional)
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.
CFG{vcs_eat_security_fips}	Scalar	Specifies that the enabled security is FIPS compliant.

Table 13-7 lists the response file variables that specify the required information to configure VCS users.

Table 13-7 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	List of encoded passwords for VCS users The value in the list can be "Administrators Operators Guests" **Note:** The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list. (Optional)
CFG{vcs_username}	List	List of names of VCS users (Optional)

Table 13-7 Response file variables specific to configuring VCS users *(continued)*

Variable	List or Scalar	Description
CFG{vcs_userpriv}	List	List of privileges for VCS users **Note:** The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list. (Optional)

Table 13-8 lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 13-8 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)
CFG{vcs_smtprsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

Table 13-9 lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 13-9 Response file variables specific to configuring VCS notifications using SNMP

Variable	List or Scalar	Description
CFG{vcs_snmpport}	Scalar	Defines the SNMP trap daemon port (default=162). (Optional)
CFG{vcs_snmpcons}	List	List of SNMP console system names (Optional)
CFG{vcs_snmpcsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

Table 13-10 lists the response file variables that specify the required information to configure SFHA global clusters.

Table 13-10 Response file variables specific to configuring SFHA global clusters

Variable	List or Scalar	Description
CFG{vcs_gconic} {system}	Scalar	Defines the NIC for the Virtual IP that the Global Cluster Option uses. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_gcovip}	Scalar	Defines the virtual IP address to that the Global Cluster Option uses. (Optional)
CFG{vcs_gconetmask}	Scalar	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. (Optional)

Sample response file for SFHA configuration

The following example shows a response file for configuring Storage Foundation High Availability.

```
##############################################
#Auto generated sfha responsefile #
##############################################



our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{rsh}=1;
$CFG{vcs_allowcomms}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vvr}=1;
$CFG{opt}{prodmode}="SF Enterprise HA";
$CFG{opt}{configure}=1;
$CFG{upi}="SF";
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw( sys1 sys2 ) ];
$CFG{vm_restore_cfg}{sys1}=0;
$CFG{vm_restore_cfg}{sys2}=0;
$CFG{vcs_clusterid}=127;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_username}=[ qw(admin operator) ];
$CFG{vcs_userenpw}=[ qw(JlmElgLimHmmKumGlj bQOsOUnVQoOUnTQsOSnUQuOUnPQtOS) ];
$CFG{vcs_userpriv}=[ qw(Administrators Operators) ];
$CFG{vcs_lltlink1}{sys1}="eth1";
$CFG{vcs_lltlink2}{sys1}="eth2";
$CFG{vcs_lltlink1}{sys2}="eth1";
$CFG{vcs_lltlink2}{sys2}="eth2";
$CFG{opt}{logpath}="/opt/VRTS/install/logs/installsf-xxxxxx/installsf-xxxxxx.response";

1;
```

Chapter 14

Performing an automated I/O fencing configuration using response files

This chapter includes the following topics:

- Configuring I/O fencing using response files
- Response file variables to configure disk-based I/O fencing
- Sample response file for configuring disk-based I/O fencing
- Response file variables to configure server-based I/O fencing
- Sample response file for configuring non-SCSI-3 I/O fencing
- Response file variables to configure non-SCSI-3 I/O fencing
- Response file variables to configure majority-based I/O fencing
- Sample response file for configuring majority-based I/O fencing

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for SFHA.

To configure I/O fencing using response files

1. Make sure that SFHA is configured.
2. Based on whether you want to configure disk-based or server-based I/O fencing, make sure you have completed the preparatory tasks.

   See " About planning to configure I/O fencing" on page 75.
3. Copy the response file to one of the cluster systems where you want to configure I/O fencing.

   See "Sample response file for configuring disk-based I/O fencing" on page 217.

   See "Sample response file for configuring server-based I/O fencing" on page 219.

   See "Sample response file for configuring non-SCSI-3 I/O fencing" on page 220.

   See "Sample response file for configuring majority-based I/O fencing" on page 222.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to configure disk-based I/O fencing" on page 214.

   See "Response file variables to configure server-based I/O fencing" on page 217.

   See "Response file variables to configure non-SCSI-3 I/O fencing" on page 220.

   See "Response file variables to configure majority-based I/O fencing" on page 222.
5. Start the configuration from the system to which you copied the response file. For example:

   ```
   # /opt/VRTS/install/installsfha<version>
   -responsefile /tmp/response_file
   ```

   Where *<version>* is the specific release version, and */tmp/response_file* is the response file's full path name.

   See "About the script-based installer" on page 69.

Response file variables to configure disk-based I/O fencing

Table 14-1 lists the response file variables that specify the required information to configure disk-based I/O fencing for SFHA.

Table 14-1 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. ■ 1—Configure Coordination Point client-based I/O fencing ■ 2—Configure disk-based I/O fencing ■ 3—Configure majority-based I/O fencing ■ 4—Configure I/O fencing in disabled mode ■ 5—Replace/Add/Remove coordination points ■ 6—Refresh keys/registrations on the existing coordination points ■ 7—Set the order of existing coordination points (Required)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) **Note:** You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.
CFG{fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) **Note:** You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.

Table 14-1 Response file variables specific to configuring disk-based I/O fencing *(continued)*

Variable	List or Scalar	Description
CFG{fencing_cpagent_monitor_freq}	Scalar	Specifies the frequency at which the Coordination Point Agent monitors for any changes to the Coordinator Disk Group constitution. **Note:** Coordination Point Agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidently deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles. If LevelTwoMonitorFreq attribute is not set, the agent will not monitor any changes to the Coordinator Disk Group. 0 means not to monitor the Coordinator Disk Group constitution.
CFG {fencing_config_cpagent}	Scalar	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.
CFG {fencing_cpagentgrp}	Scalar	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the **fencing_config_cpagent** field is given a value of '0'.

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 214.

```
#
# Configuration Values:
#
our %CFG;

$CFG{fencing_config_cpagent}=1;
$CFG{fencing_cpagent_monitor_freq}=5;
$CFG{fencing_cpagentgrp}="vxfen";
$CFG{fencing_dgname}="fencingdg1";
$CFG{fencing_newdg_disks}=[ qw(emc_clariion0_155
 emc_clariion0_162 emc_clariion0_163) ];
$CFG{fencing_option}=2;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{fencing_cpagent_monitor_freq}=5;

$CFG{prod}="SFHA62";

$CFG{systems}=[ qwsys1sys2 ];
$CFG{vcs_clusterid}=32283;
$CFG{vcs_clustername}="clus1";
1;
```

Response file variables to configure server-based I/O fencing

You can use a coordination point server-based fencing response file to configure server-based customized I/O fencing.

Table 14-2 lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table 14-2 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the `fencing_config_cpagent` field is given a value of '0'.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_reusedg}	This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks). Enter either a "1" or "0". Entering a "1" indicates reuse, and entering a "0" indicates do not reuse. When reusing an existing DG name for the mixed mode fencing configuration. you need to manually add a line of text , such as "$CFG{fencing_reusedg}=0" or "$CFG{fencing_reusedg}=1" before proceeding with a silent installation.
CFG {fencing_dgname}	The name of the disk group to be used in the customized fencing, where at least one disk is being used.
CFG {fencing_disks}	The disks being used as coordination points if any.
CFG {fencing_ncp}	Total number of coordination points being used, including both CP servers and disks.
CFG {fencing_ndisks}	The number of disks being used.

Table 14-2 Coordination point server (CP server) based fencing response file definitions *(continued)*

Response file field	Definition
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_cps_ports}	The port that the virtual IP address or the fully qualified host name of the CP server listens on.
CFG{fencing_option}	Specifies the I/O fencing configuration mode. ■ 1—Configure Coordination Point client-based I/O fencing ■ 2—Configure disk-based I/O fencing ■ 3—Configure majority-based I/O fencing ■ 4—Configure I/O fencing in disabled mode ■ 5—Replace/Add/Remove coordination points ■ 6—Refresh keys/registrations on the existing coordination points ■ 7—Set the order of existing coordination points

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing:

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cps_vips}{"10.200.117.145"}=[ qw(10.200.117.145) ];
$CFG{fencing_dgname}="vxfencoorddg";
$CFG{fencing_disks}=[ qw(emc_clariion0_37 emc_clariion0_13
emc_clariion0_12) ];
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=2;
$CFG{fencing_cps_ports}{"10.200.117.145"}=443;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Sample response file for configuring non-SCSI-3 I/O fencing

The following is a sample response file used for non-SCSI-3 I/O fencing :

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.198.89.251 10.198.89.252 10.198.89.253) ];
$CFG{fencing_cps_vips}{"10.198.89.251"}=[ qw(10.198.89.251) ];
$CFG{fencing_cps_vips}{"10.198.89.252"}=[ qw(10.198.89.252) ];
$CFG{fencing_cps_vips}{"10.198.89.253"}=[ qw(10.198.89.253) ];
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=0;
$CFG{fencing_cps_ports}{"10.198.89.251"}=443;
$CFG{fencing_cps_ports}{"10.198.89.252"}=443;
$CFG{fencing_cps_ports}{"10.198.89.253"}=443;
$CFG{non_scsi3_fencing}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Response file variables to configure non-SCSI-3 I/O fencing

Table 14-3 lists the fields in the response file that are relevant for non-SCSI-3 I/O fencing.

See "About I/O fencing for SFHA in virtual machines that do not support SCSI-3 PR" on page 30.

Table 14-3 Non-SCSI-3 I/O fencing response file definitions

Response file field	Definition
CFG{non_scsi3_fencing}	Defines whether to configure non-SCSI-3 I/O fencing. Valid values are 1 or 0. Enter 1 to configure non-SCSI-3 I/O fencing.

Table 14-3 Non-SCSI-3 I/O fencing response file definitions *(continued)*

Response file field	Definition
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent. **Note:** This variable does not apply to majority-based fencing.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the `fencing_config_cpagent` field is given a value of '0'. This variable does not apply to majority-based fencing.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers. **Note:** This variable does not apply to majority-based fencing.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server. **Note:** This variable does not apply to majority-based fencing.
CFG {fencing_ncp}	Total number of coordination points (CP servers only) being used. **Note:** This variable does not apply to majority-based fencing.
CFG {fencing_cps_ports}	The port of the CP server that is denoted by *cps* . **Note:** This variable does not apply to majority-based fencing.

Response file variables to configure majority-based I/O fencing

Table 14-4 lists the response file variables that specify the required information to configure disk-based I/O fencing for SFHA.

Table 14-4 Response file variables specific to configuring majority-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled-based fencing ■ 4—Online fencing migration ■ 5—Refresh keys/registrations on the existing coordination points ■ 6—Change the order of existing coordination points ■ 7—Majority-based fencing (Required)

Sample response file for configuring majority-based I/O fencing

```
$CFG{fencing_option}=7;
$CFG{config_majority_based_fencing}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw(sys1  sys2) ];
```

```
$CFG{vcs_clusterid}=59082;
$CFG{vcs_clustername}="clus1";
```

Section 6

Installation using operating system-specific methods

- Chapter 15. Installing SFHA using operating system-specific methods
- Chapter 16. Configuring SFHA clusters for data integrity

Chapter 15

Installing SFHA using operating system-specific methods

This chapter includes the following topics:

- About installing SFHA using operating system-specific methods
- Installing SFHA using Kickstart
- Sample Kickstart configuration file
- Installing Storage Foundation and High Availability using yum
- Installing SFHA using the Red Hat Satellite server

About installing SFHA using operating system-specific methods

On Linux, you can install SFHA using the following methods:

- You can install SFHA using Kickstart. Kickstart is supported for Red Hat Enterprise Linux 6 (RHEL6) and Red Hat Enterprise Linux 7 (RHEL7).
  See "Installing SFHA using Kickstart" on page 226.
- You can install SFHA using yum. yum is supported for Red Hat Enterprise Linux 6 (RHEL6) and Red Hat Enterprise Linux 7 (RHEL7).
  See "Installing Storage Foundation and High Availability using yum" on page 230.

- You can install SFHA using the Red Hat Satellite server. Red Hat Satellite server is supported for Red Hat Enterprise Linux 6 (RHEL6) and Red Hat Enterprise Linux 7 (RHEL7).
  See "Installing SFHA using the Red Hat Satellite server" on page 235.

Installing SFHA using Kickstart

You can install SFHA using Kickstart. Kickstart is supported for Red Hat Enterprise Linux operating system.

To install SFHA using Kickstart

1. Create a directory for the Kickstart configuration files.

   ```
   # mkdir /kickstart_files/
   ```

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

   - To generate configuration files, enter the following command:

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

     The system lists the files.

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

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

     The output includes the following:

     ```
     The kickstart script for SFHA is generated at
     /kickstart_files/kickstart_sfha62.ks
     ```

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

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

4. Copy the `rpms` directory from the installation media to the NFS location.

5. Verify the contents of the directory.

   ```
   # ls /nfs_mount_kickstart/
   ```

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

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

7 Append the entire modified contents of the Kickstart configuration file to the operating system `ks.cfg` file.

8 Launch the Kickstart installation for the operating system.

9 After the operating system installation is complete, check the file `/var/tmp/kickstart.log` for any errors that are related to the installation of RPMs and product installer scripts.

10 Verify that all the product RPMs have been installed. Enter the following command:

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

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

```
# /opt/VRTS/install/installsfha<version> -configure node1 node2
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 69.

12 Verify that all the configured `llt` links and gab ports have started.

13 Verify that the product is configured properly by entering commands such as `hastatus` and `lltstat -n`.

```
# hastatus -sum

-- SYSTEM STATE
-- System               State                Frozen

A  sys1                 RUNNING              0
A  sys2                 RUNNING              0

# lltstat -n
LLT node information:
    Node                 State     Links
    0 sys1               OPEN      2
  * 1 sys2               OPEN      2
```

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

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

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

```
# haclus -value SecureClus
1
```

Sample Kickstart configuration file

The following is a sample RedHat Enterprise Linux 6 (RHEL6) Kickstart configuration file.

```
# The packages below are required and will be installed
from OS installation media automatically # during the automated
installation of products in the DVD, if they have not been
installed yet.
%packages
systemd-libs.i686
device-mapper
device-mapper-libs
parted
libgcc.i686
ed
ksh
nss-softokn-freebl.i686
glibc.i686
libstdc++.i686
audit-libs.i686
cracklib.i686
libselinux.i686
pam.i686
libattr.i686
libacl.i686
%end

%post --nochroot
# Add necessary scripts or commands here to your need
# This generated kickstart file is only for the automated installation of pro
```

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

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

# define path varibles
ROOT=/mnt/sysimage
BUILDDIR="${ROOT}/build"
RPMDIR="${BUILDDIR}/rpms"

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

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

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


# Install the RPMs in the following order.
for RPM in VRTSperl VRTSvlic VRTSspt VRTSvxvm VRTSaslapm VRTSob
VRTSvxfs VRTSfsadv VRTSfssdk VRTSllt VRTSgab VRTSvxfen VRTSamf
VRTSvcs VRTScps VRTSvcsag VRTSvcsdr VRTSvcsea VRTSdbed VRTSodm
VRTSsfmh VRTSvbs VRTSvcswiz VRTSsfcpi62


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

umount ${BUILDDIR}

CALLED_BY=KICKSTART ${ROOT}/opt/VRTS/install/bin/
UXRT62/add_install_scripts >> ${KSLOG} 2>&1

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

exit 0
%end
```

Installing Storage Foundation and High Availability using yum

You can install SFHA using yum. yum is supported for Red Hat Enterprise operating system.

To install SFHA using yum

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

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

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

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

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

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

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

        ```
        # cp -f VRTSvlic-* VRTSperl-* ... VRTSsfcpi62-*\
        /path/to/new/repository
        ```

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

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

        Output resembles:

        ```
        27/27 - VRTSsfcpi62-6.2.0.000-GENERIC.noarch.rpm
        Saving Primary metadata
        Saving file lists metadata
        Saving other metadata
        ```

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

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

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

     ```
     # cp -f VRTSvlic-* VRTSperl-* ... VRTSsfcpi62-*\
     /path/to/existing/repository
     ```

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

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

     Output resembles:

     ```
     27/27 * VRTSsfcpi62-6.2.0.000-GENERIC.noarch.rpm
     Saving Primary metadata
     Saving file lists metadata
     Saving other metadata
     ```

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

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

1. Create an XML file, which you can name SFHA_group.xml in the repository directory. In the file specify the name, the ID, the RPM list, and other information for the group. You can generate this XML file using the installer with the option `-yumgroupxml`. An example of this XML file for SFHA is:

```
# cat SFHA_group.xml
<comps>
  <group>
    <id>SFHA62</id>
    <name>SFHA62</name>
    <default>true</default>
    <description>RPMs of SFHA 6.2.0.000</description>
    <uservisible>true</uservisible>
    <packagelist>
      <packagereq type="default">VRTSvlic</packagereq>
      <packagereq type="default">VRTSperl</packagereq>
       ... [other RPMs for SFHA]
      <packagereq type="default">VRTSsfcpi62</packagereq>
    </packagelist>
  </group>
</comps>
```

2. Create the group when the repository is created or updated.

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

Or

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

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

3 Configure a yum repository on a client system.

- Create a `.repo` file under `/etc/yum.repos.d/`. An example of this `.repo` file for SFHA is:

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

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

- Check the yum configuration. List SFHA RPMs.

```
# yum list 'VRTS*'
Available Packages
VRTSperl.x86_64        5.16.1.4-RHEL5.2           repo-SFHA
VRTSsfcpi62.noarch     6.2.0.000-GENERIC       repo-SFHA
VRTSvlic.x86_64        3.02.62.010-0              repo-SFHA
...
```

The SFHA RPMs may not be visible immediately if:

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

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

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

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

4 Install the RPMs on the target systems.

- **To install all the RPMs**

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

     ```
     # yum install VRTSvlic VRTSperl ... VRTSsfcpi62
     ```

  2. Specify all of the SFHA RPMs using its RPM glob. For example:

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

3. Specify the group name if a group is configured for SFHA's RPMs. In this example, the group name is *SFHA62*:

   ```
   # yum install @SFHA62
   ```

   Or

   ```
   # yum groupinstall SFHA62
   ```

- **To install one RPM at a time**

1. Run the `installsfha -pkginfo` command to determine RPM installation order.

2. Use the same order as the output from the `installsfha -pkginfo` command:

```
# yum install VRTSperl
# yum install VRTSvlic
  ...
# yum install VRTSsfcpi62
```

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

Where `<version>` is the specific release version.

See "About the script-based installer" on page 69.

If the `VRTSsfcpi62` RPM is installed before you use yum to install SFHA, the RPM is not upgraded or uninstalled. If the `/opt/VRTS/install/installsfha<release_version>` script is not created properly, use the `/opt/VRTS/install/bin/UXRT62/add_install_scripts` script to create the installsfha or uninstallsfha scripts after all the other SFHA RPMs are installed. For example, your output may be similar to the following, depending on the products you install:

```
# /opt/VRTS/install/bin/UXRT62/add_install_scripts
Creating install/uninstall scripts for installed products
Creating /opt/VRTS/install/installvcs62 for UXRT62
Creating /opt/VRTS/install/uninstallvcs62 for UXRT62
Creating /opt/VRTS/install/installdmp62 for UXRT62
Creating /opt/VRTS/install/uninstalldmp62 for UXRT62
Creating /opt/VRTS/install/installfs62 for UXRT62
Creating /opt/VRTS/install/uninstallfs62 for UXRT62
Creating /opt/VRTS/install/installsf62 for UXRT62
Creating /opt/VRTS/install/uninstallsf62 for UXRT62
Creating /opt/VRTS/install/installvm62 for UXRT62
Creating /opt/VRTS/install/uninstallvm62 for UXRT62
```

Installing SFHA using the Red Hat Satellite server

You can install SFHA using the Red Hat Satellite server. Red Hat Satellite is supported for Red Hat Enterprise Linux operating system. You can install RPMs and rolling patches on the systems which the Red Hat Satellite server manages.

Red Hat Satellite server is a systems management solution. It lets you:

- Inventory the hardware and the software information of your systems.
- Install and update software on systems.
- Collect and distribute custom software RPMs into manageable groups.
- Provision (Kickstart) systems.
- Manage and deploy configuration files to systems.
- Monitor your systems.
- Provision virtual guests.
- Start, stop, and configure virtual guests.

In a Red Hat Satellite server, you can manage the system by creating a channel. A Red Hat Satellite channel is a collection of software RPMs. Using channels, you can segregate the RPMs by defining some rules. For instance, a channel may contain RPMs only from a specific Red Hat distribution. You can define channels according to your own requirement. You can create a channel that contains SFHA RPMs for custom usage in your organization's network.

Channels are of two types:

- Base channel
  A base channel consists of RPMs based on a specific architecture and Red Hat Enterprise Linux release.
- Child channel
  A child channel is a channel which is associated with a base channel that contains extra custom RPMs like SFHA.

A system can subscribe to only one base channel and multiple child channels of its base channel. The subscribed system can only install or update the RPMs that are available through its satellite channels.

For more information, see the *Red Hat Satellite5.6 User Guide*.

Using Red Hat Satellite server to install SFHA products

You can use the Red Hat Satellite server to install SFHA products on your system.

To use Red Hat Satellite server to install SFHA products

1. Set the base channel, child channel, and target system by following the Red Hat Satellite documentation. You need to ensure that:
   - The base channel consists of RPMs based on RHEL6.3, RHEL6.4, RHEL6.5, or the RHEL7 release

- The child channel consists of SFHA RPMs or patches.
- The target system is registered to the Red Hat Satellite.

2 Log on to the Red Hat Satellite admin page. Select the **Systems** tab. Click on the **target system**.

3 Select **Alter Channel Subscriptions** to alter the channel subscription of the target system.

4 Select the channel which contains the repository of SFHA.

5 Enter the following command to check the YUM repository on the target system.

```
# yum repolist
```

6 Enter the following command to install the SFHA RPMs using YUM:

```
# yum install @SFHA62
```

7 Enter the following command to generate the script of the installer:

```
# /opt/VRTS/install/bin/UXRT62/add_install_scripts
```

8 Enter the following command to configure SFHA using the installer:

```
# ./installsfha62 -configure
```

For configuring SFHA, See “*Configuring SFHA*” on page 113.

Chapter 16

Configuring SFHA clusters for data integrity

This chapter includes the following topics:

- Setting up disk-based I/O fencing manually
- Setting up server-based I/O fencing manually
- Setting up non-SCSI-3 fencing in virtual environments manually
- Setting up majority-based I/O fencing manually

Setting up disk-based I/O fencing manually

Table 16-1 lists the tasks that are involved in setting up I/O fencing.

Table 16-1

Task	Reference
Initializing disks as VxVM disks	See "Initializing disks as VxVM disks" on page 139.
Identifying disks to use as coordinator disks	See "Identifying disks to use as coordinator disks" on page 239.
Checking shared disks for I/O fencing	See "Checking shared disks for I/O fencing" on page 140.
Setting up coordinator disk groups	See "Setting up coordinator disk groups" on page 239.
Creating I/O fencing configuration files	See "Creating I/O fencing configuration files" on page 240.

Table 16-1 *(continued)*

Task	Reference
Modifying SFHA configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 241.
Configuring CoordPoint agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 256.
Verifying I/O fencing configuration	See "Verifying I/O fencing configuration" on page 243.

Removing permissions for communication

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

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

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See "Initializing disks as VxVM disks" on page 139.

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

To identify the coordinator disks

1. List the disks on each node.

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

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

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

   See "Checking shared disks for I/O fencing" on page 140.

Setting up coordinator disk groups

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

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

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

The following example procedure assumes that the disks have the device names sdx, sdy, and sdz.

To create the vxfencoorddg disk group

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

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

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

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

3. Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

1. On each nodes, type:

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

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

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

2. On all cluster nodes specify the use of DMP disk policy in the `/etc/vxfenmode` file.

   - ```
     # cp /etc/vxfen.d/vxfenmode_scsi3_dmp /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. Ensure that you edit the following file on each node in the cluster to change the values of the VXFEN_START and the VXFEN_STOP environment variables to 1:

   /etc/sysconfig/vxfen

Modifying VCS configuration to use I/O fencing

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

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

To modify VCS configuration to enable I/O fencing

1. Save the existing configuration:

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

2. Stop VCS on all nodes:

   ```
   # hastop -all
   ```
   ```

3 To ensure High Availability has stopped cleanly, run `gabconfig -a`.

In the output of the commans, check that Port h is not present.

4 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

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

5 Make a backup of the main.cf file on all the nodes:

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

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

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

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

7 Save and close the file.

8 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

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

9 Using rcp or another utility, copy the VCS configuration file from a node (for example, sys1) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp sys1:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

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

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

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

- Start VCS on the node where main.cf is modified.

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

- Start VCS on all other nodes once VCS on first node reaches RUNNING state.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

1 On one of the nodes, type:

```
# vxfenadm -d
```

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

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

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

   * 0 (sys1)
   1 (sys2)

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

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

```
# vxfenconfig -l
```

Setting up server-based I/O fencing manually

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

Table 16-2 Tasks to set up server-based I/O fencing manually

Task	Reference
Preparing the CP servers for use by the SFHA cluster	See "Preparing the CP servers manually for use by the Storage Foundation High Availability" on page 244.
Generating the client key and certificates on the client nodes manually	See "Generating the client key and certificates manually on the client nodes " on page 247.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See "Configuring server-based fencing on the Storage Foundation High Availability manually" on page 249.
Modifying SFHA configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 241.
Configuring Coordination Point agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 256.
Verifying the server-based I/O fencing configuration	See "Verifying server-based I/O fencing configuration" on page 257.

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

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

Table 16-3 displays the sample values used in this procedure.

Table 16-3 Sample values in procedure

CP server configuration component	Sample name
CP server	cps1
Node #1 - Storage Foundation High Availability	sys1
Node #2 - Storage Foundation High Availability	sys2

Table 16-3 Sample values in procedure *(continued)*

CP server configuration component	Sample name
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

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

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

For example, issue the following commands on one of the Storage Foundation High Availability nodes (sys1):

```
# grep cluster /etc/VRTSvcs/conf/config/main.cf

cluster clus1

# cat /etc/vx/.uuids/clusuuid

{f0735332-1dd1-11b2-bb31-00306eea460a}
```

2 Use the `cpsadm` command to check whether the Storage Foundation High Availability and nodes are present in the CP server.

For example:

```
# cpsadm -s cps1.symantecexample.com -a list_nodes

ClusName  UUID                                    Hostname(Node ID) Registered
clus1  {f0735332-1dd1-11b2-bb31-00306eea460a} sys1(0)         0
clus1  {f0735332-1dd1-11b2-bb31-00306eea460a} sys2(1)         0
```

If the output does not show the cluster and nodes, then add them as described in the next step.

For detailed information about the `cpsadm` command, see the *Symantec Cluster Server Administrator's Guide.*

3 Add the Storage Foundation High Availability and nodes to each CP server.

For example, issue the following command on the CP server (cps1.symantecexample.com) to add the cluster:

```
# cpsadm -s cps1.symantecexample.com -a add_clus\
 -c clus1  -u {f0735332-1dd1-11b2}

Cluster clus1 added successfully
```

Issue the following command on the CP server (cps1.symantecexample.com) to add the first node:

```
# cpsadm -s cps1.symantecexample.com -a add_node\
 -c clus1 -u {f0735332-1dd1-11b2} -h sys1 -n0

Node 0 (sys1) successfully added
```

Issue the following command on the CP server (cps1.symantecexample.com) to add the second node:

```
# cpsadm -s cps1.symantecexample.com -a add_node\
 -c clus1 -u {f0735332-1dd1-11b2} -h sys2 -n1

Node 1 (sys2) successfully added
```

4 If security is to be disabled, then add the user name "cpsclient@hostname" to the server.

5 Add the users to the CP server.

Issue the following commands on the CP server (cps1.symantecexample.com):

```
# cpsadm -s cps1.symantecexample.com -a add_user -e\
cpsclient@hostname\
 -f cps_operator -g vx

User cpsclient@hostname
successfully added
```

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

For example, issue the following command on the CP server (cps1.symantecexample.com) for Storage Foundation High Availability clus1 with two nodes sys1 and sys2:

```
# cpsadm -s cps1.symantecexample.com -a\
add_clus_to_user -c clus1\
 -u {f0735332-1dd1-11b2}\
 -e cpsclient@hostname\
 -f cps_operator -g vx

Cluster successfully added to user
 cpsclient@hostname  privileges.
```

See “Generating the client key and certificates manually on the client nodes ” on page 247.

Generating the client key and certificates manually on the client nodes

The client node that wants to connect to a CP server using HTTPS must have a private key and certificates signed by the Certificate Authority (CA) on the CP server

The client uses its private key and certificates to establish connection with the CP server. The key and the certificate must be present on the node at a predefined location. Each client has one client certificate and one CA certificate for every CP server, so, the certificate files must follow a specific naming convention. Distinct certificate names help the `cpsadm` command to identify which certificates have to be used when a client node connects to a specific CP server.

The certificate names must be as follows: ca_*cps-vip*.crt and client _*cps-vip*.crt

Where, *cps-vip* is the VIP or FQHN of the CP server listed in the `/etc/vxfenmode` file. For example, for a sample VIP, *192.168.1.201*, the corresponding certificate name is ca*_192.168.1.201*.

To manually set up certificates on the client node

1 Create the directory to store certificates.

```
# mkdir -p /var/VRTSvxfen/security/keys
/var/VRTSvxfen/security/certs
```

Note: Since the openssl utility might not be available on client nodes, Symantec recommends that you access the CP server using SSH to generate the client keys or certificates on the CP server and copy the certificates to each of the nodes.

---

2 Generate the private key for the client node.

```
# /usr/bin/openssl genrsa -out client_private.key 2048
```

3 Generate the client CSR for the cluster. CN is the UUID of the client's cluster.

```
# /usr/bin/openssl req -new -key client_private.key\
-subj '/C=countryname/L=localityname/OU=COMPANY/CN=CLUS_UUID'\
-out client_192.168.1.201.csr
```

Where, *countryname* is the country code, *localityname* is the city, *COMPANY* is the name of the company, and *CLUS_UUID* is the certificate name.

4 Generate the client certificate by using the CA key and the CA certificate. Run this command from the CP server.

```
# /usr/bin/openssl x509 -req -days days -in
client_192.168.1.201.csr\
-CA /var/VRTScps/security/certs/ca.crt -CAkey\
/var/VRTScps/security/keys/ca.key -set_serial 01 -out
client_192.168.10.1.crt
```

Where, *days* is the days you want the certificate to remain valid, *192.168.1.201* is the VIP or FQHN of the CP server.

5 Copy the client key, client certificate, and CA certificate to each of the client nodes at the following location.

Copy the client key at `/var/VRTSvxfen/security/keys/client_private.key`. The client is common for all the client nodes and hence you need to generate it only once.

Copy the client certificate at `/var/VRTSvxfen/security/certs/client_192.168.1.201.crt`.

Copy the CA certificate at `/var/VRTSvxfen/security/certs/ca_192.168.1.201.crt`

Note: Copy the certificates and the key to all the nodes at the locations that are listed in this step.

6 If the client nodes need to access the CP server using the FQHN and or the host name, make a copy of the certificates you generated and replace the VIP with the FQHN or host name. Make sure that you copy these certificates to all the nodes.

7 Repeat the procedure for every CP server.

8 After you copy the key and certificates to each client node, delete the client keys and client certificates on the CP server.

Configuring server-based fencing on the Storage Foundation High Availability manually

The configuration process for the client or Storage Foundation High Availability 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)
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)
- Set the order of coordination points

Note: Whenever coordinator disks are used as coordination points in your I/O fencing configuration, you must create a disk group (vxfencoorddg). You must specify this disk group in the `/etc/vxfenmode` file.

See "Setting up coordinator disk groups" on page 239.

The customized fencing framework also generates the `/etc/vxfentab` file which has coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

To configure server-based fencing on the Storage Foundation High Availability manually

1. Use a text editor to edit the following file on each node in the cluster:

   `/etc/sysconfig/vxfen`

   You must change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1.

2. Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

   - If your server-based fencing configuration uses a single highly available CP server as its only coordination point, make sure to add the `single_cp=1` entry in the `/etc/vxfenmode` file.
   - If you want the `vxfen` module to use a specific order of coordination points during a network partition scenario, set the `vxfen_honor_cp_order` value to be **1**. By default, the parameter is disabled.

   The following sample file output displays what the `/etc/vxfenmode` file contains:

   See "Sample vxfenmode file output for server-based fencing" on page 250.

3. After editing the `/etc/vxfenmode` file, run the vxfen init script to start fencing.

   For example:

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

Sample vxfenmode file output for server-based fencing

The following is a sample vxfenmode file for server-based fencing:

```
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
```

```
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

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

#
# scsi3_disk_policy determines the way in which I/O fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
#
scsi3_disk_policy=dmp

#
# security parameter is deprecated release 6.1 onwards
# since communication with CP server will always happen
# over HTTPS which is inherently secure. In pre-6.1 releases,
# it was used to configure 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

#
# vxfen_honor_cp_order determines the order in which vxfen
# should use the coordination points specified in this file.
#
```

```
# available options:
# 0 - vxfen uses a sorted list of coordination points specified
# in this file,
# the order in which coordination points are specified does not matter.
#   (default)
# 1 - vxfen uses the coordination points in the same order they are
#    specified in this file

# Specify 3 or more odd number of coordination points in this file,
# each one in its own line. They can be all-CP servers,
# all-SCSI-3 compliant coordinator disks, or a combination of
# CP servers and SCSI-3 compliant coordinator disks.
# Please ensure that the CP server coordination points
# are numbered sequentially and in the same order
# on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as follows:
#
#  cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames
# over different subnets, all of the IPs/names can be specified
# in a comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,
...,[<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#  is the serial number of the CPS as a coordination point; must
#  start with 1.
# <vip>
#  is the virtual IP address of the CPS, must be specified in
#  square brackets ("[]").
# <vhn>
#  is the virtual hostname of the CPS, must be specified in square
#  brackets ("[]").
# <port>
#  is the port number bound to a particular <vip/vhn> of the CPS.
#  It is optional to specify a <port>. However, if specified, it
#  must follow a colon (":") after <vip/vhn>. If not specified, the
#  colon (":") must not exist after <vip/vhn>.
#
```

```
# For all the <vip/vhn>s which do not have a specified <port>,
# a default port can be specified as follows:
#
# port=<default_port>
#
#  Where <default_port> is applicable to all the <vip/vhn>s for
#  which a <port> is not specified. In other words, specifying
#  <port> with a <vip/vhn> overrides the <default_port> for that
#  <vip/vhn>. If the <default_port> is not specified, and there
#  are <vip/vhn>s for which <port> is not specified, then port
#  number 14250 will be used for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[cps1.company.com]
#  cps2=[192.168.0.25]
# cps3=[cps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [cps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
# - if default port 57777 were not specified, port 14250
# would be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
#  vxfendg=<coordinator disk group name>
# Example:
#  vxfendg=vxfencoorddg
#
# Examples of different configurations:
#  1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
```

```
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
# cps1=[cps1.company.com]
# cps2=[cps2.company.com]
# cps3=[cps3.company.com]
# port=443
```

Table 16-4 defines the vxfenmode parameters that must be edited.

Table 16-4 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 DMP devices, "dmp". **Note:** The configured disk policy is applied on all the nodes.
security	Deprecated from release 6.1 onwards. Security parameter is deprecated release 6.1 onwards as communication between CP servers and application clusters happens over the HTTPS protocol which is inherently secure. In releases prior to 6.1, the security parameter was used to configure secure communication to the CP server using the VxAT (Veritas Authentication Service) options. The options are: ■ 0 - Do not use Veritas Authentication Service for CP server communication ■ 1 - Use Veritas Authentication Service for CP server communication

Table 16-4 vxfenmode file parameters *(continued)*

vxfenmode File Parameter	Description

cps1, cps2, or vxfendg	Coordination point parameters.
	Enter either the virtual IP address or the FQHN (whichever is accessible) of the CP server.
	`cps<number>=[virtual_ip_address/virtual_host_name]:port`
	Where *port* is optional. The default port value is 443.
	If you have configured multiple virtual IP addresses or host names over different subnets, you can specify these as comma-separated values. For example:
	`cps1=[192.168.0.23],[192.168.0.24]:58888, [cps1.company.com]`
	Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfencoorddg) 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).

port	Default port for the CP server to listen on. If you have not specified port numbers for individual virtual IP addresses or host names, the default port number value that the CP server uses for those individual virtual IP addresses or host names is 443. You can change this default port value using the port parameter.
single_cp	Value 1 for single_cp parameter indicates that the server-based fencing uses a single highly available CP server as its only coordination point. Value 0 for single_cp parameter indicates that the server-based fencing uses at least three coordination points.
vxfen_honor_cp_order	Set the value to 1 for vxfen module to use a specific order of coordination points during a network partition scenario. By default the parameter is disabled. The default value is 0.

Configuring CoordPoint agent to monitor coordination points

The following procedure describes how to manually configure the CoordPoint agent to monitor coordination points.

The CoordPoint agent can monitor CP servers and SCSI-3 disks.

See the *Symantec Cluster Server Bundled Agents Reference Guide* for more information on the agent.

To configure CoordPoint agent to monitor coordination points

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

```
# haconf -makerw
# hagrp -add vxfen
# hagrp -modify vxfen SystemList sys1 0 sys2 1
# hagrp -modify vxfen AutoFailOver 0
# hagrp -modify vxfen Parallel 1
# hagrp -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 0
# hares -override coordpoint LevelTwoMonitorFreq
# hares -modify coordpoint LevelTwoMonitorFreq 5
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

3. Configure the Phantom resource for the vxfen disk group.

```
# haconf -makerw
# hares -add RES_phantom_vxfen Phantom vxfen
# hares -modify RES_phantom_vxfen Enabled 1
# haconf -dump -makero
```

4 Verify the status of the agent on the Storage Foundation High Availability using the `hares` commands. For example:

```
# hares -state coordpoint
```

The following is an example of the command and output::

```
# hares -state coordpoint

# Resource    Attribute   System   Value
coordpoint   State        sys1   ONLINE
coordpoint   State        sys2   ONLINE
```

5 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed CoordPoint agent monitoring information; including information about whether the CoordPoint agent is able to access all the coordination points, information to check on which coordination points the CoordPoint agent is reporting missing keys, etc.

To view the debug logs in the engine log, change the dbg level for that node using the following commands:

```
# haconf -makerw

# hatype -modify Coordpoint LogDbg 10

# haconf -dump -makero
```

The agent log can now be viewed at the following location:

/var/VRTSvcs/log/engine_A.log

Note: The Coordpoint agent is always in the online state when the I/O fencing is configured in the majority or the disabled mode. For both these modes the I/O fencing does not have any coordination points to monitor. Thereby, the Coordpoint agent is always in the online state.

Verifying server-based I/O fencing configuration

Follow the procedure described below to verify your server-based I/O fencing configuration.

To verify the server-based I/O fencing configuration

1 Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

Note: For troubleshooting any server-based I/O fencing configuration issues, refer to the *Symantec Cluster Server Administrator's Guide*.

2 Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```

If the output displays `single_cp=1`, it indicates that the application cluster uses a CP server as the single coordination point for server-based fencing.

Setting up non-SCSI-3 fencing in virtual environments manually

To manually set up I/O fencing in a non-SCSI-3 PR compliant setup

1 Configure I/O fencing either in majority-based fencing mode with no coordination points or in server-based fencing mode only with CP servers as coordination points.

See "Setting up server-based I/O fencing manually" on page 244.

See "Setting up majority-based I/O fencing manually " on page 264.

2 Make sure that the SFHA cluster is online and check that the fencing mode is customized mode or majority mode.

```
# vxfenadm -d
```

3 Make sure that the cluster attribute UseFence is set to SCSI-3.

```
# haclus -value UseFence
```

4 On each node, edit the /etc/vxenviron file as follows:

```
data_disk_fencing=off
```

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

```
vxfen_vxfnd_tmt=25
```

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

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

Refer to the sample /etc/vxfenmode file.

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

- Run the following command:

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

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

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

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

- Make the VCS configuration file writable:

```
# haconf -makerw
```

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

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

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

- Run the following command to verify the value:

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

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

The command should not list any resources.

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

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

- Make the VCS configuration file read-only

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

9 Make sure that the UseFence attribute in the VCS configuration file main.cf is set to SCSI-3.

10 To make these VxFEN changes take effect, stop and restart VxFEN and the dependent modules

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

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

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

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

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

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

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

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

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

```
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
#
scsi3_disk_policy=dmp

#
# Seconds for which the winning sub cluster waits to allow for the
# losing subcluster to panic & drain I/Os. Useful in the absence of
# SCSI3 based data disk fencing loser_exit_delay=55
#
# Seconds for which vxfend process wait for a customized fencing
# script to complete. Only used with vxfen_mode=customized
# vxfen_script_timeout=25

# security parameter is deprecated release 6.1 onwards since
# communication with CP server will always happen over HTTPS
# which is inherently secure. In pre-6.1 releases, it was used
# to configure 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

#
# vxfen_honor_cp_order determines the order in which vxfen
# should use the coordination points specified in this file.
#
# available options:
# 0 - vxfen uses a sorted list of coordination points specified
# in this file, the order in which coordination points are specified
# does not matter.
#   (default)
# 1 - vxfen uses the coordination points in the same order they are
#    specified in this file
```

```
# Specify 3 or more odd number of coordination points in this file,
# each one in its own line. They can be all-CP servers, all-SCSI-3
# compliant coordinator disks, or a combination of CP servers and
# SCSI-3 compliant coordinator disks.
# Please ensure that the CP server coordination points are
# numbered sequentially and in the same order on all the cluster
# nodes.
#
# Coordination Point Server(CPS) is specified as follows:
#
#  cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames
# over different subnets, all of the IPs/names can be specified
# in a comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,
# ...,[<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#  is the serial number of the CPS as a coordination point; must
#  start with 1.
# <vip>
#  is the virtual IP address of the CPS, must be specified in
#  square brackets ("[]").
# <vhn>
#  is the virtual hostname of the CPS, must be specified in square
#  brackets ("[]").
# <port>
#  is the port number bound to a particular <vip/vhn> of the CPS.
#  It is optional to specify a <port>. However, if specified, it
#  must follow a colon (":") after <vip/vhn>. If not specified, the
#  colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>,
# a default port can be specified as follows:
#
# port=<default_port>
#
#  Where <default_port> is applicable to all the <vip/vhn>s for which a
#  <port> is not specified. In other words, specifying <port> with a
```

```
#  <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[cps1.company.com]
#  cps2=[192.168.0.25]
# cps3=[cps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [cps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
# - if default port 57777 were not specified, port 14250 would be
#  used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
#  vxfendg=<coordinator disk group name>
# Example:
#  vxfendg=vxfencoorddg
#
# Examples of different configurations:
#  1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
```

```
# vxfendg=
# Note: The disk group specified in case should have three disks
# cps1=[cps1.company.com]
# cps2=[cps2.company.com]
# cps3=[cps3.company.com]
# port=443
```

Setting up majority-based I/O fencing manually

Table 16-5 lists the tasks that are involved in setting up I/O fencing.

Task	Reference
Creating I/O fencing configuration files	Creating I/O fencing configuration files
Modifying VCS configuration to use I/O fencing	Modifying VCS configuration to use I/O fencing
Verifying I/O fencing configuration	Verifying I/O fencing configuration

Creating I/O fencing configuration files

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

1. On all cluster nodes, run the following command

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

2. To check the updated /etc/vxfenmode configuration, enter the following command on one of the nodes.

   ```
   # cat /etc/vxfenmode
   ```

3. Ensure that you 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 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. To ensure High Availability has stopped cleanly, run `gabconfig -a`.

   In the output of the commans, check that Port h is not present.

4. If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

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

5. Make a backup of the main.cf file on all the nodes:

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

6. 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
   )
   ```

   For fencing configuration in any mode except the disabled mode, the value of the cluster-level attribute UseFence is set to SCSI3.

7. Save and close the file.

8. Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

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

9 Using rcp or another utility, copy the VCS configuration file from a node (for example, sys1) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp sys1:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

10 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.

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

- Start VCS on the node where main.cf is modified.

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

- Start VCS on all other nodes once VCS on first node reaches RUNNING state.

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

Verifying I/O fencing configuration

Verify from the vxfenadm output that the fencing mode reflects the configuration in the /etc/vxfenmode file.

To verify I/O fencing configuration

- On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is majority:

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

 Fencing Protocol Version: 201
 Fencing Mode: MAJORITY
 Cluster Members:

        * 0 (sys1)
          1 (sys2)

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

Section 7

Managing your Symantec deployments

- Chapter 17. Performing centralized installations using the Deployment Server

Chapter 17

Performing centralized installations using the Deployment Server

This chapter includes the following topics:

- About the Deployment Server
- Deployment Server overview
- Installing the Deployment Server
- Setting up a Deployment Server
- Setting deployment preferences
- Specifying a non-default repository location
- Downloading the most recent release information
- Loading release information and patches on to your Deployment Server
- Viewing or downloading available release images
- Viewing or removing repository images stored in your repository
- Deploying Symantec product updates to your environment
- Finding out which releases you have installed, and which upgrades or updates you may need
- Defining Install Bundles
- Creating Install Templates

- Deploying Symantec releases
- Connecting the Deployment Server to SORT using a proxy server

About the Deployment Server

The Deployment Server makes it easier to install or upgrade SFHA releases from a central location. The Deployment Server lets you store multiple release images and patches in one central location and deploy them to systems of any supported UNIX or Linux operating system (6.1 or later).

Note: The script-based installer for version 6.1 and higher supports installations from one operating system node onto a different operating system. Therefore, heterogeneous push installations are supported for 6.1 and higher releases only.

Push installations for product versions 5.1, 6.0, or 6.0.1 releases must be executed from a system that is running the same operating system as the target systems. In order to perform push installations for product versions 5.1, 6.0, or 6.0.1 releases on multiple platforms, you must have a separate Deployment Server for each operating system.

The Deployment Server lets you do the following as described in Table 17-1.

Table 17-1 Deployment Server functionality

Feature	Description
Manage repository images	■ View available SFHA releases. ■ Download maintenance and patch release images from the Symantec Operations Readiness Tools (SORT) website into a repository. ■ Load the downloaded release image files from FileConnect and SORT into the repository. ■ View and remove the release image files that are stored in the repository.
Version check systems	■ Discover RPMs and patches installed on your systems and informs you of the product and version installed ■ Identify base, maintenance, and patch level upgrades to your system and to download maintenance and patch releases. ■ Query SORT for the most recent updates.

Table 17-1 Deployment Server functionality *(continued)*

Feature	Description
Install or upgrade systems	■ Install base, maintenance, or patch level releases. ■ Install SFHA from any supported UNIX or Linux operating system to any other supported UNIX or Linux operating system. ■ Automatically load the script-based installer patches that apply to that release. ■ Install or upgrade an Install Bundle that is created from the **Define/Modify Install Bundles** menu. ■ Install an Install Template that is created from the **Create Install Templates** menu.
Define or modify Install Bundles	Define or modify Install Bundles and save them using the Deployment Server.
Create Install Templates	Discover installed components on a running system that you want to replicate on to new systems.
Update metadata	Download, load the release matrix updates, and product installer updates for systems behind a firewall. This process happens automatically when you connect the Deployment Server to the Internet, or it can be initiated manually. If the Deployment Server is not connected to the Internet, then the **Update Metadata** option is used to upload current metadata.
Set preferences	Define or reset program settings.
Connecting the Deployment Server to SORT using a proxy server	Use a proxy server, a server that acts as an intermediary for requests from clients, for connecting the Deployment Server to the Symantec Operations Readiness Tools (SORT) website.

Note: The Deployment Server is available only from the command line. The Deployment Server is not available for the web-based installer.

Note: Many of the example outputs used in this chapter are based on Red Hat Enterprise Linux.

Deployment Server overview

After obtaining and installing the Deployment Server and defining a central repository, you can begin managing your deployments from that repository. You

can load and store product images for Symantec products back to version 5.1 in your Deployment Server. The Deployment Server is a central installation server for storing and managing your product updates.

Setting up and managing your repository involves the following tasks:

- Installing the Deployment Server.
  See "Installing the Deployment Server" on page 272.
- Setting up a Deployment Server.
  See "Setting up a Deployment Server" on page 273.
- Finding out which products you have installed, and which upgrades or updates you may need.
  See "Viewing or downloading available release images" on page 280.
- Adding release images to your Deployment Server.
  See "Viewing or downloading available release images" on page 280.
- Removing release images from your Deployment Server.
  See "Viewing or removing repository images stored in your repository" on page 285.
- Defining or modifying Install Bundles to manually install or upgrade a bundle of two or more releases.
  See "Defining Install Bundles" on page 290.
- Creating Install Templates to discover installed components on a system that you want to replicate to another system.
  See "Creating Install Templates" on page 295.

Later, when your repository is set up, you can use it to deploy Symantec products to other systems in your environment.

See "Deploying Symantec product updates to your environment" on page 287.

See "Deploying Symantec releases" on page 297.

Installing the Deployment Server

You can obtain the Deployment Server by either:

- Installing the Deployment Server manually.
- Running the Deployment Server after installing at least one Symantec 6.2 product.

Note: The `VRTSperl` and the `VRTSsfcpi<version>`RPMs are included in all Storage Foundation (SF) products, so installing any Symantec 6.2 product lets you access the Deployment Server.

To install the Deployment Server manually without installing a Symantec 6.2 product

1. Log in as superuser.
2. Mount the installation media.

   See "Mounting the product disc" on page 65.
3. Move to the top-level directory on the disc.

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

   where *dist* is rhel6, rhel7, or sles11, and *arch* is x86_64 for RHEL and SLES.
4. Navigate to the following directory:

   ```
   # cd rpms
   ```
5. Run the following command to install the `VRTSperl` and the `VRTSsfcpi<version>` RPMs:

   ```
   # rpm -ivh VRTSperl*.rpm VRTSsfcpi<version>*.rpm
   ```

To run the Deployment Server

1. Log in as superuser.
2. Navigate to the following directory:

   ```
   # cd /opt/VRTS/install
   ```
3. Run the Deployment Server.

   ```
   # ./deploy_sfha
   ```

Setting up a Deployment Server

Symantec recommends that you create a dedicated Deployment Server to manage your product updates.

A Deployment Server is useful for doing the following tasks:

- Storing release images for the latest upgrades and updates from Symantec in a central repository directory.
- Installing and updating systems directly by accessing the release images that are stored within a central repository.
- Defining or modifying Install Bundles for deploying a bundle of two or more releases.
- Discovering installed components on a system that you want to replicate to another system.
- Installing Symantec products from the Deployment Server to systems running any supported platform.
- Creating a file share on the repository directory provides a convenient, central location from which systems running any supported platform can install the latest Symantec products and updates.

Create a central repository on the Deployment Server to store and manage the following types of Symantec releases:

- Base releases. These major releases and minor releases are available for all Symantec products. They contain new features, and you can download them from FileConnect.
- Maintenance releases. These releases are available for all Symantec products. They contain bug fixes and a limited number of new features, and you can download them from the Symantec Operations Readiness Tools (SORT) website.
- Patches. These releases contain fixes for specific products, and you can download them from the SORT website.

Note: All base releases and maintenance releases can be deployed using the install scripts that are included in the release. Before version 6.0.1, patches were installed manually. From the 6.0.1 release and onwards, install scripts are included with patch releases.

You can set up a Deployment Server with or without Internet access.

- If you set up a Deployment Server that has Internet access, you can download maintenance releases and patches from Symantec directly. Then, you can deploy them to your systems.
  Setting up a Deployment Server that has Internet access
- If you set up a Deployment Server that does not have Internet access, you can download maintenance releases and patches from Symantec on another system that has Internet access. Then, you can load the images onto the Deployment Server separately.

Setting up a Deployment Server that does not have Internet access

Setting up a Deployment Server that has Internet access

Figure 17-1 shows a Deployment Server that can download product images directly from Symantec using the Deployment Server.

Figure 17-1 Example Deployment Server that has Internet access

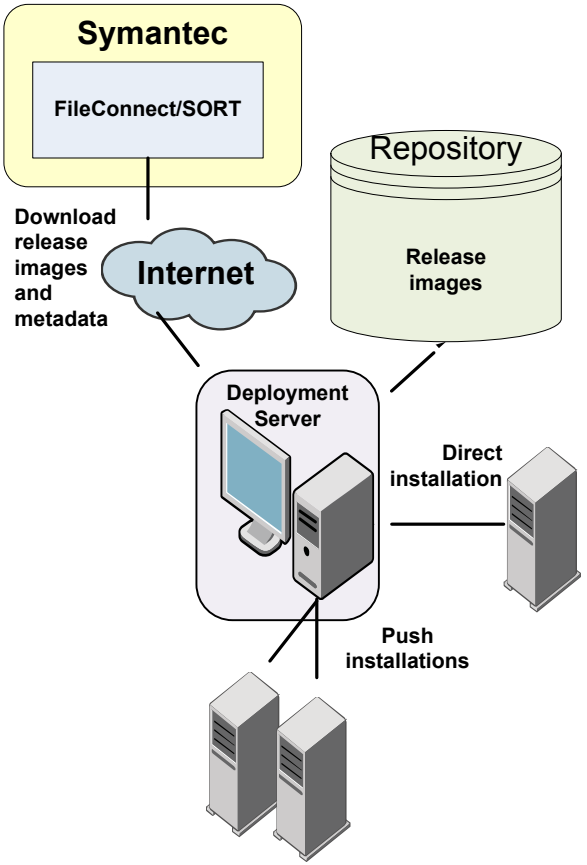

Setting up a Deployment Server that does not have Internet access

Figure 17-2 shows a Deployment Server that does not have Internet access. In this scenario, release images and metadata updates are downloaded from another system. Then, they are copied to a file location available to the Deployment Server, and loaded.

Figure 17-2 Example Deployment Server that does not have Internet access

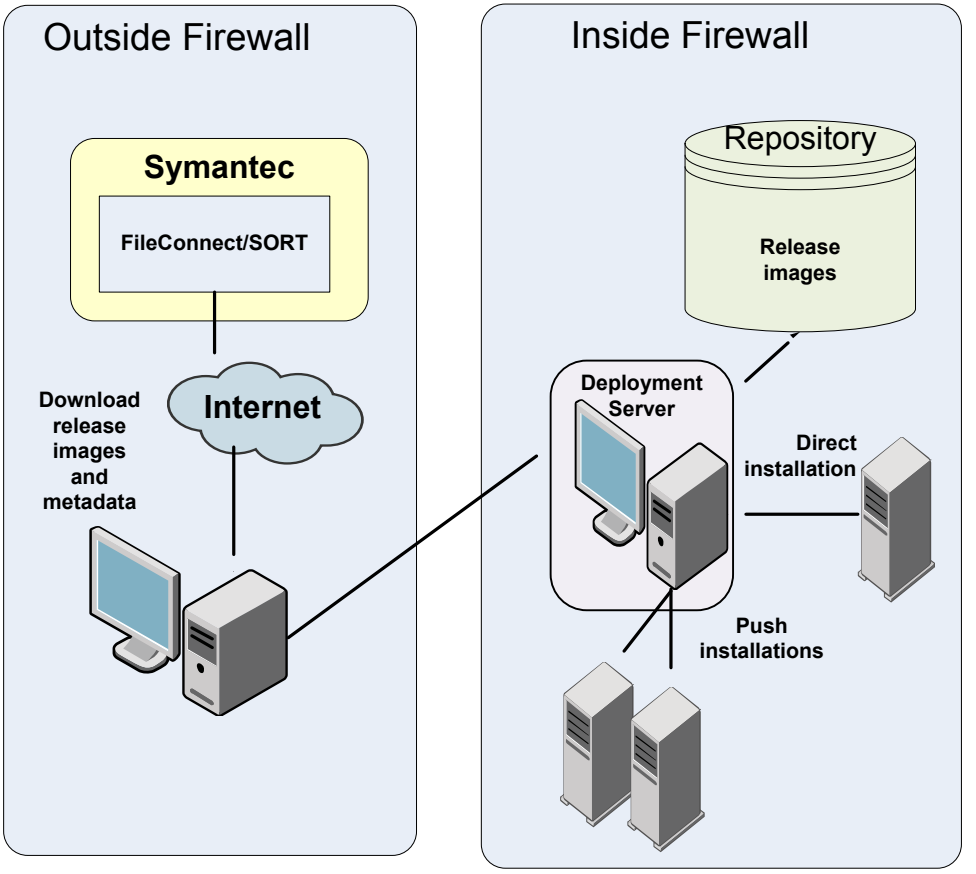

Release image files for base releases must be manually downloaded from FileConnect and loaded in a similar manner.

Setting deployment preferences

You can set preferences for managing the deployment of products dating back to version 5.1.

Note: You can select option **U (Terminology and Usage)** to obtain more information about Deployment Server terminology and usage.

To set deployment preferences

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images             M) Update Metadata
V) Version Check Systems                S) Set Preferences
I) Install/Upgrade Systems              U) Terminology and Usage
B) Define/Modify Install Bundles        ?) Help
T) Create Install Templates             Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **S**, **Set Preferences**.

You see the following output:

```
Current Preferences:

Repository            /opt/VRTS/repository
Selected Platforms    N/A
Save Tar Files        N/A

Preference List:

     1)  Repository
     2)  Selected Platforms
     3)  Save Tar Files
     b)  Back to previous menu

Select a preference to set: [1-3,b,q,?]
```

3 Do one of the following:

- To set the default repository, enter **1**. Then enter the name of the repository in which you want to store your downloads. For example, enter the following:

  /opt/VRTS/install/ProductDownloads

  If the specified repository replaces a previous repository, the installer asks if you want to move all your files into the new repository. To move your files to the new repository, enter **y**.

- To add or remove a platform, enter **2**. You are provided selections for adding or deleting a platform. When a single platform is removed, it becomes **N/A**, which means that it is not defined. By default, all platforms are chosen. Once you select to add or remove a platform, the platform is added or removed in the preferences file and the preference platforms are updated.
  If only one platform is defined, no platform, architecture, distribution, and version selection menu is displayed.
- To set the option for saving or removing tar files, enter **3**. At the prompt, if you want to save the tar files after untarring them, enter **y**. Or, if you want to remove tar files after untarring them, enter **n**.
  By default, the installer does not remove tar files after the releases have been untarred.

Specifying a non-default repository location

You can specify a repository location other than the default that has been set within the system preferences by using the command line option. The command line option is mainly used to install a release image from a different repository location. When you use the command line option, the designated repository folder is used instead of the default for the execution time of the script. Using the command line option does not override the repository preference set by the **Set Preference** menu item.

Note: When you specify a non-default repository, you are allowed only to view the repository (**View/Remove Repository**), and use the repository to install or upgrade (**Install/Upgrade Systems**) on other systems.

To use the command line option to specify a non-default repository location

- At the command line, to specify a non-default repository location, enter the following:

  ```
  # ./deploy_sfha -repository repository_path
  ```

  where *repository_path* is the location of the repository.

Downloading the most recent release information

Use one of the following methods to obtain a `.tar` file with the most recent release information:

- Download a copy from the SORT website.
- Run the Deployment Server from a system that has Internet access.

To obtain a data file by downloading a copy from the SORT website

1 Download the `.tar` file from the SORT site at:

https://sort.symantec.com/support/related_links/offline-release-updates

2 Click on **deploy_sfha.tar [Download]**, and save the file to your desktop.

To obtain a data file by running the Deployment Server from a system with Internet access

1 Run the Deployment Server. Enter the following:

```
# /opt/VRTS/install/deploy_sfha
```

2 Select option **M**, **Update Metadata**.

You see the following output:

```
The Update Metadata option is used to load release matrix updates on to systems that do
not have an Internet connection with SORT (https://sort.symantec.com).  Your system has a
connection with SORT and is able to receive updates.  No action is necessary unless you
would like to create a file to update another Deployment Server system.

     1)  Download release matrix updates and installer patches
     2)  Load an update tar file
     b)  Back to previous menu

Select the option: [1-2,b,q,?]
```

3 Select option **1**, **Download release matrix updates and installer patches**.

Loading release information and patches on to your Deployment Server

In this procedure, the Internet-enabled system is the system to which you downloaded the `deploy_sfha.tar` file.

See “Downloading the most recent release information” on page 278.

To load release information and patches on to your Deployment Server

1. On the Internet-enabled system, copy the `deploy_sfha.tar` file you downloaded to a location accessible by the Deployment Server.
2. On the Deployment Server, change to the installation directory. For example, enter the following:

   ```
   # cd /opt/VRTS/install/
   ```

3. Run the Deployment Server. Enter the following:

   ```
   # ./deploy_sfha
   ```

4. Select option **M**, **Update Metadata**, and select option **2**, **Load an update tar file**. Enter the location of the `deploy_sfha.tar` file (the installer calls it a "meta-data tar file").

   ```
   Enter the location of the meta-data tar file: [b]
   (/opt/VRTS/install/deploy_sfha.tar)
   ```

   For example, enter the location of the meta-data tar file:

   ```
   /tmp/deploy_sfha.tar
   ```

Viewing or downloading available release images

You can use the Deployment Server to conveniently view or download available release images to be deployed on other systems in your environment.

Note: If you have Internet access, communication with the Symantec Operations Readiness Tools (SORT) provides the latest release information. If you do not have Internet access, static release matrix files are referenced, and the most recent updates may not be included.

See “Loading release information and patches on to your Deployment Server” on page 279.

To view or download available release images

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

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

3 Select option **1**, **View/Download Available Releases**, to view or download what is currently installed on your system.

You see a list of platforms and release levels.

```
To view or download available releases, the platform type and release
level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  RHEL7 x86_64              8)  SLES10 x86_64
9)  SLES11 x86_64             10) Solaris 9 Sparc
11) Solaris 10 Sparc          12) Solaris 10 x64
13) Solaris 11 Sparc          14) Solaris 11 x64
b)  Back to previous menu

Select the platform of the release to view/download [1-14,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels.

```
1) Base
2) Maintenance
3) Patch
b) Back to previous menu

Select the level of the <platform> releases to view/download
[1-3,b,q,?]
```

5 Select the number corresponding to the type of release you want to view (Base, Maintenance, or Patch).

You see a list of releases available for download.

```
Available Maintenance releases for rhel6_x86_64:

release_version SORT_release_name              DL OBS AI rel_date   size_KB
================================================================================
5.1SP1PR2RP2    sfha-rhel6_x86_64-5.1SP1PR2RP2 -  Y   Y  2011-09-28 145611
5.1SP1PR2RP3    sfha-rhel6_x86_64-5.1SP1PR2RP3 -  Y   Y  2012-10-02 153924
5.1SP1PR2RP4    sfha-rhel6_x86_64-5.1SP1PR2RP4 -  -   Y  2013-08-21 186859
5.1SP1PR3RP2    sfha-rhel6_x86_64-5.1SP1PR3RP2 -  Y   Y  2011-09-28 145611
5.1SP1PR3RP3    sfha-rhel6_x86_64-5.1SP1PR3RP3 -  Y   Y  2012-10-02 153924
```

```
5.1SP1PR3RP4    sfha-rhel6_x86_64-5.1SP1PR3RP4 -  -   Y  2013-08-21 186859
6.0RP1          sfha-rhel6_x86_64-6.0RP1       -  -   Y  2012-03-22 210076
6.0.3           sfha-rhel6_x86_64-6.0.3        -  Y   Y  2013-02-01 212845
6.0.5           sfha-rhel6_x86_64-6.0.5        -  -   Y  2014-04-15 199143
6.1.1           sfha-rhel6_x86_64-6.1.1        -  -   Y  2014-07-24 208028
Enter the release_version to view details about a release or press 'Enter' to continue [b,q,?]
```

The following are the descriptions for the column headers:

- release_version: The version of the release.
- SORT_release_name: The name of the release, used when accessing SORT (https://sort.symantec.com).
- DL: An indicator that the release is present in your repository.
- OBS: An indicator that the release is obsolete by another higher release.
- AI: An indicator that the release has scripted install capabilities. All base and maintenance releases have auto-install capabilities. Patch releases with auto-install capabilities are available beginning with version 6.1. Otherwise the patch requires a manual installation.
- rel_date: The date the release is available.
- size_KB: The file size of the release in kilobytes.

6 If you are interested in viewing more details about any release, type the release version. For example, enter the following:

```
6.0.3
```

You see the following output:

```
release_version: 6.0.3
release_name: sfha-rhel6_x86_64-6.0.3
release_type: MR
release_date: 2013-02-01
downloaded: Y
install_path: rhel6_x86_64/installmr
upload_location: ftp://ftp.veritas.com/pub/support/patchcentral
/Linux/6.0.3/sfha/sfha-rhel6_x86_64-6.0.3-patches.tar.gz
obsoletes: 6.0.1.200-fs,6.0.1.200-vm,6.0.1.300-fs
obsoleted_by: None
Would you like to download this Maintenance Release? [y,n,q] (y) n

Enter the release_version to view the details about a release or press
'Enter' to continue [b,q,?]
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to download a rhel6_x86_64 Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all releases that are not currently in the repository.

```
1)  5.1SP1PR2RP2
     2)  5.1SP1PR2RP3
     3)  5.1SP1PR2RP4
     4)  5.1SP1PR3RP2
     5)  5.1SP1PR3RP3
     6)  5.1SP1PR3RP4
     7)  6.0RP1
     8)  6.0.3
     9)  6.0.5
    10)  6.1.1
    11)  All non-obsoleted releases
    12)  All releases
     b)  Back to previous menu

Select the patch release to download, 'All non-obsoleted releases' to
download all non-obsoleted releases, or 'All releases' to download
all releases [1-5,b,q] 3
```

8 Select the number corresponding to the release that you want to download. You can download a single release, all non-obsoleted releases, or all releases.

The selected release images are downloaded to the Deployment Server.

```
Downloading sfha-rhel6_x86_64-6.0RP1 from SORT - https://sort.symantec.com
Downloading 215118373 bytes (Total 215118373 bytes [205.15 MB]): 100%
Untarring sfha-rhel6_x86_64-6.0RP1 ..................................... Done

sfha-rhel6_x86_64-6.0RP1 has been downloaded successfully.
```

9 From the menu, select option **2**, **View/Remove Repository Images**, and follow the prompts to check that the release images are loaded.

See "Viewing or downloading available release images" on page 280.

Viewing or removing repository images stored in your repository

You can use the Deployment Server to conveniently view or remove the release images that are stored in your repository.

To view or remove release images stored in your repository

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images             M) Update Metadata
V) Version Check Systems                S) Set Preferences
I) Install/Upgrade Systems              U) Terminology and Usage
B) Define/Modify Install Bundles        ?) Help
T) Create Install Templates             Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

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

3 Select option **2**, **View/Remove Repository Images**, to view or remove the release images currently installed on your system.

You see a list of platforms and release levels if you have downloaded the corresponding Base, Maintenance, or Patch release on that platform.

```
To view or remove repository images, the platform type and
release level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  RHEL7 x86_64              8)  SLES10 x86_64
9)  SLES11 x86_64             10) Solaris 9 Sparc
11) Solaris 10 Sparc          12) Solaris 10 x64
13) Solaris 11 Sparc          14) Solaris 11 x64
b)  Back to previous menu

Select the platform of the release to view/remove [1-14,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels if you have downloaded the corresponding Base, Maintenance, or Patch release.

```
1) Base
2) Maintenance
3) Patch
b) Back to previous menu

Select the level of the <platform> releases to view/remove
[1-3,b,q]
```

5 Select the number corresponding to the type of release you want to view or remove (Base, Maintenance, or Patch).

You see a list of releases that are stored in your repository.

```
Stored Repository Releases:

release_version SORT_release_name        OBS AI
=============================================
6.0RP1          sfha-rhel6_x86_64-6.0RP1  -   Y
6.0.3           sfha-rhel6_x86_64-6.0.3   -   Y
```

6 If you are interested in viewing more details about a release image that is stored in your repository, type the release version. For example, enter the following:

```
6.0.3
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to remove a rhel6_x86_64 Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all the releases that are stored in your repository that match the selected platform and release level.

```
1)  6.0RP1
2)  6.0.3
b)  Back to previous menu

Select the patch release to remove [1-2,b,q] 1
```

8 Type the number corresponding to the release version you want to remove.

The release images are removed from the Deployment Server.

```
Removing sfha-rhel6_x86_64-6.0RP1-patches ..................... Done
sfha-rhel6_x86_64-6.0RP1-patches has been removed successfully.
```

Deploying Symantec product updates to your environment

You can use the Deployment Server to deploy release images to the systems in your environment as follows:

- If you are not sure what to deploy, perform a version check. A version check tells you if there are any Symantec products installed on your systems. It suggests patches and maintenance releases, and gives you the option to install updates.
  See "Finding out which releases you have installed, and which upgrades or updates you may need" on page 288.

- If you know which update you want to deploy on your systems, use the Install/Upgrade Systems script to deploy a specific Symantec release.
  See "Deploying Symantec releases" on page 297.

Finding out which releases you have installed, and which upgrades or updates you may need

Use the Version Check option to determine which Symantec product you need to deploy. The Version Check option is useful if you are not sure which releases you already have installed, or you want to know about available releases.

The Version Check option gives you the following information:

- Installed products and their versions (base, maintenance releases, and patches)
- Installed RPMs (required and optional)
- Available releases (base, maintenance releases, and patches) relative to the version which is installed on the system

To determine which Symantec product updates to deploy

1. Launch the Deployment Server. For example, enter the following:

   ```
   # /opt/VRTS/install/deploy_sfha
   ```

   You see the following output:

   ```
   Task Menu:

   R) Manage Repository Images              M) Update Metadata
   V) Version Check Systems                 S) Set Preferences
   I) Install/Upgrade Systems               U) Terminology and Usage
   B) Define/Modify Install Bundles         ?) Help
   T) Create Install Templates              Q) Quit

   Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
   ```

2. Select option **V**, **Version Check Systems**.

3 At the prompt, enter the system names for the systems you want to check. For example, enter the following:

```
sys1
```

You see output for the installed RPMs (required, optional, or missing).

You see a list of releases available for download.

```
Available Base Releases for Veritas Storage Foundation HA 6.0.1:


release_version SORT_release_name       DL rel_date   size_KB
============================================================
6.0.2           sfha-rhel6_x86_64-6.0.2 -  2012-10-22 985909
6.1             sfha-rhel6_x86_64-6.1   Y  2013-12-02 1047939

Available Maintenance Releases for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name       DL OBS AI rel_date   size_KB
=====================================================================
6.0.3           sfha-rhel6_x86_64-6.0.3 -  Y   -  2013-02-01 212845
6.0.5           sfha-rhel6_x86_64-6.0.5 Y  -   -  2014-04-15 199143
6.1.1           sfha-rhel6_x86_64-6.1.1 -  -   -  2014-07-24 208028

Available Patches for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name          DL OBS AI rel_date   size_KB
============================================================================
6.0.1.200-vm    vm-rhel6_x86_64-6.0.1.200  -  Y   -  2012-10-10 21359
6.0.1.300-fs    fs-rhel6_x86_64-6.0.1.300a -  Y   -  2012-12-20 7601

Would you like to download the available Maintenance or Patch releases that
cannot be found in the repository? [y,n,q] (n)
```

4 If you want to download any of the available maintenance releases or patches, enter **y**.

5 If you have not set a default repository for releases you download, the installer prompts you for a directory. (You can also set the default repository in **Set Preferences**).

See "Setting deployment preferences" on page 276.

6 Select an option for downloading products.

The installer downloads the releases you specified and stores them in the repository.

Defining Install Bundles

You can use Install Bundles to directly install the latest base, maintenance, and patch releases on your system. Install Bundles are a combination of base, maintenance, and patch releases that can be bundled and installed or upgraded in one operation.

Note: Install Bundles can be defined only from version 6.1 or later. The exception to this rule is base releases 6.0.1, 6.0.2, or 6.0.4 or later with maintenance release 6.0.5 or later.

To define Install Bundles

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **B**, **Define/Modify Install Bundles**.

You see the following output the first time you enter:

```
Select a Task:

     1)  Create a new Install Bundle
     b)  Back to previous menu

Select the task you would like to perform [1-1,b,q]
```

3 Select option **1**, **Create a new Install Bundle**.

You see the following output:

```
Enter the name of the Install Bundle you would like to define:
{press [Enter] to go back)
```

For example, if you entered:

```
rhel605
```

You see the following output:

```
To create an Install Bundle, the platform type must be selected:

     1)  AIX 5.3                    2)  AIX 6.1
     3)  AIX 7.1                    4)  HP-UX 11.31
     5)  RHEL5 x86_64               6)  RHEL6 x86_64
     7)  RHEL7 x86_64               8)  SLES10 x86_64
     9)  SLES11 x86_64              10) Solaris 9 Sparc
     11) Solaris 10 Sparc           12) Solaris 10 x64
     13) Solaris 11 Sparc           14) Solaris 11 x64
     b)  Back to previous menu


Select the platform of the release for the Install Bundle rhel605:
[1-14,b,q]
```

4 Select the number corresponding to the platform you want to include in the Install Bundle. For example, select the number for the **RHEL5 x86_64** release, **5**.

You see the following output:

```
Details of the Install Bundle: rhel605

Install Bundle Name    rhel605
Platform               RHEL5 x86_64
Base Release           N/A
Maintenance Release    N/A
Patch Releases         N/A

     1)  Add a Base Release
     2)  Add a Maintenance Release
     3)  Add a Patch Release
     4)  Change Install Bundle Name
     b)  Back to previous menu

Select an action to perform on the Install Bundle rhel605 [1-4,b,q]
```

5 Select option **1**, **Add a Base Release**.

You see the following output:

```
     1)  6.0.1
     2)  6.0.2
     3)  6.1
     b)  Back to previous menu

Select the Base Release version to add to the Install Bundle rhel605
[1-3,b,q]
```

6 Select option **1**, **6.0.1**.

You see the following output:

```
Symantec Storage Foundation and High Availability Solutions 6.2 Deployment Server Program
                    pilotlnx11

Details of the Install Bundle: rhel605

Install Bundle Name    rhel605
Platform               RHEL5 x86_64
Base Release           6.0.1
Maintenance Release    N/A
Patch Releases         N/A

     1)  Remove Base Release 6.0.1
     2)  Add a Maintenance Release
     3)  Add a Patch Release
     4)  Change Install Bundle Name
     b)  Back to previous menu

Select an action to perform on the Install Bundle rhel605 [1-4,b,q]
```

7 Select option **2**, **Add a Maintenance Release**.

You see the following output:

```
     1)  6.0.5
     b)  Back to previous menu

Select the Maintenance Release version to add to the Install Bundle
rhel605 [1-1,b,q]
```

8 Select option **1**, **6.0.5**.

You see the following output:

```
Symantec Storage Foundation and High Availability Solutions 6.2 Deployment Server Program
                                  pilotlnx11

Details of the Install Bundle: rhel605

Install Bundle Name    rhel605
Platform               RHEL5 x86_64
Base Release           6.0.1
Maintenance Release    6.0.5
Patch Releases         N/A

     1)  Remove Base Release 6.0.1
     2)  Remove Maintenance Release 6.0.5
     3)  Add a Patch Release
     4)  Save Install Bundle rhel605
     5)  Change Install Bundle Name
     b)  Back to previous menu

Select an action to perform on the Install Bundle rhel605
[1-5,b,q]
```

9 Select option **4**, **Save Install Bundle**.

You see the following output:

```
Install Bundle rhel605 has been saved successfully

Press [Enter] to continue:
```

If there are no releases for the option you selected, you see a prompt saying that there are no releases at this time. You are prompted to continue.

After selecting the desired base, maintenance, or patch releases, you can choose to save your Install Bundle.

The specified Install Bundle is saved on your system. The specified Install Bundle is available as an installation option when using the **I) Install/Upgrade Systems** option to perform an installation or upgrade.

Creating Install Templates

You can use Install Templates to discover installed components (RPMs, patches, products, or versions) on a system that you want to replicate. Use Install Templates to automatically install those same components on to other systems.

To create Install Templates

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

2 You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

3 Select option **T**, **Create Install Templates**.

4 You see the following output:

```
Select a Task:

     1)  Create a new Install Template
     b)  Back to previous menu

Select the task you would like to perform [1-1,b,q]
```

5 Select option **1**, **Create a new Install Template**.

You see the following output:

```
Enter the system names separated by spaces for creating an Install Template:
(press [Enter] to go back)
```

For example, if you entered **`rhel89202`** as the system name, you see the following output:

```
Enter the system names separated by spaces for version checking: rhel89202

Checking communication on rhel89202 ..... Done
Checking installed products on rhel89202 ..... Done

Platform of rhel89202:
   Linux RHEL 6.3 x86_64

Installed product(s) on rhel89202:
   Symantec Storage Foundation Cluster File System HA - 6.1.1 - license vxkeyless

Product:
   Symantec Storage Foundation Cluster File System HA - 6.1.1 - license vxkeyless

Packages:
   Installed Required packages for Symantec Storage Foundation Cluster File System HA 6.1.1:
          #PACKAGE     #VERSION
          VRTSamf      6.1.1.000
          VRTSaslapm   6.1.1.000
          ..........   .........
          ..........   .........
          VRTSvxfs     6.1.1.000
          VRTSvxvm     6.1.1.000

   Installed optional packages for Symantec Storage Foundation Cluster File System HA 6.1.1:
          #PACKAGE     #VERSION
          VRTSdbed     6.1.1.000
          VRTSgms      6.1.0.000
          ..........   .........
          ..........   .........
          VRTSvcsdr    6.1.0.000
          VRTSvcsea    6.1.1.000

   Missing optional packages for Symantec Storage Foundation Cluster File System HA 6.1.1:
          #PACKAGE
```

```
          VRTScps
          VRTSfssdk
          VRTSlvmconv

Summary:

Packages:
 17 of 17 required Symantec Storage Foundation Cluster File System HA 6.1.1 packages installed
 8 of 11 optional Symantec Storage Foundation Cluster File System HA 6.1.1 packages installed

Installed Public and Private Hot Fixes for Symantec Storage Foundation Cluster File
System HA 6.1.1:
 None

Would you like to generate a template file based on the above release information? [y,n,q] (y)

 1)  rhel89202
 b)  Back to previous menu

Select a machine list to generate the template file [1-1,b,q]
```

6 Select option **1**, **rhel89202**.

You see the following output:

```
Enter the name of the Install Template you would like to define:
(press [Enter] to go back)
```

7 Enter the name of your Install Template. For example, if you enter **MyTemplate** as the name for your Install Template, you would see the following:

```
Install Template MyTemplate has been saved successfully

Press [Enter] to continue:
```

All of the necessary information is stored in the Install Template you created.

Deploying Symantec releases

You can use the Deployment Server to deploy your licensed Symantec products dating back to version 5.1. If you know which product version you want to install, follow the steps in this section to install it.

You can use the Deployment Server to install the following:

- A single Symantec release
- Two or more releases using defined Install Bundles
  See "Defining Install Bundles" on page 290.
- Installed components on a system that you want to replicate on another system
  See "Creating Install Templates" on page 295.

To deploy a specific Symantec release

1 From the directory in which you installed your Symantec product (version 6.1 or later), launch the Deployment Server with the upgrade and install systems option. For example, enter the following:

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **I**, **Install/Upgrade Systems**.

You see the following output:

```
    1)  AIX 5.3
    2)  AIX 6.1
    3)  AIX 7.1
    4)  RHEL5 x86_64
    b)  Back to previous menu

Select the platform of the available release(s) to be
upgraded/installed [1-4,b,q,?]
```

3 Select the number corresponding to the platform for the release you want to deploy. For example, select the number for the **RHEL5 x86_64** release or the **AIX 6.1** release.

You see the following output:

```
1) Install/Upgrade systems using a single release
2) Install/Upgrade systems using an Install Bundle
3) Install systems using an Install Template
b) Back to previous menu

Select the method by which you want to Install/Upgrade your systems
[1-3,b,q]
```

4 Section option **1**, **Install/Upgrade systems using a single release** if you want to deploy a specific Symantec release.

Select a Symantec product release.

The installation script is executed and the release is deployed on the specified server.

To deploy an Install Bundle

1 Follow Steps 1 - 3.

2 Select option **2**, **Install/Upgrade systems using an Install Bundle**.

You see the following output:

```
1) <NameofInstallBundle1>
2) <NameofInstallBundle2>
b) Back to previous menu

Select the bundle to be installed/upgraded [1-2,b,q]
```

You see the following output:

```
Enter the platform target system name(s) separated by spaces:
[press [Enter] to go back)
```

3 Enter the name of the target system for which you want to install or upgrade the Install Bundle.

The installation script for the selected Install Bundle is executed, and the Install Bundle is deployed on the specified target system.

To deploy an Install Template

1. Follow Steps 1 - 3.
2. Select option **3**, **Install/Upgrade systems using an Install Template**.

   You see the following output:

   ```
   1) <NameofInstallTemplate>
   b) Back to previous menu

   Select the template to be installed [1-1,b,q] 1
   ```

   You see the following output:

   ```
   Enter the platform target system name(s) separated by spaces:
   [press [Enter] to go back)
   ```

   The installation script for the selected Install Template is executed, and the Install Template is deployed on the specified target system.

Connecting the Deployment Server to SORT using a proxy server

You can use a proxy server, a server that acts as an intermediary for requests from clients, for connecting the Deployment Server to the Symantec Operations Readiness Tools (SORT) website.

To enable the proxy access, run the following commands to set the shell environment variables before you launch Deployment Server. The shell environment variables enable Deployment Server to use the proxy server myproxy.mydomain.com which connects to port 3128.

```
http_proxy="http://myproxy.mydomain.com:3128"
export http_proxy

ftp_proxy="http://myproxy.mydomain.com:3128"
export ftp_proxy
```

The lines above can be added to the user's shell profile. For the bash shell, the profile is the `~/.bash_profile` file.

Section 8

Upgrade of SFHA

- Chapter 18. Planning to upgrade SFHA
- Chapter 19. Upgrading Storage Foundation and High Availability
- Chapter 20. Performing a rolling upgrade of SFHA
- Chapter 21. Performing a phased upgrade of SFHA
- Chapter 22. Performing an automated SFHA upgrade using response files
- Chapter 23. Performing post-upgrade tasks

Chapter 18

Planning to upgrade SFHA

This chapter includes the following topics:

- Upgrade methods for SFHA
- Supported upgrade paths for SFHA 6.2
- Considerations for upgrading SFHA to 6.2 on systems configured with an Oracle resource
- About using the installer to upgrade when the root disk is encapsulated
- Preparing to upgrade SFHA
- Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

Upgrade methods for SFHA

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

Table 18-1 Review this table to determine how you want to perform the upgrade

Upgrade types and considerations	Methods available for upgrade
Typical upgrades—use a Symantec provided tool or you can perform the upgrade manually. Requires some server downtime.	Script-based—you can use this method to upgrade for the supported upgrade paths Web-based—you can use this method to upgrade for the supported upgrade paths Response file—you can use this method to upgrade from the supported upgrade paths

Table 18-1 Review this table to determine how you want to perform the upgrade *(continued)*

Upgrade types and considerations	Methods available for upgrade
Rolling upgrade—use a Symantec provided tool or you can perform the upgrade manually. Requires the least amount of server downtime.	Script-based—you can use this method to upgrade from the previous release Web-based—you can use this method to upgrade from the previous release Response files—you can use this method to upgrade from the supported upgrade paths
Phased upgrades—use a Symantec provided tool and some manual steps. Requires a lesser server downtime than a regular upgrade.	Script-based with some manual steps—you can use this method to upgrade from the previous release Web-based —you can use this method to upgrade from the previous release
Upgrade from any supported UNIX or Linux platform to any other supported UNIX or Linux platform.	Deployment Server See "About the Deployment Server" on page 270.
Simultaneously upgrade base releases, maintenance patches, and patches.	Install Bundles See "Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches" on page 318.

Supported upgrade paths for SFHA 6.2

The following tables describe upgrading to 6.2.

Table 18-2 RHEL6 x64 upgrades using the script- or web-based installer

Symantec product versions	RHEL5	RHEL 6
5.1 5.1 RPs 5.1 SP1RP1	No upgrade path exists. Uninstall the product. Upgrade the operating system to one of the supported Linux versions, and then use the installer script to install 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.	N/A
5.1 SP1 PR2	N/A	Upgrade the operating system to one of the supported Linux versions, and then use the installer script to upgrade to 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.
5.1 SP1 PR3 5.1 SP1 RP2 5.1 SP1 RP3 5.1 SP1 RP4	No upgrade path exists. Uninstall the product. Upgrade the operating system to one of the supported Linux versions, and then use the installer script to install 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.	Upgrade the operating system to one of the supported Linux versions, and then use the installer script to upgrade to 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.
6.0 6.0 RPs 6.0.1 6.0.2 6.0.3 6.0.5	No upgrade path exists. Uninstall the product. Upgrade the operating system to one of the supported Linux versions, and then use the installer script to install 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.	Upgrade the operating system to one of the supported Linux versions, and then use the installer script to upgrade to 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.

Table 18-2 RHEL6 x64 upgrades using the script- or web-based installer *(continued)*

Symantec product versions	RHEL5	RHEL 6
6.1 6.1.1	No upgrade path exists. Uninstall the product. Upgrade the operating system to one of the supported Linux versions, and then use the installer script to install 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.	Use the installer to upgrade your Symantec product to 6.2.

Note: Symantec does not support the RHEL5 operating system with SFHA 6.2.

Note: For upgrades on RHEL 6, RHEL 6 Update 1, Update 2, use a phased upgrade instead of a rolling upgrade.

Table 18-3 SLES 11 x86-x64 upgrades using the script- or web-based installer

Symantec product versions	SLES 10	SLES 11
5.1 5.1 RPx 5.1 SP1 5.1 SP1 RPx	No upgrade path exists. Uninstall the product. Upgrade the operating system to one of the supported Linux versions, and then use the installer script to install 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.	Upgrade the operating system to one of the supported Linux versions, and then use the installer script to upgrade to 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.

Table 18-3 SLES 11 x86-x64 upgrades using the script- or web-based installer *(continued)*

Symantec product versions	SLES 10	SLES 11
6.0 6.0 RPx	No upgrade path exists. Uninstall the product. Upgrade the operating system to one of the supported Linux versions, and then use the installer script to install 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.	Upgrade the operating system to one of the supported Linux versions, and then use the installer script to upgrade to 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.
6.0.1 6.0.2 6.0.3 6.0.5	No upgrade path exists. Uninstall the product. Upgrade the operating system to one of the supported Linux versions, and then use the installer script to install 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.	Upgrade the operating system to one of the supported Linux versions, and then use the installer script to upgrade to 6.2. See the *Storage Foundation and High Availability Release Notes* for the supported Linux versions.
6.0.4 6.1 6.1.1	N/A	Use the installer to upgrade to 6.2.

Note: For upgrades on SLES 11, SLES 11 SP1, use a phased upgrade instead of a rolling upgrade.

Considerations for upgrading SFHA to 6.2 on systems configured with an Oracle resource

If you plan to upgrade SFHA running on systems configured with an Oracle resource, set the `MonitorOption` attribute to 0 (zero) before you start the upgrade. If you use the product installer for the rolling upgrade, it sets the `MonitorOption` to 0 through

its scripts. In a manual upgrade, the `MonitorOption` value must be set to 0 using the `hares` command. When the upgrade is complete, invoke the `build_oraapi.sh` script, and then set the `MonitorOption` to 1 to enable the Oracle health check.

For more information on enabling the Oracle health check, see the *Symantec Cluster Server Agent for Oracle Installation and Configuration Guide*.

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

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

Table 18-4 Upgrading using the installer when the root disk is encapsulated (SUSE Linux Enterprise 11)

Starting version	Ending version	Action required
5.1 5.1 RPx	6.2	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 SP1 5.1 SP1 RPx	6.2	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0 6.0 RPx	6.2	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0.1 6.0.3 6.0.4 6.0.5	6.2	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.1 6.1.1	6.2	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Table 18-5 Upgrading using the installer when the root disk is encapsulated (Red Hat Enterprise Linux 6)

Starting version	Ending version	Action required
5.1 SP1 PR2 5.1 SP1 RP3	6.2	You need to unencapsulate the root disk. The installer exits.
5.1 SP1 RP4	6.2	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0 6.0 RPx	6.2	You need to unencapsulate the root disk. The installer exits.
6.0.1 6.0.3 6.0.5 6.1 6.1.1	6.2	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Preparing to upgrade SFHA

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

Getting ready for the upgrade

Complete the following tasks before you perform the upgrade:

- Review the Symantec Technical Support website for additional information: http://www.symantec.com/techsupp/
- Make sure that the administrator who performs the upgrade has root access and a good knowledge of the operating system's administration.
- Make sure that all users are logged off and that all major user applications are properly shut down.
- Make sure that you have created a valid backup.
  See “Creating backups” on page 309.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the RPMs, for example `/packages/Veritas` when the root file

system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.
Do not put the files under `/tmp`, which is erased during a system restart.
Do not put the files on a file system that is inaccessible before running the upgrade script.
You can use a Symantec-supplied disc for the upgrade as long as modifications to the upgrade script are not required.

- For any startup scripts in `/etc/init.d/`, comment out any application commands or processes that are known to hang if their file systems are not present.
- Make sure that the current operating system supports version 6.2 of the product. If the operating system does not support it, plan for a staged upgrade.
- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Symantec products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/fstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/fstab` and not mounted during the upgrade. The active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure that the file systems are clean before upgrading.
- Upgrade arrays (if required).
See "Upgrading the array support" on page 317.
- To reliably save information on a mirrored disk, shut down the system and physically remove the mirrored disk. Removing the disk in this manner offers a failback point.
- Determine if the root disk is encapsulated.
See "Determining if the root disk is encapsulated" on page 310.
- If CP server-based coordination points are used in your current fencing configuration, then check that your CP servers are upgraded to 6.2 before starting the upgrade process.
- Make sure that DMP support for native stack is disabled (`dmp_native_support=off`). If DMP support for native stack is enabled (`dmp_native_support=on`), the installer may detect it and ask you to restart the system.

Creating backups

Save relevant system information before the upgrade.

To create backups

1. Log in as superuser.
2. Before the upgrade, ensure that you have made backups of all data that you want to preserve.
3. Back up information in files such as `/boot/grub/menu.lst`, `/etc/grub.conf` or `/etc/lilo.conf`, and `/etc/fstab`.
4. Installer verifies that recent backups of configuration files in VxVM private region have been saved in `/etc/vx/cbr/bk`.

   If not, a warning message is displayed.

   Warning: Backup `/etc/vx/cbr/bk` directory.

   ---

5. Copy the `fstab` file to `fstab.orig`:

   ```
   # cp /etc/fstab /etc/fstab.orig
   ```

6. Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
7. If you install the high availability version of the Symantec Storage Foundation 6.2 software, follow the guidelines that are given in the *Symantec Cluster Server Installation Guide* and *Symantec Cluster Server Release Notes* for information on preserving your VCS configuration across the installation procedure.
8. Back up the external `quotas` and `quotas.grp` files.

   If you are upgrading from 6.0.3, you must also back up the `quotas.grp.64` and `quotas.64` files.
9. If you are planning on performing a Phased or Rolling upgrade from 6.0.3 and use quotas, you need to disable them:

   ```
   # vxquotaoff -av
   ```

10. Verify that quotas are turned off on all the mounted file systems.

Determining if the root disk is encapsulated

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

```
# df -v /
```

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

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

See "About using the installer to upgrade when the root disk is encapsulated" on page 307.

Pre-upgrade tasks for migrating the SFDB repository database

Note: The Sfua_Base repository resource group will be removed from the main.cf file. It is not required as a separate service group for SFHA 6.2.

Perform the following before upgrading SFHA.

To prepare to migrate the repository database

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

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

Warning: The Database Flashsnap clone database will not be able to be carried over after upgrading. You must create a new Database Flashsnap clone database after upgrading to 6.2.

Pre-upgrade planning for Volume Replicator

Before installing or upgrading Volume Replicator (VVR):

- Confirm that your system has enough free disk space to install VVR.
- Make sure you have root permissions. You must have root permissions to perform the install and upgrade procedures.
- If replication using VVR is configured, Symantec recommends that the disk group version is at least 110 prior to upgrading.
  You can check the Disk Group version using the following command:

```
# vxdg list diskgroup
```

- If replication using VVR is configured, make sure the size of the SRL volume is greater than 110 MB.

Refer to the *Symantec Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

- If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date on all the hosts.

```
# /usr/sbin/vxrlink -g diskgroup status rlink_name
```

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

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

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

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

Planning an upgrade from the previous VVR version

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

See the *Symantec Storage Foundation Release Notes* for information regarding VVR support for replicating across Storage Foundation versions.

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

Also, if you specify TCP as the network protocol, the VVR versions on the Primary and Secondary determine whether the checksum is calculated. As shown in Table 18-6, if either the Primary or Secondary are running a version of VVR prior to 6.2, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 6.2, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 18-6 VVR versions and checksum calculations

VVR prior to 6.2 (DG version <= 140)	VVR 6.2 (DG version >= 150)	VVR calculates checksum TCP connections?
Primary	Secondary	Yes
Secondary	Primary	Yes
Primary and Secondary		Yes
	Primary and Secondary	No

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

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

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

Planning and upgrading VVR to use IPv6 as connection protocol

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

This release supports the following configurations for VVR:

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

Preparing to upgrade VVR when VCS agents are configured

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

- Freezing the service groups and stopping all the applications
- Preparing for the upgrade when VCS agents are configured

Freezing the service groups and stopping all the applications

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

To freeze the service groups and stop applications

Perform the following steps for the Primary and Secondary clusters:

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

   Note: You must also stop any remaining applications not managed by VCS.
4. On any node in the cluster, make the VCS configuration writable:

   ```
   # haconf -makerw
   ```
5. On any node in the cluster, list the groups in your configuration:

   ```
   # hagrp -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.

```
# hagrp -freeze group_name -persistent
```

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

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

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

Note: Continue only after you have performed steps 3 to step 7 for each node of the 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            sys2    ONLINE
 ORAGrp           State            sys2    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 node of the 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 315.

Determining the nodes on which disk groups are online

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

To determine the online disk groups

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

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

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

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

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

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

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

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

Preparing for the upgrade when VCS agents are configured

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

To prepare a configuration with VCS agents for an upgrade

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

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

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

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

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

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

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

   ```
   # vxrecover -bs
   ```

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

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

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

Upgrading the array support

The Storage Foundation 6.2 release includes all array support in a single RPM, `VRTSaslapm`. The array support RPM includes the array support previously included in the `VRTSvxvm` RPM. The array support RPM also includes support previously packaged as external Array Support Libraries (ASLs) and array policy modules (APMs).

See the 6.2 Hardware Compatibility List for information about supported arrays.

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

After you have installed Storage Foundation 6.2, Symantec provides support for new disk arrays through updates to the `VRTSaslapm` RPM.

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

Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

Beginning with version 6.1, Symantec offers you a method to easily install or upgrade your systems directly to a base, maintenance, patch level or a combination of multiple patches and packages together in one step using Install Bundles. With Install Bundles, the installer has the ability to merge so that customers can install or upgrade directly to maintenance or patch levels in one execution. The various scripts, RPMs, and patch components are merged, and multiple releases are installed together as if they are one combined release. You do not have to perform two or more install actions to install or upgrade systems to maintenance levels or patch levels.

Releases are divided into the following categories:

Table 18-7 Release Levels

Level	Content	Form factor	Applies to	Release types	Download location
Base	Features	RPMs	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect
Maintenance	Fixes, new features	RPMs	All products	Maintenance Release (MR), Rolling Patch (RP)	Symantec Operations Readiness Tools (SORT)

Table 18-7 Release Levels *(continued)*

Level	Content	Form factor	Applies to	Release types	Download location
Patch	Fixes	RPMs	Single product	P-Patch, Private Patch, Public patch	SORT, Support site

When you install or upgrade using Install Bundles:

- SFHA products are discovered and assigned as a single version to the maintenance level. Each system can also have one or more patches applied.
- Base releases are accessible from FileConnect that requires customer serial numbers. Maintenance and patch releases can be automatically downloaded from SORT. You can download them from the SORT website manually or use the `deploy_sfha` script.
- Patches can be installed using automated installers from the 6.0.1 version or later.
- Patches can now be detected to prevent upgrade conflict. Patch releases are not offered as a combined release. They are only available from Symantec Technical Support on a need basis.

You can use the `-base_path` and `-patch_path` options to import installation code from multiple releases. You can find RPMs and patches from different media paths, and merge RPM and patch definitions for multiple releases. You can use these options to use new task and phase functionality to correctly perform required operations for each release component. You can install the RPMs and patches in defined phases using these options, which helps you when you want to perform a single start or stop process and perform pre and post operations for all level in a single operation.

Four possible methods of integration exist. All commands must be executed from the highest base or maintenance level install script.

For example:

1. Base + maintenance:

   This integration method can be used when you install or upgrade from a lower version to 6.2.1.

   Enter the following command:

   ```
   # installmr -base_path <path_to_base>
   ```

2. Base + patch:

This integration method can be used when you install or upgrade from a lower version to 6.2.0.100.

Enter the following command:

```
# installer -patch_path <path_to_patch>
```

3. Maintenance + patch:

   This integration method can be used when you upgrade from version 6.2 to 6.2.1.100.

   Enter the following command:

   ```
   # installmr -patch_path <path_to_patch>
   ```

4. Base + maintenance + patch:

   This integration method can be used when you install or upgrade from a lower version to 6.2.1.100.

   Enter the following command:

   ```
   # installmr -base_path <path_to_base>
   -patch_path <path_to_patch>
   ```

Note: From the 6.1 or later release, you can add a maximum of five patches using *-patch_path <path_to_patch> -patch2_path <path_to_patch> ... -patch5_path <path_to_patch>*

Chapter 19

Upgrading Storage Foundation and High Availability

This chapter includes the following topics:

- Upgrading Storage Foundation and High Availability from previous versions to 6.2
- Upgrading SFHA using the web-based installer
- Upgrading Volume Replicator
- Upgrading SFDB

Upgrading Storage Foundation and High Availability from previous versions to 6.2

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

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

See "Upgrading Storage Foundation and High Availability using the script-based installer" on page 322.

If you need to upgrade your kernel with Symantec Storage Foundation 6.2 already installed, use the kernel upgrade procedure.

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

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

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

To upgrade Symantec Storage Foundation and High Availability

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

   List all service groups:

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

   For each service group listed, take it offline:

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

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

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

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

   ```
   # df -h | grep vxfs
   ```

5. Unmount all Storage Checkpoints and file systems:

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

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

   ```
   # echo "8192B.p S" | fsdb -t vxfs filesystem | grep clean
   flags 0 mod 0 clean clean_value
   ```

   A *clean_value* value of `0x5a` indicates the file system is clean, `0x3c` indicates the file system is dirty, and `0x69` indicates the file system is dusty. A dusty file system has pending extended operations.

   Perform the following steps in the order listed:

- If a file system is not clean, enter the following commands for that file system:

```
# fsck -t vxfs filesystem
# mount -t vxfs filesystem mountpoint
# umount mountpoint
```

This should complete any extended operations that were outstanding on the file system and unmount the file system cleanly.
There may be a pending large RPM clone removal extended operation if the `umount` command fails with the following error:

```
file system device busy
```

You know for certain that an extended operation is pending if the following message is generated on the console:

```
Storage Checkpoint asynchronous operation on file_system
file system still in progress.
```

- If an extended operation is pending, you must leave the file system mounted for a longer time to allow the operation to complete. Removing a very large RPM clone can take several hours.
- Repeat this step to verify that the unclean file system is now clean.

7 If a cache area is online, you must take the cache area offline before you upgrade the VxVM RPM. Use the following command to take the cache area offline:

```
# sfcache offline cachename
```

8 Stop activity to all VxVM volumes. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes.

9 Stop all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

To verify that no volumes remain open, use the following command:

```
# vxprint -Aht -e v_open
```

10 Make a record of the mount points for VxFS file systems and VxVM volumes that are defined in the `/etc/fstab` file. You will need to recreate these entries in the `/etc/fstab` file on the freshly installed system.

11 Perform any necessary preinstallation checks.

12 To invoke the common installer, run the `installer` command on the disc as shown in this example:

```
# cd /cdrom/cdrom0
# ./installer
```

13 Enter `G` to upgrade and press Return.

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

```
Enter the 64 bit RHEL5 system names separated
by spaces : [q, ?] host1
    host2
```

Depending on your existing configuration, various messages and prompts may appear. Answer the prompts appropriately.

During the system verification phase, the installer checks if the boot disk is encapsulated and the upgrade's path. If the upgrade is not supported, you need to un-encapsulate the boot disk.

15 The installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.

16 The installler lists the RPMs to install or upgrade. You are prompted to confirm that you are ready to upgrade.

17 The installer discovers if any of the systems that you are upgrading have mirrored encapsulated boot disks. You now have the option to create a backup of the systems' root disks before the upgrade proceeds. If you want to split the mirrors on the encapsulated boot disks to create the backup, answer **y**.

18 The installer then prompts you to name the backup root disk. Enter the name for the backup and mirrored boot disk or press **Enter** to accept the default.

Note: The split operation can take some time to complete.

---

19 You are prompted to start the split operation. Press **y** to continue.

20 The installer lists the RPMs that it installs or upgades.

21 Reboot the system if the boot disk is encapsulated before the upgrade.

22 If you need to re-encapsulate and mirror the root disk on each of the nodes, follow the procedures in the "Administering Disks" chapter of the *Symantec Storage Foundation Administrator's Guide*.

If you have split the mirrored root disk to back it up, then after a successful reboot, verify the upgrade and re-join the backup disk group. If the upgrade fails, revert to the backup disk group.

23 If necessary, reinstate any missing mount points in the `/etc/fstab` file on each node that you recorded in step 10.

24 If any VCS configuration files need to be restored, stop the cluster, restore the files to the `/etc/VRTSvcs/conf/config` directory, and restart the cluster.

25 Make the VCS configuration writable again from any node in the upgraded group:

```
# haconf -makerw
```

26 Enter the following command on each node in the upgraded group to unfreeze HA service group operations:

```
# hasys -unfreeze -persistent nodename
```

27 Make the configuration read-only:

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

28 Bring all of the VCS service groups, such as failover groups, online on the required node using the below command:

```
# hagrp -online groupname -sys nodename
```

29 Restart all the volumes by entering the following command for each disk group:

```
# vxvol -g diskgroup startall
```

30 Remount all VxFS file systems and Storage Checkpoints on all nodes:

```
# mount /filesystem
# mount /checkpoint_name
```

31 You can perform the following optional configuration steps:

- If you want to use features of Symantec Storage Foundation 6.2 for which you do not currently have an appropriate license installed, obtain the license and run the `vxlicinst` command to add it to your system.

- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
  See "Upgrading VxVM disk group versions" on page 381.

Upgrading SFHA using the web-based installer

This section describes upgrading SFHA with the web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems.

To upgrade SFHA

1. Perform the required steps to save any data that you want to preserve. For example, make configuration file backups.
2. If you want to upgrade 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 -any
   ```

3. Start the web-based installer.

   See "Starting the web-based installer" on page 170.
4. Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.
5. Installer detects the product that is installed on the specified system. It shows the cluster information and lets you confirm if you want to perform upgrade on the cluster. Select **Yes** and click **Next**.
6. The installer discovers if any of the systems that you want to upgrade have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the book disk group. To create the backup, check the **Split mirrors on all the systems** box. Check the appropriate box to use the same name for the backup disk group on all systems. You can use the default name or choose a new one. Check the systems where you want to create the backup. When you are ready, click the **Next** button.
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 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant more usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**.
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

9 If you are prompted to restart the systems, enter the following restart command:

```
# /sbin/shutdown -r now
```

10 After the upgrade, if the product is not configured, the web-based installer asks: "Do you want to configure this product?" If the product is already configured, it does not ask any questions.

11 If you want to upgrade application clusters that use VCS or SFHA to 6.2, make sure that you upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.2. For instructions to upgrade VCS or SFHA, see the *VCS or SFHA Installation Guide*.

12 Only perform this step if you have split the mirrored root disk to back it up. After a successful restart, verify the upgrade and re-join the backup disk group into the upgraded boot disk group. If the upgrade fails, revert the upgrade boot disk group to the backup disk group.

See "Re-joining the backup boot disk group into the current disk group" on page 369.

See "Reverting to the backup boot disk group after an unsuccessful upgrade" on page 369.

Upgrading Volume Replicator

If a previous version of Volume Replicator (VVR) is configured, the product installer upgrades VVR automatically when you upgrade the Storage Foundation products.

You have the option to upgrade without disrupting replication.

See "Upgrading VVR without disrupting replication" on page 328.

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 312.

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 sec_hostname
   ```

2. Upgrade from VVR 5.1 or later to VVR 6.2 on the Secondary.

3. Do one of the following:

   - Upgrade the disk group now. Enter the following:

     ```
     # vxdg upgrade dgname
     ```

   - Upgrade the disk group later.

If you upgrade the disk group later, be sure to pause replication before you upgrade the disk group. Also, after pausing replication, upgrade the disk group on Primary as well as Secondary.

4 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 product installer to upgrade the Primary.

To upgrade the Primary

1 Stop replication to the Primary host by initiating a Primary pause using the following command:

```
# vradmin -g diskgroup pauserep local_rvgname
```

2 Upgrade from VVR 5.1 or later to VVR 6.2 on the Secondary.

3 Do one of the following:

- Upgrade the disk group now. Enter the following:

  ```
  # vxdg upgrade dgname
  ```

- Upgrade the disk group later.
  If you upgrade the disk group later, be sure to pause replication before you upgrade the disk group. Also, after pausing replication, upgrade the disk group on Primary as well as Secondary.

4 Resume the replication from the Primary using the following command:

```
# vradmin -g diskgroup resumerep local_rvgname
    sec_hostname
```

See “Planning an upgrade from the previous VVR version” on page 312.

Upgrading SFDB

While upgrading from 6.x to 6.2 the SFDB tools are enabled by default, which implies that the vxdbd daemon is configured. You can enable the SFDB tools, if they are disabled.

To enable SFDB tools

1. Log in as root.
2. Run the following command to configure and start the vxdbd daemon.

```
# /opt/VRTS/bin/sfae_config enable
```

Note: If any SFDB installation with authentication setup is upgraded to 6.2, the commands fail with an error. To resolve the issue, setup the SFDB authentication again. For more information, see the *Storage and Availability Management for Oracle Databases* or *Storage and Availability Management for DB2 Databases*.

Chapter 20

Performing a rolling upgrade of SFHA

This chapter includes the following topics:

- About rolling upgrades
- Supported rolling upgrade paths
- Performing a rolling upgrade using the script-based installer
- Performing a rolling upgrade of SFHA using the web-based installer

About rolling upgrades

The rolling upgrade minimizes downtime for highly available clusters to the amount of time that it takes to perform a service group failover. The rolling upgrade has two main phases where the installer upgrades kernel RPMs in phase 1 and VCS agent related RPMs in phase 2.

Note: You need to perform a rolling upgrade on a completely configured cluster.

The following is an overview of the flow for a rolling upgrade:

1. The installer performs prechecks on the cluster.
2. Application downtime occurs during the first phase as the installer moves service groups to free nodes for the upgrade. The only downtime that is incurred is the normal time required for the service group to fail over. The downtime is limited to the applications that are failed over and not the entire cluster.

3. The installer performs the second phase of the upgrade on all of the nodes in the cluster. The second phase of the upgrade includes downtime of the Symantec Cluster Server (VCS) engine HAD, but does not include application downtime.

Figure 20-1 illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 20-1 Example of the installer performing a rolling upgrade

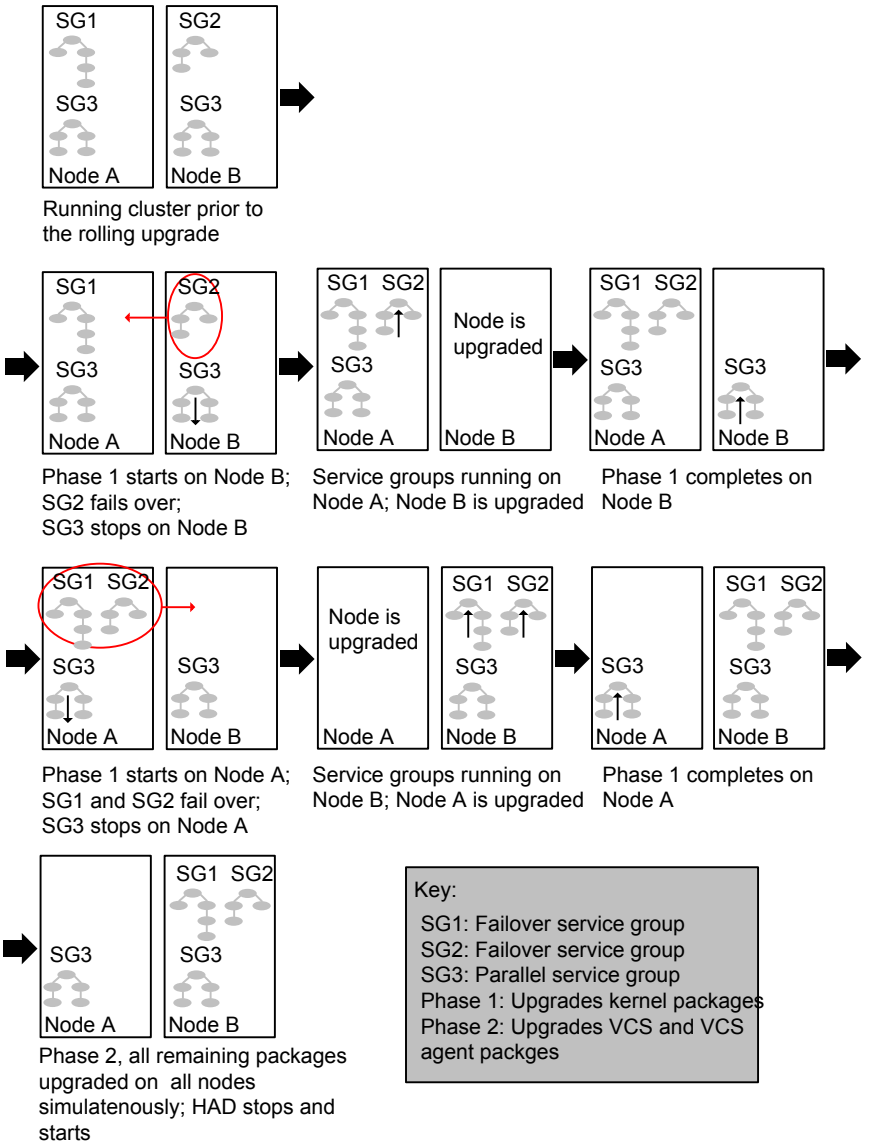

The following limitations apply to rolling upgrades:

- Rolling upgrades are not compatible with phased upgrades. Do not mix rolling upgrades and phased upgrades.
- You can perform a rolling upgrade from 5.1 and later versions.

Supported rolling upgrade paths

You can perform a rolling upgrade of SFHA with the script-based installer, the web-based installer, or manually.

The rolling upgrade procedures support only minor operating system upgrades.

Table 20-1 shows the versions of SFHA for which you can perform a rolling upgrade to Symantec Storage Foundation 6.2.

Table 20-1 Supported rolling upgrade paths

Platform	SFHA version
RHEL6:	5.1SP1RP4 6.0.1, 6.0.2, 6.0.3, 6.0.5 6.1, 6.1.1
SLES11:	6.0.1, 6.0.2, 6.0.3, 6.0.4, 6.0.5 6.1, 6.1.1

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.2, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in Patch Lookup on the SORT website.

Performing a rolling upgrade using the script-based installer

Before you start the rolling upgrade, make sure that Symantec Cluster Server (VCS) is running on all the nodes of the cluster.

Stop all activity for all the VxVM volumes that are not under VCS control. For example, stop any applications such as databases that access the volumes, and unmount any file systems that have been created on the volumes. Then stop all the volumes.

Unmount alll VxFS file systems that are not under VCS control.

To perform a rolling upgrade

1 Complete the preparatory steps on the first sub-cluster.

Unmount all VxFS file systems not under VCS control:

```
# umount mount_point
```

2 If a cache area is online, you must take the cache area offline before you upgrade the VxVM RPM. Use the following command to take the cache area offline:

```
# sfcache offline cachename
```

3 Log in as superuser and mount the SFHA 6.2 installation media.

4 From root, start the installer.

```
# ./installer
```

5 From the menu, select `Upgrade a Product` and from the sub menu, select `Rolling Upgrade`.

6 The installer suggests system names for the upgrade. Press **Enter** to upgrade the suggested systems, or enter the name of any one system in the cluster on which you want to perform a rolling upgrade and then press **Enter**.

7 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type **y** to continue.

8 The installer inventories the running service groups and determines the node or nodes to upgrade in phase 1 of the rolling upgrade. Type **y** to continue. If you choose to specify the nodes, type **n** and enter the names of the nodes.

9 The installer performs further prechecks on the nodes in the cluster and may present warnings. You can type **y** to continue or quit the installer and address the precheck's warnings.

10 Review the end-user license agreement, and type **y** if you agree to its terms.

11 If the boot disk is encapsulated and mirrored, you can create a backup boot disk.

If you choose to create a backup boot disk, type **y**. Provide a backup name for the boot disk group or accept the default name. The installer then creates a backup copy of the boot disk group.

The installer lists the RPMs to upgrade on the selected node or nodes.

12 After the installer detects the online service groups, the installer prompts the user to do one of the following:

- Manually switch service groups
- Use the CPI to automatically switch service groups

The downtime is the time that it normally takes for the service group's failover.

Note: It is recommended that you manually switch the service groups. Automatic switching of service groups does not resolve dependency issues.

13 The installer prompts you to stop the applicable processes. Type **y** to continue.

The installer evacuates all service groups to the node or nodes that are not upgraded at this time. The installer stops parallel service groups on the nodes that are to be upgraded.

14 The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs. The installer asks if you want to update your licenses to the current version. Select **Yes** or **No**. Symantec recommends that you update your licenses to fully use the new features in the current release.

15 If the cluster has configured Coordination Point Server based fencing, then during upgrade, installer may ask the user to provide the new HTTPS Coordination Point Server.

The installer performs the upgrade configuration and starts the processes. If the boot disk is encapsulated before the upgrade, installer prompts the user to reboot the node after performing the upgrade configuration.

16 Complete the preparatory steps on the nodes that you have not yet upgraded.

Unmount all VxFS file systems not under VCS control on all the nodes.

```
# umount mount_point
```

17 The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade. If the installer was invoked on the upgraded (rebooted) nodes, you must invoke the installer again.

The installer repeats step 8 through step 14.

For clusters with larger number of nodes, this process may repeat several times. Service groups come down and are brought up to accommodate the upgrade.

18 When Phase 1 of the rolling upgrade completes, mount all the VxFS file systems that are not under VCS control manually. Begin Phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.

19 The installer determines the remaining RPMs to upgrade. Press **Enter** to continue.

20 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

21 Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

22 The installer stops Symantec Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs prestop, uninstalls old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.

23 If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

24 A prompt message appears to ask if the user wants to read the summary file. You can choose **y** if you want to read the install summary file.

25 If you want to upgrade application clusters that use CP server-based fencing to 6.2, make sure that you upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.2. However, note that the CP server upgraded to 6.2 can support application clusters on 6.1 and later (HTTPS-based communication) and application clusters prior to 6.1 (IPM-based communication). When you configure the CP server, the installer asks the VIPs for HTTPS-based communication (if the clients are on release version 6.1 and later) or VIPs for IPM-based communication (if the clients are on a release version prior to 6.1).

For instructions to upgrade VCS or SFHA on the CP server systems, refer to the appropriate installation guide.

Performing a rolling upgrade of SFHA using the web-based installer

This section describes using the web-based installer to perform a rolling upgrade. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

See "About rolling upgrades" on page 331.

To start the rolling upgrade—phase 1

1 Perform the required steps to save any data that you want to preserve. For example, take backups of configuration files.

2 Start the web-based installer.

See "Starting the web-based installer" on page 170.

3 In the Task pull-down menu, select `Rolling Upgrade`.

The option `Phase-1: Upgrade Kernel packages` is displayed and selected by default.

Click **Next** to proceed.

4 Enter the name of any one system in the cluster on which you want to perform a rolling upgrade.The installer identifies the cluster information of the system and displays the information.

Click **Yes** to confirm the cluster information. The installer now displays the nodes in the cluster that will be upgraded during phase 1 of the upgrade.

5 Review the systems that the installer has chosen for phase 1 of the rolling upgrade. These systems are chosen to minimize downtime during the upgrade.

Click **Yes** to proceed.

The installer validates systems.

6 Review the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.

7 If you have online failover service groups, the installer prompts you to choose to switch these service groups either manually or automatically. Choose any option and follow the steps to switch all the failover service groups to the other subcluster.

8 The installer stops all processes. Click **Next** to proceed.

The installer removes old software and upgrades the software on the systems that you selected.

9 The installer asks if you want to update your licenses to the current version. Select **Yes** or **No**. Symantec recommends that you update your licenses to fully use the new features in the current release.

10 If the cluster has configured Coordination Point Server-based fencing, then during upgrade, installer asks the user to provide the new HTTPS Coordination Point Server. If you are prompted, restart the product.

The installer starts all the relevant processes and brings all the service groups online if the nodes do not require a restart.

11 Restart the nodes, if required.

Restart the installer.

12 Repeat step 5 through step 11 until the kernel RPMs of all the nodes are upgraded. For clusters with larger number of nodes, this process may get repeated several times. Service groups come down and are brought up to accommodate the upgrade.

13 When prompted, start phase 2. Click **Yes** to continue with the rolling upgrade.

To upgrade the non-kernel components—phase 2

1 The installer detects the information of cluster and the state of rolling upgrade.

The installer validates systems. Click **Next**. If it throws an error, address the error and return to the installer.

2 Review the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.

3 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

4 The installer stops the `HAD` and `CmdServer` processes in phase 2 of the rolling upgrade process but the applications continue to run. Click **Next** to proceed.

5 The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted. Register the software and click **Next** to proceed. The installer starts all the relevant processes and brings all the service groups online.

6 If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

7 A prompt message appears to ask if the user wants to read the summary file. You can choose **y** if you want to read the install summary file.

The upgrade is complete.

Chapter 21

Performing a phased upgrade of SFHA

This chapter includes the following topics:

- About phased upgrade
- Performing a phased upgrade using the script-based installer

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:

Table 21-1

Fail over condition	Downtime
You can fail over all your service groups to the nodes that are up.	Downtime equals the time that is taken to offline and online the service groups.
You have a service group that you cannot fail over to a node that runs during upgrade.	Downtime for that service group equals the time that is taken to perform an upgrade and restart the node.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade.

Planning for a phased upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group.

Some rough guidelines follow:

- Split the cluster into two sub-clusters of equal or near equal size.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.
- Before you start the upgrade, back up the VCS configuration files `main.cf` and `types.cf` which are in the `/etc/VRTSvcs/conf/config/` directory.
- Before you start the upgrade make sure that all the disk groups have the latest backup of configuration files in the `/etc/vx/cbr/bk` directory. If not, then run the following command to take the latest backup.

  `# /etc/vx/bin/vxconfigbackup -| [dir] [dgname|dgid]`

Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules.
- When you start the installer, only select SFHA.
- While you perform the upgrades, do not add or remove service groups to any of the nodes.
- After you upgrade the first half of your cluster (the first subcluster), you need to set up password-less SSH or RSH. Create the connection between an upgraded node in the first subcluster and a node from the other subcluster. The node from the other subcluster is where you plan to run the installer and also plan to upgrade.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.

Phased upgrade example

In this example, you have a secure cluster that you have configured to run on four nodes: node01, node02, node03, and node04. You also have four service groups:

sg1, sg2, sg3, and sg4. For the purposes of this example, the cluster is split into two subclusters. The nodes node01 and node02 are in the first subcluster, which you first upgrade. The nodes node03 and node04 are in the second subcluster, which you upgrade last.

Figure 21-1 Example of phased upgrade set up

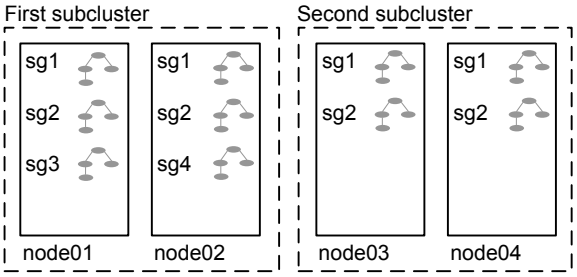

Each service group is running on the nodes as follows:

- sg1 and sg2 are parallel service groups and run on all the nodes.
- sg3 and sg4 are failover service groups. sg3 runs on node01 and sg4 runs on node02.

In your system list, you have each service group that fails over to other nodes as follows:

- sg1 and sg2 are running on all the nodes.
- sg3 and sg4 can fail over to any of the nodes in the cluster.

Phased upgrade example overview

This example's upgrade path follows:

- Move all the failover service groups from the first subcluster to the second subcluster.
- Take all the parallel service groups offline on the first subcluster.
- Upgrade the operating system on the first subcluster's nodes, if required.
- On the first subcluster, start the upgrade using the installation program.
- Get the second subcluster ready.
- Activate the first subcluster. After activating the first cluster, switch the service groups online on the second subcluster to the first subcluster.
- Upgrade the operating system on the second subcluster's nodes, if required.
- On the second subcluster, start the upgrade using the installation program.

- Activate the second subcluster.

See “Performing a phased upgrade using the script-based installer” on page 344.

Performing a phased upgrade using the script-based installer

This section explains how to perform a phased upgrade of SFHA on four nodes with four service groups. Note that in this scenario, VCS and the service groups cannot stay online on the second subcluster during the upgrade of the second subcluster. Do not add, remove, or change resources or service groups on any nodes during the upgrade. These changes are likely to get lost after the upgrade.

An example of a phased upgrade follows. It illustrates the steps to perform a phased upgrade. The example makes use of a secure SFHA cluster.

You can perform a phased upgrade from SFHA 5.1 or other supported previous versions to SFHA 6.2.

See “About phased upgrade” on page 341.

See “Phased upgrade example” on page 342.

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.

```
# hagrp -state
```

The output resembles:

```
#Group   Attribute System Value
sg1      State     node01 |ONLINE|
sg1      State     node02 |ONLINE|
sg1      State     node03 |ONLINE|
sg1      State     node04 |ONLINE|
sg2      State     node01 |ONLINE|
sg2      State     node02 |ONLINE|
sg2      State     node03 |ONLINE|
sg2      State     node04 |ONLINE|
sg3      State     node01 |ONLINE|
sg3      State     node02 |OFFLINE|
sg3      State     node03 |OFFLINE|
sg3      State     node04 |OFFLINE|
sg4      State     node01 |OFFLINE|
sg4      State     node02 |ONLINE|
sg4      State     node03 |OFFLINE|
sg4      State     node04 |OFFLINE|
```

2 Offline the parallel service groups (sg1 and sg2) from the first subcluster. Switch the failover service groups (sg3 and sg4) from the first subcluster (node01 and node02) to the nodes on the second subcluster (node03 and node04). For SFHA, vxfen sg is the parallel service group.

```
# hagrp -offline sg1 -sys node01
# hagrp -offline sg2 -sys node01
# hagrp -offline sg1 -sys node02
# hagrp -offline sg2 -sys node02
# hagrp -switch sg3 -to node03
# hagrp -switch sg4 -to node04
```

3 On the nodes in the first subcluster, unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/sda1              26G  3.3G   22G  14% /
udev                 1007M  352K 1006M   1% /dev
tmpfs                 4.0K     0  4.0K   0% /dev/vx
/dev/vx/dsk/dg2/dg2vol1
                      3.0G   18M  2.8G   1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2
                      1.0G   18M  944M   2% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3
                       10G   20M  9.4G   1% /mnt/dg2/dg2vol3
# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

4 On the nodes in the first subcluster, use the following command to take all the cache area offline:

```
# sfcache offline cachename
```

5 On the nodes in the first subcluster, stop all VxVM volumes (for each disk group) that VCS does not manage.

6 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

7 Freeze the nodes in the first subcluster.

```
# hasys -freeze -persistent node01
# hasys -freeze -persistent node02
```

8 Dump the configuration and make it read-only.

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

9 Verify that the service groups are offline on the first subcluster that you want to upgrade.

```
# hagrp -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|
```

Upgrading the operating system on the first subcluster

You can perform the operating system upgrade on the first subcluster, if required.

Before performing operating system upgrade, it is better to prevent LLT from starting automatically when the node starts. For example, you can do the following:

```
# mv /etc/llttab /etc/llttab.save
```

or you can change the `/etc/default/llt` file by setting LLT_START = **0**.

After you finish upgrading the OS, remember to change the LLT configuration to its original configuration.

Refer to the operating system's documentation for more information.

Upgrading the first subcluster

You now navigate to the installer program and start it.

To start the installer for the phased upgrade

1. Confirm that you are logged on as the superuser and you mounted the product disc.
2. Navigate to the folder that contains installsfha.

   ```
   # cd storage_foundation_high_availability
   ```

3. Make sure that you can ssh or rsh from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
4. Start the installsfha program, specify the nodes in the first subcluster (node1 and node2).

   ```
   # ./installsfha<version> -upgrade node1 node2
   ```

   Where *<version>* is the specific release version.

   See "About the script-based installer" on page 69.

   The program starts with a copyright message and specifies the directory where it creates the logs. It performs a system verification and outputs upgrade information.
5. Enter **y** to agree to the End User License Agreement (EULA).

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

6. The installer displays the list of RPMs that get removed, installed, and upgraded on the selected systems.
7. When you are prompted, reply **y** to stop appropriate processes.

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

   The installer stops processes, uninstalls RPMs, and installs RPMs.

   The upgrade is finished on the first subcluster. Do not reboot the nodes in the first subcluster until you complete the Preparing the second subcluster procedure.

Preparing the second subcluster

Perform the following steps on the second subcluster before rebooting nodes in the first subcluster.

To prepare to upgrade the second subcluster

1 Get the summary of the status of your resources.

```
# hastatus -summ
-- SYSTEM STATE
-- System               State                Frozen

A  node01               EXITED                1
A  node02               EXITED                1
A  node03               RUNNING               0
A  node04               RUNNING               0

-- GROUP STATE
-- Group           System  Probed     AutoDisabled    State

B  SG1             node01   Y          N               OFFLINE
B  SG1             node02   Y          N               OFFLINE
B  SG1             node03   Y          N               ONLINE
B  SG1             node04   Y          N               ONLINE
B  SG2             node01   Y          N               OFFLINE
B  SG2             node02   Y          N               OFFLINE
B  SG2             node03   Y          N               ONLINE
B  SG2             node04   Y          N               ONLINE
B  SG3             node01   Y          N               OFFLINE
B  SG3             node02   Y          N               OFFLINE
B  SG3             node03   Y          N               ONLINE
B  SG3             node04   Y          N               OFFLINE
B  SG4             node01   Y          N               OFFLINE
B  SG4             node02   Y          N               OFFLINE
B  SG4             node03   Y          N               OFFLINE
B  SG4             node04   Y          N               ONLINE
```

2 Unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -h
Filesystem            Size  Used Avail Use% Mounted on
/dev/sda1              26G  3.3G   22G  14% /
udev                 1007M  352K 1006M   1% /dev
tmpfs                 4.0K     0  4.0K   0% /dev/vx
/dev/vx/dsk/dg2/dg2vol1
                      3.0G   18M  2.8G   1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2
                      1.0G   18M  944M   2% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3
                       10G   20M  9.4G   1% /mnt/dg2/dg2vol3
# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

3 Make the configuration writable on the second subcluster.

```
# haconf -makerw
```

4 On the nodes in the second subcluster, use the following command to take all the cache area offline:

```
# sfcache offline cachename
```

5 Unfreeze the service groups.

```
# hagrp -unfreeze sg1 -persistent
# hagrp -unfreeze sg2 -persistent
# hagrp -unfreeze sg3 -persistent
# hagrp -unfreeze sg4 -persistent
```

6 Dump the configuration and make it read-only.

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

7 Take the service groups offline on node03 and node04.

```
# hagrp -offline sg1 -sys node03
# hagrp -offline sg1 -sys node04
# hagrp -offline sg2 -sys node03
# hagrp -offline sg2 -sys node04
# hagrp -offline sg3 -sys node03
# hagrp -offline sg4 -sys node04
```

8 Verify the state of the service groups.

```
# hagrp -state
#Group       Attribute    System     Value
SG1             State      node01    |OFFLINE|
SG1             State      node02    |OFFLINE|
SG1             State      node03    |OFFLINE|
SG1             State      node04    |OFFLINE|
SG2             State      node01    |OFFLINE|
SG2             State      node02    |OFFLINE|
SG2             State      node03    |OFFLINE|
SG2             State      node04    |OFFLINE|
SG3             State      node01    |OFFLINE|
SG3             State      node02    |OFFLINE|
SG3             State      node03    |OFFLINE|
SG3             State      node04    |OFFLINE|
```

9 Stop all VxVM volumes (for each disk group) that VCS does not manage.

10 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

```
# hastop -local
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

11 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 are not added.

```
# /etc/init.d/vxfen status
VXFEN module is not loaded

# /etc/init.d/gab status
GAB module is not loaded

# /etc/init.d/llt status
LLT module is not loaded
```

Activating the first subcluster

Get the first subcluster ready for the service groups.

To activate the first subcluster

1 Start LLT and GAB on one node in the first half of the cluster..

```
# /etc/init.d/llt start
# /etc/init.d/gab start
```

2 Seed node01 in the first subcluster.

```
# gabconfig -x
```

3 On the first half of the cluster, start SFHA:

```
# cd /opt/VRTS/install

# ./installsfha<version> -start sys1 sys2
```

Where <*version*> is the specific release version.

See “About the script-based installer” on page 69.

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.

```
# hagrp -online sg1 -sys node01
# hagrp -online sg1 -sys node02
# hagrp -online sg2 -sys node01
# hagrp -online sg2 -sys node02
# hagrp -online sg3 -sys node01
# hagrp -online sg4 -sys node02
```

Upgrading the operating system on the second subcluster

You can perform the operating system upgrade on the second subcluster, if required.

Before performing operating system upgrade, it is better to prevent LLT from starting automatically when the node starts. For example, you can do the following:

```
# mv /etc/llttab /etc/llttab.save
```

or you can change the `/etc/default/llt` file by setting LLT_START = **0**.

After you finish upgrading the OS, remember to change the LLT configuration to its original configuration.

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 installsfha.

   ```
   # cd storage_foundation_high_availability
   ```

3. Confirm that SFHA is stopped on node03 and node04. Start the installsfha program, specify the nodes in the second subcluster (node3 and node4).

   ```
   # ./installsfha -upgrade node3 node4
   ```

   The program starts with a copyright message and specifies the directory where it creates the logs.
4. Enter **y** to agree to the End User License Agreement (EULA).

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

5. The installer displays the list of RPMs that get removed, installed, and upgraded on the selected systems.
6. When you are prompted, reply **y** to stop appropriate processes.

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

   The installer stops processes, uninstalls RPMs, and installs RPMs.
7. Monitor the installer program answering questions as appropriate until the upgrade completes.

Finishing the phased upgrade

Complete the following procedure to complete the upgrade.

To finish the upgrade

1 Verify that the cluster UUID is the same on the nodes in the second subcluster and the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl
-clus -display node1 [node2 ...]
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first subcluster to the nodes in the second subcluster. For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus
-copy -from_sys node01 -to_sys node03 node04
```

2 On the second half of the cluster, start SFHA:

```
# cd /opt/VRTS/install

# ./installsfha<version> -start sys3 sys4
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 69.

3 Check to see if SFHA and High Availabiltiy and its components are up.

```
# gabconfig -a
GAB Port Memberships
===============================================================
Port a gen   nxxxnn membership 0123
Port b gen   nxxxnn membership 0123
Port h gen   nxxxnn 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  sg1      node01    Y        N              ONLINE
B  sg1      node02    Y        N              ONLINE
B  sg1      node03    Y        N              ONLINE
B  sg1      node04    Y        N              ONLINE
B  sg2      node01    Y        N              ONLINE
B  sg2      node02    Y        N              ONLINE
B  sg2      node03    Y        N              ONLINE
B  sg2      node04    Y        N              ONLINE
B  sg3      node01    Y        N              ONLINE
B  sg3      node02    Y        N              OFFLINE
B  sg3      node03    Y        N              OFFLINE
B  sg3      node04    Y        N              OFFLINE
B  sg4      node01    Y        N              OFFLINE
B  sg4      node02    Y        N              ONLINE
B  sg4      node03    Y        N              OFFLINE
B  sg4      node04    Y        N              OFFLINE
```

5 After the upgrade is complete, start the VxVM volumes (for each disk group) and mount the VxFS file systems.

In this example, you have performed a phased upgrade of SFHA. The service groups were down when you took them offline on node03 and node04, to the time SFHA brought them online on node01 or node02.

Note: If you want to upgrade application clusters that use CP server based fencing to 6.2, make sure that you first upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.2. However, note that the CP server upgraded to 6.2 can support application clusters on 6.2 (HTTPS-based communication) and application clusters prior to 6.2 (IPM-based communication). When you configure the CP server, the installer asks the VIPs for HTTPS-based communication (if the clients are on release version 6.2) or VIPs for IPM-based communication (if the clients are on a release version prior to 6.2).

For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.

Chapter 22

Performing an automated SFHA upgrade using response files

This chapter includes the following topics:

- Upgrading SFHA using response files
- Response file variables to upgrade Storage Foundation and High Availability
- Sample response file for SFHA upgrade
- Performing rolling upgrade of SFHA using response files
- Response file variables to upgrade SFHA using rolling upgrade
- Sample response file for SFHA using rolling upgrade

Upgrading SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA upgrade on one system to upgrade SFHA on other systems.

You can also create a response file using the `makeresponsefile` option of the installer.

```
# ./installer -makeresponsefile
```

To perform automated SFHA upgrade

1. Make sure the systems where you want to upgrade SFHA meet the upgrade requirements.
2. Make sure the pre-upgrade tasks are completed.
3. Copy the response file to the system where you want to upgrade SFHA.
4. Edit the values of the response file variables as necessary.
5. Mount the product disc and navigate to the folder that contains the installation program.
6. Start the upgrade from the system to which you copied the response file. For example:

   ```
   # ./installer -responsefile /tmp/response_file

   # ./installsfha -responsefile /tmp/response_file
   ```

   Where /tmp/*response_file* is the response file's full path name.

Response file variables to upgrade Storage Foundation and High Availability

Table 22-1 lists the response file variables that you can define to configure SFHA.

Table 22-1 Response file variables for upgrading SFHA

Variable	Description
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required

Table 22-1 Response file variables for upgrading SFHA *(continued)*

Variable	Description
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
CFG{opt}{upgrade}	Upgrades all RPMs installed. List or scalar: list Optional or required: optional
CFG{mirrordgname}{system}	If the root dg is encapsulated and you select split mirror is selected: Splits the target disk group name for a system. List or scalar: scalar Optional or required: optional
CFG{splitmirror}{system}	If the root dg is encapsulated and you select split mirror is selected: Indicates the system where you want a split mirror backup disk group created. List or scalar: scalar Optional or required: optional

Table 22-1 Response file variables for upgrading SFHA *(continued)*

Variable	Description
CFG{opt}{disable_dmp_native_support}	If it is set to 1, Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools is disabled after upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases RPM upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system. List or scalar: scalar Optional or required: optional
CFG{opt}{patch_path}	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed . List or scalar: scalar Optional or required: optional
CFG{opt}{patch2_path}	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{opt}{patch3_path}	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{opt}{patch4_path}	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional

Table 22-1 Response file variables for upgrading SFHA *(continued)*

Variable	Description
CFG{opt}{patch5_path}	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed. List or scalar: scalar Optional or required: optional
CFG{rootsecusrgrps}	Defines if the user chooses to grant read access to the cluster only for root and other users/usergroups which are granted explicit privileges on VCS objects. List or scalar: scalar Optional or required: optional
CFG{secusrgrps}	Defines the usergroup names that are granted read access to the cluster. List or scalar: scalar Optional or required: optional

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{client_vxfen_warning}=1;
$CFG{fencing_cps}=[ qw(10.198.92.157 10.198.92.158) ];
$CFG{fencing_cps_ports}{"10.198.92.157"}=443;
$CFG{fencing_cps_ports}{"10.198.92.158"}=443;
$CFG{fencing_cps_vips}{"10.198.92.157"}=[ qw(10.198.92.157) ];
$CFG{fencing_cps_vips}{"10.198.92.158"}=[ qw(10.198.92.158) ];
$CFG{opt}{noipc}=1;
$CFG{opt}{updatekeys}=1;
$CFG{opt}{upgrade}=1;
$CFG{systems}=[ qw(cdclab-p51a-03 cdclab-p51a-04) ];
$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.

Performing rolling upgrade of SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA upgrade on one system to upgrade SFHA on other systems.

You can also create a response file using the `makeresponsefile` option of the installer.

To perform automated SFHA rolling upgrade

1. Make sure the systems where you want to upgrade SFHA meet the upgrade requirements.
2. Make sure the pre-upgrade tasks are completed.
3. Copy the response file to the systems where you want to launch the installer.

   See "Sample response file for SFHA using rolling upgrade" on page 366.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to upgrade SFHA using rolling upgrade" on page 364.
5. Mount the product disc and navigate to the folder that contains the installation program.
6. Start the upgrade from the system to which you copied the response file. For example:

   ```
   # ./installer -responsefile /tmp/response_file

   # ./installsfha -responsefile /tmp/response_file
   ```

   Where /tmp/*response_file* is the response file's full path name.

Response file variables to upgrade SFHA using rolling upgrade

Table 22-2 lists the response file variables that you can define to upgrade SFHA using rolling upgrade.

Table 22-2 Response file variables for upgrading SFHA using rolling upgrade

Variable	Description
CFG{phase1}{0}	A series of $CFG{phase1}{N} items define sub-cluster division. The index N indicatse the order to do RU phase1. The index starts from 0. Each item has a list of node(at least 1). List or scalar: list Optional or required: conditional required Required if rolling upgrade phase1 needs to be performed.
CFG{rollingupgrade_phase2}	The `CFG{rollingupgrade_phase2}` option is used to perform rolling upgrade Phase 2. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version. List or scalar: scalar Optional or required: conditional required Required if rolling upgrade phase 2 needs to be performed.
CFG{rolling_upgrade}	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade Phase 1 or Phase 2 explicitly.
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{opt}{upgrade}	Upgrades all RPMs installed. List or scalar: scalar Optional or required: optional

Table 22-2 Response file variables for upgrading SFHA using rolling upgrade *(continued)*

Variable	Description
CFG{secusrgrps}	Defines the user groups which get read access to the cluster. List or scalar: list Optional or required: optional
CFG{rootsecusrgrps}	Defines the read access to the cluster from root users, specific users, or usergroups based on your choice. The selected users or usergroups get explicit privileges on VCS objects. 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

Sample response file for SFHA using rolling upgrade

The following example shows a response file for SFHA using Rolling Upgrade.

```
our %CFG;
$CFG{accepteula}=1;
$CFG{client_vxfen_warning}=1;
$CFG{fencing_cps}=[ qw(10.198.90.6) ];
$CFG{fencing_cps_ports}{"10.198.90.6"}=50006;
$CFG{fencing_cps_vips}{"10.198.90.6"}=[ qw(10.198.90.6) ];
$CFG{opt}{gco}=1;
$CFG{opt}{noipc}=1;
$CFG{opt}{rolling_upgrade}=1;
$CFG{opt}{rollingupgrade_phase2}=1;
$CFG{opt}{updatekeys}=1;
$CFG{opt}{upgrade}=1;
$CFG{secusrgrps}=qw{staff pilotaix218@cdc.veritas.com};
$CFG{opt}{vr}=1;
```

```
$CFG{phase1}{"0"}=[ qw(sys3 sys2) ];
$CFG{phase1}{"1"}=[ qw(sys1) ];
$CFG{systems}=[ qw(sys1 sys2 sys3) ];
$CFG{vcs_allowcomms}=1;
1;
```

Chapter 23

Performing post-upgrade tasks

This chapter includes the following topics:

- Optional configuration steps
- Re-joining the backup boot disk group into the current disk group
- Reverting to the backup boot disk group after an unsuccessful upgrade
- Post upgrade tasks for migrating the SFDB repository database
- Recovering VVR if automatic upgrade fails
- Post-upgrade tasks when VCS agents for VVR are configured
- Upgrading disk layout versions
- Upgrading VxVM disk group versions
- Updating variables
- Setting the default disk group
- About enabling LDAP authentication for clusters that run in secure mode
- Verifying the Storage Foundation and High Availability upgrade

Optional configuration steps

After the upgrade is complete, additional tasks may need to be performed.

You can perform the following optional configuration steps:

- If Volume Replicator (VVR) is configured, do the following steps in the order shown:
  - Reattach the RLINKs.
  - Associate the SRL.
- To encapsulate and mirror the boot disk, follow the procedures in the "Administering Disks" chapter of the *Symantec Storage Foundation Administrator's Guide*.
- To upgrade VxFS Disk Layout versions and VxVM Disk Group versions, follow the upgrade instructions.
  See "Upgrading VxVM disk group versions" on page 381.

Re-joining the backup boot disk group into the current disk group

Perform this procedure to rejoin the backup boot disk if you split the mirrored boot disk during upgrade. After a successful upgrade and reboot, you no longer need to keep the boot disk group backup.

To re-join the backup boot disk group

- Re-join the *backup_bootdg* disk group to the boot disk group.

  ```
  # /etc/vx/bin/vxrootadm -Y join backup_bootdg
  ```

  where the `-Y` option indicates a silent operation, and *backup_bootdg* is the name of the backup boot disk group that you created during the upgrade.

Reverting to the backup boot disk group after an unsuccessful upgrade

Perform this procedure if your upgrade was unsuccessful and you split the mirrored boot disk to back it up during upgrade. You can revert to the backup that you created when you upgraded.

To revert the backup boot disk group after an unsuccessful upgrade

1. To determine the boot disk groups, look for the *rootvol* volume in the output of the `vxprint` command.

   ```
   # vxprint
   ```

2. Use the `vxdg` command to find the boot disk group where you are currently booted.

   ```
   # vxdg bootdg
   ```

3. Boot the operating system from the backup boot disk group.
4. Join the original boot disk group to the backup disk group.

   ```
   # /etc/vx/bin/vxrootadm -Y join original_bootdg
   ```

   where the `-Y` option indicates a silent operation, and *original_bootdg* is the boot disk group that you no longer need.

Post upgrade tasks for migrating the SFDB repository database

Database Storage Checkpoints that have been created by using the SFDB tools before upgrade are visible using the `vxsfadm` CLI, and you can mount these Database Storage Checkpoints and roll back to them, if required. However, creating clones by using migrated Database Storage Checkpoints is not supported.

If you want to continue using previously created FlashSnap snapplans to take snapshots, you must validate them by using the `-o validate` option of the `vxsfadm` command.

- Rename startup script after upgrading from 5.0x and before migrating the SFDB repository
  See "Migrating SFDB from 5.0x to 6.2" on page 376.
- Migrate from a 5.0x SFDB repository database to 6.2
  See "Migrating from a 5.0 repository database to 6.2" on page 370.
- Migrate from a 5.1 or 5.1SP1 repository database to 6.2
  See "Migrating from a 5.1 or higher repository database to 6.2" on page 373.

Migrating from a 5.0 repository database to 6.2

Perform the following on one node only.

To migrate from a 5.0 repository database to 6.2

1. Rename the startup script NO_S*vxdbms3 to S*vxdbms3.

   See "Migrating SFDB from 5.0x to 6.2" on page 376.

2. As root, dump out the old Sybase ASA repository. If you are using SFHA or SF Oracle RAC, you only need to do this on one node.

   ```
   # /opt/VRTSdbed/migrate/sfua_rept_migrate
   ```

3. On the same node that you ran `sfua_rept_migrate` run the following command as Oracle user. For each Oracle instance, migrate the old repository data to the SQLite repository.

   ```
   $ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME -G \
   Oracle_service_group
   ```

4. By default, the repository is created on the file system which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:

   - Repository path has to be a directory writable by Oracle user.
   - If you are using SFHA, the repository must be accessible by all nodes. You can put it in a resource group under VCS control so it can be failed over together with the Oracle database.
   - The update commands will not be able to verify accessibility of the repository path and will fail if you have not set up the path correctly.

   Create an alternate repository path.

   ```
   $ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME \
   -G Oracle_service_group -R Alternate_path
   ```

5. If you are using Database Flashsnap for off-host processing, and if you have a repository on the secondary host that you need to migrate: perform the previous steps on the secondary host.

6 On the primary host, edit your snapplans to remove the "SNAPSHOT_DG=SNAP_*" parameter and add “SNAPSHOT_DG_PREFIX=SNAP_*". The parameter can be any PREFIX value and not necessarily "SNAP_*".

For example:

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=system1_data
SNAPSHOT_DG=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=judge_data
SNAPSHOT_DG_PREFIX=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

7 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/vxsfadm -s flashsnap \
-a oracle -c SNAPPLAN -o validate
```

This completes the migration of the repository for Database Storage Checkpoints and Database Tiered Storage parameters.

Note: While you revalidate the snapshot configuration file (`snapplan`) from an older release, use the `vxsfadm -c <configfile>` option to avoid the default values from overriding the old values.

To begin using the Storage Foundation for Databases (SFDB) tools:

see *Storage Foundation: Storage and Availability Management for Oracle Databases*.

Migrating from a 5.1 or higher repository database to 6.2

Perform the following on one node only.

To migrate from a 5.1 or higher repository database to 6.2

1 Run the following command as Oracle user. For each Oracle instance, migrate the old repository data to the SQLite repository.

```
$ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME -G \
Oracle_service_group
```

2 By default, the repository is created on the file system which contains the Oracle SYSTEM tablespace. If you need an alternative repository path, first verify the following requirements:

- Repository path has to be a directory writable by Oracle user.
- If you are using SFHA, the repository must be accessible by all nodes. You can put it in a resource group under VCS control so it can be failed over together with the Oracle database.
- The update commands will not be able to verify accessibility of the repository path and will fail if you have not set up the path correctly.

Create an alternate repository path.

```
$ /opt/VRTS/bin/dbed_update -S $ORACLE_SID -H $ORACLE_HOME \
-G Oracle_service_group -R Alternate_path
```

3 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.

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=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=system1_data
SNAPSHOT_DG=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

```
$ /usr/oracle> more SNAPPLAN1
SNAPSHOT_VERSION=5.0
PRIMARY_HOST=system1
SECONDARY_HOST=system1.example.com
PRIMARY_DG=judge_data
SNAPSHOT_DG_PREFIX=SNAP_system1_data
ORACLE_SID=HN1
ARCHIVELOG_DEST=/oracle/orahome/dbs/arch
SNAPSHOT_ARCHIVE_LOG=yes
SNAPSHOT_MODE=online
SNAPSHOT_PLAN_FOR=database
SNAPSHOT_PLEX_TAG=dbed_flashsnap
SNAPSHOT_VOL_PREFIX=SNAP_
ALLOW_REVERSE_RESYNC=no
SNAPSHOT_MIRROR=1
```

5 On the primary host, revalidate your snapshots using the following command:

```
$ /opt/VRTS/bin/vxsfadm -s flashsnap \
-a oracle -c SNAPPLAN -o validate
```

This completes the migration of the repository for Database Storage Checkpoints and Database Tiered Storage parameters.

Note: While you revalidate the snapshot configuration file (`snapplan`) from an older release, use the `vxsfadm -c <configfile>` option to avoid the default values from overriding the old values.

To begin using the Storage Foundation for Databases (SFDB) tools:

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

Migrating SFDB from 5.0x to 6.2

When upgrading from SFHA version 5.0 to SFHA 6.2 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_migrate`. Thus when `sfua_rept_migrate` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

```
/etc/rc.d/rc2.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File:  is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

To prevent S*vxdbms3 startup script error

- Rename the startup script NO_S*vxdbms3 to S*vxdbms3.

Recovering VVR if automatic upgrade fails

If the upgrade fails during the configuration phase, after displaying the VVR upgrade directory, the configuration needs to be restored before the next attempt. Run the scripts in the upgrade directory in the following order to restore the configuration:

```
# restoresrl
# adddcm
# srlprot
# attrlink
# start.rvg
```

After the configuration is restored, the current step can be retried.

Post-upgrade tasks when VCS agents for VVR are configured

The following lists post-upgrade tasks with VCS agents for VVR:

- Unfreezing the service groups
- Restoring the original configuration when VCS agents are configured

Unfreezing the service groups

This section describes how to unfreeze services groups and bring them online.

To unfreeze the service groups

1. On any node in the cluster, make the VCS configuration writable:

   ```
   # haconf -makerw
   ```

2. Edit the `/etc/VRTSvcs/conf/config/main.cf` file to remove the deprecated attributes, SRL and RLinks, in the RVG and RVGShared resources.
3. Verify the syntax of the main.cf file, using the following command:

   ```
   # hacf -verify
   ```

4. Unfreeze all service groups that you froze previously. Enter the following command on any node in the cluster:

   ```
   # hagrp -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:

   ```
   # hagrp -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 316.

Type the following command on any node in the cluster.

```
# hares -online ip_name -sys system
```

This IP is the virtual IP that is used for replication within the cluster.

8 In shared disk group environment, online the virtual IP resource on the master node.

Restoring the original configuration when VCS agents are configured

This section describes how to restore a configuration with VCS configured agents.

Note: Restore the original configuration only after you have upgraded VVR on all nodes for the Primary and Secondary cluster.

To restore the original configuration

1 Import all the disk groups in your VVR configuration.

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

Each disk group should be imported onto the same node on which it was online when the upgrade was performed. The reboot after the upgrade could result in another node being online; for example, because of the order of the nodes in the AutoStartList. In this case, switch the VCS group containing the disk groups to the node on which the disk group was online while preparing for the upgrade.

```
# hagrp -switch grpname -to system
```

2 Recover all the disk groups by typing the following command on the node on which the disk group was imported in step 1.

```
# vxrecover -bs
```

3 Upgrade all the disk groups on all the nodes on which VVR has been upgraded:

```
# vxdg upgrade diskgroup
```

4 On all nodes that are Secondary hosts of VVR, make sure the data volumes on the Secondary are the same length as the corresponding ones on the Primary. To shrink volumes that are longer on the Secondary than the Primary, use the following command on each volume on the Secondary:

```
# vxassist -g diskgroup  shrinkto volume_name volume_length
```

where *volume_length* is the length of the volume on the Primary.

Note: Do not continue until you complete this step on all the nodes in the Primary and Secondary clusters on which VVR is upgraded.

---

5 Restore the configuration according to the method you used for upgrade:

If you upgraded with the VVR upgrade scripts

Complete the upgrade by running the `vvr_upgrade_finish` script on all the nodes on which VVR was upgraded. We recommend that you first run the `vvr_upgrade_finish` script on each node that is a Secondary host of VVR.

Perform the following tasks in the order indicated:

- To run the `vvr_upgrade_finish` script, type the following command:

  ```
  # /disc_path/scripts/vvr_upgrade_finish
  ```

  where *disc_path* is the location where the Veritas software disc is mounted.

- Attach the RLINKs on the nodes on which the messages were displayed:

  ```
  # vxrlink -g diskgroup -f att rlink_name
  ```

If you upgraded with the product installer

Use the Veritas product installer and select Start an Installed Product. Or use the installation script with the `-start` option.

6 Bring online the RVGLogowner group on the master:

```
# hagrp -online RVGLogownerGrp -sys masterhost
```

7 If you plan on using IPv6, you must bring up IPv6 addresses for virtual replication IP on primary/secondary nodes and switch from using IPv4 to IPv6 host names or addresses, enter:

```
# vradmin changeip newpri=v6 newsec=v6
```

where *v6* is the IPv6 address.

8 Restart the applications that were stopped.

Upgrading disk layout versions

In this release, you can create and mount only file systems with disk layout Version 7, 8, 9, and 10. You can only local mount disk layout Version 6 only to upgrade to a later disk layout version.

Note: If you plan to use 64-bit quotas, you must upgrade to the latest disk layout Version 10. The use of 64-bit quota on earlier disk layout versions is deprecated in this release.

Disk layout Version 6 has been deprecated and you cannot cluster mount an existing file system that has disk layout Version 6. To upgrade a cluster file system with disk layout Version 6, you must local mount the file system and then upgrade the file system using the `vxupgrade` utility to a later version.

To upgrade the disk layout versions

- To get to disk layout Version 10 from Version 6. You must incrementally upgrade the disk layout of this file system. For example:

```
# vxupgrade -n 7 /mnt
# vxupgrade -n 8 /mnt
# vxupgrade -n 9 /mnt
# vxupgrade -n 10 /mnt
```

See the `vxupgrade`(1M) manual page.

Support for disk layout Version 4 has been removed. You must upgrade any existing file systems with disk layout Version 4 to disk layout Version 7 or later using the `vxfsconvert` command.

See the `vxfsconvert`(1M) manual page.

Note: Symantec recommends that you upgrade existing file systems to the highest supported disk layout version prior to upgrading to this release. Once a disk layout version has been upgraded, it is not possible to downgrade to the previous version.

You can check which disk layout version your file system has by using the following command:

```
# fstyp -v /dev/vx/dsk/dg1/vol1 | grep -i version
```

For more information about disk layout versions, see the *Symantec Storage Foundation Administrator's Guide*.

Upgrading VxVM disk group versions

All Veritas Volume Manager disk groups have an associated version number. Each VxVM release supports a specific set of disk group versions. VxVM can import and perform tasks on disk groups with those versions. Some new features and tasks work only on disk groups with the current disk group version. Before you can perform the tasks or use the features, upgrade the existing disk groups.

For 6.2, the Veritas Volume Manager disk group version is different than in previous VxVM releases. Symantec recommends that you upgrade the disk group version if you upgraded from a previous VxVM release.

After upgrading to SFHA 6.2, you must upgrade any existing disk groups that are organized by ISP. Without the version upgrade, configuration query operations continue to work fine. However, configuration change operations will not function correctly.

For more information about ISP disk groups, refer to the *Symantec Storage Foundation Administrator's Guide*.

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 *Symantec Storage Foundation Administrator's Guide*.

Updating variables

In `/etc/profile`, update the `PATH` and `MANPATH` variables as needed.

`MANPATH` can include `/opt/VRTS/man` and `PATH` can include `/opt/VRTS/bin`.

Setting the default disk group

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 *Symantec Storage Foundation Administrator's Guide*.

About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as 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.

If you have not already added VCS users during installation, you can add the users later.

See the *Symantec Cluster Server Administrator's Guide* for instructions to add VCS users.

Figure 23-1 depicts the SFHA cluster communication with the LDAP servers when clusters run in secure mode.

Figure 23-1 Client communication with LDAP servers

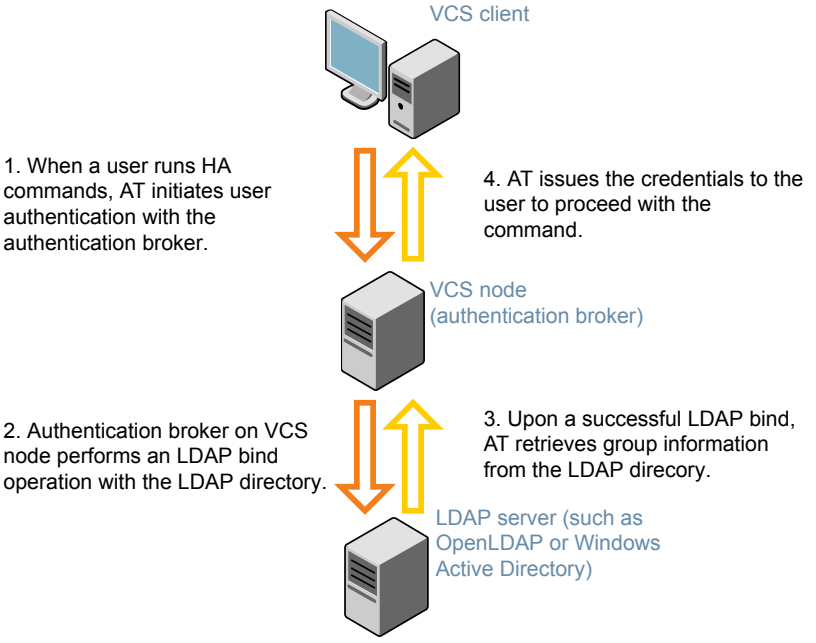

The LDAP schema and syntax for LDAP commands (such as, ldapadd, ldapmodify, and ldapsearch) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
  - UserObjectClass (the default is posixAccount)
  - UserObject Attribute (the default is uid)
  - User Group Attribute (the default is gidNumber)
  - Group Object Class (the default is posixGroup)
  - GroupObject Attribute (the default is cn)
  - Group GID Attribute (the default is gidNumber)
  - Group Membership Attribute (the default is memberUid)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)

- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 6.1.6.0 or later.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion
vssat version: 6.1.12.8
```

To enable OpenLDAP authentication for clusters that run in secure mode

1. Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-d -s domain_controller_name_or_ipaddress -u domain_user

Attribute list file name not provided, using AttributeList.txt

Attribute file created.
```

   You can use the `catatldapconf` command to view the entries in the attributes file.

2. Run the LDAP configuration tool using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-c -d LDAP_domain_name

Attribute list file not provided, using default AttributeList.txt

CLI file name not provided, using default CLI.txt

CLI for addldapdomain generated.
```

3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -x

Using default broker port 14149

CLI file not provided, using default CLI.txt

Looking for AT installation...

AT found installed at ./vssat

Successfully added LDAP domain.
```

4 Check the AT version and list the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion

vssat version: 6.1.12.8

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat listldapdomains

Domain Name : mydomain.com

Server URL : ldap://192.168.20.32:389

SSL Enabled : No

User Base DN : CN=people,DC=mydomain,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 Check the other domains in the cluster.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showdomains -p vx
```

The command output lists the number of domains that are found, with the domain names and domain types.

6 Generate credentials for the user.

```
# unset EAT_LOG

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat authenticate \
-d ldap:LDAP_domain_name -p user_name -s user_password -b \
localhost:14149
```

7 Add non-root users as applicable.

```
# useradd user1

# passwd pw1

Changing password for "user1"

user1's New password:

Re-enter user1's new password:

# su user1

# bash

# id

uid=204(user1) gid=1(staff)

# pwd

# mkdir /home/user1

# chown user1 /home/ user1
```

8 Add the non-root user to the VCS configuration.

```
# haconf -makerw
# hauser -add user1
# haconf -dump -makero
```

9 Log in as non-root user and run VCS commands as LDAP user.

```
# cd /home/user1

# ls

# cat .vcspwd

101 localhost mpise LDAP_SERVER ldap

# unset VCS_DOMAINTYPE

# unset VCS_DOMAIN

# /opt/VRTSvcs/bin/hasys -state

   #System        Attribute          Value

cluster1:sysA     SysState           FAULTED

cluster1:sysB     SysState           FAULTED

cluster2:sysC     SysState           RUNNING

cluster2:sysD     SysState           RUNNING
```

Verifying the Storage Foundation and High Availability upgrade

Refer to the section about verifying the installation to verify the upgrade.

See “Verifying that the products were installed” on page 394.

Section 9

Post-installation tasks

- Chapter 24. Performing post-installation tasks
- Chapter 25. Verifying the SFHA installation

Chapter 24

Performing post-installation tasks

This chapter includes the following topics:

- Switching on Quotas
- About configuring authentication for SFDB tools

Switching on Quotas

This turns on the group and user quotas once all the nodes are upgraded to 6.2, if it was turned off earlier.

To turn on the group and user quotas

- Switch on quotas:

```
# vxquotaon -av
```

About configuring authentication for SFDB tools

To configure authentication for Storage Foundation for Databases (SFDB) tools, perform the following tasks:

Configure the `vxdbd` daemon to require authentication	See "Configuring vxdbd for SFDB tools authentication" on page 391.
Add a node to a cluster that is using authentication for SFDB tools	See "Adding nodes to a cluster that is using authentication for SFDB tools" on page 436.

Configuring vxdbd for SFDB tools authentication

To configure vxdbd, perform the following steps as the root user

1. Run the `sfae_auth_op` command to set up the authentication services.

```
# /opt/VRTS/bin/sfae_auth_op -o setup
Setting up AT
Starting SFAE AT broker
Creating SFAE private domain
Backing up AT configuration
Creating principal for vxdbd
```

2. Stop the `vxdbd` daemon.

```
# /opt/VRTS/bin/sfae_config disable
vxdbd has been disabled and the daemon has been stopped.
```

3. Enable authentication by setting the `AUTHENTICATION` key to `yes` in the `/etc/vx/vxdbed/admin.properties` configuration file.

   If `/etc/vx/vxdbed/admin.properties` does not exist, then use`cp /opt/VRTSdbed/bin/admin.properties.example /etc/vx/vxdbed/admin.properties`.

4. Start the `vxdbd` daemon.

```
# /opt/VRTS/bin/sfae_config enable
vxdbd has been enabled and the daemon has been started.
It will start automatically on reboot.
```

The `vxdbd` daemon is now configured to require authentication.

Chapter 25

Verifying the SFHA installation

This chapter includes the following topics:

- Upgrading the disk group version
- Performing a postcheck on a node
- Verifying that the products were installed
- Installation log files
- Starting and stopping processes for the Symantec products
- Checking Veritas Volume Manager processes
- Verifying the LLT, GAB, and VCS configuration files
- Verifying LLT, GAB, and cluster operation

Upgrading the disk group version

After you upgrade from previous versions to 6.2, you have to upgrade the disk group version manually.

To upgrade disk group version, you have to first upgrade the cluster protocol version using the `vxdctl upgrade` command.

```
# vxdctl list
Volboot file
version: 3/1
seqno:   0.1
cluster protocol version: 120
```

```
hostid:  sys1
hostguid:  {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
#
# vxdctl upgrade
#

# vxdctl list

Volboot file
version: 3/1
seqno:   0.2
cluster protocol version: 140
hostid:  sys1
hostguid:  {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
```

Verify if the cluster protocol version shows 140 and disk group version is upgraded to 200.

```
# vxdctl list |grep version

version:   140
#
# vxdg upgrade dg_name
#
# vxdg list dg_name |grep version

version:   200
```

Performing a postcheck on a node

The installer's `postcheck` command can help you to determine installation-related problems and provide troubleshooting information.

See "About using the postcheck option" on page 457.

To run the postcheck command on a node

1. Run the installer with the `-postcheck` option.

   ```
   # ./installer -postcheck system_name
   ```

2. Review the output for installation-related information.

Verifying that the products were installed

Verify that the SFHA products are installed.

Use the command to check which RPMs have been installed.

```
# rpm -qa | grep VRTS
```

You can verify the version of the installed product. Use the following command:

```
# /opt/VRTS/install/installsfha<version> -version
```

Where `<version>` is the specific release version.

You can find out the about the installed RPMs and its versions by using the following command:

```
# /opt/VRTS/install/showversion
```

See "About the script-based installer" on page 69.

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 Symantec Support.

Using the summary file

The summary file contains the results of the installation by the installer or product installation scripts. The summary includes the list of the RPMs, and the status (success or failure) of each RPM. The summary also indicates which processes

were stopped or restarted during the installation. After installation, refer to the summary file to determine whether any processes need to be started.

Starting and stopping processes for the Symantec products

After the installation and configuration is complete, the Symantec product installer starts the processes that the installed products use. You can use the product installer to stop or start the processes, if required.

To stop the processes

- Use the `-stop` option to stop the product installation script.

  For example, to stop the product's processes, enter the following command:

  ```
  # ./installer -stop
  ```

  or

  ```
  # /opt/VRTS/install/installsfha<version> -stop
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 69.

To start the processes

- Use the `-start` option to start the product installation script.

  For example, to start the product's processes, enter the following command:

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

  or

  ```
  # /opt/VRTS/install/installsfha<version> -start
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 69.

Checking Veritas Volume Manager processes

Use the following procedure to verify that Volume Manager processes are running.

To confirm that key Volume Manager processes are running

- Type the following command:

  ```
  # ps -ef | grep vx
  ```

  Entries for the `vxiod`, `vxconfigd`, `vxnotify`, `vxesd`, `vxrelocd`, `vxcached` and `vxconfigbackupd` processes should appear in the output from this command. If you disable hot-relocation, the `vxrelocd` and `vxnotify` processes are not displayed.

Verifying the LLT, GAB, and VCS configuration files

Make sure that the LLT, GAB, and VCS configuration files contain the information you provided during VCS installation and configuration.

To verify the LLT, GAB, and VCS configuration files

1. Navigate to the location of the configuration files:
   - LLT
     /etc/llthosts
     /etc/lltab
   - 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 473.

   See "About the VCS configuration files" on page 477.

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 397.

4 Verify GAB operation.

5 Verify the cluster operation.

See "Verifying the cluster" on page 399.

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 sys1.

2 Run the `lltstat` command on the node sys1 to view the status of LLT.

```
lltstat -n
```

The output on sys1 resembles:

```
LLT node information:
    Node             State            Links
    *0 sys1          OPEN            2
     1 sys2          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 sys1             OPEN       2
     1 sys2             OPEN       2
     2 sys5             OPEN       1
```

3 Log in as superuser on the node sys2.

4 Run the `lltstat` command on the node sys2 to view the status of LLT.

```
lltstat -n
```

The output on sys2 resembles:

```
LLT node information:
    Node             State           Links
     0 sys1          OPEN            2
    *1 sys2          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 sys1 in a two-node cluster:

```
lltstat -nvv active
```

The output on sys1 resembles:

```
Node           State      Link    Status      Address
*0 sys1        OPEN
                          eth1 UP     08:00:20:93:0E:34
                          eth2 UP     08:00:20:93:0E:38
1 sys2         OPEN
                          eth1 UP     08:00:20:8F:D1:F2
                          eth2 DOWN
```

The command reports the status on the two active nodes in the cluster, sys1 and sys2.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- An address for each link

However, the output in the example shows different details for the node sys2. 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 sys1 in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
  Port  Usage       Cookie
  0     gab          0x0
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
        connects:    0 1
  7     gab          0x7
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
        connects:    0 1
  31    gab          0x1F
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
        connects:    0 1
```

Verifying the cluster

Verify the status of the cluster using the `hastatus` command. This command returns the system state and the group state.

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

Refer to the *Symantec 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  sys1                 RUNNING               0
A  sys2                 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 *Symantec 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 in the following procedure is for SPARC and it shows the output when the command is run on the node sys1. The list continues with similar information for sys2 (not shown) and any other nodes in the cluster.

  ```
  #System    Attribute             Value

  sys1       AgentsStopped         0

  sys1       AvailableCapacity     100

  sys1       CPUThresholdLevel     Critical 90 Warning 80 Note 70
                                   Info 60

  sys1       CPUUsage              0

  sys1       CPUUsageMonitoring    Enabled 0 ActionThreshold 0
                                   ActionTimeLimit 0 Action NONE
                                   NotifyThreshold 0 NotifyTimeLimit 0


  sys1       Capacity              100

  sys1       ConfigBlockCount      293

  sys1       ConfigCheckSum        37283

  sys1       ConfigDiskState       CURRENT

  sys1       ConfigFile            /etc/VRTSvcs/conf/config

  sys1       ConfigInfoCnt         0

  sys1       ConfigModDate         Mon Sep 03 07:14:23 CDT 2012

  sys1       ConnectorState        Up
  ```

```
sys1       CurrentLimits
sys1       DiskHbStatus
sys1       DynamicLoad             0
sys1       EngineRestarted         0
sys1       EngineVersion           6.2.00.0
sys1       FencingWeight           0
sys1       Frozen                  0
sys1       GUIIPAddr
sys1       HostUtilization         CPU 0 Swap 0
sys1       LLTNodeId               0
sys1       LicenseType             PERMANENT_SITE
sys1       Limits
sys1       LinkHbStatus            eth1 UP eth2 UP
sys1       LoadTimeCounter         0
sys1       LoadTimeThreshold       600
sys1       LoadWarningLevel        80
sys1       NoAutoDisable           0
sys1       NodeId                  0
sys1       OnGrpCnt                7
sys1       PhysicalServer
sys1       ShutdownTimeout         600
sys1       SourceFile              ./main.cf
sys1       SwapThresholdLevel      Critical 90 Warning 80 Note 70
                                   Info 60
sys1       SysInfo                 Linux:sys1,#1 SMP Fri Jul 8
                                   17:36:59 EDT
                                   2011,2.6.18-274.el5,x86_64
sys1       SysName                 sys1
```

```
sys1        SysState                RUNNING

sys1        SystemLocation

sys1        SystemOwner

sys1        SystemRecipients

sys1        TFrozen                 0

sys1        TRSE                    0

sys1        UpDownState             Up

sys1        UserInt                 0

sys1        UserStr

sys1        VCSFeatures             DR

sys1        VCSMode                 VCS
```

Section 10

Uninstallation of SFHA

- Chapter 26. Uninstalling Storage Foundation and High Availability
- Chapter 27. Uninstalling SFHA using response files

Chapter 26

Uninstalling Storage Foundation and High Availability

This chapter includes the following topics:

- Removing VxFS file systems
- Removing rootability
- Moving volumes to disk partitions
- Disabling VCS agents for VVR the agents on a system
- Removing the Replicated Data Set
- Uninstalling SFHA RPMs using the script-based installer
- Uninstalling SFHA with the web-based installer
- Removing license files (Optional)
- Removing the CP server configuration using the installer program
- Removing the Storage Foundation for Databases (SFDB) repository

Removing VxFS file systems

The VxFS RPM cannot be removed if there are any mounted VxFS file systems. Unmount all VxFS file systems before removing the RPM. After you remove the VxFS RPM, VxFS file systems are not mountable or accessible until another VxFS RPM is installed. It is advisable to back up VxFS file systems before installing a

new VxFS RPM. If VxFS will not be installed again, all VxFS file systems must be converted to a new file system type.

To remove VxFS file systems

1. Check if any VxFS file systems or Storage Checkpoints are mounted:

   ```
   # df -T | grep vxfs
   ```

2. Make backups of all data on the file systems that you wish to preserve, or recreate them as non-VxFS file systems on non-VxVM volumes or partitions.
3. Unmount all Storage Checkpoints and file systems:

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

4. Comment out or remove any VxFS file system entries from the `/etc/fstab` file.

Removing rootability

Perform this procedure if you configured rootability by encapsulating the root disk.

To remove rootability

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

   ```
   # df -v /
   ```

   The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

2. Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

   For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

   ```
   # vxplex -o rm dis mirrootvol-01 mirswapvol-01
   ```

   Warning: Do not remove the plexes that correspond to the original root disk partitions.

3. Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices:

   ```
   # /etc/vx/bin/vxunroot
   ```

   Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

Moving volumes to disk partitions

All volumes must be moved to disk partitions.

This can be done using one of the following procedures:

- Back up the system fully onto tape and then recover from it.
- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Use VxVM to move volumes incrementally onto disk partitions as described in the following section.

Moving volumes onto disk partitions using VxVM

Use the following procedure to move volumes onto disk partitions.

To move volumes onto disk partitions

1. Evacuate disks using the `vxdiskadm` program or the `vxevac` script. You should consider the amount of target disk space required for this before you begin.

   Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.

2. Remove the evacuated disks from VxVM control using the following commands:

   ```
   # vxdg -g diskgroup rmdisk disk _media_name
   # vxdisk rm disk_access_name
   ```

3. Decide which volume to move first. If the volume to be moved is mounted, unmount it.

4. If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume is synced.

5. Create a partition on free disk space of the same size as the volume. If there is not enough free space for the partition, a new disk must be added to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.

6. `Copy` the data on the volume onto the newly created disk partition using a command similar to the following:

   ```
   # dd if=/dev/vx/dsk/diskgroup/volume-name of=/dev/sdb2
   ```

   where `sdb` is the disk outside of VxVM and `2` is the newly created partition on that disk.

7. Replace the entry for that volume (if present) in `/etc/fstab` with an entry for the newly created partition.

8. Mount the disk partition if the corresponding volume was previously mounted.

9. Stop the volume and remove it from VxVM using the following commands:

   ```
   # vxvol -g diskgroup -f stop volume_name
   # vxedit -g diskgroup -rf rm volume_name
   ```

10. Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

    ```
    # vxprint -F "%sdnum" disk_media_name
    ```

11 If the output is not 0, there are still some subdisks on this disk that must be subsequently removed. If the output is 0, remove the disk from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name
# vxdisk rm disk_access_name
```

12 The free space now created can be used for adding the data in the next volume to be removed.

13 After all volumes have been converted into disk partitions successfully, reboot the system. After the reboot, none of the volumes should be open. To verify that none of the volumes are open, use the following command:

```
# vxprint -Aht -e v_open
```

14 If any volumes remain open, repeat the steps listed above.

Disabling VCS agents for VVR the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

To disable the agents

1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
# hagrp -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:

```
# hagrp -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 *Symantec Cluster Server Administrator's Guide*.

Removing the Replicated Data Set

If you use VVR, you need to perform the following steps. This section gives the steps to remove a Replicated Data Set (RDS) when the application is either active or stopped.

To remove the Replicated Data Set

1 Verify that all RLINKs are up-to-date:

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

If the Secondary is not required to be up-to-date, proceed to 2 and stop replication using the `-f` option with the `vradmin stoprep` command.

2 Stop replication to the Secondary by issuing the following command on any host in the RDS:

The `vradmin stoprep` command fails if the Primary and Secondary RLINKs are not up-to-date. Use the `-f` option to stop replication to a Secondary even when the RLINKs are not up-to-date.

```
# vradmin -g diskgroup stoprep local_rvgname sec_hostname
```

The argument `local_rvgname` is the name of the RVG on the local host and represents its RDS.

The argument `sec_hostname` is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

3 Remove the Secondary from the RDS by issuing the following command on any host in the RDS:

```
# vradmin -g diskgroup delsec local_rvgname sec_hostname
```

The argument local_rvgname is the name of the RVG on the local host and represents its RDS.

The argument sec_hostname is the name of the Secondary host as displayed in the output of the `vradmin printrvg` command.

4 Remove the Primary from the RDS by issuing the following command on the Primary:

```
# vradmin -g diskgroup delpri local_rvgname
```

When used with the `-f` option, the `vradmin delpri` command removes the Primary even when the application is running on the Primary.

The RDS is removed.

5 If you want to delete the SRLs from the Primary and Secondary hosts in the RDS, issue the following command on the Primary and all Secondaries:

```
# vxedit -r -g diskgroup rm srl_name
```

Uninstalling SFHA RPMs using the script-based installer

Use the following procedure to remove SFHA products.

Not all RPMs may be installed on your system depending on the choices that you made when you installed the software.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in SFHA 6.2 with a previous version of SFHA.

To shut down and remove the installed SFHA RPMs

1. Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/fstab`. Failing to remove these entries could result in system boot problems later.
2. Unmount all mount points for VxFS file systems.

   ```
   # umount /mount_point
   ```

3. If the VxVM RPM (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.

   See "Removing rootability" on page 405.

4. If a cache area is online, you must take the cache area offline before uninstalling the VxVM RPM. Use the following command to take the cache area offline:

   ```
   # sfcache offline cachename
   ```

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
   ```

   ```
   # ./uninstallsfha<version>
   ```

   Where `<version>` is the specific release version.

   Or, if you are using rsh, use the following:

   ```
   # ./uninstallsfha<version> -rsh
   ```

   See "About the script-based installer" on page 69.

8 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall SFHA, for example, `sys1`:

```
Enter the system names separated by spaces: [q?] sys1 sys2
```

9 The uninstall script prompts you to stop the product processes. If you respond yes, the processes are stopped and the RPMs are uninstalled.

The uninstall script creates log files and displays the location of the log files.

10 Most RPMs have kernel components. In order to ensure complete removal, a system reboot is recommended after all RPMs have been removed.

11 In case the uninstallation fails to remove any of the VRTS RPMs, check the installer logs for the reason for failure or try to remove the RPMs manually using the `pkgrm` command. For example:

```
pkgrm VRTSvxvm
```

Uninstalling SFHA with the web-based installer

This section describes how to uninstall using the web-based installer.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in SFHA 6.2 with a previous version of SFHA.

To uninstall SFHA

1 Perform the required steps to save any data that you want to preserve. For example, take backups of configuration files.

2 Start the web-based installer.

See "Starting the web-based installer" on page 170.

3 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.

4 Select **Storage Foundation High Availability** from the Product drop-down list, and click **Next**.

5 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Next**.

6 After the validation completes successfully, click **Next** to uninstall SFHA on the selected system.

7 If there are any processes running on the target system, the installer stops the processes. Click **Next**.

8 After the installer stops the processes, the installer removes the products from the specified system.

Click **Next**.

9 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.

10 Click **Finish**.

Most RPMs have kernel components. To ensure their complete removal, a system restart is recommended after all the RPMs have been removed.

Removing license files (Optional)

Optionally, you can remove the license files.

To remove the Veritas license files

1 To see what license key files you have installed on a system, enter:

```
# /sbin/vxlicrep
```

The output lists the license keys and information about their respective products.

2 Go to the directory containing the license key files and list them:

```
# cd /etc/vx/licenses/lic
# ls -a
```

3 Using the output from step 1, identify and delete the unwanted key files that are listed in step 2. Unwanted keys may be deleted by removing the license key file.

Removing the CP server configuration using the installer program

This section describes how to remove the CP server configuration from a node or a cluster that hosts the CP server.

Warning: Ensure that no SFHA cluster (application cluster) uses the CP server that you want to unconfigure. Run the `# cpsadm -s CPS_VIP -p CPS_Port -a list_nodes` to know if any application cluster is using the CP server.

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@cps1.symantecexample.com
# /opt/VRTS/install/installvcs<version>  -configcps
```

2 Select option 3 from the menu to unconfigure the CP server.

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

[2] Configure Coordination Point Server on SFHA cluster

[3] Unconfigure Coordination Point Server
```

3 Review the warning message and confirm that you want to unconfigure the CP server.

```
Unconfiguring coordination point server stops the vxcpserv process.
VCS clusters using this server for coordination purpose will have
one less coordination point.
Are you sure you want to take the CP server offline? [y,n,q] (n) y
```

4 Review the screen output as the script performs the following steps to remove the CP server configuration:

- Stops the CP server
- Removes the CP server from VCS configuration
- Removes resource dependencies
- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration
- Successfully unconfigured the Veritas Coordination Point Server

```
The CP server database is not being deleted on the shared storage.
It can be re-used if CP server is reconfigured on the cluster.
The same database location can be specified during CP server
configuration.
```

5 Decide if you want to delete the CP server configuration file.

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files
(in /var/VRTScps)? [y,n,q] (n) y

Deleting /etc/vxcps.conf and log files on sys1.... Done
Deleting /etc/vxcps.conf and log files on sys2... Done
```

6 Confirm if you want to send information about this installation to Symantec to help improve installation in the future.

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

Upload completed successfully.

Removing the Storage Foundation for Databases (SFDB) repository

After removing the product, you can remove the SFDB repository file and any backups.

Removing the SFDB repository file disables the SFDB tools.

To remove the SFDB repository

1 Identify the SFDB repositories created on the host.

Oracle:

```
# cat /var/vx/vxdba/rep_loc
```

```
{
   "sfae_rept_version" : 1,
   "oracle" : {
      "SFAEDB" : {
         "location" : "/data/sfaedb/.sfae",
         "old_location" : "",
         "alias" : [
            "sfaedb"
         ]
      }
   }
}
```

2 Remove the directory identified by the `location` key.

Oracle:

```
# rm -rf /data/sfaedb/.sfae
```

DB2 9.5 and 9.7:

```
# rm -rf /db2data/db2inst1/NODE0000/SQL00001/.sfae
```

DB2 10.1 and 10.5:

```
# rm -rf /db2data/db2inst1/NODE0000/SQL00001/MEMBER0000/.sfae
```

3 Remove the repository location file.

```
# rm -rf /var/vx/vxdba/rep_loc
```

This completes the removal of the SFDB repository.

Chapter 27

Uninstalling SFHA using response files

This chapter includes the following topics:

- Uninstalling SFHA using response files
- Response file variables to uninstall Storage Foundation and High Availability
- Sample response file for SFHA uninstallation

Uninstalling SFHA using response files

Typically, you can use the response file that the installer generates after you perform SFHA uninstallation on one cluster to uninstall SFHA on other clusters.

To perform an automated uninstallation

1. Make sure that you meet the prerequisites to uninstall SFHA.
2. Copy the response file to the system where you want to uninstall SFHA.
3. Edit the values of the response file variables as necessary.
4. Start the uninstallation from the system to which you copied the response file. For example:

   ```
   # /opt/VRTS/install/uninstallsfha<version>
    -responsefile /tmp/response_file
   ```

   Where *<version>* is the specific release version, and /tmp/*response_file* is the response file's full path name.

   See "About the script-based installer" on page 69.

Response file variables to uninstall Storage Foundation and High Availability

Table 27-1 lists the response file variables that you can define to configure SFHA.

Table 27-1 Response file variables for uninstalling SFHA

Variable	Description
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{prod}	Defines the product to be installed or uninstalled. List or scalar: scalar Optional or required: required
CFG{opt}{keyfile}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional
CFG{opt}{tmppath}	Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
CFG{opt}{uninstall}	Uninstalls SFHA RPMs. List or scalar: scalar Optional or required: optional

Sample response file for SFHA uninstallation

The following example shows a response file for uninstalling Storage Foundation High Availability.

```
our %CFG;

$CFG{opt}{redirect}=1;
$CFG{opt}{uninstall}=1;
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw(cdgv240a cdgv240b) ];

1;
```

Section 11

Adding and removing nodes

- Chapter 28. Adding a node to SFHA clusters
- Chapter 29. Removing a node from SFHA clusters

Chapter 28

Adding a node to SFHA clusters

This chapter includes the following topics:

- About adding a node to a cluster
- Before adding a node to a cluster
- Adding a node to a cluster using the SFHA installer
- Adding a node using the web-based installer
- Adding the node to a cluster manually
- Adding a node using response files
- Configuring server-based fencing on the new node
- After adding the new node
- Adding nodes to a cluster that is using authentication for SFDB tools
- Updating the Storage Foundation for Databases (SFDB) repository after adding a node

About adding a node to a cluster

After you install SFHA and create a cluster, you can add and remove nodes from the cluster.You can create clusters of up to 64 nodes.

You can add a node:

- Using the product installer
- Using the web installer

- Manually

The following table provides a summary of the tasks required to add a node to an existing SFHA cluster.

Table 28-1 Tasks for adding a node to a cluster

Step	Description
Complete the prerequisites and preparatory tasks before adding a node to the cluster.	See "Before adding a node to a cluster" on page 422.
Add a new node to the cluster.	See "Adding a node to a cluster using the SFHA installer" on page 424. See "Adding a node using the web-based installer" on page 427. See "Adding the node to a cluster manually" on page 428.
If you are using the Storage Foundation for Databases (SFDB) tools, you must update the repository database.	See "Adding nodes to a cluster that is using authentication for SFDB tools" on page 436. See "Updating the Storage Foundation for Databases (SFDB) repository after adding a node" on page 437.

The example procedures describe how to add a node to an existing cluster with two nodes.

Before adding a node to a cluster

Before preparing to add the node to an existing SFHA cluster, perform the required preparations.

- Verify hardware and software requirements are met.
- Set up the hardware.
- Prepare the new node.

To verify hardware and software requirements are met

1. Review hardware and software requirements for SFHA.
2. Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
3. Verify the existing cluster is a SFHA cluster and that SFHA is running on the cluster.

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in Figure 28-1.

Figure 28-1 Adding a node to a two-node cluster using two switches

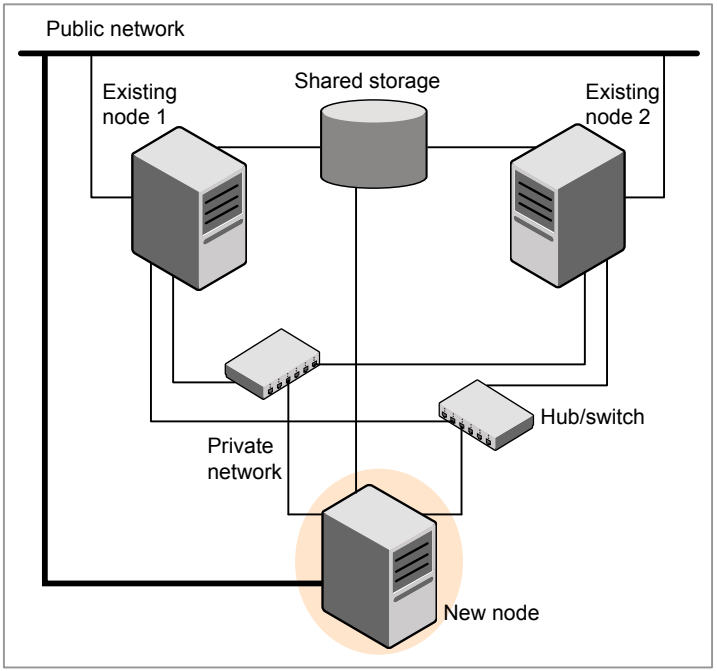

To set up the hardware

1 Connect the SFHA private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a two-node cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

Figure 28-1 illustrates a new node being added to an existing two-node cluster using two independent hubs.

2 Make sure that you meet the following requirements:

- The node must be connected to the same shared storage devices as the existing nodes.

- The node must have private network connections to two independent switches for the cluster.
  For more information, see the *Symantec Cluster Server Installation Guide*.
- The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

Complete the following preparatory steps on the new node before you add it to an existing SFHA cluster.

To prepare the new node

1. Navigate to the folder that contains the installsfha program. Verify that the new node meets installation requirements.Verify that the new node meets installation requirements.

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

   You can also use the web-based installer for the precheck.

2. Install SFHA RPMs only without configuration on the new system. Make sure all the VRTS RPMs available on the existing nodes are also available on the new node.

   ```
   # ./installsfha
   ```

   Do not configure SFHA when prompted.

   ```
   Would you like to configure SFHA on sys5? [y,n,q]? n
   ```

Adding a node to a cluster using the SFHA installer

You can add a node to a cluster using the `-addnode` option with the SFHA installer.

The SFHA installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and RPMs installed but not configured 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 and copies the following files to the new node from the existing 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/sysconfig/llt
  /etc/sysconfig/gab
  /etc/sysconfig/vxfen
- Generate security credentials on the new node if the CPS server of existing cluster is secure
- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.

At the end of the process, the new node joins the SFHA cluster.

Note: If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

See "Removing the node configuration from the CP server" on page 444.

To add the node to an existing cluster using the installer

1. Log in as the root user on one of the nodes of the existing cluster.
2. Run the SFHA installer with the -addnode option.

   ```
   # cd /opt/VRTS/install

   # ./installsfha<version> -addnode
   ```

   Where *<version>* is the specific release version.

   See "About the script-based installer" on page 69.

   The installer displays the copyright message and the location where it stores the temporary installation logs.

3 Enter the name of a node in the existing SFHA cluster.

The installer uses the node information to identify the existing cluster.

```
Enter one node of the SFHA cluster to which
you would like to add one or more new nodes: sys1
```

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: sys5
```

Confirm if the installer prompts if you want to add the node to the cluster.

The installer checks the installed products and RPMs 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.

```
Enter the NIC for the first private heartbeat
link on sys5: [b,q,?] eth1

Enter the NIC for the second private heartbeat
link on sys5: [b,q,?] eth2
```

Note: At least two private heartbeat links must be configured for high availability of the cluster.

7 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

8 Review and confirm the information.

9 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on sys5: eth3
```

10 If the existing cluster uses server-based fencing in secure mode, the installer will configure server-based fencing in secure mode on the new nodes.

The installer then starts all the required Symantec processes and joins the new node to cluster.

The installer indicates the location of the log file, summary file, and response file with details of the actions performed.

If you have enabled security on the cluster, the installer displays the following message:

```
Since the cluster is in secure mode, check the main.cf
whether you need to modify the usergroup that you would like
to grant read access. If needed, use the following commands
to modify:

# hauser -addpriv <user group> GuestGroup

# haconf -makerw

# haconf -dump -makero
```

11 Confirm that the new node has joined the SFHA cluster using `lltstat -n` and `gabconfig -a` commands.

Adding a node using the web-based installer

You can use the web-based installer to add a node to a cluster.

To add a node to a cluster using the web-based installer

1 From the Task pull-down menu, select **Add a Cluster node**.

From the product pull-down menu, select the product.

Click the **Next** button.

2 Click **OK** to confirm the prerequisites to add a node.

3 In the System Names field enter a name of a node in the cluster where you plan to add the node and click **OK**.

The installer program checks inter-system communications and compatibility. If the node fails any of the checks, review the error and fix the issue.

If prompted, review the cluster's name, ID, and its systems. Click the **Yes** button to proceed.

4 In the System Names field, enter the names of the systems that you want to add to the cluster as nodes. Separate system names with spaces. Click the **Next** button.

The installer program checks inter-system communications and compatibility. If the system fails any of the checks, review the error and fix the issue.

Click the **Next** button. If prompted, click the **Yes** button to add the system and to proceed.

5 From the heartbeat NIC pull-down menus, select the heartbeat NICs for the cluster. Click the **Next** button.

6 Once the addition is complete, review the log files. Optionally send installation information to Symantec. Click the **Finish** button to complete the node's addition to the cluster.

Adding the node to a cluster manually

Perform this procedure after you install SFHA only if you plan to add the node to the cluster manually.

Table 28-2 Procedures for adding a node to a cluster manually

Step	Description
Start the Veritas Volume Manager (VxVM) on the new node.	See "Starting Veritas Volume Manager (VxVM) on the new node" on page 429.
Configure the cluster processes on the new node.	See "Configuring cluster processes on the new node" on page 429.
If the CPS server of existing cluster is secure, generate security credentials on the new node.	See "Setting up the node to run in secure mode" on page 431.
Configure fencing for the new node to match the fencing configuration on the existing cluster. If the existing cluster is configured to use server-based I/O fencing, configure server-based I/O fencing on the new node.	See "Starting fencing on the new node" on page 432.
Start VCS.	See "To start VCS on the new node" on page 436.

Table 28-2 Procedures for adding a node to a cluster manually *(continued)*

Step	Description
If the ClusterService group is configured on the existing cluster, add the node to the group.	See "Configuring the ClusterService group for the new node" on page 432.

Starting Veritas Volume Manager (VxVM) on the new node

Veritas Volume Manager (VxVM) uses license keys to control access. As you run the `vxinstall` utility, answer **n** to prompts about licensing. You installed the appropriate license when you ran the `installsfha` program.

To start VxVM on the new node

1. To start VxVM on the new node, use the vxinstall utility:

   ```
   # vxinstall
   ```

2. Enter **n** when prompted to set up a system wide disk group for the system.

   The installation completes.

3. Verify that the daemons are up and running. Enter the command:

   ```
   # vxdisk list
   ```

   Make sure the output displays the shared disks without errors.

Configuring cluster processes on the new node

Perform the steps in the following procedure to configure cluster processes on the new node.

1. Do not apply for SUSE Linux.
2. Edit the /etc/llthosts file on the existing nodes. Using vi or another text editor, add the line for the new node to the file. The file resembles:

   ```
   0 sys1
   1 sys2
   2 sys5
   ```

3. Copy the `/etc/llthosts` file from one of the existing systems over to the new system. The `/etc/llthosts` file must be identical on all nodes in the cluster.

4 Create an `/etc/llttab` file on the new system. For example:

```
set-node sys5
set-cluster 101

link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - ether - -
```

Except for the first line that refers to the node, the file resembles the /etc/llttab files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

5 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

```
/sbin/gabconfig -c -nN
```

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

6 Edit the /etc/gabtab file on each of the existing systems, changing the content to match the file on the new system.

7 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
sys5
```

8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -rsh -clus -copy \
-from_sys sys1 -to_sys sys5
```

9 Start the LLT, GAB, and ODM drivers on the new node:

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

# /etc/init.d/gab start

# /etc/init.d/odm restart
```

10 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a
GAB Port Memberships
===============================================================
Port a gen df204 membership 012
Port b gen df20a membership 012
Port d gen df207 membership 012
```

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

Table 28-3 uses the following information for the following command examples.

Table 28-3 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
sys5	sys5.nodes.example.com	The new node that you are adding to the cluster.

Setting up SFHA related security configuration

Perform the following steps to configure SFHA related security settings.

Setting up SFHA related security configuration

1 Start /opt/VRTSat/bin/vxatd process.

2 Create HA_SERVICES domain for SFHA.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

3 Add SFHA and webserver principal to AB on node sys5.

```
# 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
```

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

1. For disk-based fencing on at least one node, copy the following files from one of the nodes in the existing cluster to the new node:

   ```
   /etc/sysconfig/vxfen
   /etc/vxfendg
   /etc/vxfenmode
   ```

   See "Configuring server-based fencing on the new node" on page 434.

2. Start fencing on the new node:

Configuring the ClusterService group for the new node

If the ClusterService group is configured on the existing cluster, add the node to the group by performing the steps in the following procedure on one of the nodes in the existing cluster.

To configure the ClusterService group for the new node

1. On an existing node, for example sys1, write-enable the configuration:

   ```
   # haconf -makerw
   ```

2. Add the node sys5 to the existing ClusterService group.

   ```
   # hagrp -modify ClusterService SystemList -add sys5 2

   # hagrp -modify ClusterService AutoStartList -add sys5
   ```

3. Modify the IP address and NIC resource in the existing group for the new node.

   ```
   # hares -modify gcoip Device eth0 -sys sys5

   # hares -modify gconic Device eth0 -sys sys5
   ```

4. Save the configuration by running the following command from any node.

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

Adding a node using response files

Typically, you can use the response file that the installer generates on one system to add nodes to an existing cluster.

To add nodes using response files

1. Make sure the systems where you want to add nodes meet the requirements.
2. Make sure all the tasks required for preparing to add a node to an existing SFHA cluster are completed.
3. Copy the response file to one of the systems where you want to add nodes.

   See "Sample response file for adding a node to a SFHA cluster" on page 434.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to add a node to a SFHA cluster" on page 433.
5. Mount the product disc and navigate to the folder that contains the installation program.
6. Start adding nodes from the system to which you copied the response file. For example:

   ```
   # ./installer -responsefile /tmp/response_file

   # ./installsfha -responsefile /tmp/response_file
   ```

   Where /tmp/*response_file* is the response file's full path name.

   Depending on the fencing configuration in the existing cluster, the installer configures fencing on the new node. The installer then starts all the required Symantec processes and joins the new node to cluster. The installer indicates the location of the log file and summary file with details of the actions performed.

Response file variables to add a node to a SFHA cluster

Table 28-4 lists the response file variables that you can define to add a node to an SFHA cluster.

Table 28-4 Response file variables for adding a node to an SFHA cluster

Variable	Description
$CFG{opt}{addnode}	Adds a node to an existing cluster. List or scalar: scalar Optional or required: required

Table 28-4 Response file variables for adding a node to an SFHA cluster *(continued)*

Variable	Description
$CFG{newnodes}	Specifies the new nodes to be added to the cluster. List or scalar: list Optional or required: required

Sample response file for adding a node to a SFHA cluster

The following example shows a response file for adding a node to a SFHA cluster.

```
our %CFG;

$CFG{clustersystems}=[ qw(sys1) ];
$CFG{newnodes}=[ qw(sys5) ];
$CFG{opt}{addnode}=1;
$CFG{opt}{configure}=1;d
$CFG{opt}{vr}=1;
$CFG{prod}="SFHA62";
d$CFG{systems}=[ qw(sys1 sys5) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_lltlink1}{sys5}="eth1";
$CFG{vcs_lltlink2}{sys5}="eth2";

1;
```

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 cps1.symantecexample.com \
-a add_node -c clus1 -h sys5 -n2

Node 2 (sys5) successfully added
```

3. Verify that the new node is added to the CP server configuration:

```
# cpsadm -s cps1.symantecexample.com \
-a list_nodes
```

   The new node must be listed in the command output.

4. Add the VCS user cpsclient@sys5 to each CP server:

```
# cpsadm -s cps1.symantecexample.com \
-a add_user -e cpsclient@sys5 \
-f cps_operator -g vx

User cpsclient@sys5 successfully added
```

To configure server-based fencing with security on the new node

1. Log in to each CP server as the root user.
2. Update each CP server configuration with the new node information:

```
# cpsadm -s cps1.symantecexample.com \
-a add_node -c clus1 -h sys5 -n2

Node 2 (sys5) successfully added
```

3. Verify that the new node is added to the CP server configuration:

```
# cpsadm -s cps1.symantecexample.com -a list_nodes
```

   The new node must be listed in the output.

Adding the new node to the vxfen service group

Perform the steps in the following procedure to add the new node to the vxfen service group.

To add the new node to the vxfen group using the CLI

1 On one of the nodes in the existing Storage Foundation High Availability, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

2 Add the node sys5 to the existing vxfen group.

```
# hagrp -modify vxfen SystemList -add sys5 2
```

3 Save the configuration by running the following command from any node in the Storage Foundation High Availability:

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

After adding the new node

Start VCS on the new node.

To start VCS on the new node

- Start VCS on the new node:

```
# hastart
```

Adding nodes to a cluster that is using authentication for SFDB tools

To add a node to a cluster that is using authentication for SFDB tools, perform the following steps as the root user

1 Export authentication data from a node in the cluster that has already been authorized, by using the `-o export_broker_config` option of the `sfae_auth_op` command.

Use the `-f` option to provide a file name in which the exported data is to be stored.

```
# /opt/VRTS/bin/sfae_auth_op \
-o export_broker_config -f exported-data
```

2 Copy the exported file to the new node by using any available copy mechanism such as `scp` or `rcp`.

3 Import the authentication data on the new node by using the `-o import_broker_config` option of the `sfae_auth_op` command.

Use the `-f` option to provide the name of the file copied in Step 2.

```
# /opt/VRTS/bin/sfae_auth_op \
 -o import_broker_config -f exported-data
Setting up AT
Importing broker configuration
Starting SFAE AT broker
```

4 Stop the `vxdbd` daemon on the new node.

```
# /opt/VRTS/bin/sfae_config disable
vxdbd has been disabled and the daemon has been stopped.
```

5 Enable authentication by setting the `AUTHENTICATION` key to `yes` in the `/etc/vx/vxdbed/admin.properties` configuration file.

If `/etc/vx/vxdbed/admin.properties` does not exist, then use `cp /opt/VRTSdbed/bin/admin.properties.example /etc/vx/vxdbed/admin.properties`

6 Start the `vxdbd` daemon.

```
# /opt/VRTS/bin/sfae_config enable
vxdbd has been enabled and the daemon has been started.
It will start automatically on reboot.
```

The new node is now authenticated to interact with the cluster to run SFDB commands.

Updating the Storage Foundation for Databases (SFDB) repository after adding a node

If you are using Database Storage Checkpoints, Database FlashSnap, or SmartTier for Oracle in your configuration, update the SFDB repository to enable access for the new node after it is added to the cluster.

To update the SFDB repository after adding a node

1. Copy the `/var/vx/vxdba/rep_loc` file from one of the nodes in the cluster to the new node.
2. If the `/var/vx/vxdba/auth/user-authorizations` file exists on the existing cluster nodes, copy it to the new node.

   If the `/var/vx/vxdba/auth/user-authorizations` file does not exist on any of the existing cluster nodes, no action is required.

   This completes the addition of the new node to the SFDB repository.

Chapter 29

Removing a node from SFHA clusters

This chapter includes the following topics:

- Removing a node from a SFHA cluster

Removing a node from a SFHA cluster

Table 29-1 specifies the tasks that are involved in removing a node from a cluster. In the example procedure, the cluster consists of nodes sys1, sys2, and sys5; node sys5 is to leave the cluster.

Table 29-1 Tasks that are involved in removing a node

Task	Reference
■ Back up the configuration file. ■ Check the status of the nodes and the service groups.	See "Verifying the status of nodes and service groups" on page 440.
■ Switch or remove any SFHA service groups on the node departing the cluster. ■ Delete the node from SFHA configuration.	See "Deleting the departing node from SFHA configuration" on page 441.
Modify the llthosts(4) and gabtab(4) files to reflect the change.	See "Modifying configuration files on each remaining node" on page 444.
For a cluster that is running in a secure mode, remove the security credentials from the leaving node.	See "Removing security credentials from the leaving node " on page 445.

Table 29-1 Tasks that are involved in removing a node *(continued)*

Task	Reference
On the node departing the cluster: ■ Modify startup scripts for LLT, GAB, and SFHA to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities.	See "Unloading LLT and GAB and removing VCS on the departing node" on page 446.

Verifying the status of nodes and service groups

Start by issuing the following commands from one of the nodes to remain in the cluster node sys1 or node sys2 in our example.

To verify the status of the nodes and the service groups

1 Make a backup copy of the current configuration file, main.cf.

```
# cp -p /etc/VRTSvcs/conf/config/main.cf\
/etc/VRTSvcs/conf/config/main.cf.goodcopy
```

2 Check the status of the systems and the service groups.

```
# hastatus -summary

-- SYSTEM STATE
-- System      State          Frozen
A  sys1    RUNNING        0
A  sys2    RUNNING        0
A  sys5      RUNNING        0

-- GROUP STATE
-- Group    System      Probed   AutoDisabled   State
B  grp1     sys1      Y        N               ONLINE
B  grp1     sys2      Y        N               OFFLINE
B  grp2     sys1      Y        N               ONLINE
B  grp3     sys2      Y        N               OFFLINE
B  grp3     sys5    Y        N              ONLINE
B  grp4     sys5    Y        N              ONLINE
```

The example output from the `hastatus` command shows that nodes sys1, sys2, and sys5 are the nodes in the cluster. Also, service group grp3 is configured to run on node sys2 and node sys5, the departing node. Service group grp4 runs only on node sys5. Service groups grp1 and grp2 do not run on node sys5.

Deleting the departing node from SFHA configuration

Before you remove a node from the cluster you need to identify the service groups that run on the node.

You then need to perform the following actions:

- Remove the service groups that other service groups depend on, or
- Switch the service groups to another node that other service groups depend on.

To remove or switch service groups from the departing node

1. Switch failover service groups from the departing node. You can switch grp3 from node sys5 to node sys2.

   ```
   # hagrp -switch grp3 -to sys2
   ```

2. Check for any dependencies involving any service groups that run on the departing node; for example, grp4 runs only on the departing node.

   ```
   # hagrp -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
   # hagrp -unlink grp4 grp1
   ```

   These commands enable you to edit the configuration and to remove the requirement grp4 has for grp1.

4. Stop SFHA on the departing node:

   ```
   # hastop -sys sys5
   ```

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  sys1    RUNNING        0
   A  sys2      RUNNING        0
   A  sys5      EXITED         0

   -- GROUP STATE
   -- Group    System      Probed   AutoDisabled   State
   B  grp1     sys1      Y        N              ONLINE
   B  grp1     sys2      Y        N              OFFLINE
   B  grp2     sys1      Y        N              ONLINE
   B  grp3     sys2      Y        N              ONLINE
   B  grp3     sys5    Y        Y              OFFLINE
   B  grp4     sys5    Y        N              OFFLINE
   ```

6 Delete the departing node from the SystemList of service groups grp3 and grp4.

```
# haconf -makerw
# hagrp -modify grp3 SystemList -delete sys5
# hagrp -modify grp4 SystemList -delete sys5
```

Note: If sys5 was in the autostart list, then you need to manually add another system in the autostart list so that after reboot, the group comes online automatically.

7 For the service groups that run only on the departing node, delete the resources from the group before you delete the group.

```
# hagrp -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.

```
# hagrp -delete grp4
```

9 Check the status.

```
# hastatus -summary
    -- SYSTEM STATE
    -- System      State          Frozen
    A  sys1      RUNNING        0
    A  sys2      RUNNING        0
    A  sys5    EXITED         0

    -- GROUP STATE
    -- Group    System      Probed   AutoDisabled   State
    B  grp1     sys1      Y        N              ONLINE
    B  grp1     sys2      Y        N              OFFLINE
    B  grp2     sys1      Y        N              ONLINE
    B  grp3     sys2      Y        N              ONLINE
```

10 Delete the node from the cluster.

```
# hasys -delete sys5
```

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 `-n`*N* option, where *N* is the number of cluster systems.

If the command has the form `/sbin/gabconfig -c -n`*N*, where *N* is the number of cluster systems, make sure that *N* is not greater than the actual number of nodes in the cluster. When *N* is greater than the number of nodes, GAB does not automatically seed.

Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`.

2 Modify /etc/llthosts file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 sys1
1 sys2
2 sys5
```

To:

```
0 sys1
1 sys2
```

Removing the node configuration from the CP server

After removing a node from a SFHA cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

Note: The `cpsadm` command is used to perform the steps in this procedure. For detailed information about the `cpsadm` command, see the *Symantec Cluster Server Administrator's Guide*.

To remove the node configuration from the CP server

1. Log into the CP server as the root user.
2. View the list of VCS users on the CP server.

   If the CP server is configured to use HTTPS-based communication, run the following command:

   ```
   # cpsadm -s cp_server -a list_users
   ```

   If the CP server is configured to use IPM-based communication, run the following command:

   ```
   # cpsadm -s cp_server -p 14250 -a list_users
   ```

   Where *cp_server* is the virtual IP/ virtual hostname of the CP server.
3. Remove the VCS user associated with the node you previously removed from the cluster.

   For CP server in non-secure mode:

   ```
   # cpsadm -s cp_server  -a rm_user \
   -e cpsclient@sys5  -f cps_operator  -g vx
   ```

4. Remove the node entry from the CP server:

   ```
   # cpsadm -s cp_server -a rm_node  -h sys5 -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 sys5. Perform the following steps.

To remove the security credentials

1. Stop the AT process.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh \
stop
```

2. Remove the credentials.

```
# rm -rf /var/VRTSvcs/vcsauth/data/
```

Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node that is departing the cluster.

You can use script-based installer to uninstall VCS on the departing node or perform the following manual steps.

If you have configured SFHA as part of the Storage Foundation and High Availability products, you may have to delete other dependent RPMs before you can delete all of the following ones.

To stop LLT and GAB and remove SFHA

1. If you had configured I/O fencing in enabled mode, then stop I/O fencing.

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

2. Stop GAB and LLT:

```
# /etc/init.d/gab stop
# /etc/init.d/llt stop
```

3. To determine the RPMs to remove, enter:

```
# rpm -qa |grep VRTS
```

4. To permanently remove the VCS RPMs from the system, use the `rpm -e` command. Start by removing the following RPMs, which may have been optionally installed, in the order shown:

```
# rpm -e VRTSfcpi62
# rpm -e VRTSvcswiz
# rpm -e VRTSvbs
# rpm -e VRTSsfmh
# rpm -e VRTSvcsea
# rpm -e VRTSvcsdr
# rpm -e VRTSvcsag
```

```
# rpm -e VRTScps
# rpm -e VRTSvcs
# rpm -e VRTSamf
# rpm -e VRTSvxfen
# rpm -e VRTSgab
# rpm -e VRTSllt
# rpm -e VRTSspt
# rpm -e VRTSvlic
# rpm -e VRTSperl
```

5 Remove the LLT and GAB configuration files.

```
# rm /etc/lltab
# rm /etc/gabtab
# rm /etc/llthosts
```

Updating the Storage Foundation for Databases (SFDB) repository after removing a node

After removing a node from a cluster, you do not need to perform any steps to update the SFDB repository.

For information on removing the SFDB repository after removing the product:

See “Removing the Storage Foundation for Databases (SFDB) repository” on page 415.

Section 12

Installation reference

- Appendix A. SFHA services and ports
- Appendix B. Installation scripts
- Appendix C. Tunable files for installation
- Appendix D. Configuration files
- Appendix E. Configuring the secure shell or the remote shell for communications
- Appendix F. Storage Foundation and High Availability components
- Appendix G. Troubleshooting installation issues
- Appendix H. Troubleshooting cluster installation
- Appendix I. Sample Storage Foundation High Availability setup diagrams for CP server-based I/O fencing
- Appendix J. Configuring LLT over UDP
- Appendix K. Using LLT over RDMA
- Appendix L. Compatibility issues when installing Storage Foundation High Availability with other products

Appendix A

SFHA services and ports

This appendix includes the following topics:

- About SFHA services and ports

About SFHA services and ports

If you have configured a firewall, ensure that the firewall settings allow access to the services and ports used by SFHA.

Table A-1 lists the services and ports used by SFHA .

Note: The port numbers that appear in bold are mandatory for configuring SFHA.

Table A-1 SFHA services and ports

Port Number	Protocol	Description	Process
4145	TCP/UDP	VVR Connection Server VCS Cluster Heartbeats	vxio
5634	HTTPS	Symantec Storage Foundation Messaging Service	xprtld
8199	TCP	Volume Replicator Administrative Service	vras
8989	TCP	VVR Resync Utility	vxreserver

Table A-1 SFHA services and ports *(continued)*

Port Number	Protocol	Description	Process
14141	TCP	Symantec High Availability Engine Veritas Cluster Manager (Java console) (ClusterManager.exe) VCS Agent driver (VCSAgDriver.exe)	had
14144	TCP/UDP	VCS Notification	Notifier
14149	TCP/UDP	VCS Authentication	vcsauthserver
14150	TCP	Veritas Command Server	CmdServer
14155	TCP/UDP	VCS Global Cluster Option (GCO)	wac
14156	TCP/UDP	VCS Steward for GCO	steward
443	TCP	Coordination Point Server	Vxcpserv
49152-65535	TCP/UDP	Volume Replicator Packets	User configurable ports created at kernel level by `vxio.sys` file

Appendix B

Installation scripts

This appendix includes the following topics:

- Installation script options
- About using the postcheck option

Installation script options

Table B-1 shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Symantec Storage Foundation product scripts, except where otherwise noted.

See "About the script-based installer" on page 69.

Table B-1 Available command line options

Command Line Option	Function
-addnode	Adds a node to a high availability cluster.
-allpkgs	Displays all RPMs required for the specified product. The RPMs are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.

Table B-1 Available command line options *(continued)*

Command Line Option	Function
-comsetup	The `-comsetup` option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.
-configure	Configures the product after installation.
-fencing	Configures I/O fencing in a running cluster.
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases RPM upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-online_upgrade	Used to perform online upgrade. Using this option, the installer upgrades the whole cluster and also supports customer's application zero down time during the upgrade procedure. Now this option only supports VCS and ApplicationHA.
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.

Table B-1 Available command line options *(continued)*

Command Line Option	Function
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-installallpkgs	The `-installallpkgs` option is used to select all RPMs.
-installrecpkgs	The `-installrecpkgs`option is used to select the recommended RPMs set.
–installminpkgs	The `-installminpkgs`option is used to select the minimum RPMs set.
-ignorepatchreqs	The `-ignorepatchreqs` option is used to allow installation or upgrading even if the prerequisite RPMs or patches are missed on the system.
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-I ssh_key_file` to every SSH invocation.
–kickstart *dir_path*	Produces a kickstart configuration file for installing with Linux RHEL Kickstart. The file contains the list of Symantec RPMs in the correct order for installing, in a format that can be used for Kickstart installations. The *dir_path* indicates the path to the directory in which to create the file.
-license	Registers or updates product licenses on the specified systems.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
-makeresponsefile	Use the `-makeresponsefile` option only to generate response files. No actual software installation occurs when you use this option.

Table B-1 Available command line options *(continued)*

Command Line Option	Function
-minpkgs	Displays the minimal RPMs required for the specified product. The RPMs are listed in correct installation order. Optional RPMs are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-noipc	Disables the installer from making outbound networking calls to Symantec Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.
-nolic	Allows installation of product RPMs without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.
–pkginfo	Displays a list of RPMs and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS RPMs.
–pkgset	Discovers and displays the RPM group (minimum, recommended, all) and RPMs that are installed on the specified systems.
-pkgtable	Displays product's RPMs in correct installation order by group.
–postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.

Table B-1 Available command line options *(continued)*

Command Line Option	Function
–recpkgs	Displays the recommended RPMs required for the specified product. The RPMs are listed in correct installation order. Optional RPMs are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer patch file.
-requirements	The `-requirements` option displays required OS version, required RPMs and patches, file system space, and other system requirements in order to install the product.
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* 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.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The `-rollingupgrade_phase1` option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs get upgraded to the latest version.
-rollingupgrade_phase2	The `-rollingupgrade_phase2` option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.

Table B-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 "About configuring secure shell or remote shell communication modes before installing products" on page 491.
–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.
-settunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the `-tunablesfile` option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the `-timeout` option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 prevents the script from timing out. The `-timeout` option does not work with the `-serial` option
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.
-tunables	Lists all supported tunables and create a tunables file template.

Table B-1 Available command line options *(continued)*

Command Line Option	Function
-tunables_file *tunables_file*	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.
-version	Checks and reports the installed products and their versions. Identifies the installed and missing RPMs and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing RPMs and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.
-yumgroupxml	The `-yumgroupxml` option is used to generate a yum group definition XML file. The `createrepo` command can use the file on Redhat Linux to create a yum group for automated installation of all RPMs for a product. An available location to store the XML file should be specified as a complete path. The `-yumgroupxml` option is supported on Redhat Linux only.

About using the postcheck option

You can use the installer's post-check to determine installation-related problems and to aid in troubleshooting.

Note: This command option requires downtime for the node.

When you use the `postcheck` option, it can help you troubleshoot the following VCS-related issues:

- The heartbeat link does not exist.
- The heartbeat link cannot communicate.
- The heartbeat link is a part of a bonded or aggregated NIC.
- A duplicated cluster ID exists (if LLT is not running at the check time).

- The VRTSllt pkg version is not consistent on the nodes.
- The llt-linkinstall value is incorrect.
- The /etc/llthosts and /etc/lltab configuration is incorrect.
- the `/etc/gabtab` file is incorrect.
- The incorrect GAB linkinstall value exists.
- The VRTSgab pkg version is not consistent on the nodes.
- The `main.cf` file or the `types.cf` file is invalid.
- The `/etc/VRTSvcs/conf/sysname` file is not consistent with the hostname.
- The cluster UUID does not exist.
- The `uuidconfig.pl` file is missing.
- The VRTSvcs pkg version is not consistent on the nodes.
- The `/etc/vxfenmode` file is missing or incorrect.
- The `/etc/vxfendg file` is invalid.
- The vxfen link-install value is incorrect.
- The VRTSvxfen pkg version is not consistent.

The `postcheck` option can help you troubleshoot the following SFHA or SFCFSHA issues:

- Volume Manager cannot start because the `/etc/vx/reconfig.d/state.d/install-db` file has not been removed.
- Volume Manager cannot start because the `volboot` file is not loaded.
- Volume Manager cannot start because no license exists.
- Cluster Volume Manager cannot start because the CVM configuration is incorrect in the `main.cf` file. For example, the Autostartlist value is missing on the nodes.
- Cluster Volume Manager cannot come online because the node ID in the `/etc/llthosts` file is not consistent.
- Cluster Volume Manager cannot come online because Vxfen is not started.
- Cluster Volume Manager cannot start because gab is not configured.
- Cluster Volume Manager cannot come online because of a CVM protocol mismatch.
- Cluster Volume Manager group name has changed from "cvm", which causes CVM to go offline.

You can use the installer's post-check option to perform the following checks:

General checks for all products:

- All the required RPMs are installed.
- The versions of the required RPMs are correct.
- There are no verification issues for the required RPMs.

Checks for Volume Manager (VM):

- Lists the daemons which are not running (`vxattachd`, `vxconfigbackupd`, `vxesd`, `vxrelocd` ...).
- Lists the disks which are not in 'online' or 'online shared' state (`vxdisk list`).
- Lists the diskgroups which are not in 'enabled' state (`vxdg list`).
- Lists the volumes which are not in 'enabled' state (`vxprint -g <dgname>`).
- Lists the volumes which are in 'Unstartable' state (`vxinfo -g <dgname>`).
- Lists the volumes which are not configured in `/etc/fstab` .

Checks for File System (FS):

- Lists the VxFS kernel modules which are not loaded (`vxfs/fdd/vxportal`.).
- Whether all VxFS file systems present in `/etc/fstab` file are mounted.
- Whether all VxFS file systems present in `/etc/fstab` are in disk layout 6 or higher.
- Whether all mounted VxFS file systems are in disk layout 6 or higher.

Checks for Cluster File System:

- Whether FS and ODM are running at the latest protocol level.
- Whether all mounted CFS file systems are managed by VCS.
- Whether cvm service group is online.

See “Performing a postcheck on a node” on page 393.

Appendix C

Tunable files for installation

This appendix includes the following topics:

- About setting tunable parameters using the installer or a response file
- Setting tunables for an installation, configuration, or upgrade
- Setting tunables with no other installer-related operations
- Setting tunables with an un-integrated response file
- Preparing the tunables file
- Setting parameters for the tunables file
- Tunables value parameter definitions

About setting tunable parameters using the installer or a response file

You can set non-default product and system tunable parameters using a tunables file. With the file, you can set tunables such as the I/O policy or toggle native multi-pathing. The tunables file passes arguments to the installer script to set tunables. With the file, you can set the tunables for the following operations:

- When you install, configure, or upgrade systems.

  ```
  # ./installer -tunablesfile tunables_file_name
  ```

  See “Setting tunables for an installation, configuration, or upgrade” on page 461.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -settunables [
sys1 sys2 ...]
```

See "Setting tunables with no other installer-related operations" on page 462.

- When you apply the tunables file with an un-integrated response file.

```
# ./installer -responsefile response_file_name -tunablesfile
tunables_file_name
```

See "Setting tunables with an un-integrated response file" on page 463.

See "About response files" on page 47.

You must select the tunables that you want to use from this guide.

See "Tunables value parameter definitions" on page 465.

Setting tunables for an installation, configuration, or upgrade

You can use a tunables file for installation procedures to set non-default tunables. You invoke the installation script with the `tunablesfile` option. The tunables file passes arguments to the script to set the selected tunables. You must select the tunables that you want to use from this guide.

See "Tunables value parameter definitions" on page 465.

Note: Certain tunables only take effect after a system reboot.

To set the non-default tunables for an installation, configuration, or upgrade

1. Prepare the tunables file.

   See "Preparing the tunables file" on page 464.

2. Make sure the systems where you want to install SFHA meet the installation requirements.
3. Complete any preinstallation tasks.
4. Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
5. Mount the product disc and navigate to the directory that contains the installation program.

6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
 -settunables [sys1 sys2 ...]
```

Where /tmp/*tunables_file* is the full path name for the tunables file.

7 Proceed with the operation. When prompted, accept the tunable parameters.

Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.

8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See "Tunables value parameter definitions" on page 465.

Note: Certain tunables only take effect after a system reboot.

To set tunables with no other installer-related operations

1 Prepare the tunables file.

See "Preparing the tunables file" on page 464.

2 Make sure the systems where you want to install SFHA meet the installation requirements.

3 Complete any preinstallation tasks.

4 Copy the tunables file to one of the systems that you want to tune.

5 Mount the product disc and navigate to the directory that contains the installation program.

6 Start the installer with the `-settunables` option.

```
# ./installer -tunablesfile tunables_file_name -settunables [
sys123 sys234 ...]
```

Where /tmp/*tunables_file* is the full path name for the tunables file.

7 Proceed with the operation. When prompted, accept the tunable parameters.

Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.

8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See "Tunables value parameter definitions" on page 465.

Note: Certain tunables only take effect after a system reboot.

To set tunables with an un-integrated response file

1 Make sure the systems where you want to install SFHA meet the installation requirements.

2 Complete any preinstallation tasks.

3 Prepare the tunables file.

See "Preparing the tunables file" on page 464.

4 Copy the tunables file to one of the systems that you want to tune.

5 Mount the product disc and navigate to the directory that contains the installation program.

6 Start the installer with the `-responsefile` and `-tunablesfile` options.

```
# ./installer -responsefile response_file_name -tunablesfile
tunables_file_name
```

Where *response_file_name* is the full path name for the response file and *tunables_file_name* is the full path name for the tunables file.

7 Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.

8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Preparing the tunables file

A tunables file is a Perl module and consists of an opening and closing statement, with the tunables defined between. Use the hash symbol at the beginning of the line to comment out the line. The tunables file opens with the line "our %TUN;" and ends with the return true "1;" line. The final return true line only needs to appear once at the end of the file. Define each tunable parameter on its own line.

You can use the installer to create a tunables file template, or manually format tunables files you create.

To create a tunables file template

- Start the installer with the `-tunables` option. Enter the following:

  ```
  # ./installer -tunables
  ```

  You see a list of all supported tunables, and the location of the tunables file template.

To manually format tunables files

- Format the tunable parameter as follows:

  ```
  $TUN{"tunable_name"}{"system_name"|"*"}=value_of_tunable;
  ```

For the *system_name*, use the name of the system, its IP address, or a wildcard symbol. The *value_of_tunable* depends on the type of tunable you are setting. End the line with a semicolon.

The following is an example of a tunables file.

```
#
# Tunable Parameter Values:
#
our %TUN;

$TUN{"tunable1"}{"*"}=1024;
$TUN{"tunable3"}{"sys123"}="SHA256";

1;
```

Setting parameters for the tunables file

Each tunables file defines different tunable parameters. The values that you can use are listed in the description of each parameter. Select the tunables that you want to add to the tunables file and then configure each parameter.

See "Tunables value parameter definitions" on page 465.

Each line for the parameter value starts with $TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{"node123"}=16;
```

In this example, you are changing the dmp_daemon_count value from its default of 10 to 16. You can use the wildcard symbol "*" for all systems. For example:

```
$TUN{"dmp_daemon_count"}{"*"}=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Symantec Storage Foundation Cluster File System High Availability Administrator's Guide* for detailed information on product tunable ranges and recommendations.

Table C-1 describes the supported tunable parameters that can be specified in a tunables file.

Table C-1 Supported tunable parameters

Tunable	Description
autoreminor	(Veritas Volume Manager) Enable reminoring in case of conflicts during disk group import.
autostartvolumes	(Veritas Volume Manager) Enable the automatic recovery of volumes.
dmp_cache_open	(Symantec Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached.
dmp_daemon_count	(Symantec Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks.
dmp_delayq_interval	(Symantec Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy.

Table C-1 Supported tunable parameters *(continued)*

Tunable	Description
dmp_fast_recovery	(Symantec Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Symantec Dynamic Multi-Pathing is started.
dmp_health_time	(Symantec Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy.
dmp_log_level	(Symantec Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed.
dmp_low_impact_probe	(Symantec Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled.
dmp_lun_retry_timeout	(Symantec Dynamic Multi-Pathing) The retry period for handling transient errors.
dmp_monitor_fabric	(Symantec Dynamic Multi-Pathing) Whether the Event Source daemon (`vxesd`) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Symantec Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Symantec Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events.
dmp_monitor_ownership	(Symantec Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored.
dmp_native_support	(Symantec Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices.
dmp_path_age	(Symantec Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy.
dmp_pathswitch_blks_shift	(Symantec Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path.

Table C-1 Supported tunable parameters *(continued)*

Tunable	Description
dmp_probe_idle_lun	(Symantec Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs.
dmp_probe_threshold	(Symantec Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon.
dmp_restore_cycles	(Symantec Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic.
dmp_restore_interval	(Symantec Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths.
dmp_restore_policy	(Symantec Dynamic Multi-Pathing) The policy used by DMP path restoration thread.
dmp_restore_state	(Symantec Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started.
dmp_retry_count	(Symantec Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed.
dmp_scsi_timeout	(Symantec Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP.
dmp_sfg_threshold	(Symantec Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature.
dmp_stat_interval	(Symantec Dynamic Multi-Pathing) The time interval between gathering DMP statistics.
fssmartmovethreshold	(Veritas Volume Manager) The file system usage threshold for SmartMove (percent). This tunable must be set after Veritas Volume Manager is started.

Table C-1 Supported tunable parameters *(continued)*

Tunable	Description
max_diskq	(Veritas File System) Specifies the maximum disk queue generated by a single file. The installer can only set the system default value of max_diskq. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_ahead	(Veritas File System) The 0 value disables read ahead functionality, the 1 value (default) retains traditional sequential read ahead behavior, and the 2 value enables enhanced read ahead for all reads. The installer can only set the system default value of read_ahead. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_nstream	(Veritas File System) The number of parallel read requests of size read_pref_io that can be outstanding at one time. The installer can only set the system default value of read_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
read_pref_io	(Veritas File System) The preferred read request size. The installer can only set the system default value of read_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
reclaim_on_delete_start_time	(Veritas Volume Manager) Time of day to start reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.
reclaim_on_delete_wait_period	(Veritas Volume Manager) Days to wait before starting reclamation for deleted volumes. This tunable must be set after Veritas Volume Manager is started.

Table C-1 Supported tunable parameters *(continued)*

Tunable	Description
same_key_for_alldgs	(Veritas Volume Manager) Use the same fencing key for all disk groups. This tunable must be set after Veritas Volume Manager is started.
sharedminorstart	(Veritas Volume Manager) Start of range to use for minor numbers for shared disk groups. This tunable must be set after Veritas Volume Manager is started.
storage_connectivity	(Veritas Volume Manager) The CVM storage connectivity type. This tunable must be set after Veritas Volume Manager is started.
usefssmartmove	(Veritas Volume Manager) Configure SmartMove feature (all, thinonly, none). This tunable must be set after Veritas Volume Manager is started.
vol_checkpt_default	(Veritas File System) Size of VxVM storage checkpoints (kBytes). This tunable requires a system reboot to take effect.
vol_cmpres_enabled	(Veritas Volume Manager) Allow enabling compression for Volume Replicator.
vol_cmpres_threads	(Veritas Volume Manager) Maximum number of compression threads for Volume Replicator.
vol_default_iodelay	(Veritas Volume Manager) Time to pause between I/O requests from VxVM utilities (10ms units). This tunable requires a system reboot to take effect.
vol_fmr_logsz	(Veritas Volume Manager) Maximum size of bitmap Fast Mirror Resync uses to track changed blocks (KBytes). This tunable requires a system reboot to take effect.
vol_max_adminio_poolsz	(Veritas Volume Manager) Maximum amount of memory used by VxVM admin I/O's (bytes). This tunable requires a system reboot to take effect.
vol_max_nmpool_sz	(Veritas Volume Manager) Maximum name pool size (bytes).

Table C-1 Supported tunable parameters *(continued)*

Tunable	Description
vol_max_rdback_sz	(Veritas Volume Manager) Storage Record readback pool maximum (bytes).
vol_max_wrspool_sz	(Veritas Volume Manager) Maximum memory used in clustered version of Volume Replicator .
vol_maxio	(Veritas Volume Manager) Maximum size of logical VxVM I/O operations (kBytes). This tunable requires a system reboot to take effect.
vol_maxioctl	(Veritas Volume Manager) Maximum size of data passed into the VxVM ioctl calls (bytes). This tunable requires a system reboot to take effect.
vol_maxparallelio	(Veritas Volume Manager) Number of I/O operations vxconfigd can request at one time. This tunable requires a system reboot to take effect.
vol_maxspecialio	(Veritas Volume Manager) Maximum size of a VxVM I/O operation issued by an ioctl call (kBytes). This tunable requires a system reboot to take effect.
vol_min_lowmem_sz	(Veritas Volume Manager) Low water mark for memory (bytes).
vol_nm_hb_timeout	(Veritas Volume Manager) Volume Replicator timeout value (ticks).
vol_rvio_maxpool_sz	(Veritas Volume Manager) Maximum memory requested by Volume Replicator (bytes).
vol_stats_enable	(Veritas Volume Manager) Enable VxVM I/O stat collection.
vol_subdisk_num	(Veritas Volume Manager) Maximum number of subdisks attached to a single VxVM plex. This tunable requires a system reboot to take effect.
voldrl_max_drtregs	(Veritas Volume Manager) Maximum number of dirty VxVM regions. This tunable requires a system reboot to take effect.

Table C-1 Supported tunable parameters *(continued)*

Tunable	Description
voldrl_max_seq_dirty	(Veritas Volume Manager) Maximum number of dirty regions in sequential mode. This tunable requires a system reboot to take effect.
voldrl_min_regionsz	(Veritas Volume Manager) Minimum size of a VxVM Dirty Region Logging (DRL) region (kBytes). This tunable requires a system reboot to take effect.
voldrl_volumemax_drtregs	(Veritas Volume Manager) Max per volume dirty regions in log-plex DRL.
voldrl_volumemax_drtregs_20	(Veritas Volume Manager) Max per volume dirty regions in DCO version 20.
voldrl_dirty_regions	(Veritas Volume Manager) Number of regions cached for DCO version 30.
voliomem_chunk_size	(Veritas Volume Manager) Size of VxVM memory allocation requests (bytes). This tunable requires a system reboot to take effect.
voliomem_maxpool_sz	(Veritas Volume Manager) Maximum amount of memory used by VxVM (bytes). This tunable requires a system reboot to take effect.
voliot_errbuf_dflt	(Veritas Volume Manager) Size of a VxVM error trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_default	(Veritas Volume Manager) Default size of a VxVM I/O trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_limit	(Veritas Volume Manager) Maximum total size of all VxVM I/O trace buffers (bytes). This tunable requires a system reboot to take effect.
voliot_iobuf_max	(Veritas Volume Manager) Maximum size of a VxVM I/O trace buffer (bytes). This tunable requires a system reboot to take effect.
voliot_max_open	(Veritas Volume Manager) Maximum number of VxVM trace channels available for vxtrace commands. This tunable requires a system reboot to take effect.

Table C-1 Supported tunable parameters *(continued)*

Tunable	Description
volpagemod_max_memsz	(Veritas Volume Manager) Maximum paging module memory used by Instant Snapshots (Kbytes).
volraid_rsrtransmax	(Veritas Volume Manager) Maximum number of VxVM RAID-5 transient reconstruct operations in parallel. This tunable requires a system reboot to take effect.
vxfs_mbuf	(Veritas File System) Maximum memory used for VxFS buffer cache. This tunable requires a system reboot to take effect.
vxfs_ninode	(Veritas File System) Number of entries in the VxFS inode table. This tunable requires a system reboot to take effect.
write_nstream	(Veritas File System) The number of parallel write requests of size write_pref_io that can be outstanding at one time. The installer can only set the system default value of write_nstream. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.
write_pref_io	(Veritas File System) The preferred write request size. The installer can only set the system default value of write_pref_io. Refer to the tunefstab(4) manual page for setting this tunable for a specified block device.

Appendix D

Configuration files

This appendix includes the following topics:

- About the LLT and GAB configuration files
- About the AMF configuration files
- About the VCS configuration files
- About I/O fencing configuration files
- Sample configuration files for CP server

About the LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires /etc/llthosts and /etc/lltab files. GAB requires /etc/gabtab file.

Table D-1 lists the LLT configuration files and the information that these files contain.

Table D-1 LLT configuration files

File	Description

Table D-1 LLT configuration files *(continued)*

File	Description
/etc/sysconfig/llt	This file stores the start and stop environment variables for LLT: ■ LLT_START—Defines the startup behavior for the LLT module after a system reboot. Valid values include: 1—Indicates that LLT is enabled to start up. 0—Indicates that LLT is disabled to start up. ■ LLT_STOP—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include: 1—Indicates that LLT is enabled to shut down. 0—Indicates that LLT is disabled to shut down. The installer sets the value of these variables to 1 at the end of SFHA configuration. If you manually configured VCS, make sure you set the values of these environment variables to 1. Assign the buffer pool memory for RDMA operations: ■ LLT_BUFPOOL_MAXMEM—Maximum assigned memory that LLT can use for the LLT buffer pool. This buffer pool is used to allocate memory for RDMA operations and packet allocation, which are delivered to the LLT clients. The default value is calculated based on the total system memory, the minimum value is 1GB, and the maximum value is 10GB. You must specify the value in GB.

/etc/llthosts	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 sys1`
	`1 sys2`

Table D-1 LLT configuration files *(continued)*

File	Description

/etc/llttab	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 sys1`
	`set-cluster 2`
	`link eth1 eth1 - ether - -`
	`link eth2 eth2 - ether - -`
	
	If you use aggregated interfaces, then the file contains the aggregated interface name instead of the eth-*MAC_address*.
	
	`set-node sys1`
	`set-cluster 2`
	`link eth1 eth-00:04:23:AC:12:C4 - ether - -`
	`link eth2 eth-00:04:23:AC:12:C5 - 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-lowpri" 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.

Table D-2 lists the GAB configuration files and the information that these files contain.

Table D-2 GAB configuration files

File	Description
/etc/sysconfig/gab	This file stores the start and stop environment variables for GAB: ■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: 1—Indicates that GAB is enabled to start up. 0—Indicates that GAB is disabled to start up. ■ GAB_STOP—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include: 1—Indicates that GAB is enabled to shut down. 0—Indicates that GAB is disabled to shut down. The installer sets the value of these variables to 1 at the end of SFHA configuration.
/etc/gabtab	After you install SFHA, the file /etc/gabtab contains a `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. Use the `-c` option for `/sbin/gabconfig` to avoid a split-brain condition. **Note:**

About the AMF configuration files

Asynchronous Monitoring Framework (AMF) kernel driver provides asynchronous event notifications to the VCS agents that are enabled for intelligent resource monitoring.

Table D-3 lists the AMF configuration files.

Table D-3 AMF configuration files

File	Description
/etc/sysconfig/amf	This file stores the start and stop environment variables for AMF: ■ AMF_START—Defines the startup behavior for the AMF module after a system reboot or when AMF is attempted to start using the init script. Valid values include: 1—Indicates that AMF is enabled to start up. (default) 0—Indicates that AMF is disabled to start up. ■ AMF_STOP—Defines the shutdown behavior for the AMF module during a system shutdown or when AMF is attempted to stop using the init script. Valid values include: 1—Indicates that AMF is enabled to shut down. (default) 0—Indicates that AMF is disabled to shut down.
/etc/amftab	After you install VCS, the file `/etc/amftab` contains a `amfconfig`(1) command that configures the AMF driver for use. The AMF init script uses this `/etc/amftab` file to configure the AMF driver. The `/etc/amftab` file contains the following line by default: `/opt/VRTSamf/bin/amfconfig -c`

About the VCS configuration files

VCS configuration files include the following:

- main.cf
  The installer creates the VCS configuration file in the /etc/VRTSvcs/conf/config folder by default during the SFHA configuration. The main.cf file contains the minimum information that defines the cluster and its nodes.
  See "Sample main.cf file for VCS clusters" on page 478.
  See "Sample main.cf file for global clusters" on page 480.
- types.cf
  The file types.cf, which is listed in the include statement in the main.cf file, defines the VCS bundled types for VCS resources. The file types.cf is also located in the folder /etc/VRTSvcs/conf/config.
  Additional files similar to types.cf may be present if agents have been added, such as OracleTypes.cf.

Note the following information about the VCS configuration file after installing and configuring VCS:

- The cluster definition includes the cluster information that you provided during the configuration. This definition includes the cluster name, cluster address, and the names of users and administrators of the cluster.

  Notice that the cluster has an attribute UserNames. The installsfha creates a user "admin" whose password is encrypted; the word "password" is the default password.
- If you set up the optional I/O fencing feature for VCS, then the UseFence = SCSI3 attribute is present.
- If you configured the cluster in secure mode, the main.cf includes "SecureClus = 1" cluster attribute.
- The installsfha creates the ClusterService service group if you configured the virtual IP, SMTP, SNMP, or global cluster options.

  The service group also has the following characteristics:

  - The group includes the IP and NIC resources.
  - The service group also includes the notifier resource configuration, which is based on your input to installsfha prompts about notification.
  - The installsfha also creates a resource dependency tree.
  - If you set up global clusters, the ClusterService service group contains an Application resource, wac (wide-area connector). This resource's attributes contain definitions for controlling the cluster in a global cluster environment. Refer to the *Symantec Cluster Server Administrator's Guide* for information about managing VCS global clusters.

Refer to the *Symantec Cluster Server Administrator's Guide* to review the configuration concepts, and descriptions of main.cf and types.cf files for Linux systems.

Sample main.cf file for VCS clusters

The following sample main.cf file is for a cluster in secure mode.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"
include "Db2udbTypes.cf"
include "SybaseTypes.cf"


cluster vcs_cluster2 (
    UserNames = { admin = cDRpdxPmHpzS, smith = dKLhKJkHLh }
```

```
    ClusterAddress = "192.168.1.16"
    Administrators = { admin, smith }
    CounterInterval = 5
    SecureClus = 1
)

    system sys1 (
   )

    system sys2 (
   )

    group ClusterService (
       SystemList = { sys1 = 0, sys2 = 1 }
       UserStrGlobal = "LocalCluster@https://10.182.2.76:8443;"
       AutoStartList = { sys1, sys2 }
       OnlineRetryLimit = 3
       OnlineRetryInterval = 120
       )

    IP webip (
       Device = eth0
       Address = "192.168.1.16"
       NetMask = "255.255.240.0"
       )

NIC csgnic (
       Device = eth0
       NetworkHosts = { "192.168.1.17", "192.168.1.18" }
       )

NotifierMngr ntfr (
   SnmpConsoles = { "sys5" = Error, "sys4" = SevereError }
   SmtpServer = "smtp.example.com"
   SmtpRecipients =  { "ozzie@example.com" = Warning,
                "harriet@example.com" = Error }
   )

webip requires csgnic
   ntfr requires csgnic

// resource dependency tree
//
//     group ClusterService
```

```
//     {
//     NotifierMngr ntfr
//         {
//         NIC csgnic
//         }
// }
```

Sample main.cf file for global clusters

If you installed SFHA with the Global Cluster option, note that the ClusterService group also contains the Application resource, wac. The wac resource is required to control the cluster in a global cluster environment.

```
    .
    .
    group ClusterService (
       SystemList = { sys1 = 0, sys2 = 1 }

       UserStrGlobal = "LocalCluster@https://10.182.2.78:8443;"

       AutoStartList = { sys1, sys2 }
       OnlineRetryLimit = 3
       OnlineRetryInterval = 120
       )

       Application wac (
          StartProgram = "/opt/VRTSvcs/bin/wacstart"
          StopProgram = "/opt/VRTSvcs/bin/wacstop"
          MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
          RestartLimit = 3
          )
    .
    .
```

In the following main.cf file example, bold text highlights global cluster specific entries.

```
include "types.cf"

cluster vcs03 (
    ClusterAddress = "10.182.13.50"
    SecureClus = 1
    )

system sysA (
```

```
    )

system sysB (
    )

system sysC (
    )

group ClusterService (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
    )

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart -secure"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac -secure" }
    RestartLimit = 3
    )

IP gcoip (
    Device = eth0
    Address = "10.182.13.50"
    NetMask = "255.255.240.0"
    )

NIC csgnic (
    Device = eth0
    NetworkHosts = { "10.182.13.1" }
    )

NotifierMngr ntfr (
   SnmpConsoles = { sys4 = SevereError }
   SmtpServer = "smtp.example.com"
   SmtpRecipients =  { "ozzie@example.com" = SevereError }
   )

gcoip requires csgnic
ntfr requires csgnic
wac requires gcoip
```

```
// resource dependency tree
//
//      group ClusterService
//      {
//      NotifierMngr ntfr
//          {
//          NIC csgnic
//          }
//      Application wac
//          {
//          IP gcoip
//              {
//              NIC csgnic
//              }
//          }
//      }
```

About I/O fencing configuration files

Table D-4 lists the I/O fencing configuration files.

Table D-4 I/O fencing configuration files

File	Description
/etc/sysconfig/vxfen	This file stores the start and stop environment variables for I/O fencing: ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. The installer sets the value of these variables to 1 at the end of SFHA configuration.
/etc/vxfendg	This file includes the coordinator disk group information. This file is not applicable for server-based fencing and majority-based fencing.

Table D-4 I/O fencing configuration files *(continued)*

File	Description
/etc/vxfenmode	This file contains the following parameters: ■ vxfen_mode ■ scsi3—For disk-based fencing. ■ customized—For server-based fencing. ■ disabled—To run the I/O fencing driver but not do any fencing operations. ■ majority— For fencing without the use of coordination points. ■ vxfen_mechanism This parameter is applicable only for server-based fencing. Set the value as cps. ■ scsi3_disk_policy ■ dmp—Configure the vxfen module to use DMP devices The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp. **Note:** You must use the same SCSI-3 disk policy on all the nodes. ■ List of coordination points This list is required only for server-based fencing configuration. Coordination points in server-based fencing can include coordinator disks, CP servers, or both. If you use coordinator disks, you must create a coordinator disk group containing the individual coordinator disks. Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points and multiple IP addresses for each CP server. ■ single_cp This parameter is applicable for server-based fencing which uses a single highly available CP server as its coordination point. Also applicable for when you use a coordinator disk group with single disk. ■ autoseed_gab_timeout This parameter enables GAB automatic seeding of the cluster even when some cluster nodes are unavailable. This feature is applicable for I/O fencing in SCSI3 and customized mode. 0—Turns the GAB auto-seed feature on. Any value greater than 0 indicates the number of seconds that GAB must delay before it automatically seeds the cluster. -1—Turns the GAB auto-seed feature off. This setting is the default.

Table D-4 I/O fencing configuration files *(continued)*

File	Description

/etc/vxfentab	When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.
	Note: The /etc/vxfentab file is a generated file; do not modify this file.
	For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk along with its unique disk identifier. A space separates the path and the unique disk identifier. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:
	■ DMP disk:
	`/dev/vx/rdmp/sdx3 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006804E795D0A3`
	`/dev/vx/rdmp/sdy3 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006814E795D0B3`
	`/dev/vx/rdmp/sdz3 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006824E795D0C3`
	For server-based fencing, the /etc/vxfentab file also includes the security settings information.
	For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.
	This file is not applicable for majority-based fencing.

Sample configuration files for CP server

The `/etc/vxcps.conf` file determines the configuration of the coordination point server (CP server.)

See "Sample CP server configuration (/etc/vxcps.conf) file output" on page 490.

The following are example main.cf files for a CP server that is hosted on a single node, and a CP server that is hosted on an SFHA cluster.

- The main.cf file for a CP server that is hosted on a single node:
  See "Sample main.cf file for CP server hosted on a single node that runs VCS" on page 485.
- The main.cf file for a CP server that is hosted on an SFHA cluster:
  See "Sample main.cf file for CP server hosted on a two-node SFHA cluster" on page 487.

Note: If you use IPM-based protocol for communication between the CP server and SFHA clusters (application clusters), the CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses). If you use HTTPS-based protocol for communication, the CP server only supports Internet Protocol version 4 (IPv4 addresses).

The example main.cf files use IPv4 addresses.

Sample main.cf file for CP server hosted on a single node that runs VCS

The following is an example of a single CP server node main.cf.

For this CP server single node main.cf, note the following values:

- Cluster name: cps1
- Node name: cps1

```
include "types.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"

// cluster name:  cps1
// CP server: cps1

cluster cps1 (
     UserNames = { admin = bMNfMHmJNiNNlVNhMK, haris = fopKojNvpHouNn,
             "cps1.symantecexample.com@root@vx" = aj,
             "root@cps1.symantecexample.com" = hq }
     Administrators = { admin, haris,
             "cps1.symantecexample.com@root@vx",
             "root@cps1.symantecexample.com" }
     SecureClus = 1
     HacliUserLevel = COMMANDROOT
     )

system cps1 (
     )

group CPSSG (
      SystemList = { cps1 = 0 }
      AutoStartList = { cps1 }
      )
```

```
        IP cpsvip1 (
              Critical = 0
              Device @cps1 = eth0
              Address = "10.209.3.1"
              NetMask = "255.255.252.0"
              )

        IP cpsvip2 (
              Critical = 0
              Device @cps1 = eth1
              Address = "10.209.3.2"
              NetMask = "255.255.252.0"
              )

        NIC cpsnic1 (
             Critical = 0
             Device @cps1 = eth0
             PingOptimize = 0
             NetworkHosts @cps1 = { "10.209.3.10" }
             )

        NIC cpsnic2 (
             Critical = 0
             Device @cps1 = eth1
             PingOptimize = 0
             )

        Process vxcpserv (
             PathName = "/opt/VRTScps/bin/vxcpserv"
             ConfInterval = 30
             RestartLimit = 3
             )

        Quorum quorum (
              QuorumResources = { cpsvip1, cpsvip2 }
              )

    cpsvip1 requires cpsnic1
    cpsvip2 requires cpsnic2
    vxcpserv requires quorum


    // resource dependency tree
```

```
//
// group CPSSG
// {
// IP cpsvip1
//     {
//     NIC cpsnic1
//     }
// IP cpsvip2
//     {
//     NIC cpsnic2
//     }
// Process vxcpserv
//     {
//     Quorum quorum
//     }
// }
```

Sample main.cf file for CP server hosted on a two-node SFHA cluster

The following is an example of a main.cf, where the CP server is hosted on an SFHA cluster.

For this CP server hosted on an SFHA cluster main.cf, note the following values:

- Cluster name: cps1
- Nodes in the cluster: cps1, cps2

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"


// cluster: cps1
// CP servers:
// cps1
// cps2

cluster cps1 (
     UserNames = { admin = ajkCjeJgkFkkIskEjh,
            "cps1.symantecexample.com@root@vx" = JK,
            "cps2.symantecexample.com@root@vx" = dl }
     Administrators = { admin, "cps1.symantecexample.com@root@vx",
            "cps2.symantecexample.com@root@vx" }
```

```
        SecureClus = 1
        )

system cps1 (
        )

system cps2 (
        )

group CPSSG (
        SystemList = { cps1 = 0, cps2 = 1 }
        AutoStartList = { cps1, cps2 } )

        DiskGroup cpsdg (
                DiskGroup = cps_dg
                )

        IP cpsvip1 (
                Critical = 0
                Device @cps1 = eth0
                Device @cps2 = eth0
                Address = "10.209.81.88"
                NetMask = "255.255.252.0"
                )

        IP cpsvip2 (
                Critical = 0
                Device @cps1 = eth1
                Device @cps2 = eth1
                Address = "10.209.81.89"
                NetMask = "255.255.252.0"
                )

        Mount cpsmount (
                MountPoint = "/etc/VRTScps/db"
                BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
                FSType = vxfs
                FsckOpt = "-y"
                )

        NIC cpsnic1 (
                Critical = 0
                Device @cps1 = eth0
```

```
            Device @cps2 = eth0
            PingOptimize = 0
            NetworkHosts @cps1 = { "10.209.81.10 }
            )

      NIC cpsnic2 (
            Critical = 0
            Device @cps1 = eth1
            Device @cps2 = eth1
            PingOptimize = 0
            )

      Process vxcpserv (
             PathName = "/opt/VRTScps/bin/vxcpserv"
             )

      Quorum quorum (
             QuorumResources = { cpsvip1, cpsvip2 }
             )

      Volume cpsvol (
             Volume = cps_volume
             DiskGroup = cps_dg
             )

cpsmount requires cpsvol
cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
cpsvol requires cpsdg
vxcpserv requires cpsmount
vxcpserv requires quorum


// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1
//     {
//     NIC cpsnic1
//     }
// IP cpsvip2
//     {
```

```
//     NIC cpsnic2
//     }
// Process vxcpserv
//     {
//     Quorum quorum
//     Mount cpsmount
//         {
//         Volume cpsvol
//             {
//             DiskGroup cpsdg
//             }
//         }
//     }
// }
```

Sample CP server configuration (/etc/vxcps.conf) file output

The following is an example of a coordination point server (CP server) configuration file `/etc/vxcps.conf` output.

```
##  The vxcps.conf file determines the
## configuration for Veritas CP Server.
cps_name=cps1
vip=[10.209.81.88]
vip=[10.209.81.89]:56789
vip_https=[10.209.81.88]:55443
vip_https=[10.209.81.89]
port=14250
port_https=443
security=1
db=/etc/VRTScps/db
ssl_conf_file=/etc/vxcps_ssl.properties
```

Appendix E

Configuring the secure shell or the remote shell for communications

This appendix includes the following topics:

- About configuring secure shell or remote shell communication modes before installing products
- Manually configuring passwordless ssh
- Setting up ssh and rsh connection using the installer -comsetup command
- Setting up ssh and rsh connection using the pwdutil.pl utility
- Restarting the ssh session
- Enabling rsh for Linux

About configuring secure shell or remote shell communication modes before installing products

Establishing communication between nodes is required to install Symantec software from a remote system, or to install and configure a cluster. The node from which the installer is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install Symantec software.

You can install products to remote systems using either secure shell (ssh) or remote shell (rsh). Symantec recommends that you use ssh as it is more secure than rsh.

You can set up ssh and rsh connections in many ways.

- You can manually set up the SSH and RSH connection with UNIX shell commands.
- You can run the `installer -comsetup` command to interactively set up SSH and RSH connection.
- You can run the password utility, `pwdutil.pl`.

This section contains an example of how to set up ssh password free communication. The example sets up ssh between a source system (sys1) that contains the installation directories, and a target system (sys2). This procedure also applies to multiple target systems.

Note: The script- and web-based installers support establishing passwordless communication for you.

Manually configuring passwordless ssh

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

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

Figure E-1 illustrates this procedure.

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

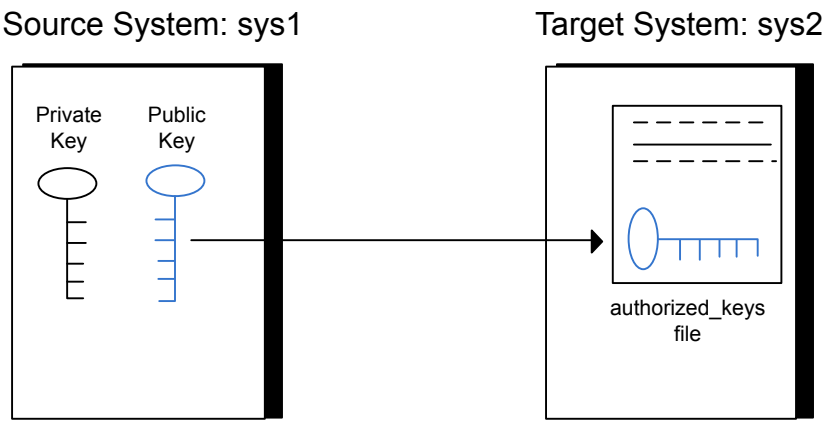

Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

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

To create the DSA key pair

1. On the source system (sys1), log in as root, and navigate to the root directory.

   ```
   sys1 # cd /root
   ```

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

   ```
   sys1 # ssh-keygen -t dsa
   ```

   System output similar to the following is displayed:

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

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

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

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

   Do not enter a passphrase. Press Enter.

   ```
   Enter same passphrase again:
   ```

   Press Enter again.

5. Output similar to the following lines appears.

   ```
   Your identification has been saved in /root/.ssh/id_dsa.
   Your public key has been saved in /root/.ssh/id_dsa.pub.
   The key fingerprint is:
   1f:00:e0:c2:9b:4e:29:b4:0b:6e:08:f8:50:de:48:d2 root@sys1
   ```

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 (sys1), move the public key to a temporary file on the target system (sys2).

   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:

   ```
   sys1 # sftp sys2
   ```

   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 sys2 ...
   The authenticity of host 'sys2 (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 'sys2,10.182.00.00'
   (DSA) to the list of known hosts.
   root@sys2 password:
   ```

3. Enter the root password of sys2.
4. At the `sftp` prompt, type the following command:

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

   The following output is displayed:

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

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

   ```
   sftp> quit
   ```

6 Add the `id_dsa.pub` keys to the `authorized_keys` file on the target system. To begin the `ssh` session on the target system (sys2 in this example), type the following command on sys1:

```
sys1 # ssh sys2
```

Enter the root password of sys2 at the prompt:

```
password:
```

Type the following commands on sys2:

```
sys2 # cat /root/id_dsa.pub >> /root/.ssh/authorized_keys
sys2 # rm  /root/id_dsa.pub
```

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

```
sys1 # exec /usr/bin/ssh-agent $SHELL
sys1 # ssh-add

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

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

To verify that you can connect to a target system

1 On the source system (sys1), enter the following command:

```
sys1 # ssh -l root sys2 uname -a
```

where sys2 is the name of the target system.

2 The command should execute from the source system (sys1) to the target system (sys2) without the system requesting a passphrase or password.

3 Repeat this procedure for each target system.

Setting up ssh and rsh connection using the installer -comsetup command

You can interactively set up the ssh and rsh connections using the `installer -comsetup` command.

Enter the following:

```
# ./installer -comsetup

Input the name of the systems to set up communication:
Enter the Solaris 10 Sparc system names separated by spaces:
[q,?] sys2
Set up communication for the system sys2:

  Checking communication on sys2 ..... Failed

CPI ERROR V-9-20-1303 ssh permission was denied on sys2. rsh
permission was denied on sys2. Either ssh or rsh is required
to be set up and ensure that it is working properly between the local
node and sys2 for communication

Either ssh or rsh needs to be set up between the local system and
sys2 for communication

Would you like the installer to setup ssh or rsh communication
automatically between the systems?
Superuser passwords for the systems will be asked. [y,n,q,?] (y) y

Enter the superuser password for system sys2:

     1)  Setup ssh between the systems
     2)  Setup rsh between the systems
     b)  Back to previous menu

Select the communication method [1-2,b,q,?] (1) 1

Setting up communication between systems. Please wait.
Re-verifying systems.

 Checking communication on sys2 ..... Done

Successfully set up communication for the system sys2
```

Setting up ssh and rsh connection using the pwdutil.pl utility

The password utility, `pwdutil.pl`, is bundled in the 6.2 release under the `scripts` directory. The users can run the utility in their script to set up the ssh and rsh connection automatically.

```
# ./pwdutil.pl -h
Usage:

Command syntax with simple format:

    pwdutil.pl check|configure|unconfigure ssh|rsh <hostname|IP addr>
    [<user>] [<password>] [<port>]

Command syntax with advanced format:

    pwdutil.pl [--action|-a 'check|configure|unconfigure']
               [--type|-t 'ssh|rsh']
               [--user|-u  '<user>']
               [--password|-p '<password>']
               [--port|-P '<port>']
               [--hostfile|-f '<hostfile>']
               [--keyfile|-k '<keyfile>']
               [-debug|-d]
               <host_URI>

    pwdutil.pl -h | -?
```

Table E-1 Options with pwdutil.pl utility

Option	Usage
--action|-a 'check|configure|unconfigure'	Specifies action type, default is 'check'.
--type|-t 'ssh|rsh'	Specifies connection type, default is 'ssh'.
--user|-u '<user>'	Specifies user id, default is the local user id.
--password|-p '<password>'	Specifies user password, default is the user id.
--port|-P '<port>'	Specifies port number for ssh connection, default is 22

Table E-1 Options with pwdutil.pl utility *(continued)*

Option	Usage
--keyfile|-k '<keyfile>'	Specifies the private key file.
--hostfile|-f '<hostfile>'	Specifies the file which list the hosts.
-debug	Prints debug information.
-h|-?	Prints help messages.
<host_URI>	Can be in the following formats: <hostname> <user>:<password>@<hostname> <user>:<password>@<hostname>: <port>

You can check, configure, and unconfigure ssh or rsh using the `pwdutil.pl`utility. For example:

- To check ssh connection for only one host:

  ```
  pwdutil.pl check ssh hostname
  ```

- To configure ssh for only one host:

  ```
  pwdutil.pl configure ssh hostname user password
  ```

- To unconfigure rsh for only one host:

  ```
  pwdutil.pl unconfigure rsh hostname
  ```

- To configure ssh for multiple hosts with same user ID and password:

  ```
  pwdutil.pl -a configure -t ssh -u user -p password hostname1
  hostname2 hostname3
  ```

- To configure ssh or rsh for different hosts with different user ID and password:

  ```
  pwdutil.pl -a configure -t ssh user1:password1@hostname1
  user2:password2@hostname2
  ```

- To check or configure ssh or rsh for multiple hosts with one configuration file:

  ```
  pwdutil.pl -a configure -t ssh --hostfile /tmp/sshrsh_hostfile
  ```

- To keep the host configuration file secret, you can use the 3rd party utility to encrypt and decrypt the host file with password.
  For example:

```
### run openssl to encrypt the host file in base64 format
# openssl aes-256-cbc -a -salt -in /hostfile -out /hostfile.enc
enter aes-256-cbc encryption password: <password>
Verifying - enter aes-256-cbc encryption password: <password>

### remove the original plain text file
# rm /hostfile

### run openssl to decrypt the encrypted host file
# pwdutil.pl -a configure -t ssh `openssl aes-256-cbc -d -a
-in /hostfile.enc`
enter aes-256-cbc decryption password: <password>
```

- To use the ssh authentication keys which are not under the default $HOME/.ssh directory, you can use --keyfile option to specify the ssh keys. For example:

```
### create a directory to host the key pairs:
# mkdir /keystore

### generate private and public key pair under the directory:
# ssh-keygen -t rsa -f /keystore/id_rsa

### setup ssh connection with the new generated key pair under
the directory:
# pwdutil.pl -a configure -t ssh --keyfile /keystore/id_rsa
user:password@hostname
```

You can see the contents of the configuration file by using the following command:

```
# cat /tmp/sshrsh_hostfile
user1:password1@hostname1
user2:password2@hostname2
user3:password3@hostname3
user4:password4@hostname4

# all default: check ssh connection with local user
hostname5
The following exit values are returned:

0     Successful completion.
```

```
1     Command syntax error.
2     Ssh or rsh binaries do not exist.
3     Ssh or rsh service is down on the remote machine.
4     Ssh or rsh command execution is denied due to password is required.
5     Invalid password is provided.
255   Other unknown error.
```

Restarting the ssh session

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

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

To restart ssh

1. On the source installation system (sys1), bring the private key into the shell environment.

   ```
   sys1 # exec /usr/bin/ssh-agent $SHELL
   ```

2. Make the key globally available for the user `root`

   ```
   sys1 # ssh-add
   ```

Enabling rsh for Linux

The following section describes how to enable remote shell.

Symantec recommends configuring a secure shell environment for Symantec product installations.

See "Manually configuring passwordless ssh" on page 492.

See the operating system documentation for more information on configuring remote shell.

To enable rsh for rhel6/sles

1. To ensure that the `rsh` and `rsh-server` RPMs are installed, type the following command:

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

   If it is not already in the file, type the following command to append the line "`rsh`" to the `/etc/securetty` file:

   ```
   # echo "rsh" >> /etc/securetty
   ```

2. Modify the line `disable = no` in the /etc/xinetd.d/rsh file.

3. In the `/etc/pam.d/rsh` file, change the "`auth`" type from "`required`" to "`sufficient`":

   ```
   auth     sufficient
   ```

4. Add the "promiscuous" flag into /etc/pam.d/rsh and /etc/pam.d/rlogin after item “pam_rhosts_auth.so”.

5. To enable the rsh server, type the following command:

   ```
   # chkconfig rsh on
   ```

6. Modify the `.rhosts` file. Each line of the `.rhosts` file contains a fully qualified domain name or IP address for each remote system. This file also contains the name of a user having access to the local system. For example, if the root user must remotely access `sys1` from `sys2`, add an entry for `sys2.`*`companyname`*`.com` to the `.rhosts` file on `sys1` by typing the following command:

   ```
   # echo "sys2.companyname.com" >> $HOME/.rhosts
   ```

7. Install the Symantec product.

To disable rsh for rhel6/sles

1. Remove the "`rsh`" entry in the `/etc/securetty` file.
2. Disable the rsh server by typing the following command:

```
# chkconfig rsh off
```

3. After you complete an installation procedure, delete the `.rhosts` file from each user's `$HOME` directory to ensure security:

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

To enable rsh for rhel7

- Run the following commands to enable rsh passwordless connection:

```
# systemctl start rsh.socket
# systemctl start rlogin.socket
# systemctl enable rsh.socket
# systemctl enable rlogin.socket
# echo rsh >> /etc/securetty
# echo rlogin >> /etc/securetty
#echo "+ +" >> /root/.rhosts
```

To disable rsh for rhel7

- Run the following commands to disable rsh passwordless connection:

```
# systemctl stop rsh.socket
# systemctl stop rlogin.socket
# systemctl disable rsh.socket
# systemctl disable rlogin.socket
```

Appendix F

Storage Foundation and High Availability components

This appendix includes the following topics:

- Storage Foundation and High Availability installation RPMs
- Symantec Cluster Server installation RPMs
- Symantec Storage Foundation obsolete and reorganized installation RPMs

Storage Foundation and High Availability installation RPMs

Table F-1 shows the RPM name and contents for each English language RPM for Storage Foundation and High Availability. The table also gives you guidelines for which RPMs to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and High Availability and Symantec Cluster Server (VCS) RPMs, the combined functionality is called Storage Foundation and High Availability and High Availability.

See “Symantec Cluster Server installation RPMs” on page 506.

Table F-1 Storage Foundation and High Availability RPMs

RPMs	Contents	Configuration
VRTSaslapm	Array Support Library (ASL) and Array Policy Module(APM) binaries Required for the support and compatibility of various storage arrays.	Minimum
VRTSperl	Perl 5.16.1 for Veritas	Minimum
VRTSvlic	Symantec License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Minimum
VRTSvxfs	Veritas File System binaries Required for VxFS file system support.	Minimum
VRTSvxvm	Veritas Volume Manager binaries, scripts, and utilities. Required for VxVM volume manager support.	Minimum
VRTSdbed	Storage Management Software for Databases	Recommended
VRTSob	Veritas Enterprise Administrator Service	Recommended
VRTSodm	Veritas Extension for Oracle Disk Manager Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle. Oracle Disk Manager enables Oracle to improve performance and manage system bandwidth.	Recommended

Table F-1 Storage Foundation and High Availability RPMs *(continued)*

RPMs	Contents	Configuration
VRTSsfcpi62	Symantec Storage Foundation Installer The Storage Foundation Common Product installer RPM contains the installer libraries and product scripts that perform the following: ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ etc. You can use these script to simplify the native operating system installations, configurations, and upgrades.	Minimum
VRTSsfmh	Veritas Operations Manager Managed Host. Discovers configuration information on a Storage Foundation managed host. If you want a central server to manage and monitor this managed host, download and install the VRTSsfmcs RPM on a server, and add this managed host to the Central Server. The VRTSsfmcs RPM is not part of this release. You can download it separately from: http://www.symantec.com/veritas-operations-manager	Recommended
VRTSspt	Veritas Software Support Tools	Recommended
VRTSvcsdr	Contains the binaries for Symantec Cluster Server disk reservation.	Recommended
VRTSfsadv	Veritas File System Advanced	Minimum

Table F-1 Storage Foundation and High Availability RPMs *(continued)*

RPMs	Contents	Configuration
VRTSfssdk	Veritas File System Software Developer Kit For VxFS APIs, the RPM contains the public Software Developer Kit (headers, libraries, and sample code). It is required if some user programs use VxFS APIs.	All

Symantec Cluster Server installation RPMs

Table F-2 shows the RPM name and contents for each English language RPM for Symantec Cluster Server (VCS). The table also gives you guidelines for which RPMs to install based whether you want the minimum, recommended, or advanced configuration.

When you install all Storage Foundation and VCS RPMs, the combined functionality is called Storage Foundation and High Availability.

See "Storage Foundation and High Availability installation RPMs" on page 503.

Table F-2 VCS installation RPMs

RPM	Contents	Configuration
VRTSgab	Symantec Cluster Server group membership and atomic broadcast services	Minimum
VRTSllt	Symantec Cluster Server low-latency transport	Minimum
VRTSamf	Symantec Cluster Server Asynchronous Monitoring Framework	Minimum
VRTSvcs	Symantec Cluster Server	Minimum
VRTSvcsag	Symantec Cluster Server Bundled Agents	Minimum
VRTSvxfen	Veritas I/O fencing	Minimum
VRTSvcsea	Consolidated database and enterprise agent RPMs	Recommended

Table F-2 VCS installation RPMs *(continued)*

RPM	Contents	Configuration
VRTScps	Veritas Coordination Point Server The Coordination Point Server is an alternate mechanism for I/O fencing. It implements I/O fencing through a client/server architecture and can provide I/O fencing for multiple VCS clusters.	Recommended

Symantec Storage Foundation obsolete and reorganized installation RPMs

Table F-3 lists the RPMs that are obsolete or reorganized for Storage Foundation and High Availability.

Table F-3 Symantec Storage Foundation obsolete and reorganized RPMs

RPM	Description
Obsolete and reorganized for 6.2	
VRTSat	Obsolete
VRTSatClient	Obsolete
VRTSatServer	Obsolete
Obsolete and reorganized for 5.1	
Infrastructure	
SYMClma	Obsolete
VRTSaa	Included in VRTSsfmh
VRTSccg	Included in VRTSsfmh
VRTSdbms3	Obsolete
VRTSicsco	Obsolete
VRTSjre	Obsolete
VRTSjre15	Obsolete

Table F-3 Symantec Storage Foundation obsolete and reorganized RPMs *(continued)*

RPM	Description
VRTSmh	Included in VRTSsfmh
VRTSobc33	Obsolete
VRTSobweb	Obsolete
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product RPMs	
VRTSacclib	Obsolete The following information is for installations, upgrades, and uninstallations using the script- or web-based installer. ■ For fresh installations VRTSacclib is not installed. ■ For upgrades, VRTSacclib is not uninstalled. ■ For uninstallation, VRTSacclib is not uninstalled.
VRTSalloc	Obsolete
VRTScmccc	Obsolete
VRTScmcm	Obsolete
VRTScmcs	Obsolete
VRTScscm	Obsolete
VRTScscw	Obsolete
VRTScsocw	Obsolete
VRTScssim	Obsolete
VRTScutil	Obsolete
VRTSd2gui-common	Included in VRTSdbed

Table F-3 Symantec Storage Foundation obsolete and reorganized RPMs *(continued)*

RPM	Description
VRTSdb2ed-common	Included in VRTSdbed
VRTSdbcom-common	Included in VRTSdbed
VRTSdbed-common	Included in VRTSdbed
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSlvmconv	Included in VRTSvxvm
VRTSdsa	Obsolete
VRTSfsman	Included in the product's main RPM.
VRTSfsmnd	Included in the product's main RPM.
VRTSfspro	Included in VRTSsfmh
VRTSmapro-common	Included in VRTSsfmh
VRTSodm-common	Included in VRTSodm
VRTSodm-platform	Included in VRTSodm
VRTSorgui-common	Obsolete
VRTSvcsdb	Included in VRTSvcsea
VRTSvcsmn	Included in VRTSvcs
VRTSvcsor	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvcs
VRTSvdid	Obsolete
VRTSvmman	Included in the product's main RPM.
VRTSvmpro	Included in VRTSsfmh
VRTSvrpro	Included in VRTSob
VRTSvrw	Obsolete
VRTSvxfs-common	Included in VRTSvxfs

Table F-3 Symantec Storage Foundation obsolete and reorganized RPMs *(continued)*

RPM	Description
VRTSvxfs-platform	Included in VRTSvxfs
VRTSvxmsa	Obsolete
VRTSvxvm-common	Included in VRTSvxvm
VRTSvxvm-platform	Included in VRTSvxvm

Appendix G

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
- Incorrect permissions for root on remote system
- Inaccessible system
- Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.2 (2184482)
- Troubleshooting the webinstaller

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 Symantec Storage
Foundation/Symantec 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. After you install the license key, you must validate the license key using the following command:

  ```
  # /opt/VRTS/bin/vxlicrep
  ```

- 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

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Failed to setup rsh communication on 10.198.89.241:
'rsh 10.198.89.241 <command>' failed
Trying to setup ssh communication on 10.198.89.241.
Failed to setup ssh communication on 10.198.89.241:
Login denied

Failed to login to remote system(s) 10.198.89.241.
Please make sure the password(s) are correct and superuser(root)
can login to the remote system(s) with the password(s).
If you want to setup rsh on remote system(s), please make sure
rsh with command argument ('rsh <host> <command>') is not
denied by remote system(s).
```

```
Either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication

Would you like the installer to setup ssh/rsh communication
automatically between the nodes?
Superuser passwords for the systems will be asked. [y,n,q] (y) n

System verification did not complete successfully

The following errors were discovered on the systems:

The ssh permission denied on 10.198.89.241
rsh exited 1 on 10.198.89.241
either ssh or rsh is needed to be setup between the local node
and 10.198.89.241 for communication
```

Suggested solution: You need to set up the systems to allow remote access using `ssh` or `rsh`.

See "About configuring secure shell or remote shell communication modes before installing products" on page 491.

Note: Remove remote shell permissions after completing the SFHA installation and configuration.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

Suggested solution: Verify that you entered the system name correctly; use the `ping`(1M) command to verify the accessibility of the host.

Upgrading Symantec Storage Foundation for Databases (SFDB) tools from 5.0.x to 6.2 (2184482)

The `sfua_rept_migrate`command results in an error message after upgrading SFHA or SF for Oracle RAC version 5.0 to SFHA or SF for Oracle RAC 6.2.

When upgrading from SFHA version 5.0 to SFHA 6.2 the S*vxdbms3 startup script is renamed to NO_S*vxdbms3. The S*vxdbms3 startup script is required by `sfua_rept_upgrade`. Thus when `sfua_rept_upgrade` is run, it is unable to find the S*vxdbms3 startup script and gives the error message:

```
/sbin/rc3.d/S*vxdbms3 not found
SFORA sfua_rept_migrate ERROR V-81-3558 File:  is missing.
SFORA sfua_rept_migrate ERROR V-81-9160 Failed to mount repository.
```

Workaround: Before running `sfua_rept_migrate`, rename the startup script NO_S*vxdbms3 to S*vxdbms3.

Troubleshooting the webinstaller

This section provides possible solutions to problems that may occur when using the `webinstaller` script:

- Issue: The `webinstaller` script may report an error.
  You may receive a similar error message when using the webinstaller:

  ```
  Error: could not get hostname and IP address
  ```

  Solution: Check whether `/etc/hosts` and `/etc/resolv.conf` file are correctly configured.

- Issue: The hostname is not a fully qualified domain name.
  You must have a fully qualified domain name for the hostname in https://<*hostname*>:<*port*>/.
  Solution: Check whether the `domain` section is defined in `/etc/resolv.conf` file.

- Issue: FireFox 3 may report an error.
  You may receive a similar error message when using FireFox 3:

  ```
  Certificate contains the same serial number as another certificate.
  ```

  Solution: Visit FireFox knowledge base website:

http://support.mozilla.com/en-US/kb/Certificate+contains+the+same+serial+number+as+another+certificate

Appendix H

Troubleshooting cluster installation

This appendix includes the following topics:

- Unmount failures
- Command failures
- Installer cannot create UUID for the cluster
- The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails
- Troubleshooting CP server
- Troubleshooting server-based fencing on the SFHA cluster nodes
- Issues during online migration of coordination points

Unmount failures

The `umount` command can fail if a reference is being held by an NFS server. Unshare the mount point and try the unmount again.

Command failures

This section describes command failures.

- Manual pages not accessible with the `man` command. Set the MANPATH environment variable appropriately.
  See "Setting environment variables" on page 63.

- The `mount`, `fsck`, and `mkfs` utilities reserve a shared volume. They fail on volumes that are in use. Be careful when accessing shared volumes with other utilities such as `dd`, it is possible for these commands to destroy data on the disk.
- Running some commands, such as `vxupgrade -n 7 /vol02`, can generate the following error message:

  ```
  vxfs vxupgrade: ERROR: not primary in a cluster file system
  ```

  This means that you can run this command only on the primary, that is, the system that mounted this file system first.

Installer cannot create UUID for the cluster

The installer displays the following error message if the installer cannot find the uuidconfig.pl script before it configures the UUID for the cluster:

```
Couldn't find uuidconfig.pl for uuid configuration,
please create uuid manually before start vcs
```

You may see the error message during SFHA configuration, upgrade, or when you add a node to the cluster using the installer.

Workaround: To start SFHA, you must run the uuidconfig.pl script manually to configure the UUID on each cluster node.

To configure the cluster UUID when you create a cluster manually

- On one node in the cluster, perform the following command to populate the cluster UUID on each node in the cluster.

  ```
  # /opt/VRTSvcs/bin/uuidconfig.pl -clus -configure nodeA
  nodeB ... nodeN
  ```

  Where nodeA, nodeB, through nodeN are the names of the cluster nodes.

The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails

While running the vxfentsthdw utility, you may see a message that resembles as follows:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node
FAILED.
```

```
Contact the storage provider to have the hardware configuration
fixed.
```

The disk array does not support returning success for a `SCSI TEST UNIT READY` command when another host has the disk reserved using SCSI-3 persistent reservations. This happens with the Hitachi Data Systems 99XX arrays if bit 186 of the system mode option is not enabled.

Troubleshooting CP server

All CP server operations and messages are logged in the /var/VRTScps/log directory in a detailed and easy to read format. The entries are sorted by date and time. The logs can be used for troubleshooting purposes or to review for any possible security issue on the system that hosts the CP server.

The following files contain logs and text files that may be useful in understanding and troubleshooting a CP server:

- /var/VRTScps/log/cpserver_[ABC].log
- /var/VRTSvcs/log/vcsauthserver.log (Security related)
- If the vxcpserv process fails on the CP server, then review the following diagnostic files:
  - /var/VRTScps/diag/FFDC_CPS_*pid*_vxcpserv.log
  - /var/VRTScps/diag/stack_*pid*_vxcpserv.txt

Note: If the vxcpserv process fails on the CP server, these files are present in addition to a core file. VCS restarts vxcpserv process automatically in such situations.

The file /var/VRTSvcs/log/vxfen/vxfend_[ABC].log contains logs that may be useful in understanding and troubleshooting fencing-related issues on a Storage Foundation High Availability (client cluster) node.

See “Troubleshooting issues related to the CP server service group” on page 518.

See “Checking the connectivity of CP server” on page 518.

See “Issues during fencing startup on Storage Foundation High Availability nodes set up for server-based fencing” on page 519.

See “Issues during online migration of coordination points” on page 520.

Troubleshooting issues related to the CP server service group

If you cannot bring up the CPSSG service group after the CP server configuration, perform the following steps:

- Verify that the CPSSG service group and its resources are valid and properly configured in the VCS configuration.
- Check the VCS engine log (`/var/VRTSvcs/log/engine_[ABC].log`) to see if any of the CPSSG service group resources are FAULTED.
- Review the sample dependency graphs to make sure the required resources are configured correctly.

Checking the connectivity of CP server

You can test the connectivity of CP server using the `cpsadm` command.

You must have set the environment variables CPS_USERNAME and CPS_DOMAINTYPE to run the `cpsadm` command on the Storage Foundation High Availability (client cluster) nodes.

To check the connectivity of CP server

- Run the following command to check whether a CP server is up and running at a process level:

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

  where *cp_server* is the virtual IP address or virtual hostname on which the CP server is listening.

Troubleshooting server-based fencing on the SFHA cluster nodes

The file /var/VRTSvcs/log/vxfen/vxfend_[ABC].log contains logs files that may be useful in understanding and troubleshooting fencing-related issues on a SFHA cluster (application cluster) node.

Issues during fencing startup on Storage Foundation High Availability nodes set up for server-based fencing

Table H-1 Fencing startup issues on Storage Foundation High Availability (client cluster) nodes

Issue	Description and resolution
`cpsadm` command on the Storage Foundation High Availability gives connection error	If you receive a connection error message after issuing the `cpsadm` command on the Storage Foundation High Availability, perform the following actions: ■ Ensure that the CP server is reachable from all the Storage Foundation High Availability nodes. ■ Check the /etc/vxfenmode file and ensure that the Storage Foundation High Availability nodes use the correct CP server virtual IP or virtual hostname and the correct port number. ■ For HTTPS communication, ensure that the virtual IP and ports listed for the server can listen to HTTPS requests.
Authorization failure	Authorization failure occurs when the nodes on the client clusters and or users are not added in the CP server configuration. Therefore, fencing on the Storage Foundation High Availability (client cluster) node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points. To resolve this issue, add the client cluster node and user in the CP server configuration and restart fencing. See "Preparing the CP servers manually for use by the Storage Foundation High Availability" on page 244.
Authentication failure	If you had configured secure communication between the CP server and the Storage Foundation High Availability (client cluster) nodes, authentication failure can occur due to the following causes: ■ The client cluster requires its own private key, a signed certificate, and a Certification Authority's (CA) certificate to establish secure communication with the CP server. If any of the files are missing or corrupt, communication fails. ■ If the client cluster certificate does not correspond to the client's private key, communication fails. ■ If the CP server and client cluster do not have a common CA in their certificate chain of trust, then communication fails.

Issues during online migration of coordination points

During online migration of coordination points using the vxfenswap utility, the operation is automatically rolled back if a failure is encountered during validation of coordination points from any of the cluster nodes.

Validation failure of the new set of coordination points can occur in the following circumstances:

- The /etc/vxfenmode.test file is not updated on all the Storage Foundation High Availability nodes, because new coordination points on the node were being picked up from an old /etc/vxfenmode.test file. The /etc/vxfenmode.test file must be updated with the current details. If the /etc/vxfenmode.test file is not present, vxfenswap copies configuration for new coordination points from the /etc/vxfenmode file.
- The coordination points listed in the /etc/vxfenmode file on the different Storage Foundation High Availability nodes are not the same. If different coordination points are listed in the /etc/vxfenmode file on the cluster nodes, then the operation fails due to failure during the coordination point snapshot check.
- There is no network connectivity from one or more Storage Foundation High Availability nodes to the CP server(s).
- Cluster, nodes, or users for the Storage Foundation High Availability nodes have not been added on the new CP servers, thereby causing authorization failure.

Vxfen service group activity after issuing the vxfenswap command

The Coordination Point agent reads the details of coordination points from the `vxfenconfig -l` output and starts monitoring the registrations on them.

Thus, during vxfenswap, when the vxfenmode file is being changed by the user, the Coordination Point agent does not move to FAULTED state but continues monitoring the old set of coordination points.

As long as the changes to vxfenmode file are not committed or the new set of coordination points are not reflected in `vxfenconfig -l` output, the Coordination Point agent continues monitoring the old set of coordination points it read from `vxfenconfig -l` output in every monitor cycle.

The status of the Coordination Point agent (either ONLINE or FAULTED) depends upon the accessibility of the coordination points, the registrations on these coordination points, and the fault tolerance value.

When the changes to vxfenmode file are committed and reflected in the `vxfenconfig -l` output, then the Coordination Point agent reads the new set of coordination points and proceeds to monitor them in its new monitor cycle.

Appendix I

Sample Storage Foundation High Availability setup diagrams for CP server-based I/O fencing

This appendix includes the following topics:

- Configuration diagrams for setting up server-based I/O fencing

Configuration diagrams for setting up server-based I/O fencing

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:

- Two unique client clusters that are served by 3 CP servers:
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
- Two node campus cluster that is served be remote CP server and 2 SCSI-3 disks:
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:

Two unique client clusters served by 3 CP servers

In the `vxfenmode` file on the client nodes, vxfenmode is set to `customized` with vxfen mechanism set to `cps`.

Client cluster served by highly available CPS and 2 SCSI-3 disks

Figure I-1 displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to customized with vxfen mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group vxfencoorddg. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure I-1 Client cluster served by highly available CP server and 2 SCSI-3 disks

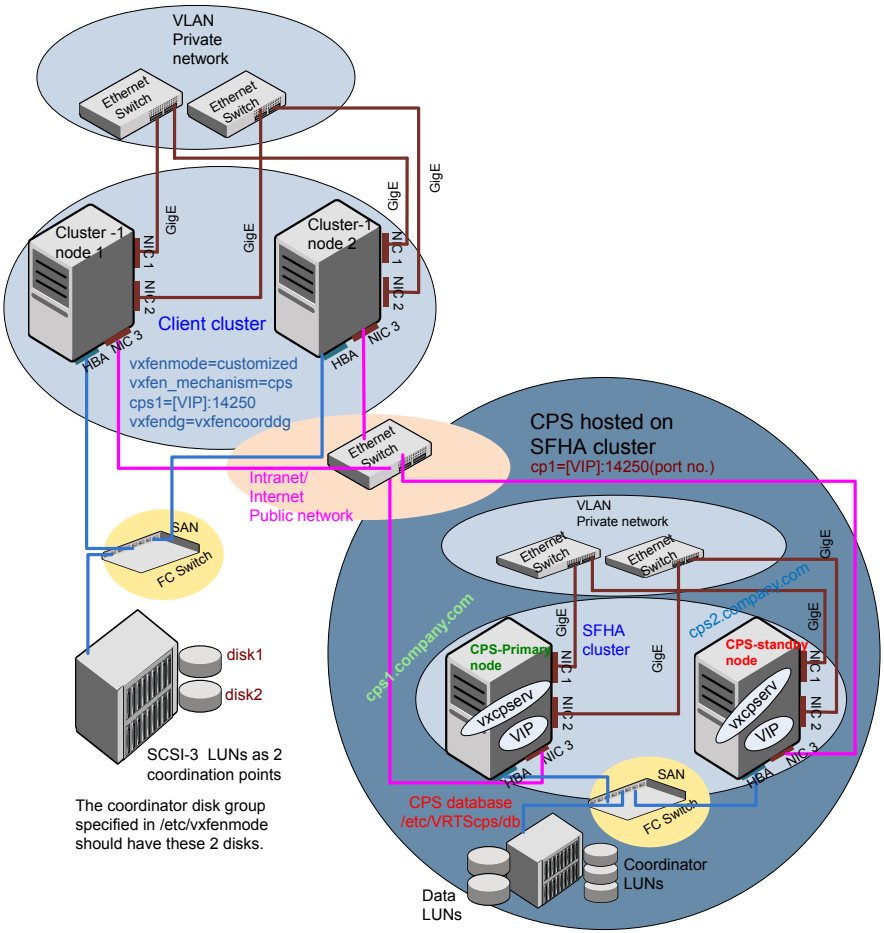

Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure I-2 displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, vxfenmode is set to `customized` with vxfen mechanism set to `cps`.

The two SCSI-3 disks (one from each site) are part of disk group vxfencoorddg. The third coordination point is a CP server on a single node VCS cluster.

Figure I-2 Two node campus cluster served by remote CP server and 2 SCSI-3

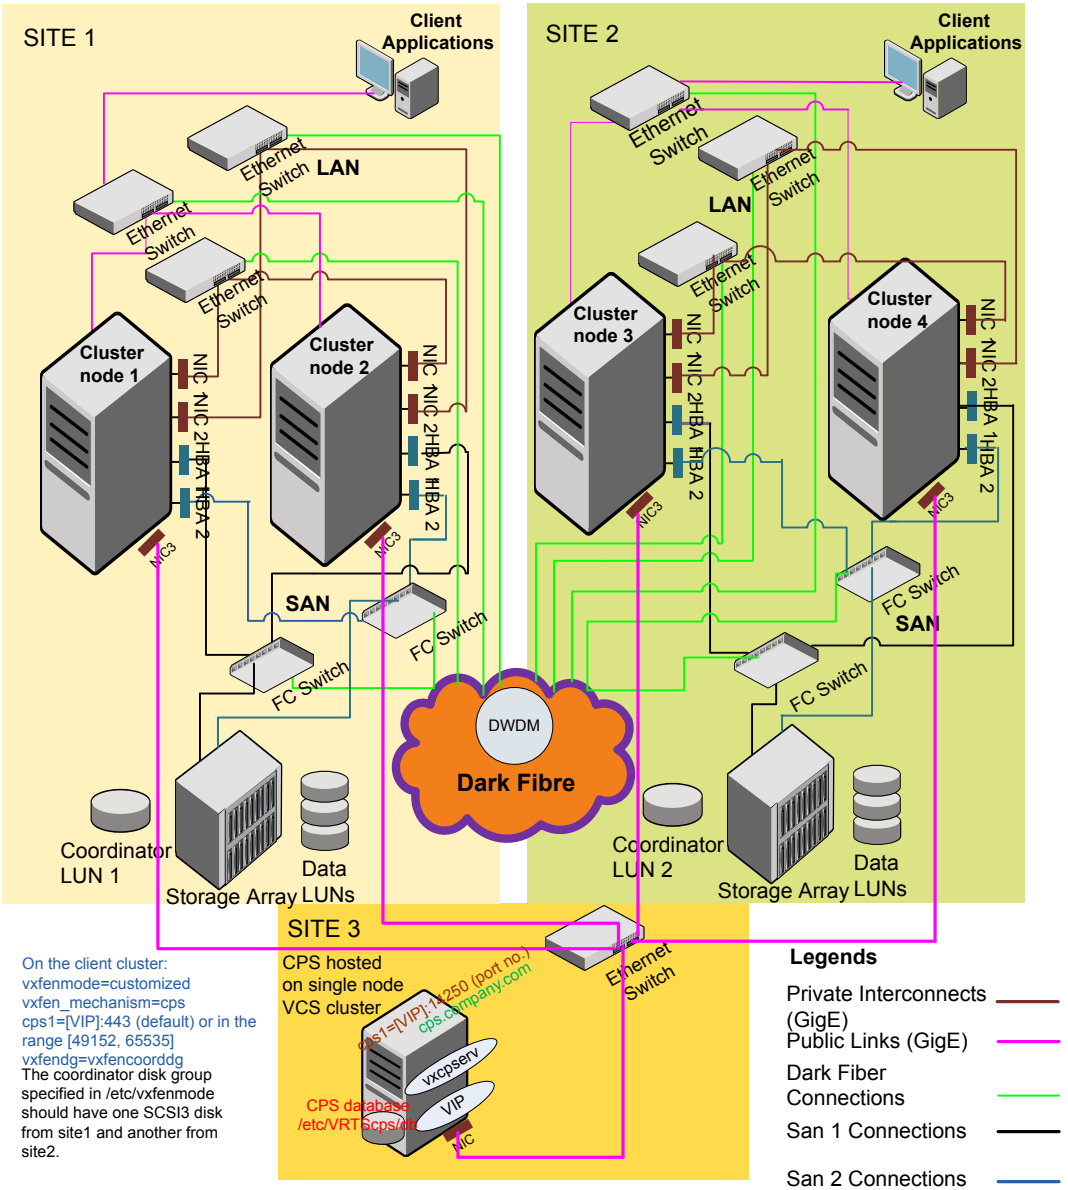

Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

In the `vxfenmode` file on the client nodes, vxfenmode is set to `customized` with vxfen mechanism set to `cps`.

The two SCSI-3 disks are are part of the disk group vxfencoorddg. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Appendix J

Configuring LLT over UDP

This appendix includes the following topics:

- Using the UDP layer for LLT
- Manually configuring LLT over UDP using IPv4
- Using the UDP layer of IPv6 for LLT
- Manually configuring LLT over UDP using IPv6

Using the UDP layer for LLT

SFHA provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

LLT over UDP is slower than LLT over Ethernet. Use LLT over UDP only when the hardware configuration makes it necessary.

Manually configuring LLT over UDP using IPv4

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on separate subnets. Set the broadcast address in /etc/lltab explicitly depending on the subnet for each link.

See "Broadcast address in the /etc/llttab file" on page 528.

- 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 530.
- Set the broadcast address correctly for direct-attached (non-routed) links.
  See "Sample configuration: direct-attached links" on page 531.
- 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 533.

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 sys1:

```
sys1 # cat /etc/llttab

set-node sys1
set-cluster 1
link link1 udp - udp  50000  -  192.168.9.1 192.168.9.255
link link2 udp - udp  50001  -  192.168.10.1 192.168.10.255
```

  Verify the subnet mask using the ifconfig command to ensure that the two links are on separate subnets.

- Display the content of the /etc/llttab file on the second node sys2:

```
sys2 # cat /etc/llttab

set-node sys2
set-cluster 1
link link1 udp - udp  50000  -  192.168.9.2 192.168.9.255
link link2 udp - udp  50001  -  192.168.10.2 192.168.10.255
```

  Verify the subnet mask using the ifconfig command to ensure that the two links are on separate subnets.

The link command in the /etc/lltab file

Review the link command information in this section for the /etc/lltab file. See the following information for sample configurations:

- See "Sample configuration: direct-attached links" on page 531.
- See "Sample configuration: links crossing IP routers" on page 533.

Table J-1 describes the fields of the link command that are shown in the /etc/lltab file examples. Note that some of the fields differ from the command for standard LLT links.

Table J-1 Field description for link command in /etc/lltab

Field	Description
tag-name	A unique string that is used as a tag by LLT; for example link1, link2,....
device	The device path of the UDP protocol; for example udp. A place holder string. On other unix platforms like Solaris or HP, this entry points to a device file (for example, /dev/udp). Linux does not have devices for protocols. So this field is ignored.
node-range	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
link-type	Type of link; must be "udp" for LLT over UDP.
udp-port	Unique UDP port in the range of 49152-65535 for the link. See "Selecting UDP ports" on page 530.
MTU	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the `lltstat -l` command to display the current value.
IP address	IP address of the link on the local node.
bcast-address	■ 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/lltab file

The `set-addr` command in the /etc/lltab 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 533.

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
node-id	The node ID of the peer node; for example, 0.
link tag-name	The string that LLT uses to identify the link; for example link1, link2,....
address	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
  - Ports from the range of well-known ports, 0 to 1023
  - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address         State
udp        0      0 *:32768                 *:*
udp        0      0 *:956                   *:*
udp        0      0 *:tftp                  *:*
udp        0      0 *:sunrpc                *:*
udp        0      0 *:ipp                   *:*
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the /etc/services file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# ifconfig interface_name netmask netmask
```

For example:

- For the first network interface on the node sys1:

  ```
  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 sys2:

  ```
  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 sys1:

  ```
  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 sys2:

  ```
  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/lltab depending on the subnet that the links are on.

An example of a typical /etc/lltab file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

```
# cat /etc/lltab
set-node nodexyz
set-cluster 100

link link1 udp - udp 50000 - 192.168.30.1 192.168.30.255
link link2 udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: direct-attached links

Figure 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

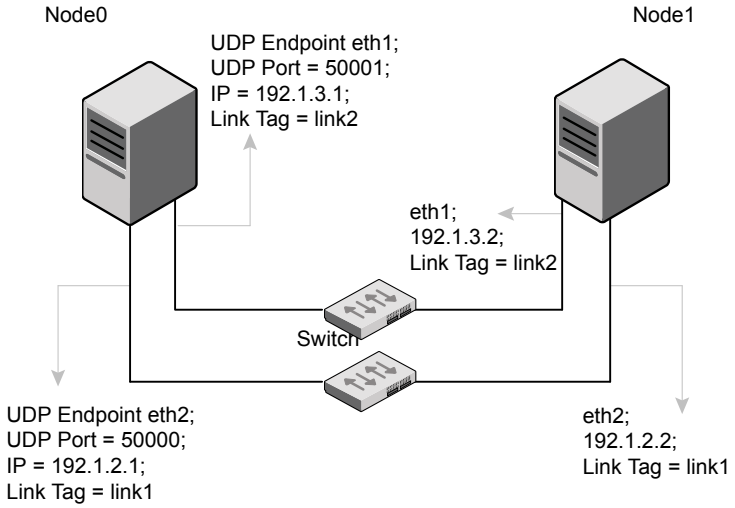

The configuration that the /etc/lltab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT sends broadcast requests to peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the /etc/lltab file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the /etc/lltab file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
```

```
link link1 udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

Figure 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

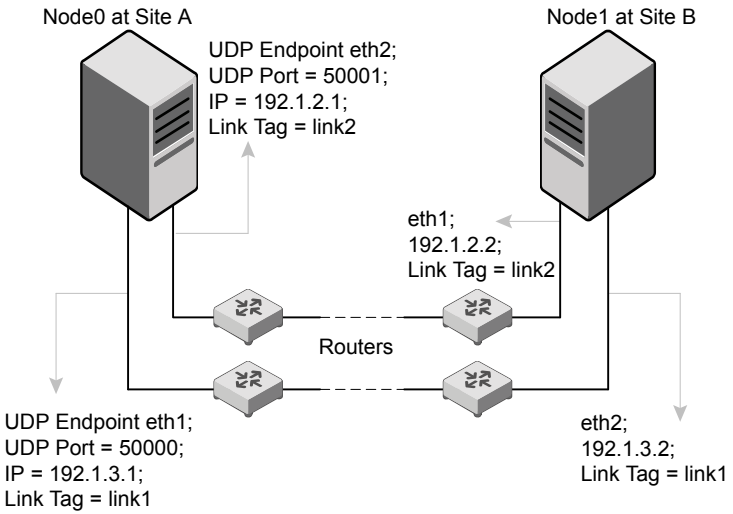

The configuration that the following `/etc/lltab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/lltab` file.

```
set-node Node1
set-cluster 1

link link1 udp - udp 50000 - 192.1.3.1 -
link link2 udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr        0 link1 192.1.1.1
set-addr        0 link2 192.1.2.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
```

```
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

The /etc/llttab file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 udp - udp 50000 - 192.1.1.1 -
link link2 udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr        1 link1 192.1.3.1
set-addr        1 link2 192.1.4.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

Using the UDP layer of IPv6 for LLT

Symantec Storage Foundation 6.2 provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

Manually configuring LLT over UDP using IPv6

The following checklist is to configure LLT over UDP:

- For UDP6, the multicast address is set to "-".
- Make sure that each NIC has an IPv6 address that is configured before configuring LLT.
- Make sure the IPv6 addresses in the /etc/llttab files are consistent with the IPv6 addresses of the network interfaces.
- Make sure that each link has a unique not well-known UDP port.
- For the links that cross an IP router, disable multicast features and specify the IPv6 address of each link manually in the /etc/llttab file.
  See "Sample configuration: links crossing IP routers" on page 536.

Sample configuration: direct-attached links

Figure J-3 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure J-3 A typical configuration of direct-attached links that use LLT over UDP

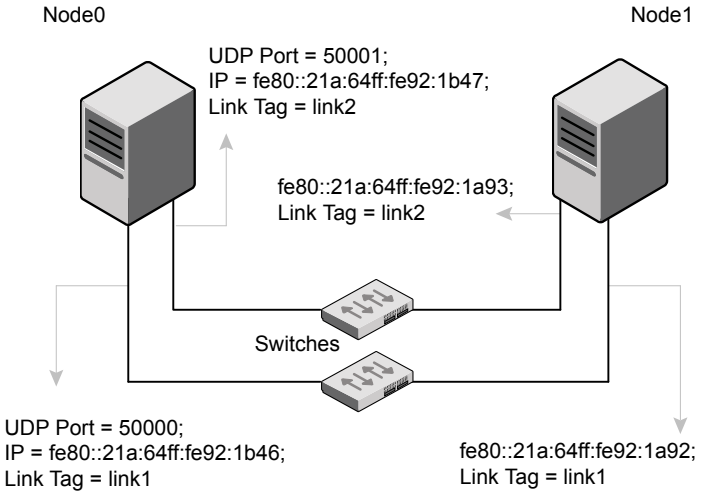

The configuration that the /etc/llttab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

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/lltab file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```

The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

Sample configuration: links crossing IP routers

Figure J-4 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-4 A typical configuration of links crossing an IP router

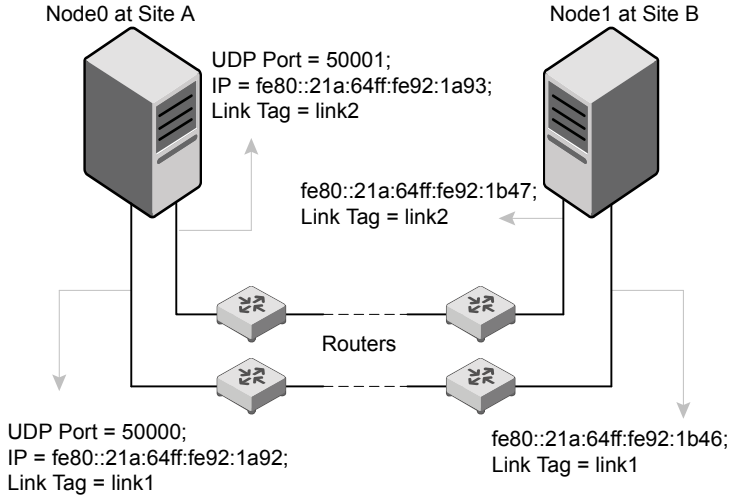

The configuration that the following `/etc/lltab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 0 link1 fe80::21a:64ff:fe92:1b46
set-addr 0 link2 fe80::21a:64ff:fe92:1b47
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb     0
set-arp         0
```

The /etc/llttab file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 1 link1 fe80::21a:64ff:fe92:1a92
set-addr 1 link2 fe80::21a:64ff:fe92:1a93
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb     0
set-arp         0
```

Appendix K

Using LLT over RDMA

This appendix includes the following topics:

- Using LLT over RDMA
- About RDMA over RoCE or InfiniBand networks in a clustering environment
- How LLT supports RDMA capability for faster interconnects between applications
- Using LLT over RDMA: supported use cases
- Configuring LLT over RDMA
- Troubleshooting LLT over RDMA

Using LLT over RDMA

This section describes how LLT works with RDMA, lists the hardware requirements for RDMA, and the procedure to configure LLT over RDMA.

About RDMA over RoCE or InfiniBand networks in a clustering environment

Remote direct memory access (RDMA) is a direct memory access capability that allows server to server data movement directly between application memories with minimal CPU involvement. Data transfer using RDMA needs RDMA-enabled network cards and switches. Networks designed with RDMA over Converged Ethernet (RoCE) and InfiniBand architecture support RDMA capability. RDMA provides fast interconnect between user-space applications or file systems between nodes over these networks. In a clustering environment, RDMA capability allows applications on separate nodes to transfer data at a faster rate with low latency and less CPU usage.

See "How LLT supports RDMA capability for faster interconnects between applications" on page 540.

How LLT supports RDMA capability for faster interconnects between applications

LLT and GAB support fast interconnect between applications using RDMA technology over InfiniBand and Ethernet media (RoCE). To leverage the RDMA capabilities of the hardware and also support the existing LLT functionalities, LLT maintains two channels (RDMA and non-RDMA) for each of the configured RDMA links. Both RDMA and non-RDMA channels are capable of transferring data between the nodes and LLT provides separate APIs to their clients, such as, CFS, CVM, to use these channels. The RDMA channel provides faster data transfer by leveraging the RDMA capabilities of the hardware. The RDMA channel is mainly used for data-transfer when the client is capable to use this channel. The non-RDMA channel is created over the UDP layer and LLT uses this channel mainly for sending and receiving heartbeats. Based on the health of the non-RDMA channel, GAB decides cluster membership for the cluster. The connection management of the RDMA channel is separate from the non-RDMA channel, but the connect and disconnect operations for the RDMA channel are triggered based on the status of the non-RDMA channel

If the non-RDMA channel is up but due to some issues in RDMA layer the RDMA channel is down, in such cases the data-transfer happens over the non-RDMA channel with a lesser performance until the RDMA channel is fixed. The system logs displays the message when the RDMA channel is up or down.

LLT uses the Open Fabrics Enterprise Distribution (OFED) layer and the drivers installed by the operating system to communicate with the hardware. LLT over RDMA allows applications running on one node to directly access the memory of an application running on another node that are connected over an RDMA-enabled network. In contrast, on nodes connected over a non-RDMA network, applications cannot directly read or write to an application running on another node. LLT clients such as, CFS and CVM, have to create intermediate copies of data before completing the read or write operation on the application, which increases the latency period and affects performance in some cases.

LLT over an RDMA network enables applications to read or write to applications on another node over the network without the need to create intermediate copies. This leads to low latency, higher throughput, and minimized CPU host usage thus improving application performance. Cluster volume manager and Cluster File Systems, which are clients of LLT and GAB, can use LLT over RDMA capability for specific use cases.

See “Using LLT over RDMA: supported use cases ” on page 541.

Using LLT over RDMA: supported use cases

You can configure the LLT over RDMA capability for the following use cases:

- Storage Foundation Smart IO feature on flash storage devices: The Smart IO feature provides file system caching on flash devices for increased application performance by reducing IO bottlenecks. It also reduces IO loads on storage controllers as the Smart IO feature meets most of the application IO needs. As the IO requirements from the storage array are much lesser, you require lesser number of servers to maintain the same IO throughput.
- Storage Foundation IO shipping feature: The IO shipping feature in Storage Foundation Cluster File System HA (SFCFSHA) provides the ability to ship IO data between applications on peer nodes without service interruption even if the IO path on one of the nodes in the cluster goes down.
- Storage Foundation Flexible Storage Sharing feature : The Flexible Storage Sharing feature in cluster volume manager allows network shared storage to co-exist with physically shared storage. It provides server administrators the ability to provision clusters for Storage Foundation Cluster File System HA (SFCFSHA) and Storage Foundation for Oracle RAC (SFRAC) or SFCFSHA applications without requiring physical shared storage.

Both Cluster File System (CFS) and Cluster Volume Manager (CVM) are clients of LLT and GAB. These clients use LLT as the transport protocol for data transfer between applications on nodes. Using LLT data transfer over an RDMA network boosts performance of file system data transfer and IO transfer between nodes.

To enable RDMA capability for faster application data transfer between nodes, you must install RDMA-capable network interface cards, RDMA-supported network switches, configure the operating system for RDMA, and configure LLT.

Ensure that you select RDMA-supported hardware and configure LLT to use RDMA functionality.

See “Choosing supported hardware for LLT over RDMA” on page 542.

See “Configuring LLT over RDMA” on page 541.

Configuring LLT over RDMA

This section describes the required hardware and configuration needed for LLT to support RDMA capability. The high-level steps to configure LLT over RDMA are as follows:

Table K-1 lists the high-level steps to configure LLT over RDMA.

Step	Action	Description
Choose supported hardware	Choose RDMA capable network interface cards (NICs), network switches, and cables.	See "Choosing supported hardware for LLT over RDMA" on page 542.
Check the supported operating system	Linux flavors only.	RHEL 6.3, RHEL 6.4, SUSE Linux Enterprise 11 SP2, SUSE Linux Enterprise 11 SP3, Oracle Linux 6.3, Oracle Linux 6.4
Install RDMA, InfiniBand or Ethernet drivers and utilities	Install the packages to access the RDMA, InfiniBand or Ethernet drivers and utilities.	See "Installing RDMA, InfiniBand or Ethernet drivers and utilities" on page 543.
Configure RDMA over an Ethernet network	Load RDMA and Ethernet drivers.	See "Configuring RDMA over an Ethernet network" on page 544.
Configuring RDMA over an InfiniBand network	Load RDMA and InfiniBand drivers.	See "Configuring RDMA over an InfiniBand network" on page 546.
Tune system performance	Tune CPU frequency and boot parameters for systems.	See "Tuning system performance" on page 550.
Configure LLT manually	Configure LLT to use RDMA capability. Alternatively, you can use the installer to automatically configure LLT to use RDMA.	See "Manually configuring LLT over RDMA" on page 552.
Verify LLT configuration	Run LLT commands to test the LLT over RDMA configuration.	See "Verifying LLT configuration" on page 556.

Choosing supported hardware for LLT over RDMA

To configure LLT over RDMA you need to use the hardware that is RDMA enabled.

Table K-2

Hardware	Supported types	Reference
Network card	Mellanox-based Host Channel Adapters (HCAs) (VPI, ConnectX, ConnectX-2 and 3)	For detailed installation information, refer to the hardware vendor documentation.
Network switch	Mellanox, InfiniBand switches Ethernet switches must be Data Center Bridging (DCB) capable	For detailed installation information, refer to the hardware vendor documentation.
Cables	Copper and Optical Cables, InfiniBand cables	For detailed installation information, refer to the hardware vendor documentation.

Warning: When you install the Mellanox NIC for using RDMA capability, do not install Mellanox drivers that come with the hardware. LLT uses the Mellanox drivers that are installed by default with the Linux operating system. LLT might not be configurable if you install Mellanox drivers provided with the hardware.

Installing RDMA, InfiniBand or Ethernet drivers and utilities

Install the following RPMs to get access to the required RDMA, InfiniBand or Ethernet drivers and utilities. Note that the rpm version of the RPMs may differ for each of the supported Linux flavors.

Symantec does not support any external Mellanox OFED packages. The supported packages are listed in this section.

Symantec recommends that you use the Yellowdog Updater Modified (yum) package management utility to install RPMs on RHEL systems and use Zypper, a command line package manager, on SUSE systems.

Note: Install the OpenSM package only if you configure an InfiniBand network. All other packages are required with both InfiniBand and Ethernet networks.

Table K-3 lists the drivers and utilities required for RDMA, InfiniBand or Ethernet network.

Packages	RHEL	SUSE
Userland device drivers for RDMA operations	▪ `libmthca` ▪ `libmlx4` ▪ `rdma` ▪ `librdmacm-utils`	▪ `libmthca-rdmav2` ▪ `libmlx4-rdmav2` ▪ `ofed` ▪ `librdmacm`
OpenSM related package (InfiniBand only)	▪ `opensm` ▪ `opensm-libs` ▪ `libibumad`	▪ `opensm` ▪ `libibumad3`
InfiniBand troubleshooting and performance tests	▪ `Ibutils` ▪ `infiniband-diags` ▪ `Perftest`	▪ `Ibutils` ▪ `infiniband-diags`
libibverbs packages for userland InfiniBand operations	▪ `libibverbs-devel` ▪ `libibverbs-utils`	▪ `libibverbs`

Configuring RDMA over an Ethernet network

Configure the RDMA and Ethernet drivers so that LLT can use the RDMA capable hardware.

See "Enable RDMA over Converged Ethernet (RoCE)" on page 544.

See "Configuring RDMA and Ethernet drivers" on page 545.

See "Configuring IP addresses over Ethernet Interfaces" on page 545.

Enable RDMA over Converged Ethernet (RoCE)

The following steps are applicable only on a system installed with RHEL Linux. On SUSE Linux, the RDMA is enabled by default.

1. Make sure that the SFHA stack is stopped and the LLT and GAB modules are not loaded.

   See "Starting and stopping processes for the Symantec products " on page 395.

2. Skip this step if you are on a RHEL 7 system. Alternatively, create or modify the `/etc/modprobe.d/mlx4.conf` configuration file and add the value `options mlx4_core hpn=1` to the file. This enables RDMA over Converged Ethernet (RoCE) in Mellanox drivers (installed by default with the operating system).

3 Verify whether the Mellanox drivers are loaded.

```
# lsmod | grep mlx4_en

# lsmod | grep mlx4_core
```

4 Unload the Mellanox drivers if the drivers are loaded.

```
# rmmod mlx4_ib

# rmmod mlx4_en

# rmmod mlx4_core
```

Configuring RDMA and Ethernet drivers

Load the Mellanox drivers that are installed by default with the operating system and enable the RDMA service.

1 (RHEL Linux only) Load the Mellanox drivers.

```
# modprobe mlx4_core

# modprobe mlx4_ib

# modprobe mlx4_en
```

2 Enable RDMA service on the Linux operating system.

On RHEL Linux: `# chkconfig --level 235 rdma on`

On SUSE Linux: `# chkconfig --level 235 openibd on`

Configuring IP addresses over Ethernet Interfaces

Perform the following steps to configure IP addresses over the network interfaces which you plan to configure under LLT. These interfaces must not be aggregated interfaces.

1. Configure IP addresses using Linux `ifconfig` command. Make sure that the IP address for each link must be from a different subnet.

   Typical private IP addresses that you can use are:

   Node0:

   link0: 192.168.1.1

   link1: 192.168.2.1

   Node1:

   link0: 192.168.1.2

   link1: 192.168.2.2

2. Run IP ping test between nodes to ensure that there is network level connectivity between nodes.

3. Configure IP addresses to start automatically after the system restarts or reboots by creating a new configuration file or by modifying the existing file.

   - On RHEL, modify the `/etc/sysconfig/network-scripts/` directory by modifying the `ifcfg-eth` (Ethernet) configuration file.
   - On SUSE, modify the `/etc/sysconfig/network/` by modifying the `ifcfg-eth` (Ethernet) configuration file.
     For example, for an Ethernet interface eth0, create the `ifcfg-eth0` file with values for the following parameters.

```
DEVICE=eth0
BOOTPROTO=static
IPADDR=192.168.27.1
NETMASK=255.255.255.0
NETWORK=192.168.27.0
BROADCAST=192.168.27.255
NM_CONTROLLED=no # This line ensures IPs are plumbed correctly after bootup
ONBOOT=yes
STARTMODE='auto' # This line is only for SUSE
```

Configuring RDMA over an InfiniBand network

While configuring RDMA over an InfiniBand network, you need to configure the InfiniBand drivers, configure the OpenSM service, and configure IP addresses for the InfiniBand interfaces.

See "Configuring RDMA and InfiniBand drivers" on page 547.

See “ Configuring the OpenSM service” on page 548.

See “Configuring IP addresses over InfiniBand Interfaces” on page 549.

Configuring RDMA and InfiniBand drivers

Configure the RDMA and InfiniBand drivers so that LLT can use the RDMA capable hardware.

1. Ensure that the following RDMA and InfiniBand drivers are loaded. Use the `lsmod` command to verify whether a driver is loaded.

   The InfiniBand interfaces are not visible by default until you load the InfiniBand drivers. This procedure is only required for initial configuration.

   ```
   # modprobe rdma_cm

   # modprobe rdma_ucm

   # modprobe mlx4_en

   # modprobe mlx4_ib

   # modprobe ib_mthca

   # modprobe ib_ipoib

   # modprobe ib_umad
   ```

2. Load the drivers at boot time by appending the configuration file on the operating system.

   On RHEL and SUSE Linux, append the `/etc/rdma/rdma.conf` and `/etc/infiniband/openib.conf` files respectively with the following values:

   ```
   ONBOOT=yes

   RDMA_UCM_LOAD=yes

   MTHCA_LOAD=yes

   IPOIB_LOAD=yes

   SDP_LOAD=yes

   MLX4_LOAD=yes

   MLX4_EN_LOAD=yes
   ```

3. Enable RDMA service on the Linux operating system.

   On RHEL Linux:

   ```
   # chkconfig --level 235 rdma on
   ```

   On SUSE Linux:

   ```
   # chkconfig --level 235 openibd on
   ```

Configuring the OpenSM service

OpenSM is an InfiniBand compliant Subnet Manager and Subnet Administrator, which is required to initialize the InfiniBand hardware. In the default mode, OpenSM

scans the IB fabric, initializes the hardware, and checks the fabric occasionally for changes.

For InfiniBand network, make sure to configure subnet manager if you have not already configured the service.

1. Modify the OpenSM configuration file if you plan to configure multiple links under LLT.

   On RHEL, update the `/etc/sysconfig/opensm` file.

2. Start OpenSM.

   On RHEL7, run **`# systemctl start opensm.service`**

   On other Linux systems, run **`# /etc/init.d/opensm start`**

3. Enable Linux service to start OpenSM automatically after restart.

   On RHEL Linux, **`# chkconfig --level 235 opensm on`**

   On SUSE Linux, **`# chkconfig --level 235 opensmd on`**

Configuring IP addresses over InfiniBand Interfaces

Perform the following steps to configure IP addresses over the network interfaces which you plan to configure under LLT. These interfaces must not be aggregated interfaces.

1. Configure IP addresses using Linux `ifconfig` command. Make sure that the IP address for each link must be from a different subnet.

   Typical private IP addresses that you can use are: **192.168.12.1, 192.168.12.2, 192.168.12.3** and so on.

2. Run the InfiniBand ping test between nodes to ensure that there is InfiniBand level connectivity between nodes.

   - On one node, start the ibping server.

     ```
     # ibping -S
     ```

   - On the node, get the GUID of an InfiniBand interface that you need to ping from another node.

     ```
     # ibstat


     CA 'mlx4_0'
     Number of ports: 2
      --
     Port 1:
     State: Active
     ```

```
---
Port GUID: 0x0002c90300a02af1
Link layer: InfiniBand
```

- Ping the peer node by using its GUID.

  ```
  # ibping -G 0x0002c90300a02af1
  ```

  Where, *0x0002c90300a02af1* is the GUID of the server.

3 Configure IP addresses automatically after restart by creating a new configuration file or by modifying the existing file.

- On RHEL, modify the `/etc/sysconfig/network-scripts/` directory by modifying the `ifcfg-ibX` (InfiniBand) configuration file.
- On SUSE, modify the `/etc/sysconfig/network/` by modifying the `ifcfg-ibX` (InfiniBand) configuration file.

```
For example, for an Infiniband interface ib0,
create ifcfg-ib0 file with values for the following parameters.

DEVICE=ib0
BOOTPROTO=static
IPADDR=192.168.27.1
NETMASK=255.255.255.0
NETWORK=192.168.27.0
BROADCAST=192.168.27.255
NM_CONTROLLED=no # This line ensures IPs are plumbed correctly
after bootup and the Network manager does not interfere
with the interfaces
ONBOOT=yes
STARTMODE='auto' # This line is only for SUSE
```

Tuning system performance

Run IP ping test to ensure that systems are tuned for the best performance. The latency should be less than 30us, if not then your system may need tuning. However, the latency may vary based on your system configuration.

To tune your system, perform the following steps. For additional tuning, follow the performance tuning guide from Mellanox.

Performance Tuning Guidelines for Mellanox Network Adapters

Tuning the CPU frequency

To tune the CPU frequency of a system, perform the following steps:

1. Verify whether the CPU frequency is already tuned.

```
# cat /proc/cpuinfo | grep Hz


model name      : Intel(R) Xeon(R) CPU E5-2643 0 @ 3.30GHz
cpu MHz         : 3300.179
```

2. If the CPU frquency displayed by the `cpu MHz` and `model name` attribute is the same, then the CPU frequency is already tuned. You can skip the next steps.

   If the CPU frequency displayed by the `cpu Mhz` and `model name` attribute is not the same, then follow the next steps to tune the frequency.

3. Go to system console and restart the system.
4. Press F11 to enter into BIOS settings.
5. Go to BIOS menu > Launch System setup > BIOS settings > System Profile Settings > System Profile > Max performance.

   The menu options might vary with system type.

Tuning the boot parameter settings

To tune the boot parameter settings, perform the following steps.

1. In the `/boot/grub/grub.conf` file or any other boot loader configuration file, ensure that the value of the `intel_iommu` is set to **off**.
2. Append the `/boot/grub/grub.conf` file or any other boot loader configuration file with the following parameters if they are not listed in the configuration file.

   `intel_idle.max_cstate=0 processor.max_cstate=1`

3. Restart the system.

On RHEL 7 systems:

1. In the `/etc/default/grub` file, append the GRUB_CMDLINE_LINUX variable with intel_idle.max_cstate=0 processor.max_cstate=1
2. After `/etc/default/grub` is modified, run the following command:

   `grub2-mkconfig -o /boot/grub2/grub.cfg`

3. Restart the system.

Manually configuring LLT over RDMA

You can automatically configure LLT to use RDMA using the installer. To manually configure LLT over RDMA follow the steps that are given in this section.

The following checklist is to configure LLT over RDMA:

- Make sure that the LLT private links are on separate subnets. Set the broadcast address in `/etc/lltab` explicitly depending on the subnet for each link.
  See "Broadcast address in the `/etc/lltab` file" on page 552.
- Make sure that each RDMA enabled NIC (RNIC) over an InfiniBand or Ethernet network has an IP address that is configured before configuring LLT.
- Make sure that the IP addresses in the `/etc/lltab` files are consistent with the IP addresses of the network interfaces (InfiniBand or Ethernet network interfaces).
- Make sure that each link has a unique and a private IP range for the UDP port.
  See "Selecting UDP ports" on page 553.
- See the sample comfiguration for direct-attached (non-routed) links.
  See "Sample configuration: direct-attached links" on page 555.

Broadcast address in the `/etc/lltab` file

The broadcast address is set explicitly for each link in the following example.

- Display the content of the `/etc/lltab` file on the first node sys1:

```
sys1 # cat /etc/lltab

set-node sys1
set-cluster 1
link link1 udp - rdma  50000  -  192.168.9.1 192.168.9.255
link link2 udp - rdma  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/lltab` file on the second node sys2:

```
sys2 # cat /etc/lltab

set-node sys2
set-cluster 1
link link1 udp - rdma  50000  -  192.168.9.2 192.168.9.255
link link2 udp - rdma  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/lltab` file

Review the link command information in this section for the `/etc/lltab` file. See the following information for sample configurations:

- See "The link command in the `/etc/lltab` file"on page 553 on page 553.

Table K-4 describes the fields of the link command that are shown in the `/etc/lltab` file examples. Note that some of the fields differ from the command for standard LLT links.

Table K-4 Field description for link command in `/etc/lltab`

Field	Description
tag-name	A unique string that is used as a tag by LLT; for example link1, link2,....
device	The device path of the UDP protocol; for example udp. A place holder string. Linux does not have devices for protocols. So this field is ignored.
node-range	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
link-type	Type of link; must be "rdma" for LLT over RDMA.
udp-port	Unique UDP port in the range of 49152-65535 for the link. See "Selecting UDP ports" on page 530.
MTU	"-" is the default, which has a value of 8192. Do not change this default value for the RDMA links.
IP address	IP address of the link on the local node.
bcast-address	Specify the value of the subnet broadcast address.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:

- Ports from the range of well-known ports, 0 to 1023
- Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address         State
udp        0      0 *:32768                 *:*
udp        0      0 *:956                   *:*
udp        0      0 *:tftp                  *:*
udp        0      0 *:sunrpc                *:*
udp        0      0 *:ipp                   *:*
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the /etc/services file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# ifconfig interface_name netmask netmask
```

For example:

- For the first network interface on the node sys1:

  ```
  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 sys2:

  ```
  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 sys1:

  ```
  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 sys2:

```
IP address=192.168.10.2, Broadcast address=192.168.10.255,
Netmask=255.255.255.0
```

Sample configuration: direct-attached links

Figure K-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure K-1 A typical configuration of direct-attached links that uses LLT over RDMA

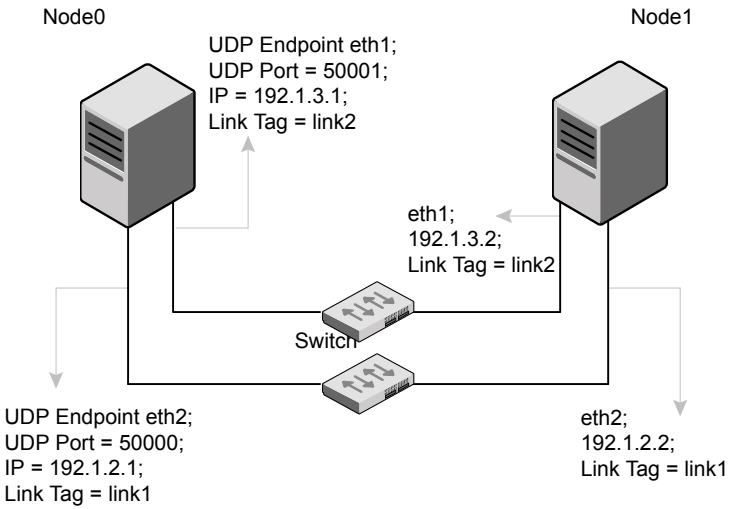

The configuration that the /etc/lltttab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT sends broadcasts to peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the /etc/lltttab file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the /etc/lltttab 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-addressbast-address
link link1 udp - rdma 50000 - 192.1.2.1 192.1.2.255
link link2 udp - rdma 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 bast-address
link link1 udp - rdma 50000 - 192.1.2.2 192.1.2.255
link link2 udp - rdma 50001 - 192.1.3.2 192.1.3.255
```

LLT over RDMA sample /etc/llttab

The following is a sample of LLT over RDMA in the etc/llttab file.

```
set-node sys1
set-cluster clus1
link eth1 udp - rdma 50000 - 192.168.10.1 - 192.168.10.255
link eth2 udp - rdma 50001 - 192.168.11.1 - 192.168.11.255
link-lowpri eth0 udp - rdma 50004 - 10.200.58.205 - 10.200.58.255
```

Verifying LLT configuration

After starting LLT, GAB and other component, run the following commands to verify the LLT configuration.

1 Run the `lltstat -l` command to view the RDMA link configuration. View the link-type configured to rdma for the RDMA links.

```
# lltstat -l


LLT link information:
link 0  link0 on rdma hipri
mtu 8192, sap 0x2345, broadcast 192.168.27.255, addrlen 4
txpkts 171  txbytes 10492
rxpkts 105  rxbytes 5124
latehb 0  badcksum 0  errors 0
```

2 Run the `lltstat -nvv -r` command to view the RDMA and non-RDMA channel connection state.

LLT internally configures each RDMA link in two modes (RDMA and non-RDMA) to allow both RDMA and non-RDMA traffic to use the same link. The GAB membership-related traffic goes over the non-RDMA channel while node to node data-transfer goes over high-speed RDMA channel for better performance.

```
# lltstat -rnvv active


LLT node information:
Node          State Link  Status TxRDMA RxRDMA Address
* 0 thorpc365 OPEN  link0   UP    UP     UP    192.168.27.1
                    link1   UP    UP     UP    192.168.28.1
                    link2   UP    N/A    N/A   00:15:17:97:91:2E

1 thorpc366   OPEN  link0   UP    UP     UP    192.168.27.2
                    link1   UP    UP     UP    192.168.28.2
                    link2   UP    N/A    N/A   00:15:17:97:A1:7C
```

Troubleshooting LLT over RDMA

This section lists the issues and their resolutions.

IP addresses associated to the RDMA NICs do not automatically plumb on node restart

If IP addresses do not plumb automatically, you might experience LLT failure.

Resolution: Assign unique IP addresses to RNICs and assign the same in the configuration script. For example, on an ethernet network, the `ifcfg-eth` script must be modified with the unique IP address of the RNIC.

See “Configuring IP addresses over InfiniBand Interfaces” on page 549.

Ping test fails for the IP addresses configured over InfiniBand interfaces.

Resolution: Check the physical configuration and configure OpenSM. If you configured multiple links, then make sure that you have configured OpenSM to monitor multiple links in the configuration file. On RHEL, configure the `/etc/sysconfig/opensm` file.

See “ Configuring the OpenSM service” on page 548.

After a node restart, by default the Mellanox card with Virtual Protocol Interconnect (VPI) gets configured in InfiniBand mode

After restart, you might expect the Mellanox VPI RNIC to get configured in the Ethernet mode. By default, the card gets configured in the InfiniBand mode.

Resolution: Update the Mellanox configuration file. On RHEL, configure the `/etc/rdma/mlx4.conf` file.

The LLT module fails to start

When you try to start LLT, it may fail to start and you may see the following message:

```
# /etc/init.d/llt start
Starting LLT:
LLT: loading module...
LLT:Error loading LLT dependency rdma_cm.
Make sure module rdma_cm is available on the system.
```

Description: Check the system log at /var/log/messages. If the log file lists the following error, the issue may be because the IPv6 module is not available on the system. In addition, the LLT module has indirect dependency on the IPv6 module.

```
ib_addr: Unknown symbol ipv6_dev_get_saddr
ib_addr: Unknown symbol ip6_route_output
ib_addr: Unknown symbol ipv6_chk_addr
```

Resolution: Load the IPv6 module. If you do not want to configure the IPv6 module on the node, then configure the IPv6 module to start in the disabled mode.

To start IPv6 in the disabled mode:

- In the `/etc/modprobe.d/` directory, create a file `ipv6.conf` and add the following line to the file

  ```
  options ipv6 disable=1
  ```

The LLT module starts up without any issues once the file loads the IPv6 module in the disabled mode.

Appendix L

Compatibility issues when installing Storage Foundation High Availability with other products

This appendix includes the following topics:

- Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present
- Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present
- Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present
- Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present

Installing Storage Foundation when other Symantec products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the VOM Central Server and Managed Host RPMs as is.
- When uninstalling Storage Foundation products where VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present

If you plan to install or upgrade Storage Foundation on systems where ApplicationHA has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not consider VCS as an installed product even though it uses the bundled VRTSvcs RPM.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not allow the installation or upgrade for products that use VCS. The following products cannot be installed or upgrade: VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SFCFSRAC or SFSYBASECE.

- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer allows the installation or upgrade of VM, FS, SF, or DMP.
- When you uninstall Storage Foundation products where ApplicationHA is present, the installer does not uninstall VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpbx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpbx, VRTSicsco, and VRTSat.

Index

A

about
- Deployment Server 270
- global clusters 27
- installation and configuration methods 45
- installation preparation 56
- installation using operating system-specific methods 225
- response files 47
- SORT 28
- Symantec product licensing 50
- Veritas Operations Manager 25
- web-based installer 169

About RDMA
- RDMA over Converged Ethernet or InfiniBand networks
  - clustering environment 539

adding
- users 130

agents
- disabling 408

applications, stopping 316

assesing system
- installation readiness 66

attributes
- UseFence 241, 264

B

backup boot disk group 369
- rejoining 369

before using
- web-based installer 170

C

cables
- cross-over Ethernet 423

checking
- installation readiness 66

checking product versions 36

cluster
- removing a node from 439
- verifying operation 399

command failures 516

commands
- hastatus 399
- hasys 400
- lltconfig 473
- lltstat 397
- vxdisksetup (initializing disks) 139
- vxlicinst 137
- vxlicrep 136

configuration
- restoring the original 378

configuring
- private network 58
- rsh 57
- ssh 57
- switches 58

configuring SFHA
- script-based installer 113

configuring VCS
- adding users 130
- event notification 131–132
- global clusters 134
- starting 115

controllers
- private Ethernet 58

coordinator disks
- DMP devices 31
- for I/O fencing 31
- setting up 239

creating
- backups 309
- Install Templates 295

D

data disks
- for I/O fencing 31

defining
- Install Bundles 290

deploying
Symantec product updates to your environment 287
Symantec releases 297
deploying using
Install Bundles 297
deploying using Install Templates
Install Templates 297
deployment preferences
setting 276
Deployment Server
about 270
downloading the most recent release information from the SORT site 278
installing 272
loading release information and patches on to 279
overview 271
proxy server 300
setting up 273
specifying a non-default repository location 278
disabling
external network connection attempts 38
disabling the agents 408
disk space requirements 35
disks
adding and initializing 139
coordinator 239
testing with vxfentsthdw 140
verifying node access 141
downloading maintenance releases and patches 36
downloading the most recent release information
by running the Deployment Server from a system with Internet access 278

E

eeprom
parameters 58
Ethernet controllers 58, 423
existing coordination points
order 192

F

freezing service groups 316

G

GAB
description 26
gabtab file
verifying after installation 473
global clusters 27
configuration 134

H

hastatus -summary command 399
hasys -display command 400
hubs 58
independent 423

I

I/O fencing
checking disks 140
setting up 238
shared storage 140
I/O fencing requirements
non-SCSI-3 44
Install Bundles
defining 290
deploying using the Deployment Server 297
integration options 318
Install Templates
creating 295
deploying using Install Templates 297
installation
Red Hat Satellite server 235
installer
about the script-based installer 69
installer patches
obtaining either manually or automatically 37
Installing
SFHA with the web-based installer 173
web-based installer 173
installing
post 135
SFHA using operating system-specific methods 225
Symantec product license keys 53
the Deployment Server 272
using Kickstart 226
using response files 196
using yum 230

K

kernel.hung_task_panic tunable 64
keyless licensing
setting or changing the product level 51

Kickstart
installing 226
sample configuration file 228

L

license keys
adding with vxlicinst 137
replacing demo key 137
licenses
information about 136
licensing 39
add-on 38
CDS 39
full 38
installing Symantec product license keys 53
setting or changing the product level for keyless licensing 51
links
private network 473
LLT
description 26
interconnects 65
verifying 397
LLT over RDMA
configure 541
faster interconnects 540
supported use cases 541
lltconfig command 473
llthosts file
verifying after installation 473
lltstat command 397
lltab file
verifying after installation 473
log files 517

M

MAC addresses 58
main.cf file
contents after installation 478
main.cf files 484
manual pages
potential problems 515
troubleshooting 515
media speed 65
optimizing 65
mounting
software disc 65

N

network switches 58
nodes
adding application nodes
configuring GAB 429
configuring LLT 429
configuring VXFEN 429
starting Volume Manager 429
non-SCSI-3 fencing
manual configuration 258
setting up 258
non-SCSI-3 I/O fencing
requirements 44
non-SCSI3 fencing
setting up 161
using installsfha 161

O

obtaining
installer patches either automatically or manually 37
security exception on Mozilla Firefox 171
optimizing
media speed 65
original configuration
restoring the 378
overview
Deployment Server 271

P

parameters
eeprom 58
PATH variable
VCS commands 396
persistent reservations
SCSI-3 60
phased 341
phased upgrade 341, 343
example 342
planning to upgrade VVR 311
post-upgrade
updating variables 382
verifying 388
prechecking
using the installer 66
preinstallation 311
preinstallation check
web-based installer 172

preparing to upgrade 308
preparing to upgrade VVR 316
private network
configuring 58
problems
accessing manual pages 515
executing file system commands 516
proxy server
connecting the Deployment Server 300

R

RDMA
Configure drivers 544, 547
Configure interfaces 549
Driver installation 543
manually configure LLT 552
OpenSM service 548
supported hardware 542
troubleshoot 558
Tune system performance 550
Verify LLT configuration 556
Red Hat Satellite server
installing 235
rejoining
backup boot disk group 369
release images
viewing or downloading available 280
release information and patches
loading using the Deployment Server 279
release notes 34
releases
finding out which releases you have, and which upgrades or updates you may need 288
removing
license files 413
the Replicated Data Set 409
removing a system from a cluster 439
Replicated Data Set
removing the 409
repository images
viewing and removing repository images stored in your repository 285
response files
about 47
installation 196
rolling upgrade 364
syntax 48
uninstalling 417
upgrading 359
restoring the original configuration 378
rolling upgrade
using response files 364
using the script-based installer 334
versions 331
rsh 116
configuration 57

S

script-based installer
about 69
SFHA configuration overview 113
SCSI-3
persistent reservations 60
SCSI-3 persistent reservations
verifying 238
service groups
freezing 316
unfreezing 377
setting
deployment preferences 276
environment variables 63
setting up
Deployment Server 273
SFDB authentication 390
adding nodes 436
configuring vxdbd 391
SFHA
configuring 113
coordinator disks 239
SFHA installation
preinstallation information 35
verifying
cluster operations 396
GAB operations 396
LLT operations 396
simulataneous install or upgrade 318
SMTP email notification 131
SNMP trap notification 132
specifying
non-default repository location 278
ssh 116
configuration 57
starting
web-based installer 170
starting configuration
installvcs program 116
product installer 115

stopping
applications 316
supported operating systems 35
switches 58
Symantec product license keys
installing 53
Symantec product updates
deploying to your environment 287
Symantec products
starting process 395
stopping process 395
Symantec releases
deploying a specific release 297
system state attribute value 399

T

troubleshooting
accessing manual pages 515
executing file system commands 516
tunables file
about setting parameters 460
parameter definitions 465
preparing 464
setting for configuration 461
setting for installation 461
setting for upgrade 461
setting parameters 464
setting with no other operations 462
setting with un-integrated response file 463

U

unfreezing service groups 377
uninstalling
using response files 417
using the web-based installer 412
unsuccessful upgrade 369
upgrade
array support 317
creating backups 309
getting ready 308
methods 302
phased 341, 343
supported upgrade paths 303
upgrades or updates
finding out which releases you have 288
upgrading
phased 341
using response files 359
upgrading *(continued)*
using the web-based installer 326
upgrading VVR
from 4.1 312
planning 311
preparing 316

V

VCS
command directory path variable 396
configuration files
main.cf 477
verifying
NIC configuration 135
product installation 394
viewing and removing repository images
stored in your repository 285
viewing or downloading
available release images 280
vradmin
delpri 410
stoprep 409
VVR 4.1
planning an upgrade from 312
vvr_upgrade_finish script 379
vxdisksetup command 139
vxlicinst command 137
vxlicrep command 136
vxplex
used to remove mirrors of root disk volumes 406

W

web-based installer 173
about 169
before using 170
installation 173
preinstallation check 172
starting 170
uninstalling 412
upgrading 326

Y

yum
installing 230