

Symantec™ VirtualStore 6.0.4 Installation and Configuration Guide - Linux

Symantec™ VirtualStore Installation and Configuration Guide

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

Product version: 6.0.4

Document version: 6.0.4 Rev 0

Legal Notice

Copyright © 2013 Symantec Corporation. All rights reserved.

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

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

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

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

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

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

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

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

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

Contacting Technical Support

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

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

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

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

- Product release level
- Hardware information

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

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

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

Support agreement resources

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

Asia-Pacific and Japan customercare_apac@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

About Symantec Connect

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

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

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>

Contents

Technical Support	4
Section 1 Installation overview and planning	18
Chapter 1 Introducing Symantec VirtualStore	19
About Symantec VirtualStore	19
About Symantec VirtualStore solution for VMware environment	19
About I/O fencing	20
About configuring SVS clusters for data integrity	21
About I/O fencing for Symantec VirtualStore in virtual machines that do not support SCSI-3 PR	22
About I/O fencing components	23
About preferred fencing	25
Chapter 2 System requirements	26
Release notes	26
Hardware compatibility list (HCL)	27
Supported operating systems	27
I/O fencing requirements	27
Coordinator disk requirements for I/O fencing	27
CP server requirements	28
Non-SCSI-3 I/O fencing requirements	31
Database requirements	32
VxVM licenses	32
Cross-Platform Data Sharing licensing	33
Disk space requirements for Storage Foundation Cluster File System	33
Discovering product versions and various requirement information	33
Chapter 3 Planning to install Symantec VirtualStore	35
About planning for SVS installation	35
About installation and configuration methods	36
About the Veritas installer	37

	Assessing the system for installation readiness	39
	Prechecking your systems using the Veritas installer	39
	Downloading the Symantec VirtualStore software	40
	Setting environment variables	41
	Optimizing LLT media speed settings on private NICs	42
	Guidelines for setting the media speed of the LLT interconnects	42
	About using ssh or rsh with the Veritas installer	43
	Setting up shared storage	43
	Setting up shared storage: SCSI	43
	Setting up shared storage: Fibre Channel	45
	Prerequisites for installing Symantec VirtualStore	46
	Hardware overview and requirements for Symantec VirtualStore	47
	Shared storage	48
	Fibre Channel switch	48
	Cluster platforms	49
	Disk space requirements for Storage Foundation Cluster File System	49
Chapter 4	Licensing Symantec VirtualStore	50
	About Veritas product licensing	50
	Setting or changing the product level for keyless licensing	51
	Installing Veritas product license keys	53
Section 2	Installation using the script-based installer	54
Chapter 5	Installing Symantec VirtualStore using the script-based installer	55
	Installing Symantec VirtualStore using the installer	55
	Configuring Symantec VirtualStore in secure mode	58
Chapter 6	Preparing to configure Symantec VirtualStore	59
	About planning to configure I/O fencing	59
	Typical SVS cluster configuration with server-based I/O fencing	63
	Recommended CP server configurations	64
	Setting up the CP server	67
	Planning your CP server setup	67
	Installing the CP server using the installer	68
	Configuring the CP server using the Web-based installer	69

	Configuring the CP server cluster in secure mode	70
	Setting up shared storage for the CP server database	71
	Configuring the CP server using the installer program	72
	Configuring the CP server manually	81
	Configuring CP server using response files	82
	Response file variables to configure CP server	83
	Sample response file for configuring the CP server on single node	
	VCS cluster	85
	Sample response file for configuring the CP server on SFHA	
	cluster	86
	Verifying the CP server configuration	86
Chapter 7	Configuring Symantec VirtualStore	88
	Overview of tasks to configure Symantec VirtualStore using the	
	script-based installer	88
	Configuring SVS using the script-based installer	89
	Starting the software configuration	89
	Specifying systems for configuration	90
	Configuring the cluster name	91
	Configuring private heartbeat links	91
	Configuring the virtual IP of the cluster	95
	Configuring a secure cluster node by node	97
	Adding VCS users	101
	Configuring SMTP email notification	102
	Configuring SNMP trap notification	103
	Configuring global clusters	105
	Completing the SVS configuration	106
	Verifying and updating licenses on the system	107
Chapter 8	Configuring SVS for data integrity	110
	Setting up disk-based I/O fencing using installsvs	110
	Configuring disk-based I/O fencing using installsvs	110
	Initializing disks as VxVM disks	113
	Checking shared disks for I/O fencing	114
	Setting up disk-based I/O fencing manually	118
	Setting up coordinator disk groups	119
	Creating I/O fencing configuration files	120
	Modifying VCS configuration to use I/O fencing	121
	Verifying I/O fencing configuration	122
	Identifying disks to use as coordinator disks	123
	Setting up server-based I/O fencing using installsvs	124

	Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsvs	132
	Setting up server-based I/O fencing manually	134
	Preparing the CP servers manually for use by the SVS cluster	134
	Configuring server-based fencing on the SVS cluster manually	137
	Configuring CoordPoint agent to monitor coordination points	143
	Verifying server-based I/O fencing configuration	145
	Setting up non-SCSI-3 fencing in virtual environments manually	146
	Sample /etc/vxfenmode file for non-SCSI-3 fencing	148
	Enabling or disabling the preferred fencing policy	151
Section 3	Installation using the Web-based installer	154
Chapter 9	Installing Symantec VirtualStore using the Web-based installer	155
	About the Web-based installer	155
	Before using the Veritas Web-based installer	156
	Starting the Veritas Web-based installer	156
	Obtaining a security exception on Mozilla Firefox	157
	Performing a pre-installation check with the Veritas Web-based installer	158
	Setting installer options with the Web-based installer	158
	Installing SVS with the Web-based installer	159
Chapter 10	Configuring Symantec VirtualStore	161
	Configuring Symantec VirtualStore using the Web-based installer	161
	Configuring Symantec VirtualStore for data integrity using the Web-based installer	166
Section 4	Automated installation using response files	174
Chapter 11	Performing an automated Symantec VirtualStore installation	175
	Installing SVS using response files	175

Chapter 12	Performing an automated Symantec VirtualStore configuration	177
	Configuring SVS using response files	177
	Response file variables to configure Symantec VirtualStore	178
Chapter 13	Performing an automated I/O fencing configuration using response files	188
	Configuring I/O fencing using response files	188
	Response file variables to configure disk-based I/O fencing	189
	Sample response file for configuring disk-based I/O fencing	192
	Response file variables to configure server-based I/O fencing	192
	Sample response file for configuring server-based I/O fencing	194
	Response file variables to configure non-SCSI-3 server-based I/O fencing	194
	Sample response file for configuring non-SCSI-3 server-based I/O fencing	195
Section 5	Upgrading Symantec VirtualStore	197
Chapter 14	Preparing to upgrade	198
	Upgrade methods for SVS	198
	Supported upgrade paths	199
	About using the installer to upgrade when the root disk is encapsulated	199
	Preparing to upgrade SVS	200
	Getting ready for the upgrade	200
	Creating backups	201
	Determining if the root disk is encapsulated	202
	Pre-upgrade planning for Veritas Volume Replicator	202
	Preparing to upgrade VVR when VCS agents are configured	205
	Upgrading the array support	209
Chapter 15	Performing a full upgrade using the installer	211
	Performing a full upgrade	211
	Ensuring the file systems are clean	211
	Performing the upgrade	212

Chapter 16	Performing a phased upgrade	216
	Performing a phased upgrade of SVS	216
	Prerequisites for a phased upgrade	216
	Planning for a phased upgrade	217
	Phased upgrade limitations	217
	Moving the service groups to the second subcluster	217
	Upgrading the SVS stack on the first subcluster	220
	Preparing the second subcluster	221
	Activating the first subcluster	222
	Upgrading the operating system on the second subcluster	223
	Upgrading the second subcluster	223
	Completing the phased upgrade	224
Chapter 17	Performing a rolling upgrade	225
	Performing a rolling upgrade using the installer	225
	About rolling upgrades	225
	Supported rolling upgrade paths	228
	Performing a rolling upgrade using the script-based installer	228
	Performing a rolling upgrade of SVS using the Web-based installer	230
Section 6	Post-installation tasks	233
Chapter 18	Performing post-upgrade tasks	234
	Re-joining the backup boot disk group into the current disk group	234
	Reverting to the backup boot disk group after an unsuccessful upgrade	235
	Register the new plugin	235
Section 7	Setting up Symantec VirtualStore	236
Chapter 19	Setting up VirtualStore	237
	About setting up VirtualStore	238
	Setting up Clustered NFS	238
	Setting up the VirtualStore vSphere Service Console	241
	Registering the VirtualStore cluster	242
	Verifying the VirtualStore cluster	243
	Virtual machine customization pre-requisites	243
	VMware View Integration pre-requisites	244
	Unregistering the VirtualStore cluster	244

	Displays all the VMware vCenter Servers registered	245
	Useful links from VMware on NFS support and customization while cloning virtual machines	245
Section 8	Using the extensions for Symantec VirtualStore	246
Chapter 20	Using the VirtualStore Plug-in for VMware vSphere	247
	About the VMware vSphere Plug-in for VirtualStore	247
	Adding storage to ESX servers	248
	Creating a virtual machine	249
	Creating virtual machine clones using Symantec FileSnap	249
	Specifying the number of virtual machine clones to be created	251
	Specifying where to create the virtual machine clones	251
	Specifying the guest operating system if customizing for Linux or Windows	252
	Specifying network customization parameters for guest operating systems if using DHCP or Static IP	254
	Configuring the VMware View	255
	Configuring Citrix XenDesktop	257
	Verifying the virtual machine clones	257
Chapter 21	Using the Citrix XenDesktop extension for Symantec VirtualStore	258
	About adding virtual machine clones to Citrix XenDesktop	258
	Connecting Citrix XenDesktop to a vCenter Server	259
	Adding a connection	260
	Creating virtual machine clones using the Citrix XenDesktop extension	263
	Importing virtual machines to Citrix Desktop Studio	264
	Creating a desktop group and testing the deployment	268
Section 9	Adding and removing nodes	272
Chapter 22	Adding a node to a cluster	273
	About adding a node to a cluster	273
	Before adding a node to a cluster	274
	Adding a node to a cluster using the SVS installer	276

	Adding a node using the Web-based installer	280
	Adding the node to a cluster manually	280
	Starting Veritas Volume Manager (VxVM) on the new node	281
	Configuring cluster processes on the new node	282
	Setting up the node to run in secure mode	284
	Configuring server-based fencing on the new node	287
	Starting fencing on the new node	289
	Configuring the ClusterService group for the new node	290
	Starting Cluster Volume Manager (CVM) and Cluster File System (CFS)	290
	After adding the new node	291
Chapter 23	Removing a node from a cluster	292
	Removing a node from a SVS cluster	292
	Verifying the status of nodes and service groups	293
	Deleting the departing node from SVS configuration	294
	Removing security credentials from the leaving node	297
	Removing the node configuration from the CP server	297
	Unloading LLT and GAB and removing VCS on the departing node	298
Section 10	Uninstallation of Symantec VirtualStore	300
Chapter 24	Uninstalling Symantec VirtualStore	301
	Shutting down cluster operations	301
	Removing VxFS file systems	302
	Removing rootability	302
	Moving volumes to disk partitions	303
	Moving volumes onto disk partitions using VxVM	303
	Disabling the agents on a system	305
	Removing the Replicated Data Set	306
	Uninstalling SVS RPMs using the script-based installer	307
	Uninstalling SVS with the Veritas Web-based installer	308
	Removing license files (Optional)	309
	Removing the CP server configuration using the installer program	310

Section 11	Installation reference	312
Appendix A	Tunable files for installation	313
	About setting tunable parameters using the installer or a response file	313
	Setting tunables for an installation, configuration, or upgrade	314
	Setting tunables with no other installer-related operations	315
	Setting tunables with an un-integrated response file	316
	Preparing the tunables file	317
	Setting parameters for the tunables file	317
	Tunables value parameter definitions	318
Appendix B	Installation scripts	322
	Installation script options	322
	Restarting the installer after a failed connection	327
	Starting and stopping processes for the Veritas products	327
	Installation program has improved failure handling	328
Appendix C	Configuration files	329
	About I/O fencing configuration files	329
Appendix D	Response files	332
	About response files	332
	Upgrading SVS using response files	333
	Uninstalling SVS using response files	333
	Syntax in the response file	334
	Response file variables to install, upgrade, or uninstall Symantec VirtualStore	334
	Sample response file for Symantec VirtualStore installation	336
	Sample response file for Symantec VirtualStore configuration	337
Appendix E	Configuring the secure shell or the remote shell for communications	339
	About configuring secure shell or remote shell communication modes before installing products	339
Appendix F	SVS components	340
	Symantec VirtualStore installation RPMs	340
	Veritas Cluster Server installation RPMs	343

	Veritas Cluster File System installation RPMs	344
	Veritas Storage Foundation obsolete and reorganized installation RPMs	345
Appendix G	Sample SVS cluster setup diagrams for CP server-based I/O fencing	348
	Configuration diagrams for setting up server-based I/O fencing	348
	Two unique client clusters served by 3 CP servers	349
	Client cluster served by highly available CPS and 2 SCSI-3 disks	350
	Client cluster served by highly available CPS and 2 SCSI-3 disks	351
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	353
	Two node campus cluster served by remote CP server and 2 SCSI-3 disks	355
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	357
	Multiple client clusters served by highly available CP server and 2 SCSI-3 disks	358
Appendix H	Configuring LLT over UDP using IPv6	360
	Using the UDP layer of IPv6 for LLT	360
	When to use LLT over UDP	360
	Manually configuring LLT over UDP using IPv6	360
	The link command in the /etc/lldtab file	361
	The set-addr command in the /etc/lldtab file	362
	Selecting UDP ports	362
	Sample configuration: direct-attached links	363
	Sample configuration: links crossing IP routers	364
Appendix I	Configuring LLT over UDP using IPv4	366
	Using the UDP layer for LLT	366
	When to use LLT over UDP	366
	Manually configuring LLT over UDP using IPv4	366
	Broadcast address in the /etc/lldtab file	367
	The link command in the /etc/lldtab file	367
	The set-addr command in the /etc/lldtab file	368
	Selecting UDP ports	369
	Configuring the netmask for LLT	369
	Configuring the broadcast address for LLT	370

	Sample configuration: direct-attached links	370
	Sample configuration: links crossing IP routers	372
Appendix J	Compatibility issues when installing Symantec VirtualStore with other products	374
	Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present	374
	Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present	375
	Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present	375
	Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present	376
Index		377

Installation overview and planning

- [Chapter 1. Introducing Symantec VirtualStore](#)
- [Chapter 2. System requirements](#)
- [Chapter 3. Planning to install Symantec VirtualStore](#)
- [Chapter 4. Licensing Symantec VirtualStore](#)

Introducing Symantec VirtualStore

This chapter includes the following topics:

- [About Symantec VirtualStore](#)
- [About Symantec VirtualStore solution for VMware environment](#)
- [About I/O fencing](#)
- [About configuring SVS clusters for data integrity](#)
- [About preferred fencing](#)

About Symantec VirtualStore

Symantec VirtualStore (SVS) powered by Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) serves as a highly scalable, highly available NAS solution optimized for deploying and hosting virtual machine. VirtualStore is built on top of Cluster File System (CFS), which provides high availability and linear scalability across the cluster.

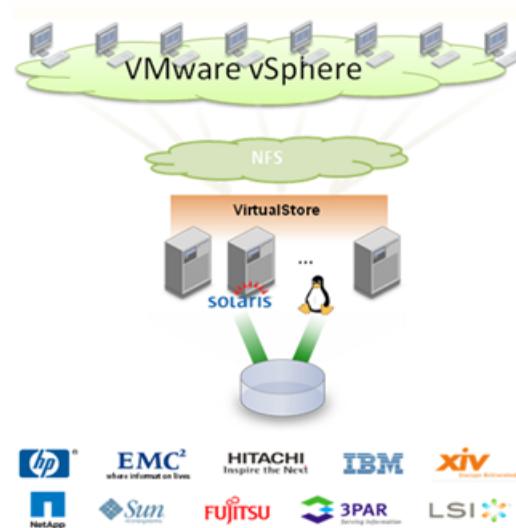
About Symantec VirtualStore solution for VMware environment

Symantec VirtualStore powered by Veritas Storage Foundation Cluster File System High Availability (SFCFSHA) is an optimized and customized NAS solution for your VMware environments. It provides the benefits of a highly scalable NAS and NFS solution for the VMware ESX workloads. With seamless integration with VMware Virtual Center, VirtualStore gives you a complete solution to efficiently manage the

VMware ESX NAS storage and leverage some of the key underlying benefits of the SFCFSHA.

- Provide standard NAS/NFS interfaces to interact with the VMware Virtual Infrastructure.
- Compatibility with VMotion (and storage VMotion).
- Balance datastore based on SVS Cluster Virtual IP addresses for optimized IO.
- Leverage key capabilities within the SFCFSHA (export same file system through multiple NAS heads, SmartTier, and high availability plus others...).
- Leverage SFCFSHA File snapshot technology that would let you take multiple copies of a large file.
- Page file caching for optimized read-IO access to shared pages for virtual machine disk access.

Figure 1-1 VirtualStore VMware Vsphere environment



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. When you install Symantec VirtualStore, the installer installs the `VRTSVxFeN` RPM, which includes the I/O fencing driver. To protect data on shared disks, you must configure I/O fencing after you install and configure Symantec VirtualStore.

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

I/O fencing coordination points can be coordinator disks or coordination point servers (CP servers) or both. You can configure disk-based or server-based I/O fencing:

Disk-based I/O fencing	<p>I/O fencing that uses coordinator disks is referred to as disk-based I/O fencing.</p> <p>Disk-based I/O fencing ensures data integrity in a single cluster.</p>
Server-based I/O fencing	<p>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.</p> <p>Server-based I/O fencing ensures data integrity in clusters.</p> <p>In virtualized environments that do not support SCSI-3 PR, Symantec VirtualStore supports non-SCSI-3 server-based I/O fencing.</p> <p>See “About I/O fencing for Symantec VirtualStore in virtual machines that do not support SCSI-3 PR” on page 22.</p>

See [“About planning to configure I/O fencing”](#) on page 59.

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

See the *Symantec VirtualStore Administrator's Guide*.

About configuring SVS clusters for data integrity

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

See [“About planning to configure I/O fencing”](#) on page 59.

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

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

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

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

See [“Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsvs”](#) on page 132.

See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 146.

About I/O fencing components

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

- Data disks—Store shared data
See [“About data disks”](#) on page 23.
- Coordination points—Act as a global lock during membership changes
See [“About coordination points”](#) on page 23.
See [“About coordination points”](#) on page 24.

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

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. SVS 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 SVS configuration.

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

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

- Coordination point servers

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

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

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

About coordination points

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

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

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 SVS configuration.

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

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

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.
- Disable preferred fencing policy to use the default node count-based race policy.

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

See [“Enabling or disabling the preferred fencing policy”](#) on page 151.

System requirements

This chapter includes the following topics:

- [Release notes](#)
- [Hardware compatibility list \(HCL\)](#)
- [Supported operating systems](#)
- [I/O fencing requirements](#)
- [Database requirements](#)
- [VxVM licenses](#)
- [Cross-Platform Data Sharing licensing](#)
- [Disk space requirements for Storage Foundation Cluster File System](#)
- [Discovering product versions and various requirement information](#)

Release notes

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

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

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

Hardware compatibility list (HCL)

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

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

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

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

Supported operating systems

For information on supported operating systems, see the *Symantec VirtualStore Release Notes*.

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 27.
- CP servers
See “[CP server requirements](#)” on page 28.

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

Coordinator disk requirements for I/O fencing

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

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

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

CP server requirements

Symantec VirtualStore 6.0.4 clusters (application clusters) support coordination point servers (CP servers) which are hosted on the following VCS and SFHA versions:

- VCS 6.0.1, VCS 6.0, VCS 6.0 PR1, VCS 6.0 RP1, VCS 5.1SP1, or VCS 5.1 single-node cluster
Single-node VCS clusters with VCS 5.1 SP1 RP1 and later or VCS 6.0 and later that hosts CP server does not require LLT and GAB to be configured.
- SFHA 6.0.1, SFHA 6.0, SFHA 6.0 PR1, SFHA 6.0 RP1, 5.1SP1, or 5.1 cluster

Warning: Before you upgrade 5.1 CP server nodes to use VCS or SFHA 6.0.4, you must upgrade all the application clusters that use this CP server to version 6.0.4. Application clusters at version 5.1 cannot communicate with CP server that runs VCS or SFHA 5.1 SP1 or later.

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

See the *Veritas Cluster Server Installation Guide* or the *Veritas Storage Foundation High Availability 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 2-1 lists additional requirements for hosting the CP server.

Table 2-1 CP server hardware requirements

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

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

Table 2-2 CP server supported operating systems and versions

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

Following are the CP server networking requirements and recommendations:

- Symantec recommends that network access from the application clusters to the CP servers should be made highly-available and redundant. The network connections require either a secure LAN or VPN.
- The CP server uses the TCP/IP protocol to connect to and communicate with the application clusters by these network paths. The CP server listens for messages from the application clusters using TCP port 14250. This is the default port that can be changed during a CP server configuration. Symantec recommends that you configure multiple network paths to access a CP server. If a network path fails, CP server does not require a restart and continues to listen on all the other available virtual IP addresses.
- The CP server supports either Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses) when communicating with the application clusters. If the CP server is configured to use an IPv6 virtual IP address, then the application clusters should also be on the IPv6 network where the CP server is hosted.
- When placing the CP servers within a specific network configuration, you must take into consideration the number of hops from the different application cluster nodes to the CP servers. As a best practice, Symantec recommends that the number of hops and network latency from the different application cluster nodes to the CP servers should be equal. This ensures that if an event occurs that

results in an I/O fencing scenario, there is no bias in the race due to difference in number of hops or network latency between the CPS and various nodes.

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

Communication mode	CP server in secure mode	CP server in non-secure mode
SVS cluster in secure mode	Yes	Yes
SVS cluster in non-secure mode	Yes	Yes

For secure communications between the SVS and CP server, consider the following requirements and suggestions:

- In a secure communication environment, all CP servers that are used by the application cluster must be configured with security enabled. A configuration where the application cluster uses some CP servers running with security enabled and other CP servers running with security disabled is not supported.
- For non-secure communication between CP server and application clusters, there is no need to configure Symantec Product Authentication Service. In non-secure mode, authorization is still provided by CP server for the application cluster users. The authorization that is performed only ensures that authorized users can perform appropriate actions as per their user privileges on the CP server.

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

Non-SCSI-3 I/O fencing requirements

Supported virtual environment for non-SCSI-3 fencing:

- VMware Server ESX 3.5, 4.0, 5.0, and 5.1 on AMD Opteron or Intel Xeon EM64T (x86_64)
Guest operating system: See the *Symantec VirtualStore Release Notes* for the list of supported Linux operating systems.

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

- Symantec VirtualStore must be configured with Cluster attribute UseFence set to SCSI3

- All coordination points must be CP servers

Database requirements

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

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

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

SVS does not support running SFDB tools with DB2 and Oracle.

VxVM licenses

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

Table 2-3 describes the levels of licensing in Veritas Volume Manager and supported features.

Table 2-3 Levels of licensing in Veritas Volume Manager and supported features

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

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

To see the license features that are enabled in VxVM

- ◆ Enter the following command:

```
# vxdctl license
```

Cross-Platform Data Sharing licensing

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

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

Disk space requirements for Storage Foundation Cluster File System

Table 2-4 shows the approximate disk space used by SFCFS for the required and optional packages.

Table 2-4 Approximate disk space use for SFCFS

English	/root	/opt	/usr	/var
Required Packages	78 MB	718 MB	85 MB	1 MB
All Packages	78 MB	740 MB	86 MB	1 MB

Discovering product versions and various requirement information

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

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

The information that the `version` option or the `showversion` script discovers on systems includes the following:

- The installed version of all released Storage Foundation and High Availability Suite of products
- The required RPMs or patches (if applicable) that are missing
- The available updates (including patches or hotfixes) from Symantec Operations Readiness Tools (SORT) for the installed products

To run the version checker

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

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

Planning to install Symantec VirtualStore

This chapter includes the following topics:

- [About planning for SVS installation](#)
- [About installation and configuration methods](#)
- [About the Veritas installer](#)
- [Assessing the system for installation readiness](#)
- [Downloading the Symantec VirtualStore software](#)
- [Setting environment variables](#)
- [Optimizing LLT media speed settings on private NICs](#)
- [Guidelines for setting the media speed of the LLT interconnects](#)
- [About using ssh or rsh with the Veritas installer](#)
- [Setting up shared storage](#)
- [Prerequisites for installing Symantec VirtualStore](#)
- [Hardware overview and requirements for Symantec VirtualStore](#)

About planning for SVS installation

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

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

Document version: 6.0.4 Rev 0.

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

Follow the preinstallation instructions if you are installing Symantec VirtualStore. The following Symantec VirtualStore is installed with these instructions:

- Symantec VirtualStore

About installation and configuration methods

You can use one of the following methods to install and configure SVS.

Table 3-1 Installation and configuration methods

Method	Description
<p>Interactive installation and configuration using the script-based installer</p> <p>Note: If you obtained SVS from an electronic download site, you must use the <code>installsvs</code> script instead of the <code>installer</code> script.</p>	<p>You can use one of the following script-based installers:</p> <ul style="list-style-type: none"> ■ Common product installer script: <code>installer</code> The common product installer script provides a menu that simplifies the selection of installation and configuration options. ■ Product-specific installation script: <code>installsvs</code> ■ The product-specific installation script provides command-line interface options. Installing and configuring with the <code>installsvs</code> script is identical to specifying SVS from the <code>installer</code> script. Use this method to install or configure only SVS.
<p>Silent installation using the response file</p>	<p>The response file automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. You can use the script-based installers with the response file to install silently on one or more systems.</p> <p>See “About response files” on page 332.</p>

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

Method	Description
Web-based installer	<p>The Web-based installer provides an interface to manage the installation and configuration from a remote site using a standard Web browser.</p> <p><code>webinstaller</code></p> <p>See “About the Web-based installer” on page 155.</p>
Manual installation and configuration	<p>Manual installation uses the Linux commands to install SVS. To retrieve a list of all RPMs and patches required for all products in the correct installation order, enter:</p> <pre># installsvs -allpkgs</pre> <p>Use the Linux commands to install SVS. Then use a manual or an interactive method with <code>installsvs</code> or <code>installer</code> script to configure the SVS stack.</p>

About the Veritas installer

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

- The general product installer. The general product installer enables you to install and configure the product, verify preinstallation requirements, and view the product’s description. You perform the installation from a disc, and you are prompted to choose a product to install.
See [“Installing Symantec VirtualStore using the installer”](#) on page 55.
- Product-specific installation scripts. If you obtained a standalone Veritas product from an electronic download site, the single product download files do not contain the general product installer. Use the product installation script to install the individual products. You can find these scripts at the root of the product media in the scripts directory. These scripts are also installed with the product.

[Table 3-2](#) lists all the SFHA Solutions product installation scripts. The list of product installation scripts that you find on your system depends on the product that you install on your system.

Note: The name of the script is different depending on whether you run the script from the install media or from a system on which the product software is installed.

Table 3-2 Product installation scripts

Veritas product name	Product installation script (When running the script from the install media)	Product installation script (When running the script from a system on which the SFHA Solutions product is installed)
Veritas Cluster Server (VCS)	<code>installvcs</code>	<code>installvcs<version></code>
Veritas Storage Foundation (SF)	<code>installsf</code>	<code>installsf<version></code>
Veritas Storage Foundation and High Availability (SFHA)	<code>installsfha</code>	<code>installsfha<version></code>
Veritas Storage Foundation Cluster File System High Availability (SFCFSHA)	<code>installsfcfsha</code>	<code>installsfcfsha<version></code>
Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)	<code>installsfrac</code>	<code>installsfrac<version></code>
Veritas Dynamic Multi-Pathing	<code>installdmp</code>	<code>installdmp<version></code>
Symantec VirtualStore	<code>installsvs</code>	<code>installsvs<version></code>

The scripts that are installed on the system include the product version in the script name. For example, to install the SVS script from the install media, run the `installsvs` command. However, to run the script from the installed binaries, run the `installsvs<version>` command.

For example, for the 6.0.4 version:

```
# /opt/VRTS/install/installsvs604 -configure
```

Note: Do not include the release version if you use the general product installer to install the product.

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

- Use `b` (back) to return to a previous section of the installation procedure. The back feature of the installation scripts is context-sensitive, so it returns to the beginning of a grouped section of questions.

- Use `Control+c` to stop and exit the program if an installation procedure hangs. After a short delay, the script exits.
- Use `q` to quit the installer.
- Use `?` to display help information.
- Use the Enter button to accept a default response.

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 6.0.4.

Prechecking your systems using the installer Performs a pre-installation check on the specified systems. The Veritas product installer reports whether the specified systems meet the minimum requirements for installing 6.0.4.

See [“Prechecking your systems using the Veritas installer”](#) on page 39.

Prechecking your systems using the Veritas installer

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

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

To use the precheck option

- 1 Start the script-based or Web-based installer.
 See [“Installing SVS with the Web-based installer”](#) on page 159.
- 2 Select the precheck option:
 - From the Web-based installer, select the **Perform a Pre-Installation Check** from the Task pull-down menu.
 - In the script-based installer, from root on the system where you want to perform the check, start the installer.

```
# ./installer
```

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

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

Downloading the Symantec VirtualStore software

One method of obtaining the Symantec VirtualStore software is to download it to your local system from the Symantec Web site.

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

To download the trialware version of the software

- 1 Open the following link in your browser:
<http://www.symantec.com/index.jsp>
- 2 In Products and Solutions section, click the **Trialware & Downloads** link.
- 3 On the next page near the bottom of the page, click **Business Continuity**.
- 4 Under Cluster Server, click **Download Now**.
- 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 Veritas product, the single product download files do not contain the product installer. Use the installation script for the specific product to install the product.

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

See [“About the Veritas installer”](#) on page 37.

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.

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

```
# df -k filesystem
```

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

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

Setting environment variables

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

After installation, SVS commands are in `/opt/VRTS/bin`. SVS 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 are using a shell such as `sh` or `bash`, enter the following:

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

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

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

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

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

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

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

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

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

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.

About using ssh or rsh with the Veritas installer

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

See [“Installation script options”](#) on page 322.

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

- When you perform installer sessions using a response file.

See [“About configuring secure shell or remote shell communication modes before installing products”](#) on page 339.

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: Symantec VirtualStore 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 59.

See also the *Symantec VirtualStore 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/rdisk/disk-group/volume
```

For example, enter the following command:

```
# mkfs -t vxfs /dev/vx/rdisk/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 115.

Prerequisites for installing Symantec VirtualStore

Each cluster node must be connected to the public network and each must have a unique host name by which it can be addressed on the public network. The local node from which you install does not have to be part of the cluster.

Provide the following information when installing SVS:

- The cluster name, beginning with a letter (a-z, A-Z).
- A unique ID from 0-65535 for the cluster. Within the public subnet, a new cluster using a duplicate cluster ID can cause existing clusters to fail.
- The host names of the cluster nodes.
- The device names of the network interface cards (NICs) used for the private networks among nodes.

- Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities as root on all cluster nodes or remote systems.
- Symantec recommends configuring the cluster with I/O fencing enabled. I/O fencing requires shared devices to support SCSI-3 Persistent Reservations (PR). Enabling I/O fencing prevents data corruption caused by a split brain scenario.

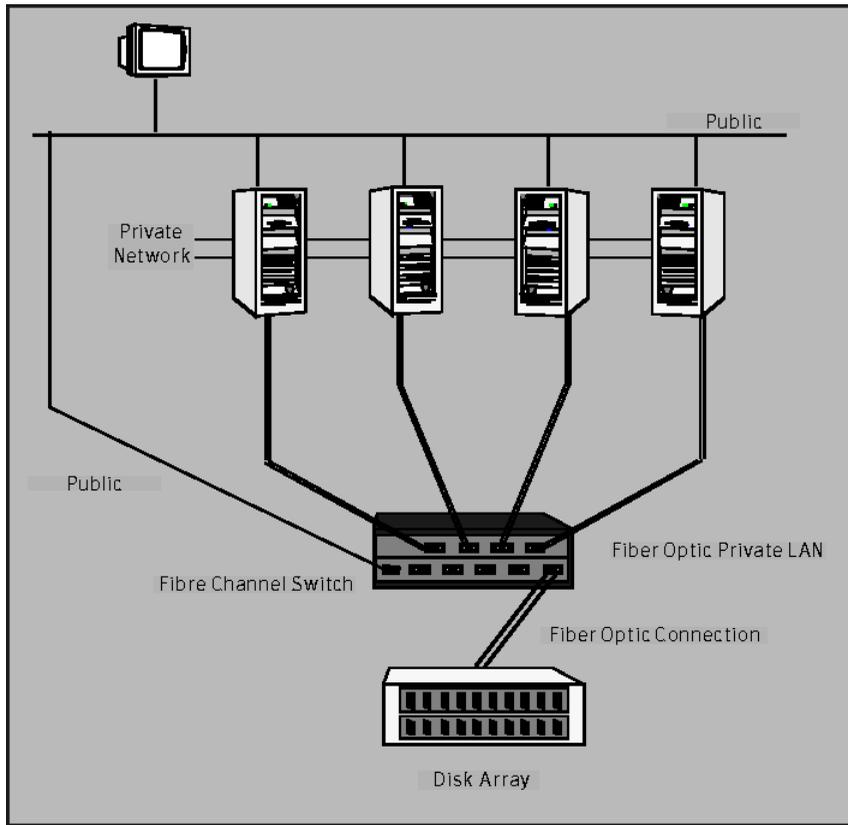
The Symantec VirtualStore is supported without I/O fencing enabled. However, without I/O fencing enabled, split brain scenarios can result in data corruption.

Hardware overview and requirements for Symantec VirtualStore

VxFS cluster functionality runs optimally on a Fibre Channel fabric. Fibre Channel technology provides the fastest, most reliable, and highest bandwidth connectivity currently available. By employing Fibre Channel technology, SVS can be used in conjunction with the latest Veritas Storage Area Network (SAN) applications to provide a complete data storage and retrieval solution.

[Figure 3-1](#) shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

Figure 3-1 Four Node SVS Cluster Built on Fibre Channel Fabric



Shared storage

Shared storage can be one or more shared disks or a disk array connected either directly to the nodes of the cluster or through a Fibre Channel Switch. Nodes can also have non-shared or local devices on a local I/O channel. It is advisable to have `/`, `/usr`, `/var` and other system partitions on local devices.

Fibre Channel switch

Each node in the cluster must have a Fibre Channel I/O channel to access shared storage devices. The primary component of the Fibre Channel fabric is the Fibre Channel switch.

Cluster platforms

There are several hardware platforms that can function as nodes in a Symantec VirtualStore (SVS) cluster.

See the *Symantec VirtualStore Release Notes*.

Note: For a cluster to work correctly, all nodes must have the same time. If you are not running the Network Time Protocol (NTP) daemon, make sure the time on all the systems comprising your cluster is synchronized.

Disk space requirements for Storage Foundation Cluster File System

[Table 3-3](#) shows the approximate disk space used by SFCFS for the required and optional packages.

Table 3-3 Approximate disk space use for SFCFS

English	/root	/opt	/usr	/var
Required Packages	58 MB	634 MB	50 MB	0.02 MB
All Packages	58 MB	689 MB	50 MB	0.68 MB

Licensing Symantec VirtualStore

This chapter includes the following topics:

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

About Veritas product licensing

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

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

www.symantec.com/techsupp/

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

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

The installer prompts for the product modes and options that you want to install, and then sets the required product level.

Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled. If you do not comply with the above terms, continuing to use the Symantec product is a violation of your end user license agreement, and results in warning messages.

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

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

If you upgrade to this release from a prior release of the Veritas software, the 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 Veritas 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: In order to change from one product group to another, you may need to perform additional steps.

Setting or changing the product level for keyless licensing

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

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.

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

To set or change the product level

- 1 Change your current working directory:

```
# cd /opt/VRTSvlic/bin
```

- 2 View the current setting for the product level.

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

- 3 View the possible settings for the product level.

```
# ./vxkeyless displayall
```

- 4 Set the desired product level.

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

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

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

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

To clear the product license level

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

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

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

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

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

Installing Veritas product license keys

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

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

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

To install a new license

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

```
# cd /opt/VRTS/bin  
  
# ./vxlicinst -k license key
```

To see a list of your `vxkeyless` keys, enter the following command:

```
# ./vxkeyless display
```

After you upgrade from a previous release, the output you see when you run the `vxkeyless display` command includes the previous release's `vxkeyless` keys. Each `vxkeyless` key name includes the suffix `_<previous_release_version>`. For example, `DMP_6.0`, or `SFENT_VR_5.1SP1`, or `VCS_GCO_5.1`. During the upgrade process, the CPI installer prompts you to update the `vxkeyless` keys to the current release level. If you update the `vxkeyless` keys during the upgrade process, you no longer see the `_<previous_release_number>` suffix after the keys are updated.

Installation using the script-based installer

- [Chapter 5. Installing Symantec VirtualStore using the script-based installer](#)
- [Chapter 6. Preparing to configure Symantec VirtualStore](#)
- [Chapter 7. Configuring Symantec VirtualStore](#)
- [Chapter 8. Configuring SVS for data integrity](#)

Installing Symantec VirtualStore using the script-based installer

This chapter includes the following topics:

- [Installing Symantec VirtualStore using the installer](#)
- [Configuring Symantec VirtualStore in secure mode](#)

Installing Symantec VirtualStore using the installer

The following walkthrough is to install Symantec VirtualStore on two systems: "host1" and "host2".

To install Symantec VirtualStore

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

Note: The `installer` program can configure both SSH and RSH communications.

- 2 Load and mount the software disc.
- 3 Move to the top-level directory on the disc.

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

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

```
# ./installer
```

- 5 Enter `I` to install and press Return.
- 6 When the list of available products appears, select Symantec VirtualStore (SVS). Enter the corresponding number, and press Return.
- 7 At the prompt, specify whether you accept the terms of the End User License Agreement (EULA).

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

- 8 Choose the mode you want to install:

- 1) Typical
- 2) Custom

```
Choose the mode you would like to install SVS: [1-2,q,?] (1)
```

If you chose the **Typical** install mode, the recommended RPMs will be installed by default. Skip to step 10.

If you chose the **Custom** install mode, proceed to step 9.

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

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

where *platform* indicates the operating system.

- 11 You have the option to let the installer configure SSH or RSH communications between the systems. If you choose to allow this configuration, select the shell and provide the root passwords for each system.
- 12 After the system checks complete, the installer displays a list of the patches and RPMs that will be installed. For **Custom** install mode, press Enter to continue with the installation.
- 13 After installation completes, choose your licensing method.

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

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

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

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

If you have a valid license key, select 1 and enter the license key at the prompt. Skip to step 17.

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

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

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

If you chose the **Custom** install mode, proceed to step 14.

If you chose the **Typical** install mode, the `installer` takes you directly to the configuration.

- 14 If you are going to use the Veritas Volume Replicator, enter `y` at the following prompt:

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

- 15 For typical install mode, SVS is automatically configured and started.
- 16 At the prompt, specify whether you want to send your installation information to Symantec.

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

- 17 View the log file, if needed, to confirm the installation.

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

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

Configuring Symantec VirtualStore in secure mode

Configuring SVS in secure mode ensures that all the communication between the systems is encrypted and users are verified against security credentials. SVS user names and passwords are not used when a cluster is running in secure mode. You can select the secure mode to be FIPS compliant while configuring the secure mode.

To configure SVS in secure mode

- 1 Enter appropriate choices when the installer prompts you:

```
Would you like to configure the VCS cluster in
secure mode [y,n,q] (n) y
1. Configure the cluster in secure mode without FIPS
2. Configure the cluster in secure mode with FIPS
3. Back to previous menu
Select the option you would like to perform [1-2,b,q] (1) 2
```

- 2 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.

Preparing to configure Symantec VirtualStore

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

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

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

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

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

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

See [Figure 6-2](#) on page 62.

[Figure 6-1](#) illustrates a high-level flowchart to configure I/O fencing for the Symantec VirtualStore cluster.

Figure 6-1 Workflow to configure I/O fencing

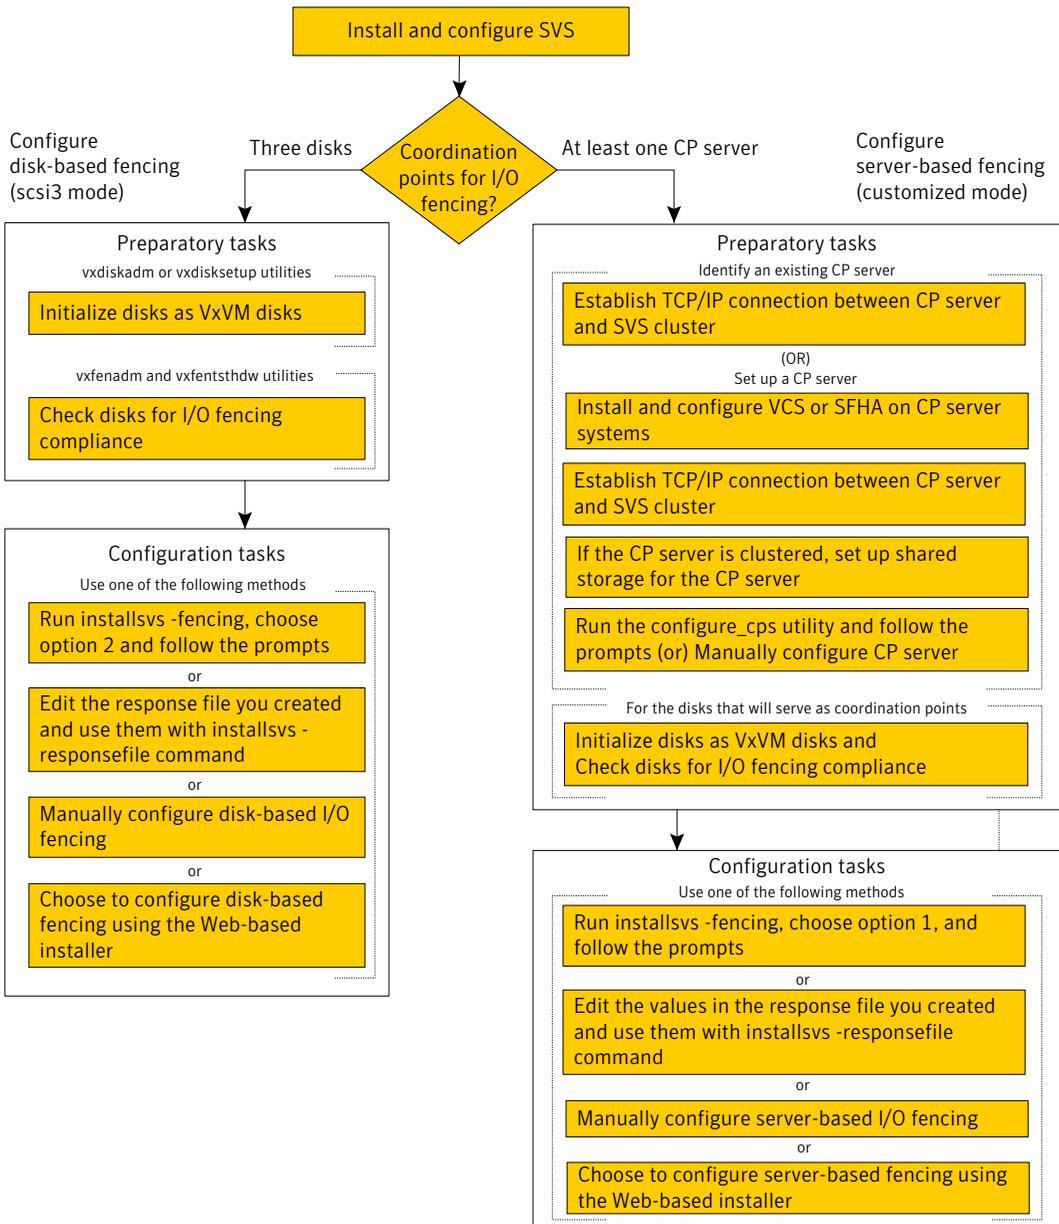
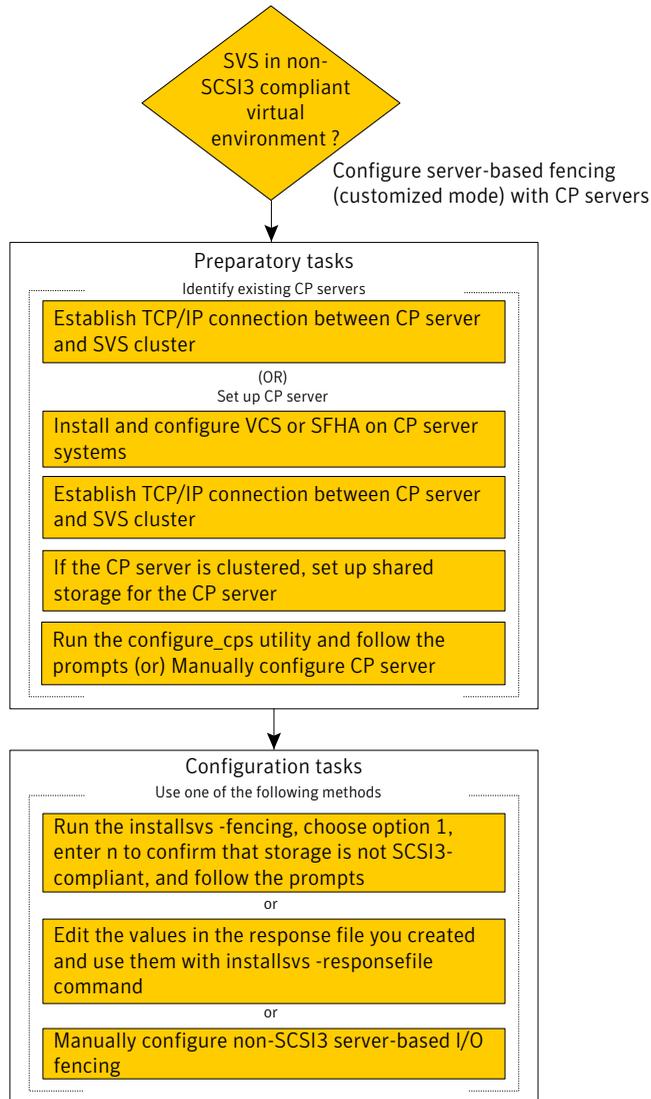


Figure 6-2 illustrates a high-level flowchart to configure non-SCSI-3 server-based I/O fencing for the Symantec VirtualStore cluster in virtual environments that do not support SCSI-3 PR.

Figure 6-2 Workflow to configure non-SCSI-3 server-based I/O fencing



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

- Using the installsvs
 - See [“Setting up disk-based I/O fencing using installsvs”](#) on page 110.
 - See [“Setting up server-based I/O fencing using installsvs”](#) on page 124.
 - See [“Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsvs”](#) on page 132.
- Using the Web-based installer
 - See [“Configuring Symantec VirtualStore for data integrity using the Web-based installer”](#) on page 166.
- Using response files
 - See [“Response file variables to configure disk-based I/O fencing”](#) on page 189.
 - See [“Response file variables to configure server-based I/O fencing”](#) on page 192.
 - See [“Response file variables to configure non-SCSI-3 server-based I/O fencing”](#) on page 194.
 - See [“Configuring I/O fencing using response files”](#) on page 188.
- Manually editing configuration files
 - See [“Setting up disk-based I/O fencing manually”](#) on page 118.
 - See [“Setting up server-based I/O fencing manually”](#) on page 134.
 - See [“Setting up non-SCSI-3 fencing in virtual environments manually”](#) on page 146.

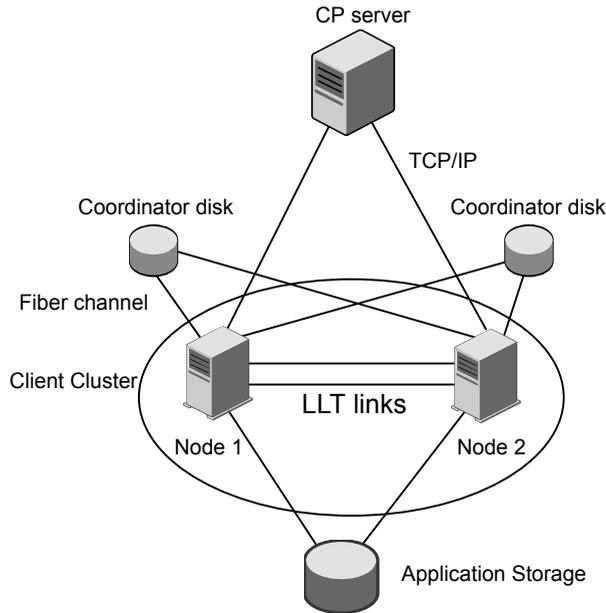
You can also migrate from one I/O fencing configuration to another.

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

Typical SVS cluster configuration with server-based I/O fencing

[Figure 6-3](#) displays a configuration using a SVS cluster (with two nodes), a single CP server, and two coordinator disks. The nodes within the SVS cluster are connected to and communicate with each other using LLT links.

Figure 6-3 CP server, SVS cluster, and coordinator disks



Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points
 See [Figure 6-4](#) on page 65.
- Multiple application clusters use a single CP server and single or multiple pairs of coordinator disks (two) as their coordination points
 See [Figure 6-5](#) on page 66.
- 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 6-6](#) on page 66.

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 6-4 displays a configuration using three CP servers that are connected to multiple application clusters.

Figure 6-4 Three CP servers connecting to multiple application clusters

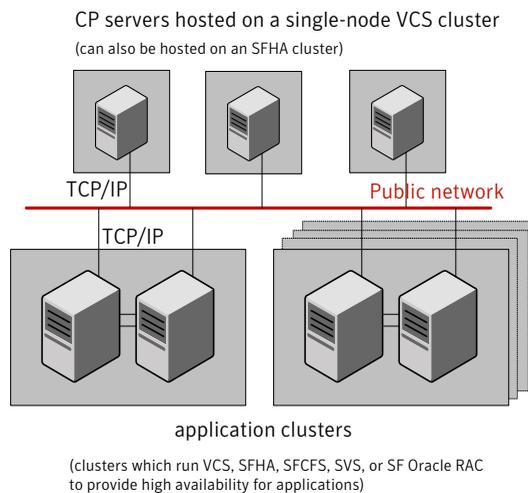


Figure 6-5 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.

Figure 6-5 Single CP server with two coordinator disks for each application cluster

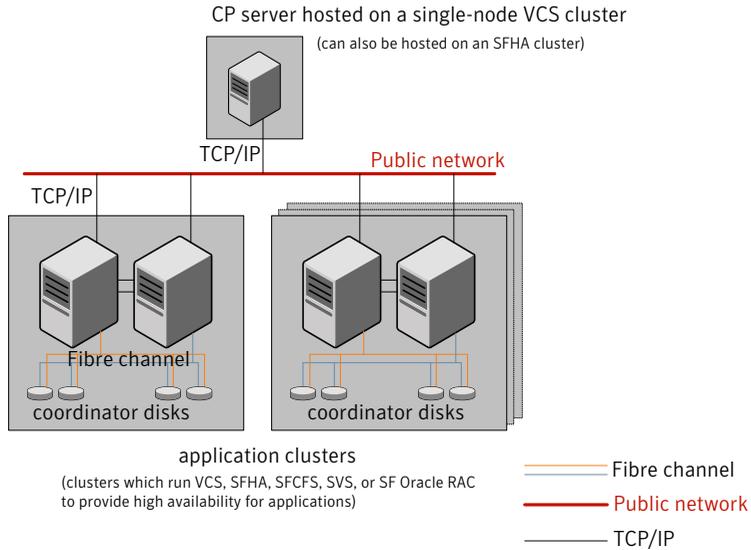
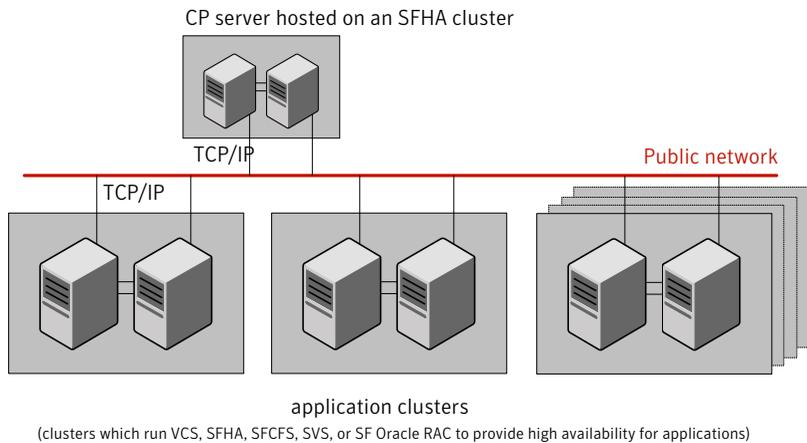


Figure 6-6 displays a configuration using a single CP server that is connected to multiple application clusters.

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



See [“Configuration diagrams for setting up server-based I/O fencing”](#) on page 348.

Setting up the CP server

[Table 6-1](#) lists the tasks to set up the CP server for server-based I/O fencing.

Table 6-1 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 67.
Install the CP server	See “Installing the CP server using the installer” on page 68.
Configure the CP server cluster in secure mode	See “Configuring the CP server cluster in secure mode” on page 70.
Set up shared storage for the CP server database	See “Setting up shared storage for the CP server database” on page 71.
Configure the CP server	See “Configuring the CP server using the installer program” on page 72. See “Configuring the CP server using the Web-based installer” on page 69. See “Configuring the CP server manually” on page 81. See “Configuring CP server using response files” on page 82.
Verify the CP server configuration	See “Verifying the CP server configuration” on page 86.

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 SVS cluster (application cluster) with a single CP server as coordination point or with at least three coordination points.
Symantec recommends using at least three coordination points.
- 3 Decide whether you want to configure the CP server cluster in secure mode.
Symantec recommends configuring the CP server cluster in secure mode to secure the communication between the CP server and its clients (SVS clusters). It also secures the HAD communication on the CP server cluster.
- 4 Set up the hardware and network for your CP server.
See [“CP server requirements”](#) on page 28.
- 5 Have the following information handy for CP server configuration:
- Name for the CP server
The CP server name should not contain any special characters. CP server name can include alphanumeric characters, underscore, and hyphen.
 - Port number for the CP server
Allocate a TCP/IP port for use by the CP server.
Valid port range is between 49152 and 65535. The default port number is 14250.
 - Virtual IP address, network interface, netmask, and networkhosts for the CP server
You can configure multiple virtual IP addresses for the CP server.

Installing the CP server using the installer

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

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

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

- CP server setup uses a single system
- Install and configure VCS to create a single-node VCS cluster.
- During installation, make sure to select all RPMs for installation. The VRTScps RPM is installed only if you select to install all RPMs.
- See the *Veritas Cluster Server Installation Guide* for instructions on installing and configuring VCS.
- Proceed to configure the CP server.
- See “[Configuring the CP server using the installer program](#)” on page 72.
- See “[Configuring the CP server manually](#)” on page 81.
- CP server setup uses multiple systems
- Install and configure SFHA to create an SFHA cluster. This makes the CP server highly available.
- Meet the following requirements for CP server:
- During installation, make sure to select all RPMs for installation. The VRTScps RPM is installed only if you select to install all RPMs.
 - During configuration, configure disk-based fencing (scsi3 mode).
- See the *Veritas Storage Foundation and High Availability Installation Guide* for instructions on installing and configuring SFHA.
- Proceed to set up shared storage for the CP server database.

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 Symantec VirtualStore on a cluster

- 1 Start the Web-based installer.
- 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 VirtualStore

Click **Next**.
- 3 On the Select Cluster page, enter the system names where you want to configure Symantec VirtualStore and click **Next**.
- 4 In the Confirmation dialog box, verify cluster information is correct and choose whether or not to configure I/O fencing.
 - To configure I/O fencing, click **Yes**.

- To configure I/O fencing later, click **No**.
- 5 On the Select Option page, select Configure CP Server on VCS 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.

Configuring the CP server cluster in secure mode

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

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

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

To configure the CP server cluster in secure mode

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

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

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

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

See [“About the Veritas installer”](#) on page 37.

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 Veritas installer”](#) on page 37.

Setting up shared storage for the CP server database

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

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

To set up shared storage for the CP server database

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

For example:

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

- 2 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/rdisk/cps_dg/cps_volume
```

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

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

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

Configuring the CP server using the 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 73.

For CP servers on an SFHA cluster: See ["To configure the CP server on an SFHA cluster"](#) on page 77.

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

1 Verify that the `VRTScps` package is installed on the node.

2 Run the `installvcs<version>` program with the `configcps` option.

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

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

See [“About the Veritas installer”](#) on page 37.

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

6 Enter the name of the CP Server.

```
Enter the name of the CP Server: [b] mycpserver1
```

- 7 Enter valid virtual IP addresses for the CP Server. A CP Server can be configured with more than one virtual IP address. You can also use IPv6 address.

Enter valid IP addresses for Virtual IPs for the CP Server, separated by space [b] 10.200.58.231 10.200.58.232

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.

- 8 Enter the corresponding CP server port number for each virtual IP address or press Enter to accept the default value (14250).

Enter corresponding port number for each Virtual IP address in the range [49152, 65535], separated by space, or simply accept the default port suggested: [b] (14250) 65535

- 9 Choose whether the communication between the CP server and the VCS clusters has to be made secure. If you have not configured the CP server cluster in secure mode, enter **n** at the prompt.

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

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**

- 10 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)

11 Verify and confirm the CP server configuration information.

```
CP Server configuration verification:
-----
CP Server Name: mycpserver1
CP Server Virtual IP(s): 10.200.58.231, 10.200.58.232
CP Server Port(s): 65535, 14250
CP Server Security: 0
CP Server Database Dir: /etc/VRTScps/db
-----
```

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

12 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
```

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

```
Enter the number of NIC resources that you want to configure.
You must use a public NIC.
Enter how many NIC resources you want to configure (1 to 2): 2
```

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

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

```
Enter a valid network interface on linux92216 for NIC resource - 1: eth0
Enter a valid network interface on linux92216 for NIC resource - 2: eth1
```

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

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

16 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 linux92216? [y,n,q] y
Enter a valid IP address to configure NetworkHosts for NIC eth0
on system linux92216: 10.200.56.22
```

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

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

```
Enter the netmask for virtual IP 10.200.58.231: (255.255.252.0)
Enter the netmask for virtual IP 10.200.58.232: (255.255.252.0)
```

18 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 Veritas Coordination Point Server is ONLINE
```

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

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

```
For example:
# hagrps -state CPSSG
#Group Attribute System Value
CPSSG State... |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 *Veritas Cluster Server Administrator's Guide*.

To configure the CP server on an SFHA cluster

- 1 Verify that the `VRTScps` package 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 Veritas installer”](#) on page 37.

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

- 7 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. You can also use IPv6 address.

```
Enter valid IP addresses for Virtual IPs for the CP Server,
separated by space [b] 10.200.58.231 10.200.58.232
```

- 9 Enter the corresponding CP server port number for each virtual IP address or press Enter to accept the default value (14250).

Enter corresponding port number for each Virtual IP address in the range [49152, 65535], separated by space, or simply accept the default port suggested: [b] **(14250) 65535**

- 10 Choose whether the communication between the CP server and the VCS clusters has to be made secure. If you have not configured the CP server cluster in secure mode, enter **n** at the prompt.

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

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)**

- 11 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**

- 12 Verify and confirm the CP server configuration information.

CP Server configuration verification:

```
CP Server Name: cps1
CP Server Virtual IP(s): 10.200.58.231, 10.200.58.232
CP Server Port(s): 65535, 14250
CP Server Security: 1
CP Server Database Dir: /cpsdb
```

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

13 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.
```

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

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

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

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

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

Enter a valid network interface on linux92216 for NIC resource - 1: eth0

Enter a valid network interface on linux92216 for NIC resource - 2: eth1

16 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

17 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 linux92216? [y,n,q] **y**

Enter a valid IP address to configure NetworkHosts for NIC eth0 on system linux92216: **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)**

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

```
Enter the netmask for virtual IP 10.200.58.231: (255.255.252.0)
Enter the netmask for virtual IP 10.200.58.232: (255.255.252.0)
```

- 19** 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
```

- 20** 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) cpsdgl
3) newcpsdg
```

- 21** 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

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

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

```
1) newcpsvol
```

- 24** 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.
```

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

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

- 26** 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 Veritas Coordination Point Server is ONLINE
The Veritas Coordination Point Server has been configured on your system.
```

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

```
For example:
# hagrps -state CPSSG
#Group Attribute System Value
CPSSG State 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 *Veritas Cluster Server Administrator's Guide*.

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

To manually configure the CP server

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

```
# hstop -local
```

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

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

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

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

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

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

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

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

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

- 5 Start VCS on all the cluster nodes.

```
# hstart
```

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

```
# hagrps -state CPSSG
```

Output similar to the following appears:

```
# Group Attribute System Value
CPSSG State 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 Veritas installer”](#) on page 37.

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 Veritas installer”](#) on page 37.

Response file variables to configure CP server

[Table 6-2](#) shows the response file variables required to configure CP server.

Table 6-2 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

Table 6-2 describes response file variables to configure CP server (*continued*)

Variable	List or Scalar	Description
CFG{cps_db_dir}	Scalar	This variable describes the absolute path of CP server database
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_vips}	List	This variable describes the virtual IP addresses for the CP server
CFG{cps_ports}	List	This variable describes the port number for the virtual IP addresses for the CP server
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
CFG{cps_newdgd_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

Table 6-2 describes response file variables to configure CP server (*continued*)

Variable	List or Scalar	Description
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

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

Review the response file variables and their definitions.

See [Table 6-2](#) on page 83.

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

$CFG{cps_db_dir}="/etc/VRTScps/db";
$CFG{cps_netmasks}=[ qw(255.255.252.0) ];
$CFG{cps_network_hosts}{cpsnic1}=[ qw(10.200.56.22) ];
$CFG{cps_nic_list}{cpsvip1}=[ qw(en0) ];
$CFG{cps_ports}=[ qw(14250) ];
$CFG{cps_security}=0;
$CFG{cps_singlenode_config}=1;
$CFG{cps_vip2nicres_map}{"10.200.58.233"}=1;
$CFG{cps_vips}=[ qw(10.200.58.233) ];
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{prod}="VCS60";
$CFG{systems}=[ qw(cps1) ];
$CFG{vcs_clusterid}=18523;
$CFG{vcs_clustername}="vcs92216";
```

1;

Sample response file for configuring the CP server on SFHA cluster

Review the response file variables and their definitions.

See [Table 6-2](#) on page 83.

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

$CFG{cps_db_dir}="/cpsdb";
$CFG{cps_diskgroup}="mycpsdg";
$CFG{cps_netmasks}=[ qw(255.255.252.0 255.255.252.0) ];
$CFG{cps_network_hosts}{cpsnic1}=[ qw(10.200.56.22) ];
$CFG{cps_network_hosts}{cpsnic2}=[ qw(10.200.56.22) ];
$CFG{cps_nic_list}{cpsvip1}=[ qw(eth0 eth1) ];
$CFG{cps_nic_list}{cpsvip2}=[ qw(eth0 eth1) ];
$CFG{cps_ports}=[ qw(65533 14250) ];
$CFG{cps_security}=1;
$CFG{cps_fips_mode}=0;
$CFG{cps_sfha_config}=1;
$CFG{cps_vip2nicres_map}{"10.200.58.231"}=1;
$CFG{cps_vip2nicres_map}{"10.200.58.232"}=2;
$CFG{cps_vips}=[ qw(10.200.58.231 10.200.58.232) ];
$CFG{cps_volume}="mycpsvol";
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{prod}="SFHA60";
$CFG{systems}=[ qw(cps1 cps2) ];
$CFG{vcs_clusterid}=46707;
$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)

2 Run the `cpsadm` command to check if the `vxcpserv` process is listening on the configured Virtual IP.

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

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

Configuring Symantec VirtualStore

This chapter includes the following topics:

- [Overview of tasks to configure Symantec VirtualStore using the script-based installer](#)
- [Configuring SVS using the script-based installer](#)

Overview of tasks to configure Symantec VirtualStore using the script-based installer

[Table 7-1](#) lists the tasks that are involved in configuring Symantec VirtualStore using the script-based installer.

Table 7-1 Tasks to configure Symantec VirtualStore using the script-based installer

Task	Reference
Start the software configuration	See “Starting the software configuration” on page 89.
Specify the systems where you want to configure Symantec VirtualStore	See “Specifying systems for configuration” on page 90.
Configure the basic cluster	See “Configuring private heartbeat links” on page 91.
Configure virtual IP address of the cluster (optional)	See “Configuring the virtual IP of the cluster” on page 95.

Table 7-1 Tasks to configure Symantec VirtualStore using the script-based installer (*continued*)

Task	Reference
Configure the cluster in secure mode (optional)	See " Configuring Symantec VirtualStore in secure mode " on page 58.
Add VCS users (required if you did not configure the cluster in secure mode)	See " Adding VCS users " on page 101.
Configure SMTP email notification (optional)	See " Configuring SMTP email notification " on page 102.
Configure SNMP email notification (optional)	See " Configuring SNMP trap notification " on page 103.
Complete the software configuration	See " Completing the SVS configuration " on page 106.

Configuring SVS using the script-based installer

Starting the software configuration

You can configure Symantec VirtualStore using the Veritas product installer or the `installsvs` command.

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

To configure Symantec VirtualStore 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:

Symantec VirtualStore

To configure Symantec VirtualStore using the `installsvs` program

- 1 Confirm that you are logged in as the superuser.
- 2 Start the `installsvs` program, with the configure option.

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

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

See ["About the Veritas installer"](#) on page 37.

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 Symantec VirtualStore. 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 Symantec VirtualStore.

```
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 remsh 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

- Verifies that Symantec VirtualStore is installed
 - Exits if 6.0.4 is not installed
- 3 Choose the mode you want to configure Symantec VirtualStore:

- 1) Typical
- 2) Custom

Choose the mode you would like to configure SVS: [1-2,q,?] (1)

- If you choose the **Typical** mode, the configuration automatically generates the cluster ID and automatically configures the network.
 - If you choose the **Custom** mode, then continue to follow the steps.
- 4 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 59.

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 the unique cluster name.

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

Configuring private heartbeat links

You now configure the private heartbeat links that LLT uses. VCS provides the option to use LLT over Ethernet or over UDP (User Datagram Protocol). Symantec recommends that you configure heartbeat links that use LLT over Ethernet 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.

See [“Using the UDP layer for LLT”](#) on page 366.

The following procedure helps you configure LLT over Ethernet.

To configure private heartbeat links

- 1 Choose one of the following options at the installer prompt based on whether you want to configure LLT over Ethernet or UDP.
 - Option 1: LLT over Ethernet (answer installer questions)
Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.
Skip to step 2.
 - Option 2: LLT over UDP (answer installer questions)
Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.
Skip to step 3.
 - Option 3: Automatically detect configuration for LLT over Ethernet
Allow the installer to automatically detect the heartbeat link details to configure LLT over Ethernet. The installer tries to detect all connected links between all systems.
Make sure that you activated the NICs for the installer to be able to detect and automatically configure the heartbeat links.
Skip to step 5.

Note: Option 3 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** 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.

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

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

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

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

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

- 6 Enter a unique cluster ID:

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

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

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

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

Configuring the virtual IP of the cluster

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

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

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

To configure the virtual IP of the cluster

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

```
Active NIC devices discovered on 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)
```

5 Enter the virtual IP address for the cluster.

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

For IPv4: ■ Enter the virtual IP address.

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

- Confirm the default netmask or enter another one:

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

- Verify and confirm the Cluster Virtual IP information.

Cluster Virtual IP verification:

```
NIC: eth0
IP: 192.168.1.16
Netmask: 255.255.240.0
```

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

For IPv6

- Enter the virtual IP address.

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

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

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

- Verify and confirm the Cluster Virtual IP information.

```
Cluster Virtual IP verification:
```

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

```
Is this information correct? [y,n,q] (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 97.

Configuring a secure cluster node by node

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

[Table 7-2](#) lists the tasks that you must perform to configure a secure cluster.

Table 7-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See “Configuring the first node” on page 98.
Configure security on the remaining nodes	See “Configuring the remaining nodes” on page 98.
Complete the manual configuration steps	See “Completing the secure cluster configuration” on page 99.

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/installsvs<version> -securityonenode
```

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

See [“About the Veritas installer”](#) on page 37.

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/installsvs<version> -securityonnode
```

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

See [“About the Veritas installer”](#) on page 37.

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

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

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

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

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

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

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

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

```
# /opt/VRTSvcs/bin/haconf -makerw
# /opt/VRTSvcs/bin/hagrp -list Frozen=0
# /opt/VRTSvcs/bin/hagrp -freeze groupname -persistent
# /opt/VRTSvcs/bin/haconf -dump -makero
```

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

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

- 3 On all nodes, stop the CmdServer.

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

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

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

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

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

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

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

- 7 On all nodes, start CmdServer.

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

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

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

```
# /opt/VRTSvcs/bin/hagrp -list Frozen=1
```

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

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

Adding VCS users

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

To add VCS users

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

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

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

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

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

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

```
Enter Again:*****
```

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

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

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

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

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP email services. You need to provide the SMTP server name and email addresses of people to be notified. Note that you can also configure the notification after installation.

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

To configure SMTP email notification

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

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

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

See [“Configuring SNMP trap notification”](#) on page 103.

- 3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on sys1: eth0
Enter the NIC for the VCS Notifier to use on sys1:
[b,q,?] (eth0)
Is eth0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SMTP server's host name.

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

- Enter the email address of each recipient.

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

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

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

4 Add more SMTP recipients, if necessary.

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

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

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

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

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

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

5 Verify and confirm the SMTP notification information.

```
NIC: eth0
```

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

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

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

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

Configuring SNMP trap notification

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

Note that you can also configure the notification after installation.

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

To configure the SNMP trap notification

- 1 Review the required information to configure the SNMP notification feature of VCS.

- 2 Specify whether you want to configure the SNMP notification.

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

See [“Configuring global clusters”](#) on page 105.

- 3 Provide information to configure SNMP trap notification.

Provide the following information:

- Enter the NIC information.

```
Active NIC devices discovered on sys1: eth0
Enter the NIC for the VCS Notifier to use on sys1:
[b,q,?] (eth0)
Is eth0 to be the public NIC used by all systems?
[y,n,q,b,?] (y)
```

- Enter the SNMP trap daemon port.

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

- Enter the SNMP console system name.

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

```
NIC: eth0
```

```
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.

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 VirtualStore Installation and Configuration Guide* for instructions to set up Symantec VirtualStore global clusters.

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

To configure the global cluster option

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

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

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

3 Provide information to configure this cluster as global cluster.

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

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

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

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

For IPv4: Global Cluster Option configuration verification:

```
NIC: eth0
IP: 10.198.89.22
Netmask: 255.255.240.0
```

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

For IPv6: Global Cluster Option configuration verification:

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

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

Completing the SVS configuration

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

To complete the SVS configuration

1 If prompted, press Enter at the following prompt.

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

2 Review the output as the installer stops various processes and performs the configuration. The installer then restarts SVS 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 Symantec VirtualStore 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 SVS using response files”](#) on page 177.

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 Symantec VirtualStore, 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 108.

See [“Updating product licenses”](#) on page 108.

Checking licensing information on the system

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

To check licensing information

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

```
# vxlicrep
```

- 2 Review the 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.

Updating product licenses

You can use the `./installer -license` command or the `vxlicinst -k` to add the Symantec VirtualStore license key on each node. If you have Symantec VirtualStore already installed and configured and you use a demo license, you can replace the demo license.

See [“Replacing a Symantec VirtualStore demo license with a permanent license”](#) on page 108.

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 Symantec VirtualStore demo license with a permanent license

When a Symantec VirtualStore 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 Symantec VirtualStore on all nodes in the cluster:

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

This command does not shut down any running applications.

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

```
# vxlicinst -k license key
```

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

```
# vxlicrep
```

- 5 Start Symantec VirtualStore on each node:

```
# hstart
```

Configuring SVS for data integrity

This chapter includes the following topics:

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

Setting up disk-based I/O fencing using installsvs

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

Configuring disk-based I/O fencing using installsvs

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

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

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

```
# /opt/VRTS/install/installsvs<version> -fencing
```

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

See [“About the Veritas installer”](#) on page 37.

The installsvs 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 Symantec VirtualStore 6.0.4 is configured properly.

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

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

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

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

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

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

- To use an existing disk group, enter the number corresponding to the disk group at the prompt.
 The program verifies whether the disk group you chose has an odd number of disks and that the disk group has a minimum of three disks.
- To create a new disk group, perform the following steps:
 - Enter the number corresponding to the **Create a new disk group** option.

The program lists the available disks that are in the CDS disk format in the cluster and asks you to choose an odd number of disks with at least three disks to be used as coordinator disks.

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

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

- Enter the numbers corresponding to the disks that you want to use as coordinator disks.
 - Enter the disk group name.
- 6 Verify that the coordinator disks you chose meet the I/O fencing requirements. You must verify that the disks are SCSI-3 PR compatible using the `vxfsntsthdw` utility and then return to this configuration program.
- See [“Checking shared disks for I/O fencing”](#) on page 114.
- 7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.
- 8 Enter the I/O fencing disk policy that you chose to use. For example:

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

The program also does the following:

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

- Starts VCS on each node to make sure that the Symantec VirtualStore is cleanly configured to use the I/O fencing feature.
- 11 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.
 - 12 Configure the Coordination Point Agent.

```
Do you want to configure Coordination Point Agent on
the client cluster? [y,n,q] (y)
```

- 13 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
```

- 14 Set the level two monitor frequency.

```
Do you want to set LevelTwoMonitorFreq? [y,n,q] (y)
```

- 15 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 143.

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

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

```
# fdisk -l
```

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

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

```
# vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i sdr
```

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

Checking shared disks for I/O fencing

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

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

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

See the *Symantec VirtualStore 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 114.
- Verifying that nodes have access to the same disk
See [“Verifying that the nodes have access to the same disk”](#) on page 115.
- Testing the shared disks for SCSI-3
See [“Testing the disks using vxfcntlshdw utility”](#) on page 116.

Verifying Array Support Library (ASL)

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

To verify Array Support Library (ASL)

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

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

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

The following output is a sample:

```
# vxddladm listsupport all
```

LIBNAME	VID	PID
libvxhitachi.so	HITACHI	DF350, DF400, DF400F, DF500, DF500F
libvxxp1281024.so	HP	All
libvxxp12k.so	HP	All
libvxdds2a.so	DDN	S2A 9550, S2A 9900, S2A 9700
libvxpurple.so	SUN	T300
libvxxiotechE5k.so	XIOTECH	ISE1400
libvxcopan.so	COPANSYS	8814, 8818
libvxibmids8k.so	IBM	2107

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

```
# vxdisk scandisks
```

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

Verifying that the nodes have access to the same disk

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

To verify that the nodes have access to the same disk

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

```
# vxfenadm -i diskpath
```

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

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

From node A, enter:

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

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

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

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

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

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

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

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

Testing the disks using `vxfcntlsthdw` utility

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

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

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

```
The disk /dev/sdx is ready to be configured for I/O Fencing on
node sys1
```

For more information on how to replace coordinator disks, refer to the *Symantec VirtualStore Administrator's Guide*.

To test the disks using `vxfcntlsthdw` utility

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

2 From one node, start the utility.

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

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

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

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

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

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

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

```
Enter the first node of the cluster: sys1
```

```
Enter the second node of the cluster: sys2
```

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

Enter the disk name to be checked for SCSI-3 PGR on node sys1 in the format:

for dmp: /dev/vx/rdmp/sdx

for raw: /dev/sdx

Make sure it's the same disk as seen by nodes sys1 and sys2

/dev/sdr

Enter the disk name to be checked for SCSI-3 PGR on node sys2 in the format:

for dmp: /dev/vx/rdmp/sdx

for raw: /dev/sdx

Make sure it's the same disk as seen by nodes sys1 and sys2

/dev/sdr

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

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

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

Setting up disk-based I/O fencing manually

Table 8-1 lists the tasks that are involved in setting up I/O fencing.

Table 8-1 Tasks to set up I/O fencing manually

Task	Reference
Initializing disks as VxVM disks	See "Initializing disks as VxVM disks" on page 113.
Identifying disks to use as coordinator disks	See "Identifying disks to use as coordinator disks" on page 123.

Table 8-1 Tasks to set up I/O fencing manually (*continued*)

Task	Reference
Checking shared disks for I/O fencing	See "Checking shared disks for I/O fencing" on page 114.
Setting up coordinator disk groups	See "Setting up coordinator disk groups" on page 119.
Creating I/O fencing configuration files	See "Creating I/O fencing configuration files" on page 120.
Modifying Symantec VirtualStore configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 121.
Configuring CoordPoint agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 143.
Verifying I/O fencing configuration	See "Verifying I/O fencing configuration" on page 122.

Setting up coordinator disk groups

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

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

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

The following example procedure assumes that the disks have the device names `sdx`, `sdz`, and `sdz`.

To create the `vxfencoorddg` disk group

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

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

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

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

3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

1 On each nodes, type:

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

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

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

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

- For DMP configuration:

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

- For raw device configuration:

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

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

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

- 4 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 commands, 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 copy of the `main.cf` file:

```
# 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.

```
# /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
```

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See [“Initializing disks as VxVM disks”](#) on page 113.

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 114.

Setting up server-based I/O fencing using installsvs

You can configure server-based I/O fencing for the Symantec VirtualStore cluster using the installsvs.

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 59.

See [“Recommended CP server configurations”](#) on page 64.

This section covers the following example procedures:

Mix of CP servers and coordinator disks

See [“To configure server-based fencing for the Symantec VirtualStore cluster \(one CP server and two coordinator disks\)”](#) on page 124.

Single CP server

See [“To configure server-based fencing for the Symantec VirtualStore cluster \(single CP server\)”](#) on page 129.

To configure server-based fencing for the Symantec VirtualStore 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 Symantec VirtualStore cluster. The Symantec VirtualStore cluster is also referred to as the application cluster or the client cluster.
See [“Setting up the CP server”](#) on page 67.
 - The coordination disks are verified for SCSI3-PR compliance.

See [“Checking shared disks for I/O fencing”](#) on page 114.

- 2 Start the installsvs with the `-fencing` option.

```
# /opt/VRTS/install/installsvs<version> -fencing
```

Where `<version>` is the specific release version. The installsvs starts with a copyright message and verifies the cluster information.

See [“About the Veritas installer”](#) on page 37.

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 Symantec VirtualStore 6.0.4 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 [1-4,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.

```
Enter the total number of Virtual IP addresses or fully
qualified host name for the
Coordination Point Server #1: [b,q,?] (1) 2
```

- 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 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] (14250)
```

8 Provide the following coordinator disks-related details at the installer prompt:

- Enter the I/O fencing disk policy for the coordinator disks.

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

- Choose the coordinator disks from the list of available disks that the installer displays. Ensure that the disk you choose is available from all the Symantec VirtualStore (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. 10.109.80.197 ([10.109.80.197]:14250)
SCSI-3 disks:
  1. sdx
  2. sdy
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk policy used for customized fencing: raw
```

The installer initializes the disks and the disk group and depots the disk group on the Symantec VirtualStore (application cluster) node.

10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the Symantec VirtualStore (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.

```
Updating client cluster information on Coordination Point Server 10.210.80.197

Adding the client cluster to the Coordination Point Server 10.210.80.197 ..... Done

Registering client node sys1 with Coordination Point Server 10.210.80.197..... Done
Adding CPClient user for communicating to Coordination Point Server 10.210.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.210.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.210.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.210.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.210.80.197 ..Done

Updating /etc/vxfenmode file on sys1 ..... Done
Updating /etc/vxfenmode file on sys2 ..... Done
```

See [“About I/O fencing configuration files”](#) on page 329.

- 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 Symantec VirtualStore (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)
```

- 15** Additionally the coordination point agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidentally 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
```

- 16** Note the location of the configuration log files, summary files, and response files that the installer displays for later use.

To configure server-based fencing for the Symantec VirtualStore cluster (single CP server)

- 1** Make sure that the CP server is configured and is reachable from the Symantec VirtualStore cluster. The Symantec VirtualStore cluster is also referred to as the application cluster or the client cluster.
- 2** See [“Setting up the CP server”](#) on page 67.
- 3** Start the `installsvs` with `-fencing` option.

```
# /opt/VRTS/install/installsvs<version> -fencing
```

Where `<version>` is the specific release version. The `installsvs` starts with a copyright message and verifies the cluster information.

See [“About the Veritas installer”](#) on page 37.

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 Symantec VirtualStore 6.0.4 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 [1-4,b,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.

```
Enter the total number of Virtual IP addresses or fully
qualified host name for the
Coordination Point Server #1: [b,q,?] (1) 2
```

- 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] (14250)
```

- 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]:14250)
```

- 10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the Symantec VirtualStore (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 10.210.80.197

Adding the client cluster to the Coordination Point Server 10.210.80.197 ..... Done

Registering client node sys1 with Coordination Point Server 10.210.80.197..... Done
Adding CPClient user for communicating to Coordination Point Server 10.210.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.210.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.210.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.210.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.210.80.197 .. Done

Updating /etc/vxfenmode file on sys1 ..... Done
Updating /etc/vxfenmode file on sys2 ..... Done
```

See [“About I/O fencing configuration files”](#) on page 329.

- 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.

Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsvs**14** Configure the CP agent on the Symantec VirtualStore (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.

Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsvs

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 installsvs in a non-SCSI-3 PR-compliant setup**1** Start the installsvs with `-fencing` option.

```
# /opt/VRTS/install/installsvs<version> -fencing
```

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

See [“About the Veritas installer”](#) on page 37.

The installsvs 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 Symantec VirtualStore 6.0.4 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type **1** to configure server-based I/O fencing.

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

Setting up non-SCSI-3 server-based I/O fencing in virtual environments using installsvs

- 4 Enter **n** to confirm that your storage environment does not support SCSI-3 PR.

```
Does your storage environment support SCSI3 PR?
[y,n,q] (y) n
```

- 5 Confirm that you want to proceed with the non-SCSI-3 I/O fencing configuration at the prompt.

- 6 Enter the number of CP server coordination points you want to use in your setup.

- 7 Enter the following details for each CP server:

- Enter the virtual IP address or the fully qualified host name.
- Enter the port address on which the CP server listens for connections. The default value is 14250. You can enter a different port address. Valid values are between 49152 and 65535.

The installer assumes that these values are identical from the view of the SVS cluster nodes that host the applications for high availability.

- 8 Verify and confirm the CP server information that you provided.

- 9 Verify and confirm the SVS cluster configuration information.

Review the output as the installer performs the following tasks:

- Updates the CP server configuration files on each CP server with the following details:
 - Registers each node of the SVS cluster with the CP server.
 - Adds CP server user to the CP server.
 - Adds SVS cluster to the CP server user.
- Updates the following configuration files on each node of the SVS cluster
 - `/etc/vxfenmode` file
 - `/etc/vxenviron` file
 - `/etc/sysconfig/vxfen` file
 - `/etc/llttab` file
 - `/etc/vxfentab` file

- 10 Review the output as the installer stops Symantec VirtualStore on each node, starts I/O fencing on each node, updates the VCS configuration file `main.cf`, and restarts Symantec VirtualStore with non-SCSI-3 server-based fencing.

Confirm to configure the CP agent on the SVS cluster.

- 11 Confirm whether you want to send the installation information to Symantec.
- 12 After the installer configures I/O fencing successfully, note the location of summary, log, and response files that installer creates.

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

Setting up server-based I/O fencing manually

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

Table 8-2 Tasks to set up server-based I/O fencing manually

Task	Reference
Preparing the CP servers for use by the Symantec VirtualStore cluster	See “Preparing the CP servers manually for use by the SVS cluster” on page 134.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See “Configuring server-based fencing on the SVS cluster manually” on page 137.
Modifying Symantec VirtualStore configuration to use I/O fencing	See “Modifying VCS configuration to use I/O fencing” on page 121.
Configuring Coordination Point agent to monitor coordination points	See “Configuring CoordPoint agent to monitor coordination points” on page 143.
Verifying the server-based I/O fencing configuration	See “Verifying server-based I/O fencing configuration” on page 145.

Preparing the CP servers manually for use by the SVS cluster

Use this procedure to manually prepare the CP server for use by the SVS cluster or clusters.

[Table 8-3](#) displays the sample values used in this procedure.

Table 8-3 Sample values in procedure

CP server configuration component	Sample name
CP server	cps1

Table 8-3 Sample values in procedure (*continued*)

CP server configuration component	Sample name
Node #1 - SVS cluster	sys1
Node #2 - SVS cluster	sys2
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the SVS cluster

- 1 Determine the cluster name and uuid on the SVS cluster.

For example, issue the following commands on one of the SVS cluster 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 SVS cluster 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 VirtualStore Administrator's Guide*.

3 Add the SVS cluster 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 enabled, check whether the CPSADM@VCS_SERVICES@cluster_uuid users are created in the CP server.

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

```
# cpsadm -s cps1.symantecexample.com -a list_users
```

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

If security is to be disabled, then add the user name "cpsclient@hostname" to the server instead of the CPSADM@VCS_SERVICES@cluster_uuid (for example, cpsclient@sys1).

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

5 Add the users to the CP server.

Issue the following commands on the CP server (cps1.symantecexample.com):

```
# cpsadm -s cps1.symantecexample.com -a add_user -e\  
CPSADM@VCS_SERVICES@cluster_uid\  
-f cps_operator -g vx
```

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

6 Authorize the CP server user to administer the SVS cluster. You must perform this task for the CP server users corresponding to each node in the SVS cluster.

For example, issue the following command on the CP server (cps1.symantecexample.com) for SVS cluster clus1 with two nodes sys1 and sys2:

```
# cpsadm -s cps1.symantecexample.com -a\  
add_clus_to_user -c clus1\  
-u {f0735332-1dd1-11b2}\  
-e CPSADM@VCS_SERVICES@cluster_uid\  
-f cps_operator -g vx
```

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

Configuring server-based fencing on the SVS cluster manually

The configuration process for the client or SVS cluster to use CP server as a coordination point requires editing the `/etc/vxfenmode` file.

You need to edit this file to specify the following information for your configuration:

- Fencing mode
- Fencing mechanism
- Fencing disk policy (if applicable to your I/O fencing configuration)
- Appropriate value for the security configuration
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)

Note: Whenever coordinator disks are used as coordination points in your I/O fencing configuration, you must create a disk group (vxencoordg). You must specify this disk group in the `/etc/vxfenmode` file.

See [“Setting up coordinator disk groups”](#) on page 119.

The customized fencing framework also generates the `/etc/vxfentab` file which has security setting and the coordination points (all the CP servers and disks from disk group specified in `/etc/vxfenmode` file).

To configure server-based fencing on the SVS cluster 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.

The following sample file output displays what the `/etc/vxfenmode` file contains:

See [“Sample vxfenmode file output for server-based fencing”](#) on page 138.

- 3 After editing the `/etc/vxfenmode` file, run the `vxfen` init script to start fencing.

For example:

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

- 4 Make sure that `/etc/vxfenmode` file contains the value of security is set to 1.

Make sure that following command displays the certificate being used by `cpsadm` client,

```
EAT_DATA_DIR=/vat/VRTSvc/vcsauth/data/CPSADM cpsat showcred
```

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
# security - 1
# security - 0

vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O Fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
# raw - connect to disks using the native interface
#
scsi3_disk_policy=dmp

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

#
# Specify 3 or more odd number of coordination points in this file,
# one in its own line. They can be all-CP servers, all-SCSI-3
# compliant coordinator disks, or a combination of CP servers and
# SCSI-3 compliant coordinator disks. Please ensure that the CP

```

```

# server coordination points are numbered sequentially and in the
# same order on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as:
#
# cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames over
# different subnets, all of the IPs/names can be specified in a
# comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,...,
# [<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
# is the serial number of the CPS as a coordination point; must
# start with 1.
# <vip>
# is the virtual IP address of the CPS, must be specified in
# square brackets ("[]").
# <vhn>
# is the virtual hostname of the CPS, must be specified in square
# brackets ("[]").
# <port>
# is the port number bound to a particular <vip/vhn> of the CPS.
# It is optional to specify a <port>. However, if specified, it
# must follow a colon (":") after <vip/vhn>. If not specified, the
# colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>, a
# default port can be specified as follows:
#
# port=<default_port>
#
# Where <default_port> is applicable to all the <vip/vhn>s for
# which a <port> is not specified. In other words, specifying <port>
# with a <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:

```

```
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[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
#
```

Table 8-4 defines the vxfenmode parameters that must be edited.

Table 8-4 vxfenmode file parameters

vxfenmode File Parameter	Description
vxfen_mode	Fencing mode of operation. This parameter must be set to "customized".
vxfen_mechanism	Fencing mechanism. This parameter defines the mechanism that is used for fencing. If one of the three coordination points is a CP server, then this parameter must be set to "cps".
scsi3_disk_policy	Configure the vxfen module to use either DMP devices, "dmp" or the underlying raw character devices, "raw". Note: The configured disk policy is applied on all the nodes.
security	Security parameter 1 indicates that secure mode is used for CP server communications. Security parameter 0 indicates that communication with the CP server is made in non-secure mode. The default security value is 1.
fips_mode	[For future use] Set the value to 0.
cps1, cps2, or vxfendg	Coordination point parameters. Enter either the virtual IP address or the FQHN (whichever is accessible) of the CP server. <i>cps<number>=[virtual_ip_address/virtual_host_name]:port</i> Where <i>port</i> is optional. The default port value is 14250. If you have configured multiple virtual IP addresses or host names over different subnets, you can specify these as comma-separated values. For example: <pre>cps1=[192.168.0.23], [192.168.0.24]:58888, [cps1.company.com]</pre> Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfencoordg) 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).

Table 8-4 vxfenmode file parameters (*continued*)

vxfenmode File Parameter	Description
port	<p>Default port for the CP server to listen on.</p> <p>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 14250. You can change this default port value using the port parameter.</p>
single_cp	<p>Value 1 for single_cp parameter indicates that the server-based fencing uses a single highly available CP server as its only coordination point.</p> <p>Value 0 for single_cp parameter indicates that the server-based fencing uses at least three coordination points.</p>

Configuring CoordPoint agent to monitor coordination points

The following procedure describes how to manually configure the CoordPoint agent to monitor coordination points.

The CoordPoint agent can monitor CP servers and SCSI-3 disks.

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

To configure CoordPoint agent to monitor coordination points

- 1 Ensure that your SVS cluster has been properly installed and configured with fencing enabled.
- 2 Create a parallel service group vxfen and add a coordpoint resource to the vxfen service group using the following commands:

```
# haconf -makerw
# hagrps -add vxfen
# hagrps -modify vxfen SystemList sys1 0 sys2 1
# hagrps -modify vxfen AutoFailOver 0
# hagrps -modify vxfen Parallel 1
# hagrps -modify vxfen SourceFile "./main.cf"
# hares -add coordpoint CoordPoint vxfen
# hares -modify coordpoint FaultTolerance 0
# hares -override coordpoint LevelTwoMonitorFreq
# hares -modify coordpoint LevelTwoMonitorFreq 5
# hares -modify coordpoint Enabled 1
# haconf -dump -makero
```

- 3 Verify the status of the agent on the SVS cluster using the `hares` commands. For example:

```
# hares -state coordpoint
```

The following is an example of the command and output::

```
# hares -state coordpoint

# Resource      Attribute  System  Value
coordpoint     State     sys1    ONLINE
coordpoint     State     sys2    ONLINE
```

- 4 Access the engine log to view the agent log. The agent log is written to the engine log.

The agent log contains detailed CoordPoint agent monitoring information; including information about whether the CoordPoint agent is able to access all the coordination points, information to check on which coordination points the CoordPoint agent is reporting missing keys, etc.

To view 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
```

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 VirtualStore Administrator's Guide*.

- 2 Verify that I/O fencing is using the specified coordination points by running the `vxfenconfig` command. For example, run the following command:

```
# vxfenconfig -l
```

If the output displays `single_cp=1`, it indicates that the application cluster uses a CP server as the single coordination point for server-based fencing.

Setting up non-SCSI-3 fencing in virtual environments manually

To manually set up I/O fencing in a non-SCSI-3 PR compliant setup

- 1 Configure I/O fencing in customized mode with only CP servers as coordination points.

See [“Setting up server-based I/O fencing manually”](#) on page 134.

- 2 Make sure that the Symantec VirtualStore cluster is online and check that the fencing mode is customized.

```
# vxfenadm -d
```

- 3 Make sure that the cluster attribute `UseFence` is set to `SCSI3`.

```
# haclus -value UseFence
```

- 4 On each node, edit the `/etc/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 SCSI3.
- 10 To make these VxFEN changes take effect, stop and restart VxFEN and the dependent modules

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

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

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

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

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

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

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

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

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

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

```

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

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

#
# Seconds for which vxmfend process wait for a customized fencing
# script to complete. Only used with vxmfen_mode=customized
vxmfen_script_timeout=25

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

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

```

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

Enabling or disabling the preferred fencing policy

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

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

See [“About preferred fencing”](#) on page 25.

To enable preferred fencing for the I/O fencing configuration

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

```
# vxfsadm -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 Veritas Cluster Server (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
```

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

- Make the VCS configuration writable.

```
# haconf -makerw
```

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

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

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

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

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

- Save the VCS configuration.

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

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

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

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

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

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

```
# haconf -makerw
```

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

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

Installation using the Web-based installer

- [Chapter 9. Installing Symantec VirtualStore using the Web-based installer](#)
- [Chapter 10. Configuring Symantec VirtualStore](#)

Installing Symantec VirtualStore using the Web-based installer

This chapter includes the following topics:

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

About the Web-based installer

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

You use the `webinstaller` script to start and stop the Veritas XPortal Server `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 is operating. While the installation processes are operating, the log files are located in a session-based directory under the `/var/tmp` directory. After the install process completes, the log files are located in the `/opt/VRTS/install/logs` directory. It is recommended that you keep these files for auditing, debugging, and future use.

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

See [“Before using the Veritas Web-based installer”](#) on page 156.

Before using the Veritas Web-based installer

The Veritas Web-based installer requires the following configuration.

Table 9-1 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Veritas products.	Must be a supported platform for 6.0.4.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must use the same operating system as the target systems and must be at one of the supported operating system update levels.
Administrative system	The system where you run the Web browser to perform the installation.	Must have a Web browser. Supported browsers: <ul style="list-style-type: none">■ Internet Explorer 6, 7, and 8■ Firefox 3.x and later

Starting the Veritas Web-based installer

This section describes starting the Veritas Web-based installer.

To start the Web-based installer

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

```
# ./webinstaller start
```

The webinstaller script displays a URL. Note this URL.

Note: If you do not see the URL, 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.

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

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

- 5 Log in as superuser.

Obtaining a security exception on Mozilla Firefox

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

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

To obtain a security exception

- 1 Click **Or you can add an exception** link.
- 2 Click **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 pre-installation check with the Veritas Web-based installer

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

To perform a pre-installation check

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

Setting installer options with the Web-based installer

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

The supported options follow:

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

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 SVS with the Web-based installer

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

To install SVS using the Web-based installer

- 1 Perform preliminary steps.
See [“Performing a pre-installation check with the Veritas Web-based installer”](#) on page 158.
- 2 Start the Web-based installer.

Note: The installer displays a warning and asks to install VMware vSphere Perl SDK on all nodes before proceeding to install. You can ignore the warning message and proceed with this step.

- 3 Select **Install a Product** from the **Task** drop-down list.
- 4 Select **Symantec VirtualStore** from the Product drop-down list, and click **Next**.
- 5 On the License agreement page, read the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 6 Select the installation mode: **Typical** or **Custom**.
Typical: automatically configures your SVS after installation with typical settings.
Custom: shows GUI for user inputs for each configuration steps.
- 7 Indicate the systems where you want to install. Separate multiple system names with spaces. Click **Next**.

- 8 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.
- 9 After the validation completes successfully, click **Next** to install SVS on the selected system.
- 10 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 SVS.

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

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

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

Click **Finish**. The installer asks if you would like to read the summary file. Select **Yes** to read the summary file. If you select **No**, the installer prompts you for another task.

Configuring Symantec VirtualStore

This chapter includes the following topics:

- [Configuring Symantec VirtualStore using the Web-based installer](#)

Configuring Symantec VirtualStore using the Web-based installer

Before you begin to configure Symantec VirtualStore 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 Symantec VirtualStore on a cluster

- 1 Start the Web-based installer.
- 2 On the Select a task and a product page, select the task and the product as follows:

Task	Configure a Product
Product	Symantec VirtualStore

Click **Next**.

- 3 On the Select Systems page, enter the system names where you want to configure Symantec VirtualStore, 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 Choose the mode you would like to configure your product.

If you choose the **Typical** mode, the configuration prompts you only one question about the cluster ID and automatically configuring networks. The cluster ID is the unique cluster ID that you entered during installation.

If you choose the **Custom** mode, then follow the steps below.

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

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

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

See [“Configuring Symantec VirtualStore for data integrity using the Web-based installer”](#) on page 166.

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

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

See [“About the Veritas installer”](#) on page 37.

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

- | | |
|--|---|
| Cluster Name | Enter a unique cluster name. |
| Cluster ID | Enter a unique cluster ID.

Note that you can have the installer check to see if the cluster ID is unique. Symantec recommends that you use the installer to check for duplicate cluster IDs in multi-cluster environments. |
| Check duplicate cluster ID | Select the check box if you want the installer to verify if the given cluster ID is unique in your private network. The verification is performed after you specify the heartbeat details in the following pages. The verification takes some time to complete. |
| LLT Type | Select an LLT type from the list. You can choose to configure LLT over UDP or over Ethernet. |
| Number of Heartbeats | Choose the number of heartbeat links you want to configure. |
| Additional Low Priority Heartbeat NIC | Select the check box if you want to configure a low priority link. The installer configures one heartbeat link as low priority link. |
| Unique Heartbeat NICs per system | For LLT over Ethernet, select the check box if you do not want to use the same NIC details to configure private heartbeat links on other systems.

For LLT over UDP, this check box is selected by default. |

Click **Next**.

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

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

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

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

Click **Next**.

- 8 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 SVS 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 `installsvs`.

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.

- 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**.

- 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 5, 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.

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

See [“Configuring Symantec VirtualStore for data integrity using the Web-based installer”](#) on page 166.

- 12 Click **Next** to complete the process of configuring Symantec VirtualStore.

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 Symantec VirtualStore for data integrity using the Web-based installer

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

See [“Configuring Symantec VirtualStore using the Web-based installer”](#) on page 161.

See [“About planning to configure I/O fencing”](#) on page 59.

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 167.
- See [“Configuring server-based fencing for data integrity using the Web-based installer”](#) on page 169.
- See [“Configuring fencing in disabled mode using the Web-based installer”](#) on page 170.
- See [“Online fencing migration mode using the Web-based installer”](#) on page 172.

Configuring disk-based fencing for data integrity using the Web-based installer

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

See [“Configuring Symantec VirtualStore using the Web-based installer”](#) on page 161.

See [“About planning to configure I/O fencing”](#) on page 59.

To configure Symantec VirtualStore for data integrity

- 1 Start the Web-based installer.
- 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 VirtualStore

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 server-based fencing in a virtual environment that is not SCSI-3 PR compliant.
- 7 On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.

Click **Next**.
- 8 On the Configure Fencing page, specify the following information:

- Select a Disk Group** Select the **Create a new disk group** option or select one of the disk groups from the list.
- If you selected one of the disk groups that is listed, choose the fencing disk policy for the disk group.
 - 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.
- Fencing Disk Policy** Choose the fencing disk policy for the disk group.

10 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.

11 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.

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

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

Configuring server-based fencing for data integrity using the Web-based installer

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

See [“Configuring Symantec VirtualStore using the Web-based installer”](#) on page 161.

See [“About planning to configure I/O fencing”](#) on page 59.

To configure Symantec VirtualStore for data integrity

- 1 Start the Web-based installer.
- 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 VirtualStore

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 server-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 server-based fencing in a virtual environment that is not SCSI-3 PR compliant.
- 7 On the Configure Fencing page, the installer prompts for details based on the fencing type you chose to configure. Specify the coordination points details.

Click **Next**.
- 8 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 Symantec VirtualStore, you must configure the cluster for data integrity. Review the configuration requirements.

See [“Configuring Symantec VirtualStore using the Web-based installer”](#) on page 161.

See [“About planning to configure I/O fencing”](#) on page 59.

To configure Symantec VirtualStore for data integrity

- 1 Start the Web-based installer.
- 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 VirtualStore

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.

Online fencing migration mode using the Web-based installer

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

See [“Configuring Symantec VirtualStore using the Web-based installer”](#) on page 161.

See [“About planning to configure I/O fencing”](#) on page 59.

To configure Symantec VirtualStore for data integrity

- 1 Start the Web-based installer.
- 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 VirtualStore

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 `Online fencing migration` 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**.

Note: If the coordination point servers are configured in secure mode, then the communication between coordination point servers and client servers happen in secure mode.

- 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 checkbox to specify whether you want to send your installation information to Symantec.
Click **Finish**. The installer prompts you for another task.

Automated installation using response files

- [Chapter 11. Performing an automated Symantec VirtualStore installation](#)
- [Chapter 12. Performing an automated Symantec VirtualStore configuration](#)
- [Chapter 13. Performing an automated I/O fencing configuration using response files](#)

Performing an automated Symantec VirtualStore installation

This chapter includes the following topics:

- [Installing SVS using response files](#)

Installing SVS using response files

Typically, you can use the response file that the installer generates after you perform SVS installation on one cluster to install SVS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To install SVS using response files

- 1 Make sure the systems where you want to install SVS meet the installation requirements.
- 2 Make sure the preinstallation tasks are completed.
- 3 Copy the response file to one of the cluster systems where you want to install SVS.
- 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  
# ./installsvs -responsefile /tmp/response_file
```

Where */tmp/response_file* is the response file's full path name.

See [“About the Veritas installer”](#) on page 37.

Performing an automated Symantec VirtualStore configuration

This chapter includes the following topics:

- [Configuring SVS using response files](#)
- [Response file variables to configure Symantec VirtualStore](#)

Configuring SVS using response files

Typically, you can use the response file that the installer generates after you perform SVS configuration on one cluster to configure SVS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

To configure SVS using response files

- 1 Make sure the SVS RPMs are installed on the systems where you want to configure SVS.
- 2 Copy the response file to one of the cluster systems where you want to configure SVS.

- 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 Symantec VirtualStore”](#) on page 178.

- 4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsvs<version>
-responsefile /tmp/response_file
```

Where *<version>* is the specific release version, and */tmp/response_file* is the response file’s full path name.

Response file variables to configure Symantec VirtualStore

[Table 12-1](#) lists the response file variables that you can define to configure SVS.

Table 12-1 Response file variables specific to configuring Symantec VirtualStore

Variable	List or Scalar	Description
\$CFG{config_cfs}	Scalar	Performs the Cluster File System configuration for SVS. (Required) Set the value to 1 to configure Cluster File System for SVS.
CFG{opt}{configure}	Scalar	Performs the configuration if the RPMs are already installed. (Required) Set the value to 1 to configure SVS.
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 12-1 Response file variables specific to configuring Symantec VirtualStore
(continued)

Variable	List or Scalar	Description
CFG{prod}	Scalar	Defines the product to be configured. The value is VCS60 for VCS. (Required)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{rsh}	Scalar	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is <i>/opt/VRTS/install/logs</i> . Note: The installer copies the response files and summary files also to the specified <i>logpath</i> location. (Optional)
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec Web site. The value 0 indicates that the installation logs are not uploaded to the Symantec Web site. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (*csgnic*, *csgvip*, and *csgnetmask*) must be defined if any are defined. The same is true for the SMTP notification (*smtpserver*, *smtprecp*, and *smtpsev*), the SNMP

trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Table 12-2 lists the response file variables that specify the required information to configure a basic Symantec VirtualStore cluster.

Table 12-2 Response file variables specific to configuring a basic Symantec VirtualStore cluster

Variable	List or Scalar	Description
CFG{vcs_clusterid}	Scalar	An integer between 0 and 65535 that uniquely identifies the cluster. (Required)
CFG{vcs_clustername}	Scalar	Defines the name of the cluster. (Required)
CFG{vcs_allowcomms}	Scalar	Indicates whether or not to start LLT and GAB when you set up a single-node cluster. The value can be 0 (do not start) or 1 (start). (Required)
CFG{fencingenabled}	Scalar	In a Symantec VirtualStore configuration, defines if fencing is enabled. Valid values are 0 or 1. (Required)

Table 12-3 lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 12-3 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. Two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. You must enclose the system name within double quotes. (Required)
CFG{vcs_lltlinklowpri#} {"system"}	Scalar	Defines a low priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication. If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on. You must enclose the system name within double quotes. (Optional)

Table 12-4 lists the response file variables that specify the required information to configure LLT over UDP.

Table 12-4 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)

Table 12-4 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG {vcs_udplinklowpri<n>_address} {<system1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_port} {<system1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 12-4 Response file variables specific to configuring LLT over UDP
(continued)

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_netmask} {<system1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 12-5 lists the response file variables that specify the required information to configure virtual IP for Symantec VirtualStore cluster.

Table 12-5 Response file variables specific to configuring virtual IP for Symantec VirtualStore 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 12-6 lists the response file variables that specify the required information to configure the Symantec VirtualStore cluster in secure mode.

Table 12-6 Response file variables specific to configuring Symantec VirtualStore cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonnode}	Scalar	Specifies that the securityonnode option is being used.
CFG{securityonnode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. <ul style="list-style-type: none"> ■ 1—Configure the first node ■ 2—Configure the other node
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.
CFG{opt}{fips}	Scalar	Specifies that the FIPS option is being used.
CFG{vcs_eat_security_fips}	Scalar	Specifies that the enabled security is FIPS compliant.

Table 12-7 lists the response file variables that specify the required information to configure VCS users.

Table 12-7 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)
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 12-8](#) lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 12-8 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_smtprecip}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)

Table 12-8 Response file variables specific to configuring VCS notifications using SMTP (*continued*)

Variable	List or Scalar	Description
CFG{vcs_smtpsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

Table 12-9 lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 12-9 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_snmpsev}	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 12-10 lists the response file variables that specify the required information to configure Symantec VirtualStore global clusters.

Table 12-10 Response file variables specific to configuring Symantec VirtualStore 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)

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 server-based I/O fencing](#)
- [Response file variables to configure non-SCSI-3 server-based I/O fencing](#)
- [Sample response file for configuring non-SCSI-3 server-based I/O fencing](#)

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for Symantec VirtualStore.

To configure I/O fencing using response files

- 1 Make sure that Symantec VirtualStore 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 59.

- 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 192.
 See [“Sample response file for configuring server-based I/O fencing”](#) on page 194.
- 4 Edit the values of the response file variables as necessary.
 See [“Response file variables to configure disk-based I/O fencing”](#) on page 189.
 See [“Response file variables to configure server-based I/O fencing”](#) on page 192.
- 5 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installsvs<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 Veritas installer”](#) on page 37.

Response file variables to configure disk-based I/O fencing

[Table 13-1](#) lists the response file variables that specify the required information to configure disk-based I/O fencing for SVS.

Table 13-1 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)

Table 13-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG{fencing_option}	Scalar	<p>Specifies the I/O fencing configuration mode.</p> <ul style="list-style-type: none"> ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled mode ■ 4—Fencing migration when the cluster is online <p>(Required)</p>
CFG {fencing_scsi3_disk_policy}	Scalar	<p>Specifies the I/O fencing mechanism.</p> <p>This variable is not required if you had configured fencing in disabled mode. For disk-based fencing, you must configure the <code>fencing_scsi3_disk_policy</code> variable and either the <code>fencing_dgname</code> variable or the <code>fencing_newdg_disks</code> variable.</p> <p>(Optional)</p>
CFG{fencing_dgname}	Scalar	<p>Specifies the disk group for I/O fencing.</p> <p>(Optional)</p> <p>Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.</p>

Table 13-1 Response file variables specific to configuring disk-based I/O fencing
(continued)

Variable	List or Scalar	Description
CFG{fencing_newdg_disks}	List	<p>Specifies the disks to use to create a new disk group for I/O fencing.</p> <p>(Optional)</p> <p>Note: You must define the <code>fencing_dgname</code> variable to use an existing disk group. If you want to create a new disk group, you must use both the <code>fencing_dgname</code> variable and the <code>fencing_newdg_disks</code> variable.</p>
CFG{fencing_cpagent_monitor_freq}	Scalar	<p>Specifies the frequency at which the Coordination Point Agent monitors for any changes to the Coordinator Disk Group constitution.</p> <p>Note: Coordination Point Agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidentally deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the <code>LevelTwoMonitorFreq</code> attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles. If <code>LevelTwoMonitorFreq</code> 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.</p>
CFG {fencing_config_cpagent}	Scalar	<p>Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not.</p> <p>Enter "0" if you do not want to configure the Coordination Point agent using the installer.</p> <p>Enter "1" if you want to use the installer to configure the Coordination Point agent.</p>

Table 13-1 Response file variables specific to configuring disk-based I/O fencing (continued)

Variable	List or Scalar	Description
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 189.

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 13-2](#) lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table 13-2 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.

Table 13-2 Coordination point server (CP server) based fencing response file definitions (*continued*)

Response file field	Definition
CFG {fencing_cpagentgrp}	<p>Name of the service group which will have the Coordination Point agent resource as part of it.</p> <p>Note: This field is obsolete if the <code>fencing_config_cpagent</code> field is given a value of '0'.</p>
CFG {fencing_cps}	<p>Virtual IP address or Virtual hostname of the CP servers.</p>
CFG {fencing_reusedg}	<p>This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks). Enter either a "1" or "0".</p> <p>Entering a "1" indicates reuse, and entering a "0" indicates do not reuse.</p> <p>When reusing an existing DG name for the mixed mode fencing configuration, you need to manually add a line of text, such as "<code>\$CFG{fencing_reusedg}=0</code>" or "<code>\$CFG{fencing_reusedg}=1</code>" before proceeding with a silent installation.</p>
CFG {fencing_dgname}	<p>The name of the disk group to be used in the customized fencing, where at least one disk is being used.</p>
CFG {fencing_disks}	<p>The disks being used as coordination points if any.</p>
CFG {fencing_ncp}	<p>Total number of coordination points being used, including both CP servers and disks.</p>
CFG {fencing_ndisks}	<p>The number of disks being used.</p>
CFG {fencing_cps_vips}	<p>The virtual IP addresses or the fully qualified host names of the CP server.</p>
CFG {fencing_ports}	<p>The port that the virtual IP address or the fully qualified host name of the CP server listens on.</p>
CFG {fencing_scsi3_disk_policy}	<p>The disk policy that the customized fencing uses.</p> <p>The value for this field is either "raw" or "dmp"</p>

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing:

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cps_vips>{"10.200.117.145"}=[ qw(10.200.117.145) ];
$CFG{fencing_dgname}="vxfencoorddg";
$CFG{fencing_disks}=[ qw(emc_clariion0_37 emc_clariion0_13) ];
$CFG{fencing_scsi3_disk_policy}="raw";
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=2;
$CFG{fencing_ports>{"10.200.117.145"}=14250;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="SFCFSHA60";
$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 server-based I/O fencing

[Table 13-3](#) lists the fields in the response file that are relevant for non-SCSI-3 server-based customized I/O fencing.

See [“About I/O fencing for Symantec VirtualStore in virtual machines that do not support SCSI-3 PR”](#) on page 22.

Table 13-3 Non-SCSI-3 server-based I/O fencing response file definitions

Response file field	Definition
CFG{non_scsi3_fencing}	<p>Defines whether to configure non-SCSI-3 server-based I/O fencing.</p> <p>Valid values are 1 or 0. Enter 1 to configure non-SCSI-3 server-based I/O fencing.</p>

Table 13-3 Non-SCSI-3 server-based 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.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. Note: This field is obsolete if the <code>fencing_config_cpagent</code> field is given a value of '0'.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_ncp}	Total number of coordination points (CP servers only) being used.
CFG {fencing_ports}	The port of the CP server that is denoted by <i>cps</i> .

Sample response file for configuring non-SCSI-3 server-based I/O fencing

The following is a sample response file used for non-SCSI-3 server-based I/O fencing :

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.198.89.251 10.198.89.252 10.198.89.253) ];
$CFG{fencing_cps_vips}{"10.198.89.251"}=[ qw(10.198.89.251) ];
$CFG{fencing_cps_vips}{"10.198.89.252"}=[ qw(10.198.89.252) ];
$CFG{fencing_cps_vips}{"10.198.89.253"}=[ qw(10.198.89.253) ];
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=0;
$CFG{fencing_ports}{"10.198.89.251"}=14250;
```

```
$CFG{fencing_ports}{"10.198.89.252"}=14250;  
$CFG{fencing_ports}{"10.198.89.253"}=14250;  
$CFG{non_scsi3_fencing}=1;  
$CFG{opt}{configure}=1;  
$CFG{opt}{fencing}=1;  
$CFG{prod}="SVS60";  
$CFG{systems}=[ qw(sys1 sys2) ];  
$CFG{vcs_clusterid}=1256;  
$CFG{vcs_clustername}="clus1";  
$CFG{fencing_option}=1;
```

Upgrading Symantec VirtualStore

- [Chapter 14. Preparing to upgrade](#)
- [Chapter 15. Performing a full upgrade using the installer](#)
- [Chapter 16. Performing a phased upgrade](#)
- [Chapter 17. Performing a rolling upgrade](#)

Preparing to upgrade

This chapter includes the following topics:

- [Upgrade methods for SVS](#)
- [Supported upgrade paths](#)
- [About using the installer to upgrade when the root disk is encapsulated](#)
- [Preparing to upgrade SVS](#)

Upgrade methods for SVS

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 14-1 Review this table to determine how you want to perform the upgrade

Upgrade types and considerations	Methods available for upgrade
Typical upgrades—use a Veritas provided tool or you can perform the upgrade manually. Requires some server downtime.	Script-based—you can use this to upgrade for the supported upgrade paths Web-based—you can use this to upgrade for the supported upgrade paths Manual—you can use this to upgrade from the previous release Response file—you can use this to upgrade from the supported upgrade paths

Table 14-1 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 Veritas provided tool or you can perform the upgrade manually. Requires least amount of server downtime.	Script-based—you can use this to upgrade from the previous release Web-based—you can use this to upgrade from the previous release
Phased upgrades—use a Veritas provided tool and some manual steps. Requires less server downtime than a regular upgrade.	Script-based with some manual steps—you can use this to upgrade from the previous release
Native operating system upgrade—use the upgrade software that comes with the operating system. Note that not all operating systems support native upgrades.	Operating system specific methods Operating system upgrades

Supported upgrade paths

The following tables describe upgrading to 6.0.1

Table 14-2 SLES 11 x64 upgrades using the script- or Web-based installer

Veritas software versions	SLES 10	SLES 11
6.0 6.0 RP1 6.0.1 6.0.3	N/A	Upgrade to SLES 11 SP2 or SLES 11 SP3. Use the installer to upgrade to 6.0.4.
New installation	N/A	Install on SLES 11 SP2 or SLES 11 SP3. Use the installer to install 6.0.4.

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

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

Table 14-3 Upgrading using installer when the root disk is encapsulated

Starting version	Ending version	Action required
5.0 MP3	6.0.4	You need to unencapsulate the root disk. The installer exits.
5.1 or 5.1 RPx	6.0.4	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
5.1 SP1 or 5.1 SP1 RPx	6.0.4	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0 or 6.0 RP1	6.0.4	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0.1	6.0.4	Do not unencapsulate. The installer runs normally. Reboot after upgrade.
6.0.3	6.0.4	Do not unencapsulate. The installer runs normally. Reboot after upgrade.

Preparing to upgrade SVS

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/>
- If you are upgrading from 5.1SP1, you need to unregister the plugin before you upgrade and after the upgrade, you need to register the new plugin.
 See “[Register the new plugin](#)” on page 235.
- 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 201.
- Ensure that you have enough file system space to upgrade. Identify where you want to copy the RPMs, for example `/packages/Veritas` when the root file system has enough space or `/var/tmp/packages` if the `/var` file system has enough space.
Do not put the files under `/tmp`, which is erased during a system reboot. Do not put the files on a file system that is inaccessible prior to running the upgrade script.
You can use a Veritas-supplied disc for the upgrade as long as modifications to the upgrade script are not required. If `/usr/local` was originally created as a slice, modifications are required.
- For any startup scripts in `/sbin/rcS.d`, comment out any application commands or processes that are known to hang if their file systems are not present.
- Make sure that the current operating system supports version 6.0.4 of the product. If the operating system does not support it, plan for a staged upgrade.
- Schedule sufficient outage time and downtime for the upgrade and any applications that use the Veritas products. Depending on the configuration, the outage can take several hours.
- Any swap partitions not in `rootdg` must be commented out of `/etc/fstab`. If possible, swap partitions other than those on the root disk should be commented out of `/etc/fstab` and not mounted during the upgrade. Active swap partitions that are not in `rootdg` cause `upgrade_start` to fail.
- Make sure the file systems are clean before upgrading.
- Upgrade arrays (if required).
See [“Upgrading the array support”](#) on page 209.
- 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.

Creating backups

Save relevant system information before the upgrade.

Note: Unconfigure any FLR jobs for any file systems before unmounting specific file systems.

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 Run the `vxlicrep`, `vxdisk list`, and `vxprint -ht` commands and record the output. Use this information to reconfigure your system after the upgrade.
- 6 If you are installing the high availability version of the Veritas Storage Foundation 6.0.4 software, follow the guidelines given in the *Veritas Cluster Server Installation Guide* and *Veritas Cluster Server Release Notes* for information on preserving your VCS configuration across the installation procedure.

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 199.

Pre-upgrade planning for Veritas Volume Replicator

Before installing or upgrading Veritas Volume Replicator (VVR):

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

- 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 *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.
- If replication using VVR is configured, verify that all the Primary RLINKs are up-to-date on all the hosts.

```
# /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 *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide* for more information.

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

VEA installation planning

The Veritas Enterprise Administrator (VEA) GUI consists of several RPMs. Follow these planning guidelines to install VEA for use with VVR.

- The VEA server packages must be installed on the hosts on which VVR is installed, not the client. If you install using the product installer, these RPMs are installed when you install Storage Foundation products.

The VEA server packages include the following:

- The Veritas Volume Replicator Management Services Provider RPM, `VRTSvrpro`, must be installed on all hosts in the Replicated Data Set (RDS).
- For `VRTSvrpro` to function, the Veritas Volume Manager Management Services Provider RPM, `VRTSvmpro`, must be installed on each system.
- To use the functionality for receiving SNMP notifications and email notifications, the Veritas Action Agent RPM, `VRTSaa` must be installed.

- The VEA client can be installed on a host on which VVR is installed, or a separate host that is used to administer the VVR hosts. To use the VEA client on a system, the `VRTSobgui` RPM must be installed on that system.

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 VirtualStore 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 14-4](#), if either the Primary or Secondary are running a version of VVR prior to 6.0.4, and you use the TCP protocol, VVR calculates the checksum for every data packet it replicates. If the Primary and Secondary are at VVR 6.0.4, VVR does not calculate the checksum. Instead, it relies on the TCP checksum mechanism.

Table 14-4 VVR versions and checksum calculations

VVR prior to 6.0.4 (DG version <= 140)	VVR 6.0.4 (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

Symantec VirtualStore supports using IPv6 as the connection protocol.

This release supports the following configurations for VVR:

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

Preparing to upgrade VVR when VCS agents are configured

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

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

Freezing the service groups and stopping all the applications

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

To freeze the service groups and stop applications

Perform the following steps for the Primary and Secondary clusters:

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

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

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

```
# haconf -makerw
```

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

```
# hagrps -list
```

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

```
# hagrps -freeze group_name -persistent<sys_name>
```

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
VVRGp         State          system02    ONLINE
ORAGrp        State          system02    ONLINE
```

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

- 9 Repeat step 8 for each 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 207.

- 11 For shared disk groups, run the following command on any node in the CVM cluster:

```
# vxdctl -c mode
```

Note the master and record it for future use.

Determining the nodes on which disk groups are online

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

To determine the online disk groups

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

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

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

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

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

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

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

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

Preparing for the upgrade when VCS agents are configured

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

To prepare a configuration with VCS agents for an upgrade

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

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

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

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

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

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

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

```
# vxrecover -bs
```

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

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

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

Upgrading the array support

The Storage Foundation 6.0.4 release includes all array support in a single RPM, VRTSaslapm. The array support RPM includes the array support previously included in the VRTSvxvm RPM. The array support RPM also includes support previously packaged as external array support libraries (ASLs) and array policy modules (APMs).

See the 6.0.4 Hardware Compatibility List for information about supported arrays.

See [“Hardware compatibility list \(HCL\)”](#) on page 27.

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.

See [“Removing external ASL or APM RPMs before upgrading Storage Foundation with manual steps”](#) on page 210.

After you have installed Storage Foundation 6.0.4, Symantec provides support for new disk arrays through updates to the VRTSaslapm RPM.

For more information about array support, see the *Veritas Storage Foundation Cluster File System High Availability Administrator's Guide*.

For more information about File Level Replication (FLR), see the *Veritas Storage Foundation and High Availability Solutions Replication Administrator's Guide*.

Removing external ASL or APM RPMs before upgrading Storage Foundation with manual steps

If you are upgrading with manual steps, remove all external ASLs or APMs before upgrading Storage Foundation. This procedure is not required if you upgrade with the product installer.

To remove external ASL or APM RPMs for a manual installation

- 1 Before you remove any RPMs, make sure you are not running anything in VxVM volumes, make sure that no volumes are mounted etc.

These steps ensure that VxVM does not access the data in any disks that were claimed by these ASLs or APMs after the RPMs are removed. Attempting to access the data could lead to data corruption if the disks are not claimed correctly.

- 2 Remove any external ASLs.

```
# rpm -e VRTSIBM-DS4xxx-2.0-1.0
```

- 3 Determine which external APM RPMs are installed:

```
# rpm -qf /etc/vx/apmkey.d | grep -v "^VRTSvxvm-platform"
```

This command lists the RPMs which installed any APM keys.

Any RPM other than VRTSvxvm is an external APM RPM.

- 4 Remove these APMs before upgrading Storage Foundation:

```
# rpm -e VRTSIBM-DS4xxx-2.0-1.0
```

Performing a full upgrade using the installer

This chapter includes the following topics:

- [Performing a full upgrade](#)

Performing a full upgrade

Performing a full upgrade involves the following tasks:

- Ensuring that the file systems are clean
- Performing the upgrade
- Updating the configuration and confirming startup

Ensuring the file systems are clean

Before upgrading to SVS 6.0.4, ensure that the file systems are clean. To ensure that the logs have been replayed and the file systems are marked clean:

To ensure the file systems are clean

- 1 Log in as superuser onto any node in the cluster.
- 2 Take the service group offline on each node of the cluster, which contains VxFS and CFS resources:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFSSMount resource.

Repeat this step for each SVS service group.

Note: This unmounts the CFS file systems.

- 3 Unmount all VxFS file systems not under VCS control:

```
# umount mount_point
```

- 4 Check and repair each VxFS file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

The `fsck` command in `/opt/VRTS/bin` accepts either the block or character device (`/dev/vx/dsk/dg/vol`) or (`/dev/vx/rdisk/dg/vol`). The operating system version of `fsck` may limit the device types it accepts.

For more information, see the `fsck` and `fsck_vxfs` man pages.

Repeat this step for each file system.

Performing the upgrade

To perform the upgrade

- 1 Log in as superuser.
- 2 Insert the appropriate media disc per your distribution and architecture into your system's DVD-ROM drive.

- 3 If volume management software is running on your system, the software disc automatically mounts as `/mnt/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

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

- 4 Change to the top-level directory on the disc:

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

- 5 Verify there are no VxFS file systems mounted on the nodes being upgraded:

```
# mount -t vxfs
```

If any VxFS file systems are mounted, offline the group on each node of the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFSSMount resource.

If VxFS are not managed by VCS then unmount them manually:

```
# umount mount_point
```

Repeat this step for each SVS service group.

- 6 Start the upgrade from any node in the cluster. Enter the following command, and then press **y** to upgrade the cluster configuration.

```
# ./installsvs -upgrade
```

- 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_cluster_file_system_ha/EULA/
en/EULA_productname_Ux_version.pdf file present on media? [y,n,q,?] y
```

- 8 The installer discovers if any of the systems that you are upgrading have mirrored and encapsulated boot disks. For each system that has a mirrored boot disk, you have the option to create a backup of the system's boot disk group before the upgrade proceeds. If you want to split the boot disk group to create a backup, answer **y**.
- 9 The installer then prompts you to name the backup boot disk group. Enter the name for it or press **Enter** to accept the default.
- 10 If you are prompted to start the split operation. Press **y** to continue.

Note: The split operation can take some time to complete.

- 11 You are prompted to enter the system names (in the following example, "sys1" and "sys2") on which the software is to be upgraded. Enter the system name or names and then press Return.

Enter the system names separated by spaces: **sys1 sys2**

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

If the installer hangs or asks for a login password, stop the installer and set up ssh or rsh. Then run the installer again.
- 13 After you accept EULA and the system checks complete, the installer displays a list of the RPMs that will be upgraded. Press Enter to continue with the upgrade.
- 14 Output shows information that SVS must be stopped on a running system. Enter **y** to continue.
- 15 Press **Return** to begin removing the previous RPMs and installing the new.
- 16 Press **Return** again for summary information about logs and reboots.

Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.
- 17 Update the configuration.

- 18 Only perform this step if you have split the mirrored root disk to back it up. After a successful reboot, verify the upgrade and re-join the backup disk group. If the upgrade fails, revert to the backup disk group.

See [“Re-joining the backup boot disk group into the current disk group”](#) on page 234.

See [“Reverting to the backup boot disk group after an unsuccessful upgrade”](#) on page 235.

Performing a phased upgrade

This chapter includes the following topics:

- [Performing a phased upgrade of SVS](#)

Performing a phased upgrade of SVS

Performing a phased upgrade involves the following tasks:

- Moving the service groups to the second subcluster
- Upgrading the SVS stack on the first subcluster
- Preparing the second subcluster
- Activating the first subcluster
- Upgrading the operating system on the second subcluster
- Upgrading the second subcluster
- Finishing the phased upgrade

Before you start the upgrade on the first half of the cluster, back up the VCS configuration files `main.cf` and `types.cf` which are in the directory `/etc/VRTSvcs/conf/config/`.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade. Note that your applications have downtime during this procedure.

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 basic guidelines follow:

- Split the cluster in half. If the cluster has an odd number of nodes, calculate $(n+1)/2$, and start the upgrade with the even number of nodes.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.

Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules. Also, do not add or remove service groups to any of the nodes.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.
- You can perform a phased upgrade when the root disk is encapsulated.

Moving the service groups to the second subcluster

To move the service groups to the second subcluster

- 1 Switch failover groups from the first half of the cluster to one of the nodes in the second half of the cluster. In this procedure, `sys1` is a node in the first half of the cluster and `sys4` is a node in the second half of the cluster. Enter the following:

```
# hagrps -switch failover_group -to sys4
```

- 2 On the first half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the applications.

- 3 On the first half of the cluster, unmount the VxFS or CFS file systems that are not managed by VCS.

```
# mount | grep vxfs
```

Verify that no processes use the VxFS or CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS or CFS mount point with the mechanism provided by the application.

Unmount the VxFS or CFS file system. Enter the following:

```
# umount /mount_point
```

- 4 On the first half of the cluster, bring all the VCS service groups offline including CVM group. Enter the following:

```
# hagrps -offline group_name -sys sys1
```

When the CVM group becomes OFFLINE, all the parallel service groups such as the CFS file system will also become OFFLINE on the first half of the cluster nodes.

- 5 Verify that the VCS service groups are offline on all the nodes in first half of the cluster. Enter the following:

```
# hagrps -state group_name
```

- 6 Freeze the nodes in the first half of the cluster. Enter the following:

```
# haconf -makerw
```

```
# hasys -freeze -persistent sys1
```

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

- 7 If the cluster-wide attribute UseFence is set to SCSI3, then reset the value to NONE in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

- 8 Verify that only GAB ports a, b, d and h are open. Enter the following:

```
# gabconfig -a
GAB Port Memberships
=====
Port a gen 6b5901 membership 01
Port b gen 6b5904 membership 01
Port d gen 6b5907 membership 01
Port h gen ada40f membership 01
```

Do not stop VCS. Port h should be up and running.

- 9 In the first half of the cluster, stop all VxVM and CVM volumes. Enter the following command for each disk group:

```
# vxvol -g diskgroup stopall
```

Verify that no volumes remain open. Enter the following:

```
# vxprint -Aht -e v_open
```

- 10 On first half of the cluster, upgrade the operating system on all the nodes, if applicable. For instructions, see the upgrade paths for the operating system.

Upgrading the SVS stack on the first subcluster

To upgrade the SVS stack on the first subcluster

- ◆ **Note:** This procedure is based on an "in-place" upgrade path; that is, if the operating system is upgraded, the release will be the same, but only the path level will change. If you are moving from major operating system release to another, you must uninstall the SVS stack before you upgrade the operating system. After the operating system is upgraded, you must reinstall SVS.
-

On the first half of the cluster, upgrade SVS by using the `installsvs` script. For example use the `installsvs` script as shown below:

```
# ./installsvs -upgrade sys1
```

where `<sys1>` is the node on the first subcluster.

After the upgrade for first half of the cluster is complete, no GAB ports will be shown in `gabconfig -a` output.

To upgrade your operating system, follow the normal procedures for your platform.

Note: After the installation completes, you can safely ignore any instructions that the installer displays.

Preparing the second subcluster

To prepare the second subcluster

- 1 On the second half of the cluster, stop all applications that are not configured under VCS. Use native application commands to stop the application. [Downtime starts now.]

- 2 On the second half of the cluster, unmount the VxFS and CFS file systems that are not managed by VCS. Enter the following:

```
# mount | grep vxfs
```

Verify that no processes use the VxFS and CFS mount point. Enter the following:

```
# fuser -c mount_point
```

Stop any processes using a VxFS and CFS mount point with the mechanism provided by the application.

Unmount the VxFS and CFS file system. Enter the following:

```
# umount /mount_point
```

- 3 On the second half of the cluster, unfreeze all the VCS service groups on all the nodes using the following commands:

```
# haconf -makerw  
# hagr -unfreeze group_name -persistent  
# haconf -dump -makero
```

- 4 On the second half of the cluster, bring all the VCS service groups offline, including CVM group. Enter the following:

```
# hagr -offline group_name -sys sys4
```

- 5 On the second half of the cluster, verify that the VCS service groups are offline. Enter the following:

```
# hagr -state group_name
```

- 6 Stop VCS on the second half of the cluster. Enter the following:

```
# hastop -local
```

- 7 If the cluster-wide attribute UseFence is set to SCSI3, reset the value to NONE in the `/etc/VRTSvcs/conf/config/main.cf` file, in second half of the cluster.
- 8 On the second half on cluster, stop the following SVS modules: VCS, VxFEN, ODM, GAB, and LLT. Enter the following:

```
# /etc/init.d/vxglm stop
# /etc/init.d/vxodm stop
# /etc/init.d/vxgms stop
# /etc/init.d/vxfen stop
# /etc/init.d/gab stop
# /etc/init.d/llt stop

# ./installer/installsvs/installsvs - stop
```

- 9 If the cluster-wide attribute UseFence is set to NONE, reset the value to SCSI3 in the `/etc/VRTSvcs/conf/config/main.cf` file, in first half of the cluster.

Activating the first subcluster

To activate the first subcluster

- 1 Restart the upgraded nodes in the first half of the cluster:

```
# /sbin/shutdown -r now
```

When the first half of the cluster nodes come up, no GAB ports are OPEN. The following command does not show any GAB ports:

```
# /sbin/gabconfig -a
GAB Port Memberships
=====
```

- 2 If required, force gab to form a cluster after the upgraded nodes are rebooted in the first half of the cluster.

```
# /sbin/gabconfig -x
```

GAB ports a, b, d and h appear in `gabconfig -a` command output.

Note: If port b and h are not up, you need to bring fencing and VCS manually online.

- 3 Unfreeze the nodes in the first half of the cluster. Enter the following:

```
# haconf -makerw
# hasys -unfreeze -persistent sys1
# haconf -dump -makero
```

- 4 Manually mount the VxFS and CFS file systems that are not managed by VCS.
[Downtime ends now.]

Upgrading the operating system on the second subcluster

To upgrade the operating system on the second subcluster

- ◆ Enter the following.

```
# chkconfig vcs off
# chkconfig vxfs off
# chkconfig gab off
# chkconfig llt off
```

On the second half of the cluster, upgrade the operating system, if applicable.
For instructions, see the upgrade paths for the operating system.

Upgrading the second subcluster

To upgrade the second subcluster

- ◆ Enter the following:

```
# ./installsvs -upgrade node_name
```

Completing the phased upgrade

To complete the phased upgrade

- 1 Verify that the cluster UUID on the nodes in the second subcluster is the same as the cluster UUID on the nodes in the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus -display nodename
```

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 Enter the following before reboot of nodes:

```
# chkconfig vcs on  
# chkconfig vxfs on  
# chkconfig gab on  
# chkconfig lltd on
```

- 3 Restart the upgraded nodes in the second half of the cluster:

```
# /sbin/shutdown -r now
```

When second half of the nodes come up, all the GAB ports a, b, d, h, u, v, w and f are ONLINE. Also all the CFS mounts service groups come online automatically.

- 4 Manually mount the VxFS and CFS file systems that are not managed by VCS in the second half of the cluster.
- 5 Find out which node is the CVM master. Enter the following:

```
# vxdctl -c mode
```

- 6 On the CVM master node, upgrade the CVM protocol. Enter the following:

```
# vxdctl upgrade
```

Performing a rolling upgrade

This chapter includes the following topics:

- [Performing a rolling upgrade using the installer](#)
- [Performing a rolling upgrade of SVS using the Web-based installer](#)

Performing a rolling upgrade using the installer

Use a rolling upgrade to upgrade Symantec VirtualStore to the latest release with minimal application downtime.

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 RPMs in phase 2.

Note: You need to perform a rolling upgrade on a completely configured cluster.

The following is an overview of the flow for a rolling upgrade:

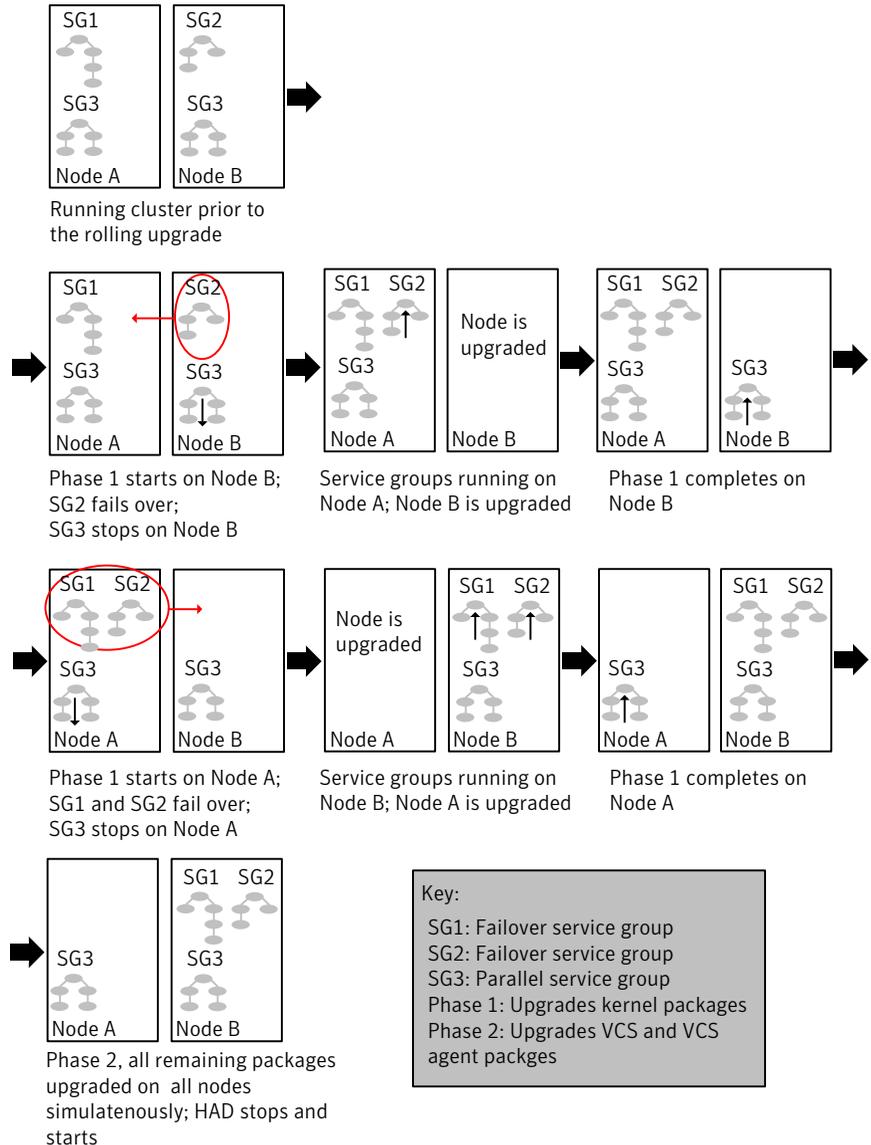
1. The installer performs prechecks on the cluster.

2. The installer moves service groups to free nodes for the first phase of the upgrade as is needed.

Application downtime occurs during the first phase as the installer moves service groups to free nodes for the upgrade. The only downtime that is incurred is the normal time required for the service group to fail over. 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 Veritas Cluster Server (VCS) engine HAD, but does not include application downtime.

[Figure 17-1](#) illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 17-1 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 SVS with the script-based installer, the Web-based installer, or manually.

The rolling upgrade procedures support both major and minor operating system upgrades.

[Table 17-1](#) shows the versions of SVS for which you can perform a rolling upgrade to 6.0.4.

Table 17-1 Supported rolling upgrade paths

Platform	SVS version
Linux	6.0.1, 6.0.2, 6.0.3

Performing a rolling upgrade using the script-based installer

Before you start the rolling upgrade, make sure that Veritas Cluster Server (VCS) is running.

To perform a rolling upgrade

- 1 Complete the preparatory steps on the first sub-cluster.
- 2 Log in as superuser and mount the SVS 6.0.4 installation media.
- 3 From root, start the installer.


```
# ./installer
```
- 4 From the menu, select `Upgrade` and from the sub menu, select `Rolling Upgrade`.
- 5 The installer suggests system names for the upgrade. Enter `Yes` to upgrade the suggested systems, or enter `No`, and then enter the name of any one system in the cluster on which you want to perform a rolling upgrade.
- 6 The installer checks system communications, release compatibility, version information, and lists the cluster name, ID, and cluster nodes. Type `y` to continue.
- 7 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.

- 8 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.
- 9 Review the end-user license agreement, and type **y** if you agree to its terms.
- 10 After the installer shows the package list, it detects if there are online failover service groups on the nodes to be upgraded. If there are online failover service groups, the installer prompts you 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.
- 11 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.
- 12 The installer stops relevant processes, uninstalls old kernel RPMs, and installs the new RPMs. When prompted, enable replication or global cluster capabilities, if required, and register the software.

The installer performs the upgrade configuration and re-starts processes.

If some processes fail to start, you may need to reboot the nodes and manually check the cluster's status.
- 13 Complete the preparatory steps on the nodes that you have not yet upgraded.
- 14 The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade.

If the installer prompts to reboot nodes, reboot the nodes.

Restart the installer.

The installer repeats step 7 through step 12.

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.
- 15 When phase 1 of the rolling upgrade completes, begin phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.
- 16 The installer determines the remaining RPMs to upgrade. Press **Enter** to continue.

- 17 The installer stops Veritas Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.
The installer performs prechecks, uninstalls old RPMs, and installs the new RPMs. It performs post-installation tasks, and the configuration for the upgrade.
- 18 Type **y** or **n** to help Symantec improve the automated installation.
- 19 If you have network connection to the Internet, the installer checks for updates. If updates are discovered, you can apply them now.
- 20 A prompt message appears to ask if the user would like to read the summary file. You can choose **y** if you want to read the install summary file.
- 21 To upgrade VCS or Storage Foundation High Availability (SFHA) on the Coordination Point (CP) server systems to version 6.0.4, upgrade VCS or SFHA on the CP server systems. You then upgrade all the application clusters to 6.0.4.
For instructions to upgrade VCS or SFHA on the CP server systems, refer to the appropriate installation guide.

Performing a rolling upgrade of SVS using the Web-based installer

This section describes using the Veritas Web-based installer to perform a rolling upgrade. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

See [“About rolling upgrades”](#) on page 225.

To start the rolling upgrade—phase 1

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.
- 3 In the Task pull-down menu, select `Rolling Upgrade`.
Click the **Next** button 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. If it throws an error, address the error and return to the installer.
- 6 Review the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 7 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 If you want to enable volume or file replication or global cluster capabilities, select from the following options:
 - Veritas Volume Replicator
 - Veritas File Replicator
 - Global Cluster Option
 Click **Register** to register the software. Click the **Next** button.
The installer starts all the relevant processes and brings all the service groups online.
- 10 When prompted by the installer, reboot the nodes on the first half of the cluster.
Restart the installer.
- 11 Repeat step 5 through step 10 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.
- 12 When prompted, perform step 3 through step 10 on the nodes that you have not yet upgraded.
- 13 When prompted, start phase 2. Click **Yes** to continue with the rolling upgrade.
You may need to restart the Web-based installer to perform phase 2.

To upgrade the non-kernel components—phase 2

- 1 In the Task pull-down menu, make sure that **Rolling Upgrade** is selected.
Click the **Next** button to proceed.
- 2 The installer detects the information of cluster and the state of rolling upgrade.
The installer validates systems. Click **Next**. If it throws an error, address the error and return to the installer.
- 3 Review the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.
- 4 The installer stops the `HAD` and `CmdServer` processes in phase 2 of the rolling upgrade process. 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 would like to read the summary file. You can choose **y** if you want to read the install summary file.

The upgrade is complete.

Post-installation tasks

- [Chapter 18. Performing post-upgrade tasks](#)

Performing post-upgrade tasks

This chapter includes the following topics:

- [Re-joining the backup boot disk group into the current disk group](#)
- [Reverting to the backup boot disk group after an unsuccessful upgrade](#)
- [Register the new plugin](#)

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.

See [“Performing a rolling upgrade using the installer”](#) on page 225.

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.

See [“Performing a rolling upgrade using the installer”](#) on page 225.

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 `vx dg` command to find the boot disk group where you are currently booted.

```
# vx dg 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.

Register the new plugin

After upgrading, need to register the new plugin.

Setting up Symantec VirtualStore

- [Chapter 19. Setting up VirtualStore](#)

Setting up VirtualStore

This chapter includes the following topics:

- [About setting up VirtualStore](#)
- [Setting up Clustered NFS](#)
- [Setting up the VirtualStore vSphere Service Console](#)
- [Registering the VirtualStore cluster](#)
- [Verifying the VirtualStore cluster](#)
- [Virtual machine customization pre-requisites](#)
- [VMware View Integration pre-requisites](#)
- [Unregistering the VirtualStore cluster](#)
- [Displays all the VMware vCenter Servers registered](#)
- [Useful links from VMware on NFS support and customization while cloning virtual machines](#)

About setting up VirtualStore

Table 19-1 VirtualStore tasks

To	Tasks
Set up VirtualStore	See “Setting up Clustered NFS” on page 238. See “Setting up the VirtualStore vSphere Service Console” on page 241. See “Registering the VirtualStore cluster” on page 242. See “Verifying the VirtualStore cluster” on page 243.
Unregister the VirtualStore cluster	See “Unregistering the VirtualStore cluster” on page 244. See “Verifying the VirtualStore cluster” on page 243.
Virtual machine customization pre-requisites	See “Virtual machine customization pre-requisites” on page 243.
VMware View Integration pre-requisites	See “VMware View Integration pre-requisites” on page 244.
View all the VMware vCenter Servers registered	See “Displays all the VMware vCenter Servers registered” on page 245.
View informational links	See “Useful links from VMware on NFS support and customization while cloning virtual machines” on page 245.

Setting up Clustered NFS

To set up Cluster NFS (CNFS), create a metadata NFS share for VirtualStore.

Note: The file system must have disk layout Version 8 or later to use the FileSnaps of virtual machines.

See the [vxupgrade \(1M\)](#) and [mkfs_vxfs \(1M\)](#) manual pages.

To set up CNFS

- 1 Create a shared disk group from the Cluster Volume Manager (CVM) master:

```
# vxdbg -s init shared_disk_group disk1 disk2 disk3 ... diskn
```

For example:

```
# vxdbg -s init svsdg xiv0_0699 xiv0_0700 ...
```

- 2 Create a shared volume (around 100 MB) on the shared disk group to store the CNFS metadata, create a disk layout Version 8 or later file system on top, and add the shared volume to the configuration:

```
# vxassist -g shared_disk_group make svv_meta_vol 100m
# mkfs -t vxfs -o version=8 \
/dev/vx/rdisk/shared_disk_group/svv_meta_vol
# cfsshare config -p nfs shared_disk_group svv_meta_vol \
/svv_meta_mntpt
```

For example:

```
# vxassist -g svsdg make svv_meta_vol 100m
# mkfs -t vxfs -o version=8 /dev/vx/rdisk/svsdg/svv_meta_vol
# cfsshare config -p nfs svsdg svv_meta_vol /svv_meta_mntpt
```

- 3 Create another shared volume to store the VMware images, create a disk layout Version 8 or later file system on it, and add the shared volume to the configuration:

```
# vxassist -g shared_disk_group make svsdata size
# mkfs -t vxfs -o version=8 \
/dev/vx/rdisk/shared_disk_group/svs_data_vol
# cfsshare add -p nfs -N nfs_share_options shared_disk_group \
svs_data_vol /svs_data_mntpt all=[mount_options]
# cfsshare display
```

The example below assumes that you have no mount options:

```
# vxassist -g svsdg make svs_data_vol 2048g layout=striped ncol=120
# mkfs -t vxfs -o version=8 ,bsize=8192 \
/dev/vx/rdisk/svsdg/svs_data_vol
# cfsshare add -p nfs -N "rw,no_root_squash" svsdg svs_meta_vol \
/svs_data_mntpt all=
# cfsshare display
```

See the `mount_vxfs(1M)` manual page for all possible `mount_options`.

- 4 Add a virtual IP (VIP) address for NFS share on the Public Network or Bonded Interface:

```
# cfsshare addvip device vip_address subnet_mask
```

For example 1:

```
# cfsshare addvip eth0 192.168.1.60 255.255.255.0
```

For example 2:

```
# cfsshare addvip bond0 192.168.1.60 255.255.255.0
```

Symantec recommends creating a bond between multiple Network interface cards to increase throughput.

- 5 Add the CNFS share created and the VIP added in the above steps to configure an NFS datastore on all your ESX machines.
 - Create VMKernel Port to allow ESX to access NFS Datastores. For example, make ESX server an NFS client this requires an IP address per Server Chassis.
 - If you add an NFS datastore and it was unsuccessful, try using the `vmkping` command to ensure basic network connectivity.

- Add the NFS Datastore to the ESX Server using the **add storage Wizard** and provide the following parameters:
 - Storage type=Network File System
 - Server: *IP_address_of_NFS_share*
 - Folder: Export mount point (For example, the */svs_data_mntpt*)
 - Data Store name: A unique name for Datastore

Setting up the VirtualStore vSphere Service Console

This section describes how to set up the VirtualStore vSphere Service Console to host the VMware vSphere Plug-ins.

The vSphere Service Console consists of the following components:

- A VirtualStore node running svswab.
svswab is the Web Server which hosts the vSphere Plug-in.
- A VIP address that communicates between the vSphere Client and the Plug-in.

To set up VirtualStore vSphere Service Console

1. You can either add a new VIP in the cluster or use the same VIP address that you defined when you set up clustered NFS.

- To set up a new VIP:

```
# cfsshare addvip device vip_address subnet_mask
```

Example 1:

```
# cfsshare addvip eth0 192.168.1.60 255.255.255.0
```

Example 2:

```
# cfsshare addvip bond0 192.168.1.60 255.255.255.0
```

Symantec recommends that you create a bond between multiple Network interface cards to increase throughput.

- To use the VIP address that you set up earlier, skip this step.

See “[Setting up Clustered NFS](#)” on page 238.

- 2 On the svswab server node, configure the VirtualStore Web Server for cluster failover:

```
$ svsvmwadm -a setup -s vip_address
```

The following example uses the virtual IP (VIP) address added in step 1:

```
$ svsvmwadm -a setup -s 192.168.1.60
```

- 3 Verify that the VirtualStore vSphere Service Console is Online:

```
# /opt/VRTS/bin/hagrp -state ` /opt/VRTS/bin/hares -value svswab Group`
```

where svswab is the name of the Symantec VirtualStore web process.

- 4 Ensure port 14191 is open. It is required to serve the vCenter Plug-in.

Registering the VirtualStore cluster

This section describes how to register the VirtualStore cluster with VMware vCenter Server.

Registering the VirtualStore cluster with VMware vCenter Server enables you the following:

- To start using the storage (NFS exported from this cluster) as a VMware datastore.
- From the vSphere Client to invoke the VirtualStore FileSnap Wizard.

To register the VirtualStore cluster

- 1 Log in to the cluster using the virtual IP address used to set up the SVS vSphere Service Console.
- 2 Register the VirtualStore cluster with VMware vCenter Server:

```
$ svsvmwadm -a register -v vcenter_server_ip -u user \  
[-p password]
```

For example:

```
$ svsvmwadm -a register -v 192.168.1.55 -u admin -p xxxxxx
```

- 3 List registered Plug-ins.

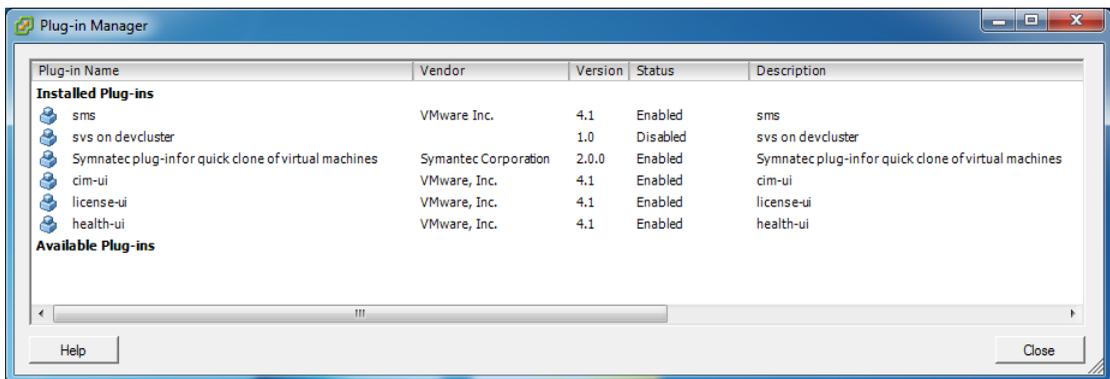
```
$ svsvmwadm -a list
```

Verifying the VirtualStore cluster

This section describes how to verify the VirtualStore cluster is installed and enabled.

To verify the VirtualStore cluster is installed and enabled

- 1 Go to vCenter VI client.
- 2 Click on **Plug-ins** and a **Plug-in Manager** window displays.
- 3 Ensure that you see **svs on *cluster_name*** in the list of installed Plug-ins.
- 4 Verify that the Plug-in is in the enabled state.



Virtual machine customization pre-requisites

Prior knowledge of VMware vSphere is required.

For more information on "Customizing Guest Operating Systems", see the *VMware vSphere Basic System Administration Guide*.

Virtual machine customization pre-requisites

- 1 VMware Tools must be installed on the virtual machine. For detailed instructions, refer to <http://kb.vmware.com/kb/1014294>.
- 2 For Windows XP, 2000 Server and 2003 Server, Microsoft Sysprep needs to be available on the vCenter Server. For detailed instructions, refer to <http://kb.vmware.com/kb/1005593>.
- 3 If customization has been applied to the virtual machine, it must be powered on in order to commit the customization changes to the guest operating system. This must be done before recustomizing.

VMware View Integration pre-requisites

Prior knowledge of VMware View is required.

For more information, see the *VMware View Administrator's Guide*.

VMware View Integration pre-requisites

- 1 VMware View Agent needs to be installed on the virtual machine.
- 2 The vCenter Server must be added to the View Connection Server. It must be added using the same address (IP or FQDN) used to login to the vCenter with the vSphere Client.

Note: Domain users which are View Administrators should not exist as local users with the same password on the View Connection Server or Active Directory Server as this may cause authentication issues.

- 3 SASL DIGEST-MD5 libraries must be installed on the VirtualStore node. The default installation of most operating systems already includes this library. If it is not installed with the OS, it should be available on the OS installation media.

Package name:
RHEL: `cyrus-sasl-md5`
SLES: `cyrus-sasl-digestmd5`

- 4 OpenLDAP Tools must be installed on the VirtualStore node. The default installation of most operating systems already includes this library. If it is not installed with the OS, it should be available on the OS installation media.

Package name: `openldap-clients` or `openldap2-clients`

Unregistering the VirtualStore cluster

This section describes how to unregister the VirtualStore cluster from VMware vCenter Server.

To unregister the VirtualStore cluster

- 1 Log in to the cluster using the virtual IP address used to set up the SVS vSphere Service Console.
- 2 Unregister the VirtualStore cluster from VMware vCenter Server:

```
$ svsvmwaadm -a unregister -v vcenter_server_ip -u user \  
[-p password]
```

For example:

```
$ svsvmwaadm -a unregister -v 192.168.1.55 -u admin -p xxxxxx
```

Displays all the VMware vCenter Servers registered

This section describes how to display all the VMware vCenter Servers registered.

To display all the VMware vCenter Servers registered

- ◆ Displays all the VMware vCenter Servers registered:

```
$ svsvmwaadm -a list
```

Useful links from VMware on NFS support and customization while cloning virtual machines

- Best Practices for running VMware vSphere on Network Attached Storage (White paper):
http://vmware.com/files/pdf/VMware_NFS_BestPractices_WP_EN.pdf
- VirtualCenter2 templates usage and best practices (White paper)
Best practices for setting up templates and guest customization:
http://www.vmware.com/pdf/vc_2_templates_usage_best_practices_wp.pdf

8

Section

Using the extensions for Symantec VirtualStore

- [Chapter 20. Using the VirtualStore Plug-in for VMware vSphere](#)
- [Chapter 21. Using the Citrix XenDesktop extension for Symantec VirtualStore](#)

Using the VirtualStore Plug-in for VMware vSphere

This chapter includes the following topics:

- [About the VMware vSphere Plug-in for VirtualStore](#)
- [Adding storage to ESX servers](#)
- [Creating a virtual machine](#)
- [Creating virtual machine clones using Symantec FileSnap](#)
- [Specifying the number of virtual machine clones to be created](#)
- [Specifying where to create the virtual machine clones](#)
- [Specifying the guest operating system if customizing for Linux or Windows](#)
- [Specifying network customization parameters for guest operating systems if using DHCP or Static IP](#)
- [Configuring the VMware View](#)
- [Configuring Citrix XenDesktop](#)
- [Verifying the virtual machine clones](#)

About the VMware vSphere Plug-in for VirtualStore

VirtualStore provides the benefits of a highly-scalable NAS and NFS solution for VMware ESX workloads.

The VirtualStore Plug-in for VMware vSphere enables an administrator to quickly clone hundreds of virtual machines.

Adding storage to ESX servers

To add storage to ESX servers

- 1 Right-click on a datacenter, cluster, folder, or an ESX server in the **vSphere Client** window. Choose **Add storage from *clustername***, where *clustername* is the SVS cluster that has been registered with this vCenter Server, and from which you want to add storage.
- 2 In the **Add Symantec Storage** window, specify ESX server information, NFS storage, and datastore information in the appropriate fields.

ESX Servers In the **ESX Servers** table, check the ESX server(s) to which you want to add storage.

The ESX servers that are displayed under the **ESX Host Name** column depend on what entity you selected in step 1. If the entity selected is a datacenter, all the ESX servers in the datacenter are displayed. If the entity selected is a folder, all the ESX servers in the folder are displayed. If the entity selected is a cluster, then all the ESX servers in that cluster are displayed.

Basically, a datacenter, a cluster, or a folder are containers of ESX servers.

Virtual IP Address In the **NFS Storage Properties** table, in the drop-down menu, select a virtual IP address.

Share In the **NFS Storage Properties** table, in the drop-down menu, select the directory path of the NFS share.

Mount NFS Read Only In the **NFS Storage Properties** table, check if you want the mounted NFS share to be read-only.

Datastore Name A datastore name is automatically generated based on the virtual IP address and the NFS share that you selected.

You can change the datastore name to whatever you choose.

Note: The datastore name must be unique. If the datastore name is not unique, you receive an error message after clicking **Submit**.

- 3 Click **Submit**.

Creating a virtual machine

To create a virtual machine

- ◆ Create a virtual machine.

See the *VMware vSphere Virtual Machine Administration Guide*.

http://www.vmware.com/pdf/vsphere4/r41/vsp_41_vm_admin_guide.pdf

If you plan on using or integrating the virtual machine with VMware View:

See the *VMware View Administrator's Guide*.

http://www.vmware.com/pdf/view45_admin_guide.pdf

Note: If you plan to use the vApp and OVF templates, then you must disable the vApp.

See the *Symantec VirtualStore Release Notes* for more information.

MAC address conflicts may arise if static addresses are used.

See the VMware Knowledge Base Article on changing the MAC address of a hosted virtual machine for more information.

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=507

Creating virtual machine clones using Symantec FileSnap

This section describes how to create virtual machine clones using Symantec FileSnap. FileSnap is a component of Symantec VirtualStore.

Before being able to create virtual machine clones using Symantec FileSnap, you must have registered your SVS cluster with a vCenter Server using the SVS CLI. You must also have created a virtual machine.

See “[About the VMware vSphere Plug-in for VirtualStore](#)” on page 247.

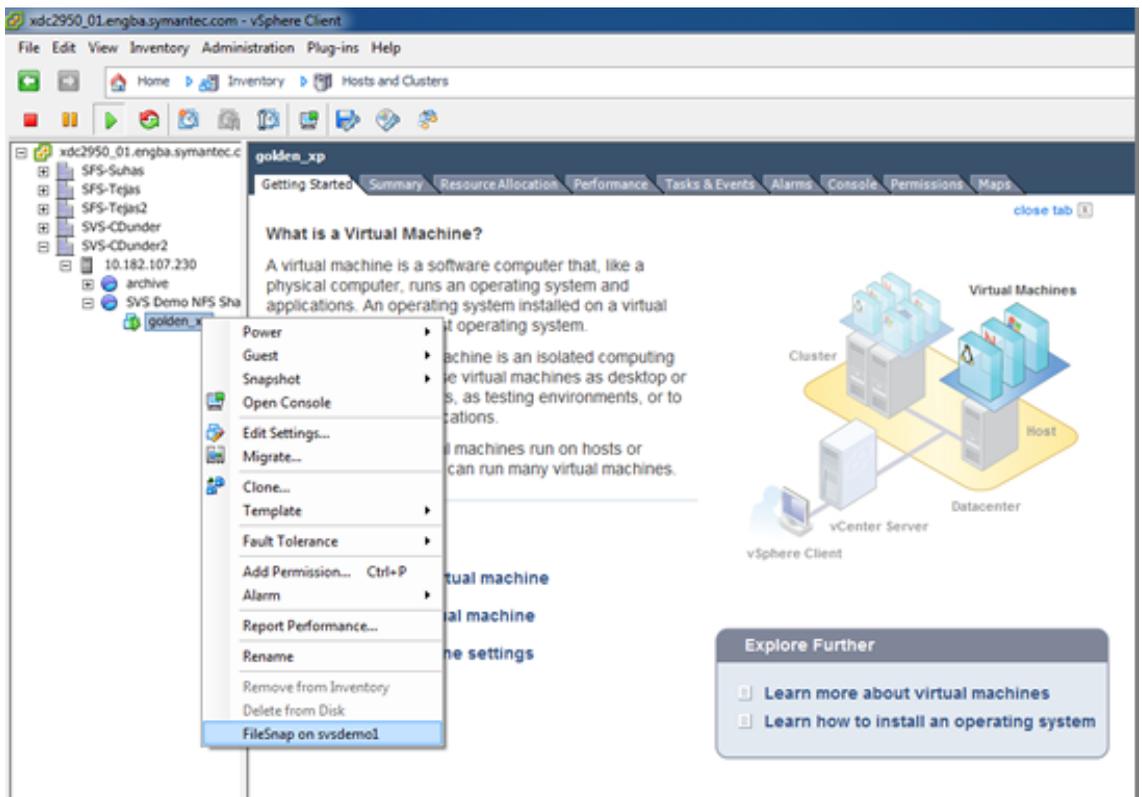
See “[Creating a virtual machine](#)” on page 249.

Note: During the use of the FileSnap wizard, if you “close” the status window this only closes the window, and does not cancel any running jobs.

To create virtual machine clones using Symantec FileSnap

- 1 Right-click on the virtual machine you want to clone, and select **FileSnap** on *clustername* in the **vSphere Client** window.

Depending upon whether the virtual disks belong to a Symantec datastore (CNFS/VirtualStore), the GUI allows or disallows the cloning operation of the virtual machine. To allow cloning, at least one virtual disk must be from a Symantec datastore. If the **NEXT** button is not displayed, then the GUI displays a message that the cloning operation is not allowed on this virtual machine.



- 2 In the **Security alert** window, click **Yes**.

Specifying the number of virtual machine clones to be created

To specify the number of virtual machine clones to be created

- 1 In the **Clone Name Pattern** window of the **Symantec Quick Clone Virtual Machine Wizard**, enter virtual machine clone information in the appropriate fields.

Number of Clones Specify the number of virtual machine clones you want to create.

Clone Name Specify the name of the virtual machine clone.
The clone name should only contain alpha-numeric characters, underscores, and or hyphens.

Start From Specify the number of the virtual machine clone that you want to start creating clones from.

- 2 Click **Next**.

Specifying where to create the virtual machine clones

To specify where to create the virtual machine clones

- 1 In the **Clone Target** window of the **Symantec Quick Clone Virtual Machine Wizard**, enter the datacenter, ESX host/cluster, and the resource pool for creating virtual machine clones in the appropriate fields.

Datacenter (required) From the drop-down menu, select the datacenter where you want the virtual machine clones to reside.

ESX Host/Cluster (required) From the drop-down menu, select the ESX host/cluster where you want the virtual machine clones to reside.

Resource Pool From the drop-down menu, select the resource pool for running the virtual machine clones.

A resource pool is a VMware term; it is a division of computing resources that are used to manage allocations between virtual machines.

- | | |
|---|--|
| Customize the guest OS and network for the virtual machine clones | If you want to customize the guest OS and network for the virtual machine clones, check the Customize the guest OS and network for the virtual machine clones checkbox. |
| Export virtual machines to Virtual Desktop Infrastructure (VDI) | <p>If you want to export virtual machines to Virtual Desktop Infrastructure (VDI), check the Export virtual machines to Virtual Desktop Infrastructure (VDI) checkbox.</p> <p>If you check the Export virtual machines to Virtual Desktop Infrastructure (VDI) checkbox, you can select from one of the following options:</p> <ul style="list-style-type: none"> ■ VMware View - allows you to configure the VMware View. ■ Citrix XenDesktop (CSV file) - allows you to export a CSV file from the vCenter Server, and then download the exported CSV file to Citrix XenDesktop. |

2 Click **Next**.

Specifying the guest operating system if customizing for Linux or Windows

To specify the guest operating system if customizing the guest operating system for Linux

1 In the **Guest OS Customization** window of the **Symantec Quick Clone Virtual Machine Wizard**, specify the guest operating system parameters for virtual machine clones if you want to customize the guest operating system for Linux.

- | | |
|---------------|--|
| Guest OS Name | Specify the guest operating system name that you want to customize. |
| Start From | Specify the number of the virtual machine clone that you want to start from. |
| Domain | Enter the active directory domain name. |

2 Click **Next**.

To specify the guest operating system if customizing the guest operating system for Windows

- 1 In the **Guest OS Customization** window of the **Symantec Quick Clone Virtual Machine Wizard**, specify the guest operating system parameters for virtual machine clones if you want to customize the guest operating system for Windows.

Guest OS Name	Specify the guest operating system name that you want to customize. The guest OS name should only contain alpha-numeric characters and or hyphens. The full host name (concatenated Guest OS Name and Start From) cannot be greater than 63 characters.
Start From	Specify the number of the virtual machine clone that you want to start from.
Workgroup	Enter the workgroup name. If you select a domain computer (Domain), the Workgroup field is not available.
Domain	Enter the active directory domain name.
Username	Enter the domain username.
Password	Enter the domain user password.
Confirm Password	Re-enter the domain user password.
Product Key	Specify the product key for the virtual machine clone. The product key should only contain alpha-numeric characters, underscores, and or hyphens.
User Name	Specify the user name for the virtual machine clone. The user name should only contain alpha-numeric characters, underscores, and or hyphens.
Company	Specify the name of the company you are from.
Timezone	From the drop-down menu, specify the time zone for creating the virtual machine clone.

- 2 Click **Next**.

Specifying network customization parameters for guest operating systems if using DHCP or Static IP

To specify network customization parameters for guest operating systems if using DHCP

- 1 If you checked the **Customize the guest OS and network for the virtual machine clones** checkbox on the **Clone Target** window of the **Symantec Quick Clone Virtual Machine Wizard**, the **Network Customization** window appears.

Each virtual machine may have multiple network adapters.

- 2 In the **Network Customization** window, if you select **DHCP** as the **IP Assignment**, you do not need to specify values for the **IP Address**, **Netmask**, and the **Gateway** fields.
- 3 Click **Next**.

To specify network customization parameters for guest operating systems if using Static IP

- 1 If you checked the **Customize the guest OS and network for the virtual machine clones** checkbox on the **Clone Target** window of the **Symantec Quick Clone Virtual Machine Wizard**, the **Network Customization** window appears.
- 2 In the **Network Customization** window, if you select **Static IP** as the **IP Assignment**, enter information in the appropriate fields.

IP Assignment If you select **Static IP** from the drop-down menu, you need to enter values for **IP Address**, **Netmask**, and **Gateway** fields.

IP Address Specify a valid IP address for the network adapter card.
Each virtual machine may have multiple network adapters.

Netmask Specify a valid netmask address for the network adapter.

Gateway Specify a valid gateway IP address for the network adapter.

Primary DNS Enter the IP address for your primary Domain Name System (DNS) server.

Secondary DNS Enter the IP address for your secondary Domain Name System (DNS) server.

- 3 Click **Next**.

Configuring the VMware View

To configure the VMware View

- 1 If you checked the **Export virtual machines to Virtual Desktop Infrastructure (VDI)** checkbox, and then selected the **VMware View** radio button on the **Clone Target** window of the **Symantec Quick Clone Virtual Machine Wizard**, the **VMware View Configuration** window appears.

- 2 In the **VMware View Configuration** window, enter information in the appropriate fields to import cloned virtual machines as a new Desktop Pool in the VMware View.

Server Name/IP Address	<p>Specify the name of the server or IP address for the server for importing cloned virtual machines.</p> <p>Note: You must log on to the vCenter Server using the same IP address or a fully-qualified domain name (FQDN) that the VMware View Server uses to access the vCenter Server for successful integration of the cloned virtual machines into the VMware View Server.</p>
Username	<p>Enter a VMware View administrator username.</p> <p>You must enter a username.</p>
Password	<p>Enter the VMware View administrator password.</p> <p>You must enter a password.</p>
Domain	<p>Enter the VMware View administrator domain name.</p> <p>Note: To be accessed by the VMware View Server, virtual machines must be joined to a domain. The Symantec Quick Clone Virtual Machine Wizard can automatically join virtual machine clones to the domain entered on the VMware View Configuration window.</p>
Pool ID	<p>Specify the pool ID for importing the cloned virtual machines.</p> <p>You must enter a pool ID.</p>
Pool Persistence	<p>Specify if you want persistent pools by checking this checkbox.</p> <p>Pool persistence indicates if users are allocated a dedicated desktop. This option statically assigns the desktop to a user on the first connection. All user documents, applications, and settings are retained between sessions.</p>

- 3 Click **Next**.

Configuring Citrix XenDesktop

To configure Citrix XenDesktop

- 1 If you checked the **Customize the guest OS and network for the virtual machine clones** radio button and the **Export virtual machines to Virtual Desktop Infrastructure (VDI) - Citrix XenDesktop (CSV file)** option on the **Clone Target** window of the **Symantec Quick Clone Virtual Machine Wizard**, the **Configure Citrix XenDesktop** window appears.

- 2 In the **Configure Citrix XenDesktop** window for **Connection Name**, provide a meaningful name for this connection.

This connection name has to be same as the one created in Citrix XenDesktop.

If you click the **Details** link, you will see additional information about **Connection Name**. **This is the VMware vCenter connection name as provided during the XenDesktop configuration for VMware environments.**

- 3 Click **Next**.

Verifying the virtual machine clones

To verify the list of requested virtual machine clones

- 1 In the **Summary** window of the **Symantec Quick Clone Virtual Machine Wizard**, verify the list of requested virtual machine clones.
- 2 If you specified a **Connection Name** on the **Configure Citrix XenDesktop** window, you will see a message saying **The CSV file for Citrix XenDesktop will be generated.**, and the hyperlink for accessing the CSV file.
- 3 If you want to turn on the virtual machine clones after they are created, check the **Power on clones after creation** checkbox.
- 4 Click **Submit**.

Using the Citrix XenDesktop extension for Symantec VirtualStore

This chapter includes the following topics:

- [About adding virtual machine clones to Citrix XenDesktop](#)
- [Connecting Citrix XenDesktop to a vCenter Server](#)
- [Adding a connection](#)
- [Creating virtual machine clones using the Citrix XenDesktop extension](#)
- [Importing virtual machines to Citrix Desktop Studio](#)
- [Creating a desktop group and testing the deployment](#)

About adding virtual machine clones to Citrix XenDesktop

You can add virtual machine clones created on SVS to Citrix XenDesktop.

Prior to being able to use this feature, you must have already configured your VMware vCenter Server and registered your SVS cluster with the vCenter Server.

You must also have successfully installed and configured the Citrix XenDesktop controller.

Connecting Citrix XenDesktop to a vCenter Server

By default, the vCenter Server is installed with a self-signed certificate. If you are not able to replace it with a certificate issued from a certification authority, you can use the following workaround to use the self-signed certificate of the vCenter Server. On the Windows server where you installed Citrix XenDesktop, do the following.

To connect Citrix XenDesktop to a vCenter Server

- 1 Open an Internet Explorer browser, and enter the hostname or IP address of the machine running the vCenter Server as **https://<vcenter_server>**.

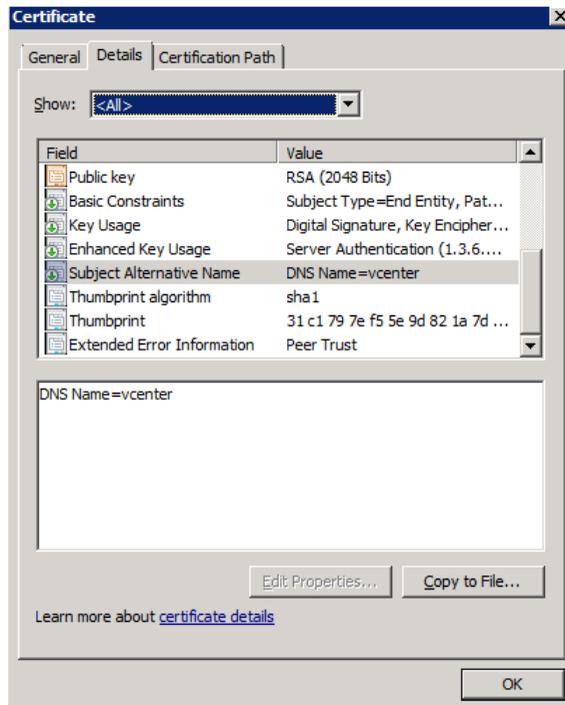
You will see a certificate error.

- 2 Click the certificate error, and select **View certificate**, and click **Install certificate**.

- 3 Click the **Details** tab on the **View Certificate** dialog.

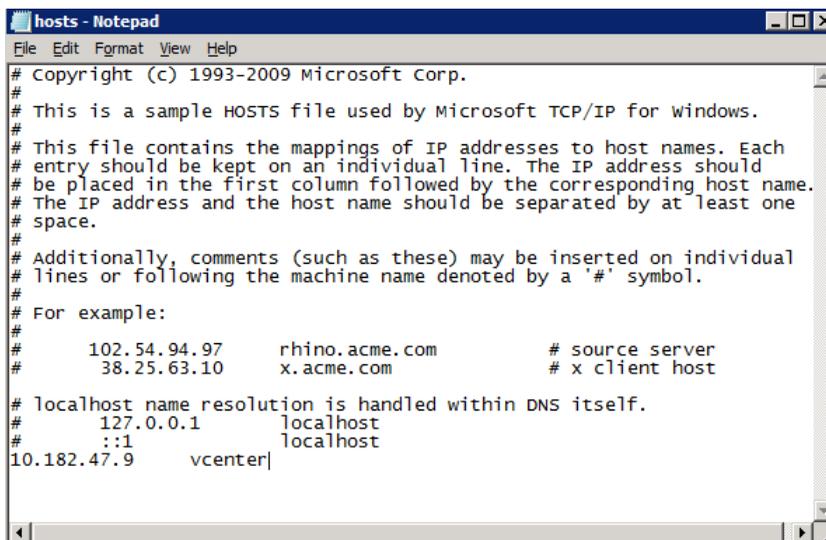
Check the DNS name in the certificate. For example, the DNS name of the vCenter Server is displayed as **vcenter**.

Make a note of the DNS name.



- 4 Open the hosts file in %SystemRoot%/WINDOWS/system32/Drivers/etc/.

Add a line with the IP address of your vCenter Server and the DNS name as displayed in the **Certificate** dialog, and save the hosts file.



```
hosts - Notepad
File Edit Format View Help
# Copyright (c) 1993-2009 Microsoft Corp.
#
# This is a sample HOSTS file used by Microsoft TCP/IP for Windows.
#
# This file contains the mappings of IP addresses to host names. Each
# entry should be kept on an individual line. The IP address should
# be placed in the first column followed by the corresponding host name.
# The IP address and the host name should be separated by at least one
# space.
#
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name denoted by a '#' symbol.
#
# For example:
#
#       102.54.94.97       rhino.acme.com       # source server
#       38.25.63.10      x.acme.com         # x client host
#
# localhost name resolution is handled within DNS itself.
#       127.0.0.1        localhost
#       ::1             localhost
10.182.47.9      vcenter|
```

- 5 Add a line with the IP address of your vCenter Server and the DNS name in the certificate and save the hosts file. The following is a sample.

- 6 Open a new Internet Explorer browser, and enter
`https://<vCenter_DNS_name_in_certificate>`.

If the certification error is gone, you are done, and you can move to the next step.

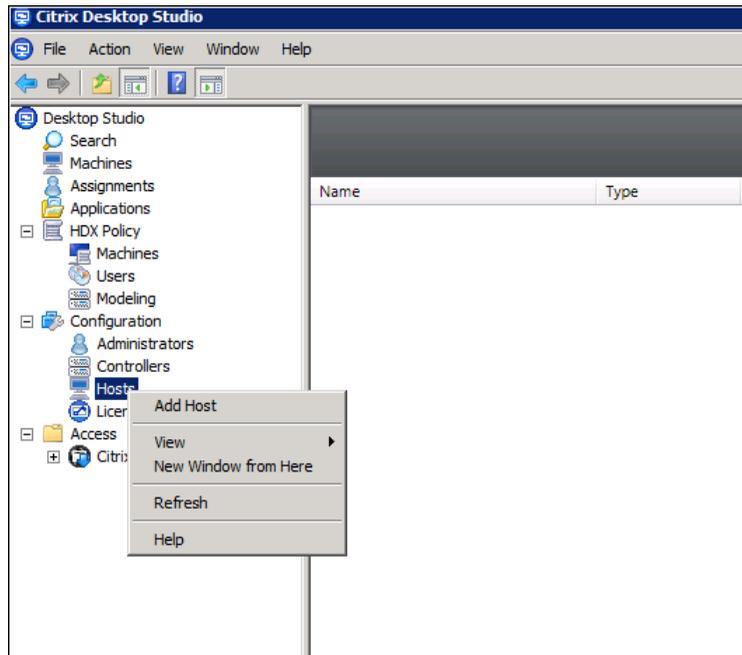
See [“Adding a connection”](#) on page 260.

Adding a connection

To add a connection

- 1 Launch **Citrix Desktop Studio** on the Windows server where you installed Citrix XenDesktop.
- 2 Select **Configuration > Hosts**.

- 3 Right-click on **Hosts**, and then select **Add Host**.



- 4 Choose **VMware virtualization** as the **Host type**.
- 5 Enter `https://<vcenter_DNS_name_in_certificate>/sdk` in the **Address** text box, and enter the vCenter Server administrator's Username and Password.

- 6 Choose a **Connection name** that you like, and select **Manually create virtual machines**.

Add Host

Steps

- Host
- Summary

Host type: VMware virtualization

Address: https://vcenter/sdk

Username: administrator

Password: ●●●●●●

The Connection name will be displayed in Desktop Studio. Consider using a name that will help administrators to identify the host type and address of the deployment to which the connection relates.

Connection name: vCenter

Virtual machines:

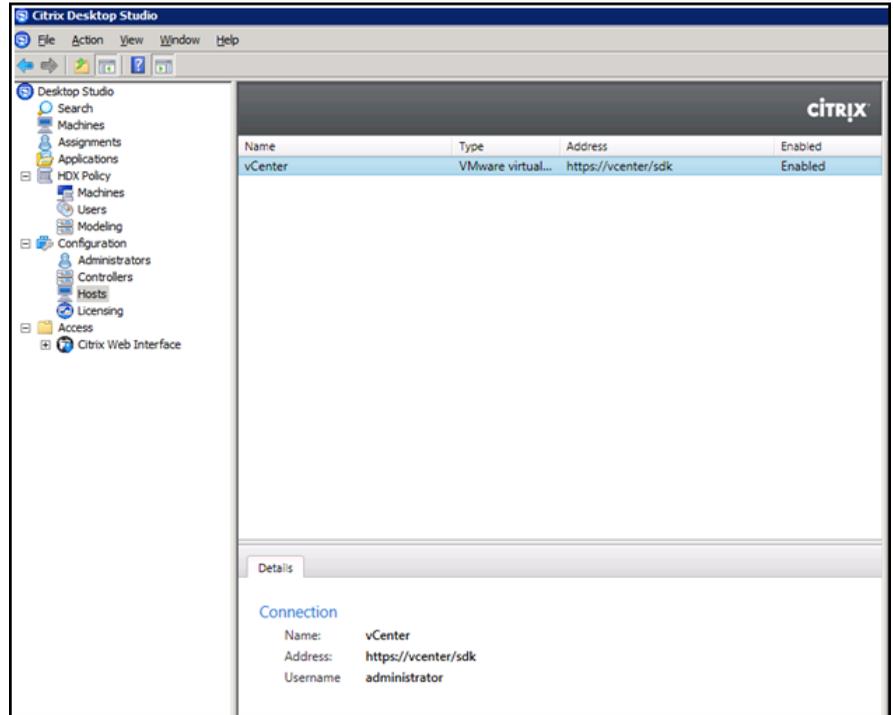
- Use XenDesktop to create virtual machines
- Manually create virtual machines

CITRIX

Back Next Cancel

7 Click **Next**.

You can see the connection name in Citrix Desktop Studio.



Creating virtual machine clones using the Citrix XenDesktop extension

This section describes how to create virtual machine clones using the Citrix XenDesktop extension with the Symantec Quick Clone Virtual Machine Wizard.

To create virtual machine clones using the Citrix XenDesktop extension

- 1 Prepare a virtual machine golden image and install Citrix XenDesktop Virtual Desktop Agent on the virtual machine.

Refer to the Citrix XenDesktop documentation for how to install the Virtual Desktop Agent.

<http://support.citrix.com/proddocs/topic/xendesktop-snma/cds-install-vda.html>

- 2 Follow the steps for creating virtual machine clones using Symantec FileSnap.

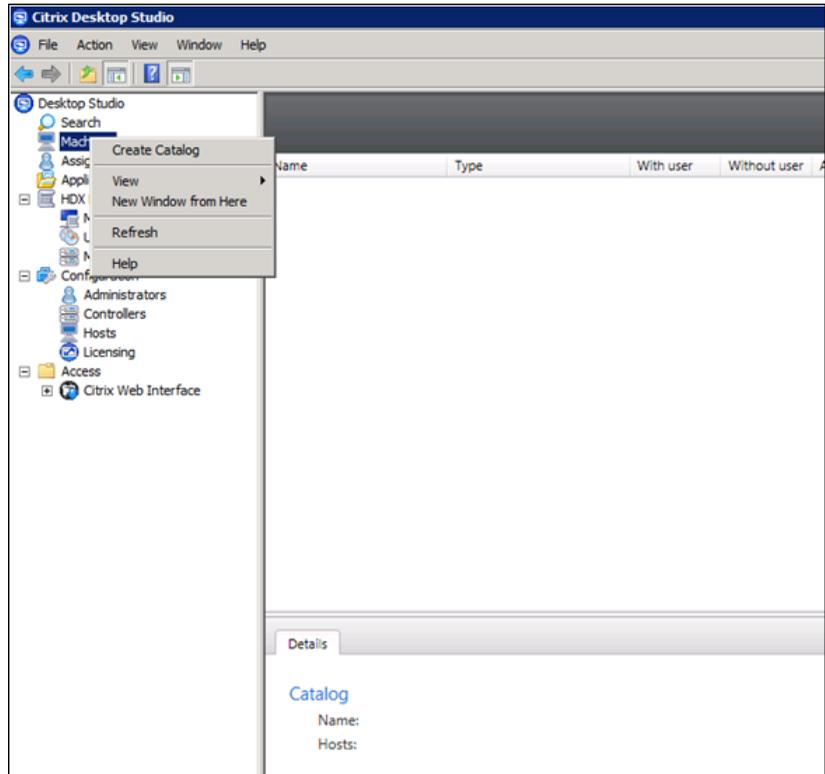
- 3 Once you have finished using the Symantec Quick Clone Virtual Machine Wizard, power on all virtual machines.
- 4 Check if the virtual machines are properly configured.

Importing virtual machines to Citrix Desktop Studio

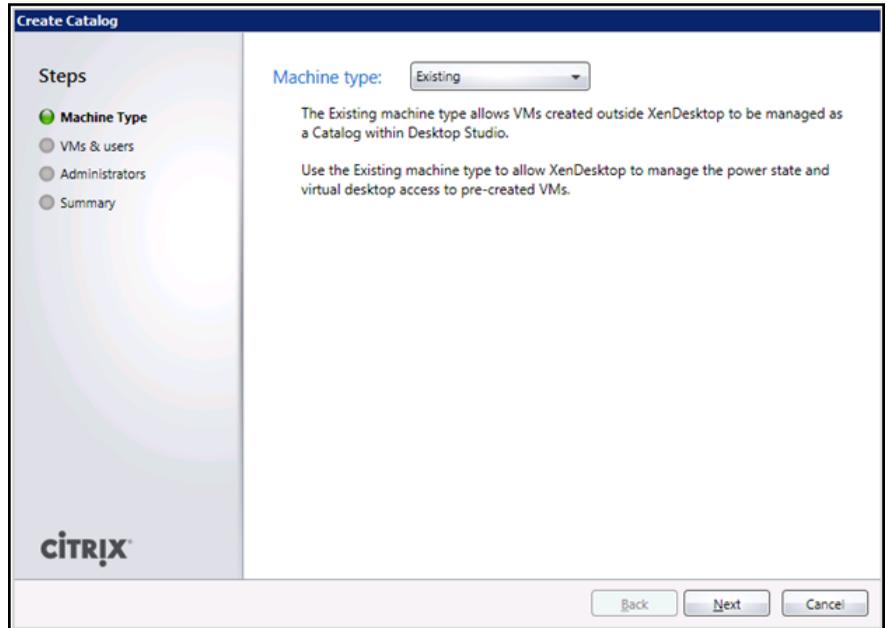
This section describes how to import virtual machines to Citrix Desktop Studio.

To import the virtual machine clones into Citrix XenDesktop

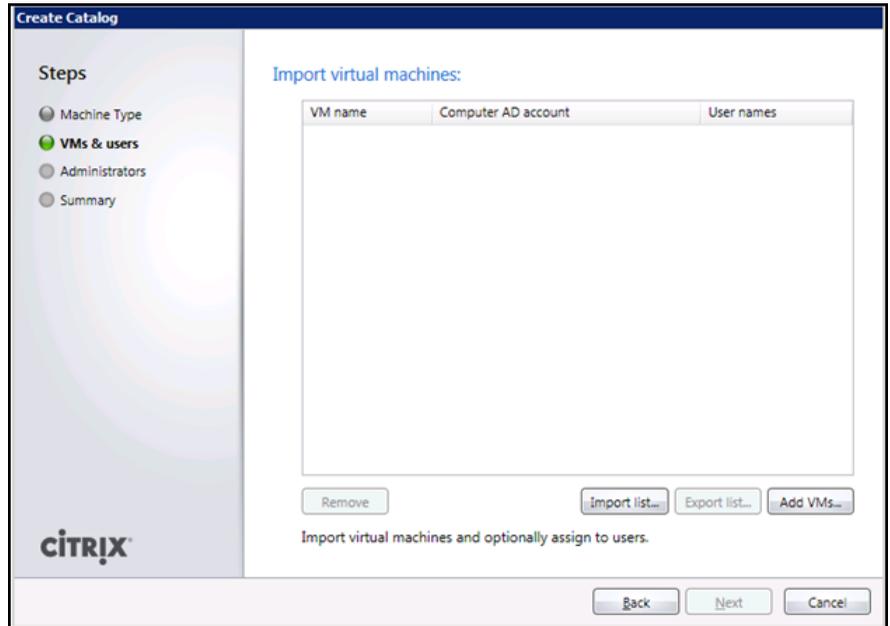
- 1 Launch Citrix Desktop Studio, and right-click **Machines**, and then select **Create Catalog**.



- 2 Select **Existing** from the **Machine type** drop-down menu.



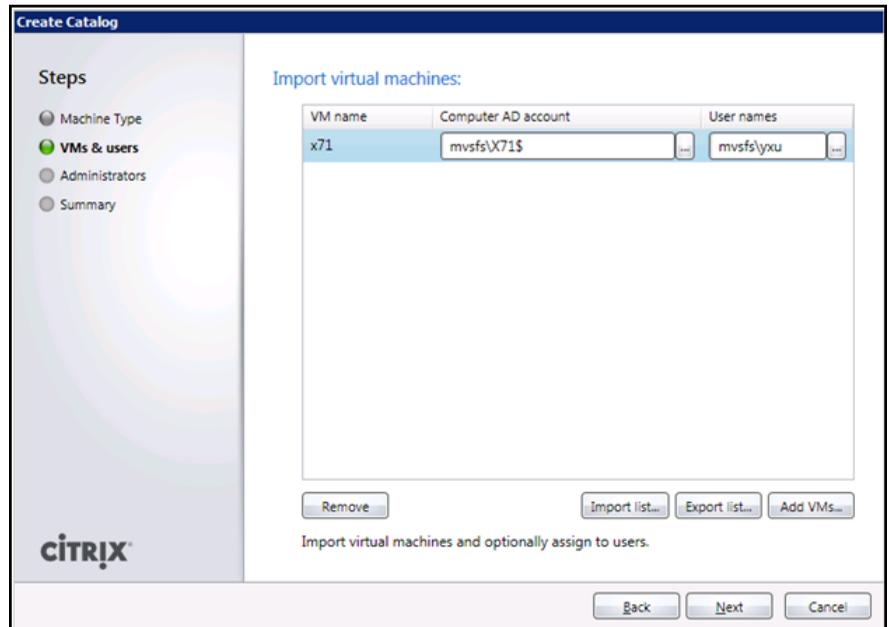
- 3 Click the **Import list** button.



- 4 Choose the CSV file that you saved, and click **Open**.

You will see the virtual machines listed in the CSV file.

You can change the **Computer AD** account or **User names** if you want.



- 5 Click **Next** to finish the operation.
- 6 Log into the virtual machines with the user account that you were assigned.
Verify that the virtual machines are working correctly.

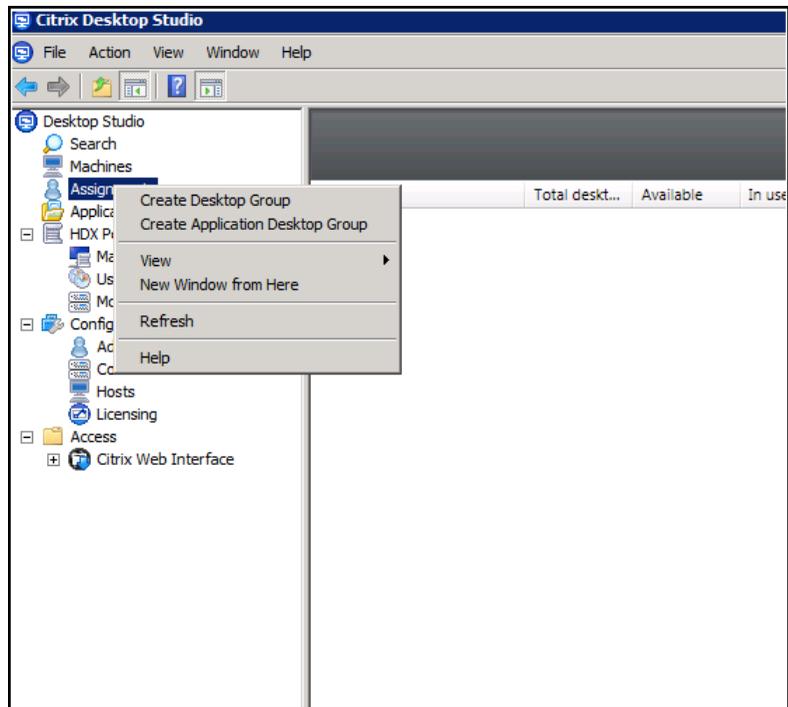
Creating a desktop group and testing the deployment

This section describes how to create a desktop group for the catalog of virtual machines that you just created.

To create a desktop group and test the deployment

- 1 At the **Citrix Desktop Studio** dialog, right-click **Assignments**, and select **Create Desktop Group**.

Follow the Wizard dialogs step-by-step.

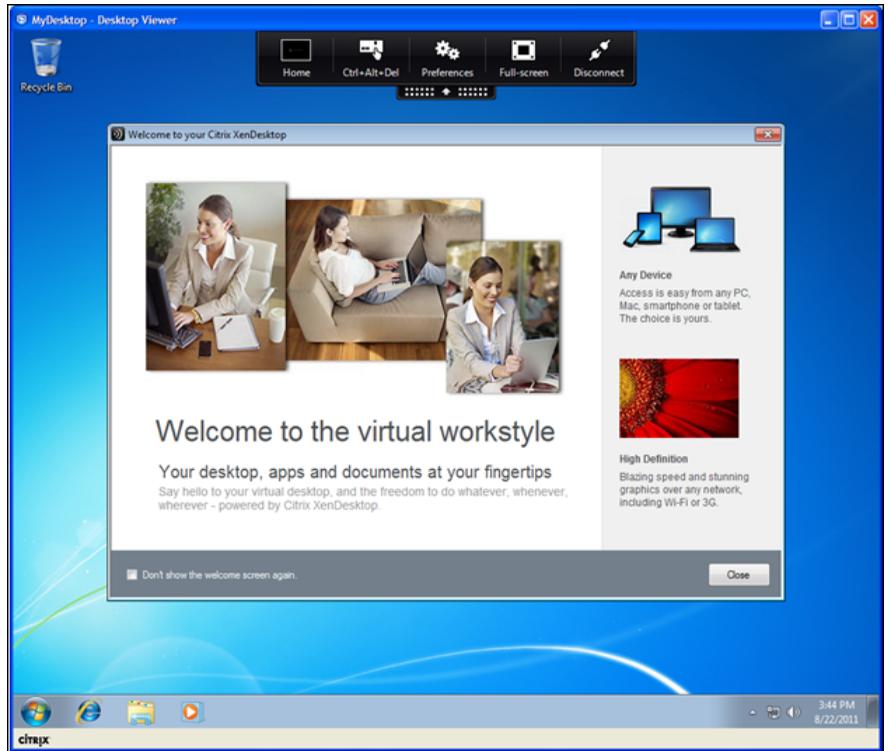


- 2 Ensure that you have installed and configured Citrix XenDesktop Virtual Desktop Agent.

See [“Creating virtual machine clones using the Citrix XenDesktop extension”](#) on page 263.

- 3 From any computer, point your browser at the controller IP address, or its fully-qualified DNS name, and authenticate it using one of the users supplied.





Adding and removing nodes

- [Chapter 22. Adding a node to a cluster](#)
- [Chapter 23. Removing a node from a cluster](#)

Adding a node to a cluster

This chapter includes the following topics:

- [About adding a node to a cluster](#)
- [Before adding a node to a cluster](#)
- [Adding a node to a cluster using the SVS installer](#)
- [Adding a node using the Web-based installer](#)
- [Adding the node to a cluster manually](#)
- [Starting Cluster Volume Manager \(CVM\) and Cluster File System \(CFS\)](#)
- [After adding the new node](#)

About adding a node to a cluster

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 SVS cluster.

Table 22-1 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 274.

Table 22-1 Tasks for adding a node to a cluster (*continued*)

Step	Description
Add a new node to the cluster.	See “Adding a node to a cluster using the SVS installer” on page 276. See “Adding the node to a cluster manually” on page 280.
Complete the configuration of the new node after adding it to the cluster.	See “Starting Cluster Volume Manager (CVM) and Cluster File System (CFS)” on page 290.

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 SVS 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 SVS.
- 2 Verify the new system has the same identical operating system versions and patch levels as that of the existing cluster
- 3 If the cluster is upgraded from the previous version, you must check the cluster protocol version to make sure it has the same version as the node to be added. If there is a protocol mismatch, the node is unable to join the existing cluster.

Check the cluster protocol version using:

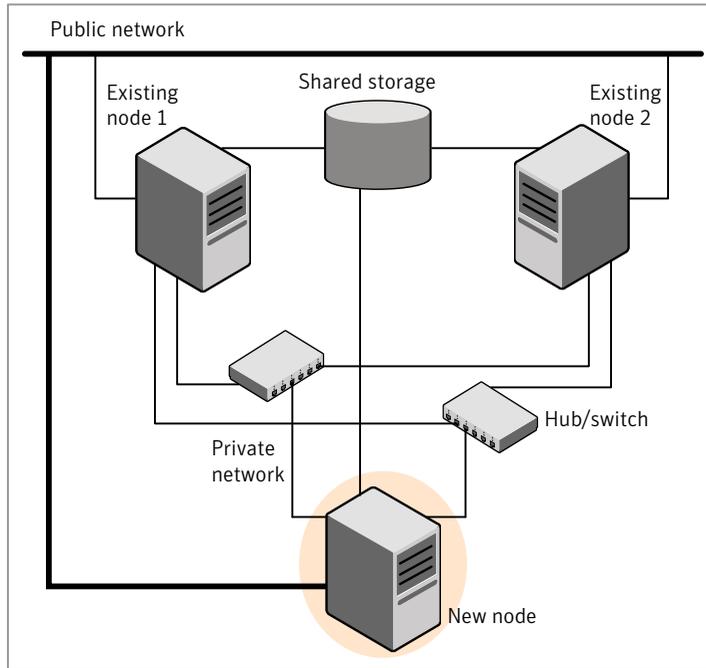
```
# vxdctl protocolversion
Cluster running at protocol 120
```

- 4 If the cluster protocol on the master node is below 120, upgrade it using:

```
# vxdctl upgrade [version]
```

Before you configure a new system on an existing cluster, you must physically add the system to the cluster as illustrated in [Figure 22-1](#).

Figure 22-1 Adding a node to a two-node cluster using two switches



To set up the hardware

- 1 Connect the SVS 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 22-1 illustrates a new node being added to an existing two-node cluster using two independent hubs.

- 2 Make sure that you meet the following requirements:

- The node must be connected to the same shared storage devices as the existing nodes.
- The node must have private network connections to two independent switches for the cluster.

For more information, see the *Veritas Cluster Server Installation Guide*.

- The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster.

Complete the following preparatory steps on the new node before you add it to an existing SVS cluster.

To prepare the new node

- 1 Verify that the new node meets installation requirements.

```
# ./installsvs -precheck
```

You can also use the Web-based installer for the precheck.

- 2 Install SVS on the new system. Make sure all the VRTS RPMs available on the existing nodes are also available on the new node.

```
# cd /opt/VRTS/install
```

```
# ./installsvs<version>
```

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

Do not configure SVS when prompted.

Adding a node to a cluster using the SVS installer

You can add a node to a cluster using the `-addnode` option with the SVS installer.

The SVS 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`
- Copies the following files on the new node:
 - `/etc/llthosts`
 - `/etc/gabtab`
 - `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files from the existing cluster to the new node:
 - `/etc/vxfenmode`

```

/etc/vxfendg
/etc/vx/.uuids/clusuuid
/etc/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.
- Adds the new node to the CVM, ClusterService, and VxSS service groups in the VCS configuration.

Note: For other service groups configured under VCS, update the configuration for the new node manually.

- Adds the new node to the CVM and ClusterService service groups in the VCS configuration.

Note: For other service groups configured under VCS, update the configuration for the new node manually.

- Starts SVS processes and configures CVM and CFS on the new node.
- Configure clustered NFS service group on the new node.

At the end of the process, the new node joins the SVS cluster.

Note: If you have configured server-based fencing on the existing cluster, make sure that the CP server does not contain entries for the new node. If the CP server already contains entries for the new node, remove these entries before adding the node to the cluster, otherwise the process may fail with an error.

To add the node to an existing cluster using the installer

- 1 Log in as the root user on one of the nodes of the existing cluster.
- 2 Run the SVS installer with the `-addnode` option.

```
# cd /opt/VRTS/install  
# ./installsvs<version> -addnode
```

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

See [“About the Veritas installer”](#) on page 37.

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 SVS cluster.

The installer uses the node information to identify the existing cluster.

```
Enter one node of the SVS 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: saturn
```

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.

Note: The network interface names used for the private interconnects on the new node must be the same as that of the existing nodes in the cluster. The LLT configuration for the new node must be the same as that of the existing cluster.

```
Enter the NIC for the first private heartbeat
link on saturn: [b,q,?] eth1
```

```
Enter the NIC for the second private heartbeat
link on saturn: [b,q,?] eth2
```

Note: At least two private heartbeat links must be configured for high availability of the cluster.

- 7 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.
 The installer verifies the network interface settings and displays the information.
- 8 Review and confirm the information.
- 9 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on saturn: eth3
```

- 10 If the existing cluster uses server-based fencing in secure mode, the installer will configure server-based fencing in secure mode on the new nodes.
 The installer then starts all the required Veritas processes and joins the new node to cluster.
 The installer indicates the location of the log file, summary file, and response file with details of the actions performed.
- 11 Confirm that the new node has joined the SVS cluster using `lltstat -n` and `gabconfig -a` commands.
 If the new node has not joined the cluster, verify if it has been added to the `SystemList`.

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 SVS only if you plan to add the node to the cluster manually.

Table 22-2 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 281.

Table 22-2 Procedures for adding a node to a cluster manually (*continued*)

Step	Description
Configure the cluster processes on the new node.	See “Configuring cluster processes on the new node” on page 282.
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 284.
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 289.
Start VCS.	See “To start VCS on the new node” on page 291.
Configure CVM and CFS.	
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 290.

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 `installsvs` 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 For Red Hat Linux, modify the file `/etc/sysctl.conf` on the new system to set the shared memory and other parameter required by your application; refer to the your application documentation for details. The value of the shared memory parameter is put to effect when the system restarts.

Do not apply 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 saturn
set-cluster 101

link eth1 eth-[MACID for eth1] - ether - -
link eth2 eth-[MACID for eth2] - ether - -
```

Except for the first line that refers to the node, the file resembles the `/etc/llttab` files on the existing nodes. The second line, the cluster ID, must be the same as in the existing nodes.

- 5 Use vi or another text editor to create the file `/etc/gabtab` on the new node. This file must contain a line that resembles the following example:

```
/sbin/gabconfig -c -nN
```

Where N represents the number of systems in the cluster including the new node. For a three-system cluster, N would equal 3.

- 6 Edit the `/etc/gabtab` file on each of the existing systems, changing the content to match the file on the new system.
- 7 Use vi or another text editor to create the file `/etc/VRTSvcs/conf/sysname` on the new node. This file must contain the name of the new node added to the cluster.

For example:

```
saturn
```

- 8 Create the Unique Universal Identifier file `/etc/vx/.uuids/clusuuid` on the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -rsh -clus -copy \  
-from_sys sys1 -to_sys saturn
```

- 9 Start the LLT, GAB, and ODM drivers on the new node:

```
# /etc/init.d/llt start  
  
# /etc/init.d/gab start  
  
# /etc/init.d/odm restart
```

- 10 On the new node, verify that the GAB port memberships are a and d:

```
# gabconfig -a  
GAB Port Memberships  
=====
```

```
Port a gen df204 membership 012  
Port 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 22-3](#) uses the following information for the following command examples.

Table 22-3 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.

Configuring the authentication broker on node sys5

To configure the authentication broker on node sys5

- 1 Extract the embedded authentication files and copy them to temporary directory:

```
# mkdir -p /var/VRTSvcs/vcsauth/bkup  
# cd /tmp; gunzip -c /opt/VRTSvcs/bin/VxAT.tar.gz | tar xvf -
```

- 2 Edit the setup file manually:

```
# cat /etc/vx/.uuids/clusuuid 2>&1
```

The output is a string denoting the UUID. This UUID (without { and }) is used as the ClusterName for the setup file.

```
{UUID}
```

```
# cat /tmp/eat_setup 2>&1
```

The file content must resemble the following example:

```
AcceptorMode=IP_ONLY  
  
BrokerExeName=vcsauthserver  
  
ClusterName=UUID  
  
DataDir=/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER  
  
DestDir=/opt/VRTSvcs/bin/vcsauth/vcsauthserver  
  
FipsMode=0  
  
IPPort=14149  
  
RootBrokerName=vcsroot_uuid  
  
SetToRBPlusABorNot=0  
  
SetupPDRs=1  
  
SourceDir=/tmp/VxAT/version
```

3 Set up the embedded authentication file:

```
# cd /tmp/VxAT/version/bin/edition_number; \
./broker_setup.sh/tmp/eat_setup

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssregctl -s -f
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER/root/.VRTSat/profile \
/VRTSatlocal.conf -b 'Security\Authentication \
\Authentication Broker' -k UpdatedDebugLogFileName \
-v /var/VRTSvcs/log/vcsauthserver.log -t string
```

4 Copy the broker credentials from one node in the cluster to sys5 by copying the entire `bkup` directory.

The `bkup` directory content resembles the following example:

```
# cd /var/VRTSvcs/vcsauth/bkup/

# ls

CMDSERVER  CPSADM  CPSEVER  HAD  VCS_SERVICES  WAC
```

5 Import the `VCS_SERVICES` domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCSAUTHSERVER -f /var/VRTSvcs/vcsauth/bkup \
/VCS_SERVICES -p password
```

6 Import the credentials for `HAD`, `CMDSERVER`, `CPSADM`, `CPSEVER`, and `WAC`.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atutil import -z \
/var/VRTSvcs/vcsauth/data/VCS_SERVICES -f /var/VRTSvcs/vcsauth/bkup \
/HAD -p password
```

7 Start the `vcsauthserver` process on `sys5`.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh
```

8 Perform the following tasks:

```
# mkdir /var/VRTSvcs/vcsauth/data/CLIENT
# mkdir /var/VRTSvcs/vcsauth/data/TRUST
# export EAT_DATA_DIR='/var/VRTSvcs/vcsauth/data/TRUST'
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat setuptrust -b \
localhost:14149 -s high
```

9 Create the `/etc/VRTSvcs/conf/config/.secure` file:

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

Setting up SVS related security configuration

Perform the following steps to configure SVS related security settings.

Setting up SVS related security configuration

1 Start `/opt/VRTSat/bin/vxatd` process.**2** Create `HA_SERVICES` domain for SVS.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

3 Add SVS 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
```

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 SVS cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

- 2 Add the node sys5 to the existing vxfen group.

```
# hagrps -modify vxfen SystemList -add sys5 2
```

- 3 Save the configuration by running the following command from any node in the SVS cluster:

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

Starting fencing on the new node

Perform the following steps to start fencing on the new node.

To start fencing on the new node

- 1 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
```

If you are using pure CP server-based fencing on the existing cluster, then only the /etc/vxfenmode file needs to be copied on the new node.

- 2 Start fencing on the new node:
- 3 On the new node, verify that the GAB port memberships are a, b, and d:

```
# gabconfig -a
```

```
GAB Port Memberships
```

```
=====
Port a gen      df204 membership 012
Port b gen      df20d membership 012
Port d gen      df20a membership 012
```

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 saturn to the existing ClusterService group.

```
# hagrps -modify ClusterService SystemList -add saturn 2
```

```
# hagrps -modify ClusterService AutoStartList -add saturn
```

- 3 Modify the IP address and NIC resource in the existing group for the new node.

```
# hares -modify gcoip Device eth0 -sys saturn
```

```
# hares -modify gconic Device eth0 -sys saturn
```

- 4 Save the configuration by running the following command from any node.

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

Starting Cluster Volume Manager (CVM) and Cluster File System (CFS)

Bring the CVM group online to start CVM and CFS on the new node.

To start CVM and CFS

- 1 Log onto one of the nodes in the existing cluster.
- 2 Bring the CVM group online on the new node:

```
# hagrps -online cvm -sys saturn
```

- 3 Verify that the CVM and CFS groups are online:

```
# hagrps -state
```

- 4 Configure cluster NFS on the new nodes. On each new node, run the command:

```
# cfsshare addnode <saturn>
```

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
```

Removing a node from a cluster

This chapter includes the following topics:

- [Removing a node from a SVS cluster](#)

Removing a node from a SVS cluster

[Table 23-1](#) 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 23-1 Tasks that are involved in removing a node

Task	Reference
<ul style="list-style-type: none"> ■ Back up the configuration file. ■ Check the status of the nodes and the service groups. 	See “Verifying the status of nodes and service groups” on page 293.
<ul style="list-style-type: none"> ■ Switch or remove any SVS service groups on the node departing the cluster. ■ Delete the node from SVS configuration. 	See “Deleting the departing node from SVS configuration” on page 294.
Modify the llthosts(4) and gabtab(4) files to reflect the change.	
If the existing cluster is configured to use server-based I/O fencing, remove the node configuration from the CP server.	See “Removing the node configuration from the CP server” on page 297.

Table 23-1 Tasks that are involved in removing a node (*continued*)

Task	Reference
For a cluster that is running in a secure mode, remove the security credentials from the leaving node.	See “Removing security credentials from the leaving node” on page 297.
On the node departing the cluster: <ul style="list-style-type: none"> ■ Modify startup scripts for LLT, GAB, and SVS to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. 	

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 SVS 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.

```
# hagrps -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.

```
# hagrps -dep
```

- 3 If the service group on the departing node requires other service groups—if it is a parent to service groups on other nodes—unlink the service groups.

```
# haconf -makerw
# hagrps -unlink grp4 grp1
```

These commands enable you to edit the configuration and to remove the requirement grp4 has for grp1.

- 4 Stop SVS 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
# hagr -modify grp3 SystemList -delete sys5
# hagr -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.

```
# hagr -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.

```
# hagr -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
```

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/
```

Removing the node configuration from the CP server

After removing a node from a SVS 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 VirtualStore Administrator's Guide*.

To remove the node configuration from the CP server

- 1 Log into the CP server as the root user.
- 2 View the list of VCS users on the CP server, using the following command:

```
# cpsadm -s cp_server -a list_users
```

Where `cp_server` is the virtual IP/ virtual hostname of the CP server.

- 3 Remove the VCS user associated with the node you previously removed from the cluster.

For CP server in non-secure mode:

```
# cpsadm -s cp_server -a rm_user \  
-e cpsclient@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
```

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 Symantec VirtualStore 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 SVS

- 1 If you had configured I/O fencing in enabled mode, then stop I/O fencing.

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

- 2 Stop GAB and LLT:

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

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

- 3 To determine the RPMs to remove, enter:

```
# rpm -qa | grep VRTS
```

- 4 To permanently remove the VCS RPMs from the system, use the `rpm -e` command. Start by removing the following RPMs, which may have been optionally installed, in the order shown:

```
# rpm -e VRTSvcsea
# rpm -e VRTSatServer
# rpm -e VRTSatClient
# rpm -e VRTSvcshr
# rpm -e VRTSvcsg
# rpm -e VRTScps
# rpm -e VRTSvcsl
# rpm -e VRTSamsf
# rpm -e VRTSvxfen
# rpm -e VRTSgab
# rpm -e VRTSllt
# rpm -e VRTSspt
# rpm -e VRTSsfcp604
# rpm -e VRTSperl
# rpm -e VRTSvlic
```

- 5 Remove the LLT and GAB configuration files.

```
# rm /etc/llttab
# rm /etc/gabtab
# rm /etc/llthosts
```

Uninstallation of Symantec VirtualStore

- [Chapter 24. Uninstalling Symantec VirtualStore](#)

Uninstalling Symantec VirtualStore

This chapter includes the following topics:

- [Shutting down cluster operations](#)
- [Removing VxFS file systems](#)
- [Removing rootability](#)
- [Moving volumes to disk partitions](#)
- [Disabling the agents on a system](#)
- [Removing the Replicated Data Set](#)
- [Uninstalling SVS RPMs using the script-based installer](#)
- [Uninstalling SVS with the Veritas Web-based installer](#)
- [Removing license files \(Optional\)](#)
- [Removing the CP server configuration using the installer program](#)

Shutting down cluster operations

If the systems are running as an HA cluster, you have to take all service groups offline and shutdown VCS.

To take all service groups offline and shutdown VCS

- ◆ Use the `hastop` command as follows:

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

Warning: Do not use the `-force` option when executing `hastop`. This will leave all service groups online and shut down VCS, causing undesired results during uninstallation of the RPMs.

Removing VxFS file systems

The VxFS RPM cannot be removed if there are any mounted VxFS file systems. Unmount all VxFS file systems before removing the RPM. After you remove the VxFS RPM, VxFS file systems are not mountable or accessible until another VxFS RPM is installed. It is advisable to back up VxFS file systems before installing a new VxFS RPM. If VxFS will not be installed again, all VxFS file systems must be converted to a new file system type.

To remove VxFS file systems

- 1 Check if any VxFS file systems or Storage Checkpoints are mounted:

```
# df -T | grep vxfs
```

- 2 Make backups of all data on the file systems that you wish to preserve, or recreate them as non-VxFS file systems on non-VxVM volumes or partitions.
- 3 Unmount all Storage Checkpoints and file systems:

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

- 4 Comment out or remove any VxFS file system entries from the `/etc/fstab` file.

Removing rootability

Perform this procedure if you configured rootability by encapsulating the root disk.

To remove rootability

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

```
# df -v /
```

The root disk is under VxVM control if `/dev/vx/dsk/rootdg/rootvol` is listed as being mounted as the root (`/`) file system. If so, unmirror and unencapsulate the root disk as described in the following steps:

- 2 Use the `vxplex` command to remove all the plexes of the volumes `rootvol`, `swapvol`, `usr`, `var`, `opt` and `home` that are on disks other than the root disk.

For example, the following command removes the plexes `mirrootvol-01`, and `mirswapvol-01` that are configured on a disk other than the root disk:

```
# vxplex -o rm dis mirrootvol-01 mirswapvol-01
```

Warning: Do not remove the plexes that correspond to the original root disk partitions.

- 3 Enter the following command to convert all the encapsulated volumes in the root disk back to being accessible directly through disk partitions instead of through volume devices:

```
# /etc/vx/bin/vxunroot
```

Following the removal of encapsulation, the system is rebooted from the unencapsulated root disk.

Moving volumes to disk partitions

All volumes must be moved to disk partitions.

This can be done using one of the following procedures:

- Back up the system fully onto tape and then recover from it.
- Back up each file system individually and then recover them all after creating new file systems on disk partitions.
- Use VxVM to move volumes incrementally onto disk partitions as described in the following section.

Moving volumes onto disk partitions using VxVM

Use the following procedure to move volumes onto disk partitions.

To move volumes onto disk partitions

- 1 Evacuate disks using the `vxdiskadm` program, VEA, or the `vxevac` script. You should consider the amount of target disk space required for this before you begin.

Evacuation moves subdisks from the specified disks to target disks. The evacuated disks provide the initial free disk space for volumes to be moved to disk partitions.

- 2 Remove the evacuated disks from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk _media_name
# vxdisk rm disk_access_name
```

- 3 Decide which volume to move first. If the volume to be moved is mounted, unmount it.
- 4 If the volume is being used as a raw partition for database applications, make sure that the application is not updating the volume and that data on the volume is synced.
- 5 Create a partition on free disk space of the same size as the volume. If there is not enough free space for the partition, a new disk must be added to the system for the first volume removed. Subsequent volumes can use the free space generated by the removal of this volume.
- 6 Copy the data on the volume onto the newly created disk partition using a command similar to the following:

```
# dd if=/dev/vx/dsk/diskgroup/volume-name of=/dev/sdb2
```

where `sdb` is the disk outside of VxVM and `2` is the newly created partition on that disk.

- 7 Replace the entry for that volume (if present) in `/etc/fstab` with an entry for the newly created partition.
- 8 Mount the disk partition if the corresponding volume was previously mounted.
- 9 Stop the volume and remove it from VxVM using the following commands:

```
# vxvol -g diskgroup -f stop volume_name
# vxedit -g diskgroup -rf rm volume_name
```

- 10 Remove any disks that have become free (have no subdisks defined on them) by removing volumes from VxVM control. To check if there are still some subdisks remaining on a particular disk, use the following command:

```
# vxprint -F "%sdnum" disk_media_name
```

- 11 If the output is not 0, there are still some subdisks on this disk that must be subsequently removed. If the output is 0, remove the disk from VxVM control using the following commands:

```
# vxdg -g diskgroup rmdisk disk_media_name  
# vxdisk rm disk_access_name
```

- 12 The free space now created can be used for adding the data in the next volume to be removed.
- 13 After all volumes have been converted into disk partitions successfully, reboot the system. After the reboot, none of the volumes should be open. To verify that none of the volumes are open, use the following command:

```
# vxprint -Aht -e v_open
```

- 14 If any volumes remain open, repeat the steps listed above.

Disabling the agents on a system

This section explains how to disable a VCS agent for VVR on a system. To disable an agent, you must change the service group containing the resource type of the agent to an OFFLINE state. Then, you can stop the application or switch the application to another system.

To disable the agents

- 1 Check whether any service group containing the resource type of the agent is online by typing the following command:

```
# hagrps -state service_group -sys system_name
```

If none of the service groups is online, skip to 3.

- 2 If the service group is online, take it offline.

To take the service group offline without bringing it online on any other system in the cluster, enter:

```
# hagrps -offline service_group -sys system_name
```

- 3 Stop the agent on the system by entering:

```
# haagent -stop agent_name -sys system_name
```

When you get the message `Please look for messages in the log file`, check the file `/var/VRTSvcs/log/engine_A.log` for a message confirming that each agent has stopped.

You can also use the `ps` command to confirm that the agent is stopped.

- 4 Remove the system from the `SystemList` of the service group. If you disable the agent on all the systems in the `SystemList`, you can also remove the service groups and resource types from the VCS configuration.

Read information on administering VCS from the command line.

Refer to the *Veritas Cluster Server User's Guide*.

Removing the Replicated Data Set

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 SVS RPMs using the script-based installer

Use the following procedure to remove SVS 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 SVS 6.0.4 with a previous version of SVS.

To shut down and remove the installed SVS RPMs

- 1 Comment out or remove any Veritas File System (VxFS) entries from the file system table `/etc/fstab`. Failing to remove these entries could result in system boot problems later.
- 2 Unmount all mount points for VxFS file systems.

```
# umount /mount_point
```

- 3 If the VxVM RPM (`VRTSvxvm`) is installed, read and follow the uninstallation procedures for VxVM.
- 4 Make sure you have performed all of the prerequisite steps.
- 5 Move to the `/opt/VRTS/install` directory and run the uninstall script.

```
# cd /opt/VRTS/install
# ./uninstallsvs<version>
```

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

Or, if you are using `ssh` or `rsh`, use one of the following:

```
■ # ./uninstallsvs<version> -rsh
■ # ./uninstallsvs<version> -ssh
```

See [“About the Veritas installer”](#) on page 37.

- 6 The uninstall script prompts for the system name. Enter one or more system names, separated by a space, from which to uninstall SVS, for example, `sys1`:

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

- 7 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.
- 8 Most RPMs have kernel components. In order to ensure complete removal, a system reboot is recommended after all RPMs have been removed.

Uninstalling SVS with the Veritas Web-based installer

This section describes how to uninstall using the Veritas Web-based installer.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout Version in SVS 6.0.4 with a previous version of SVS.

To uninstall SVS

- 1 Perform the required steps to save any data that you wish to preserve. For example, take back-ups of configuration files.
- 2 Start the Web-based installer.

- 3 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.
- 4 Select **Symantec VirtualStore** 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 SVS 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. In order to ensure their complete removal, a system reboot 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 unwanted key files listed in step 2. Unwanted keys may be deleted by removing the license key file.

Removing the CP server configuration using the 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 SVS cluster (application cluster) uses the CP server that you want to unconfigure.

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/installvcsversion -configcps
```

- 2 Select option 3 from the menu to unconfigure the CP server.

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

Select one of the following:

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

- 3 Review the warning message and confirm that you want to unconfigure the CP server.

```
WARNING: Unconfiguring Coordination Point Server stops the
vxcpserv process. VCS clusters using this server for
coordination purpose will have one less coordination point.
```

```
Are you sure you want to bring down the cp server? (y/n)
(Default:n) :y
```

- 4 Review the screen output as the script performs the following steps to remove the CP server configuration:

- Stops the CP server
- Removes the CP server from VCS configuration
- Removes resource dependencies
- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration
- 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.

Installation reference

- [Appendix A. Tunable files for installation](#)
- [Appendix B. Installation scripts](#)
- [Appendix C. Configuration files](#)
- [Appendix D. Response files](#)
- [Appendix E. Configuring the secure shell or the remote shell for communications](#)
- [Appendix F. SVS components](#)
- [Appendix G. Sample SVS cluster setup diagrams for CP server-based I/O fencing](#)
- [Appendix H. Configuring LLT over UDP using IPv6](#)
- [Appendix I. Configuring LLT over UDP using IPv4](#)
- [Appendix J. Compatibility issues when installing Symantec VirtualStore with other products](#)

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 314.

- When you apply the tunables file with no other installer-related operations.

```
# ./installer -tunablesfile tunables_file_name -settunables [  
system1 system2 ...]
```

See [“Setting tunables with no other installer-related operations”](#) on page 315.

- 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 316.

You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 318.

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 318.

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 317.
- 2 Make sure the systems where you want to install SVS meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems where you want to install, configure, or upgrade the product.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer for the installation, configuration, or upgrade. For example:

```
# ./installer -tunablesfile /tmp/tunables_file
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with no other installer-related operations

You can use the installer to set tunable parameters without any other installer-related operations. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 318.

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 317.
- 2 Make sure the systems where you want to install SVS meet the installation requirements.
- 3 Complete any preinstallation tasks.
- 4 Copy the tunables file to one of the systems that you want to tune.
- 5 Mount the product disc and navigate to the directory that contains the installation program.
- 6 Start the installer with the `-setttunables` option.

```
# ./installer -tunablesfile tunables_file_name -setttunables [
sys123 sys234 ...]
```

Where `/tmp/tunables_file` is the full path name for the tunables file.

- 7 Proceed with the operation. When prompted, accept the tunable parameters.
Certain tunables are only activated after a reboot. Review the output carefully to determine if the system requires a reboot to set the tunable value.
- 8 The installer validates the tunables. If an error occurs, exit the installer and check the tunables file.

Setting tunables with an un-integrated response file

You can use the installer to set tunable parameters with an un-integrated response file. You must use the parameters described in this guide. Note that many of the parameters are product-specific. You must select the tunables that you want to use from this guide.

See [“Tunables value parameter definitions”](#) on page 318.

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 SVS meet the installation requirements.
- 2 Complete any preinstallation tasks.
- 3 Prepare the tunables file.
See [“Preparing the tunables file”](#) on page 317.
- 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 318.

Each line for the parameter value starts with \$TUN. The name of the tunable is in curly brackets and double-quotes. The system name is enclosed in curly brackets and double-quotes. Finally define the value and end the line with a semicolon, for example:

```
$TUN{"dmp_daemon_count"}{"node123"}=16;
```

In this example, you are changing the dmp_daemon_count value from its default of 10 to 16. You can use the wildcard symbol "*" for all systems. For example:

```
$TUN{"dmp_daemon_count"}{"*"}=16;
```

Tunables value parameter definitions

When you create a tunables file for the installer you can only use the parameters in the following list.

Prior to making any updates to the tunables, refer to the *Veritas Storage Foundation and High Availability Solutions Tuning Guide* for detailed information on product tunable ranges and recommendations .

[Table A-1](#) describes the supported tunable parameters that can be specified in a tunables file.

Table A-1 Supported tunable parameters

Tunable	Description
dmp_cache_open	(Veritas Dynamic Multi-Pathing) Whether the first open on a device performed by an array support library (ASL) is cached. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_daemon_count	(Veritas Dynamic Multi-Pathing) The number of kernel threads for DMP administrative tasks. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_delayq_interval	(Veritas Dynamic Multi-Pathing) The time interval for which DMP delays the error processing if the device is busy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_fast_recovery	(Veritas Dynamic Multi-Pathing) Whether DMP should attempt to obtain SCSI error information directly from the HBA interface. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table A-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_health_time	(Veritas Dynamic Multi-Pathing) The time in seconds for which a path must stay healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_log_level	(Veritas Dynamic Multi-Pathing) The level of detail to which DMP console messages are displayed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_low_impact_probe	(Veritas Dynamic Multi-Pathing) Whether the low impact path probing feature is enabled. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_lun_retry_timeout	(Veritas Dynamic Multi-Pathing) The retry period for handling transient errors. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_fabric	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) uses the Storage Networking Industry Association (SNIA) HBA API. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_osevent	(Veritas Dynamic Multi-Pathing) Whether the Event Source daemon (vxesd) monitors operating system events. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_monitor_ownership	(Veritas Dynamic Multi-Pathing) Whether the dynamic change in LUN ownership is monitored. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_native_support	(Veritas Dynamic Multi-Pathing) Whether DMP does multi-pathing for native devices. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_path_age	(Veritas Dynamic Multi-Pathing) The time for which an intermittently failing path needs to be monitored before DMP marks it as healthy. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table A-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_pathswitch_blks_shift	(Veritas Dynamic Multi-Pathing) The default number of contiguous I/O blocks sent along a DMP path to an array before switching to the next available path. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_idle_lun	(Veritas Dynamic Multi-Pathing) Whether the path restoration kernel thread probes idle LUNs. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_probe_threshold	(Veritas Dynamic Multi-Pathing) The number of paths will be probed by the restore daemon. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_cycles	(Veritas Dynamic Multi-Pathing) The number of cycles between running the check_all policy when the restore policy is check_periodic. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_interval	(Veritas Dynamic Multi-Pathing) The time interval in seconds the restore daemon analyzes the condition of paths. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_policy	(Veritas Dynamic Multi-Pathing) The policy used by DMP path restoration thread. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_restore_state	(Veritas Dynamic Multi-Pathing) Whether kernel thread for DMP path restoration is started. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_retry_count	(Veritas Dynamic Multi-Pathing) The number of times a path reports a path busy error consecutively before DMP marks the path as failed. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_scsi_timeout	(Veritas Dynamic Multi-Pathing) The timeout value for any SCSI command sent via DMP. This tunable must be set after Veritas Dynamic Multi-Pathing is started.
dmp_sfg_threshold	(Veritas Dynamic Multi-Pathing) The status of the subpaths failover group (SFG) feature. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Table A-1 Supported tunable parameters (*continued*)

Tunable	Description
dmp_stat_interval	(Veritas Dynamic Multi-Pathing) The time interval between gathering DMP statistics. This tunable must be set after Veritas Dynamic Multi-Pathing is started.

Installation scripts

This appendix includes the following topics:

- [Installation script options](#)
- [Restarting the installer after a failed connection](#)
- [Starting and stopping processes for the Veritas products](#)
- [Installation program has improved failure handling](#)

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 Veritas Storage Foundation product scripts, except where otherwise noted.

See [Table B-1](#) on page 322.

Table B-1 Available command line options

Commandline 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.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-comcleanup	The <code>-comcleanup</code> option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-configcps	The <code>-configcps</code> option is used to configure CP server on a running system or cluster.
-configure	Configures the product after installation.
-fencing	Configures I/O fencing in a running cluster.
-fips	The <code>-fips</code> option is used to enable or disable security with fips mode on a running VCS cluster. It could only be used together with <code>-security</code> or <code>-securityonnode</code> option.
-hostfile <i>full_path_to_file</i>	Specifies the location of a file that contains a list of hostnames on which to install.
-installallpkgs	The <code>-installallpkgs</code> option is used to select all RPMs.
-installrecpkgs	The <code>-installrecpkgs</code> option is used to select the recommended RPMs set.
-installminpkgs	The <code>-installminpkgs</code> option is used to select the minimum RPMs set.
-ignorepatchreqs	The <code>-ignorepatchreqs</code> option is used to allow installation or upgrading even if the prerequisite packages or patches are missed on the system.
-keyfile <i>ssh_key_file</i>	Specifies a key file for secure shell (SSH) installs. This option passes <code>-i ssh_key_file</code> to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
-logpath <i>log_path</i>	Specifies a directory other than <code>/opt/VRTS/install/logs</code> as the location where installer log files, summary files, and response files are saved.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-makeresponsefile	Use the <code>-makeresponsefile</code> option only to generate response files. No actual software installation occurs when you use this option.
-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 <code>allpkgs</code> option.
-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 <code>-pkginfo</code> option with the <code>installvcs</code> script to display VCS RPMs.
-pkgpath <i>package_path</i>	Designates the path of a directory that contains all RPMs to install. The directory is typically an NFS-mounted location and must be accessible by all specified installation systems.
-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.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
<code>-recpkgs</code>	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 <code>allpkgs</code> option.
<code>-redirect</code>	Displays progress details without showing the progress bar.
<code>-requirements</code>	The <code>-requirements</code> option displays required OS version, required RPMs and patches, file system space, and other system requirements in order to install the product.
<code>-responsefile <i>response_file</i></code>	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The <i>response_file</i> must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
<code>-rolling_upgrade</code>	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
<code>-rollingupgrade_phase1</code>	The <code>-rollingupgrade_phase1</code> option is used to perform rolling upgrade Phase-I. In the phase, the product kernel RPMs get upgraded to the latest version.
<code>-rollingupgrade_phase2</code>	The <code>-rollingupgrade_phase2</code> option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent RPMs upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
<code>-rsh</code>	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-security	The <code>-security</code> option is used to convert a running VCS cluster between secure and non-secure modes of operation.
-securityonenode	The <code>-securityonenode</code> option is used to configure a secure cluster node by node.
-securitytrust	The <code>-securitytrust</code> option is used to setup trust with another broker.
-serial	Specifies that the installation script performs install, uninstall, start, and stop operations on each system in a serial fashion. If this option is not specified, these operations are performed simultaneously on all systems.
-setunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the <code>-tunablesfile</code> option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The <code>-timeout</code> option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the <code>-timeout</code> option overrides the default value of 1200 seconds. Setting the <code>-timeout</code> option to 0 prevents the script from timing out. The <code>-timeout</code> option does not work with the <code>-serial</code> option.
-tmppath <i>tmp_path</i>	Specifies a directory other than <code>/var/tmp</code> as the working directory for the installation scripts. This destination is where initial logging is performed and where RPMs are copied on remote systems before installation.

Table B-1 Available command line options (*continued*)

Commandline Option	Function
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file <i>tunables_file</i>	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.
-upgrade_kernelpkgs	The <code>-upgrade_kernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase1</code> .
-upgrade_nonkernelpkgs	The <code>-upgrade_nonkernelpkgs</code> option has been renamed to <code>-rollingupgrade_phase2</code> .
-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, hotfixes, and available updates for the installed product if an Internet connection is available.

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.

Starting and stopping processes for the Veritas products

After the installation and configuration is complete, the Veritas product installer starts the processes that are used by the installed products. You can use the product installer to stop or start the processes, if required.

To stop the processes

- ◆ Use the `-stop` option to stop the product installation script.

For example, to stop the product's processes, enter the following command:

```
# ./installer -stop
```

or

```
# /opt/VRTS/install/installsvs<version> -stop
```

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

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/installsvs<version> -start
```

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

Installation program has improved failure handling

The product installer has improved ability to recover from failed installations, as follows:

- A recovery file is created if an installation fails due to a failed network connection. This file enables the install program to resume from the point where the installation failed.
- New options are available to start or stop the Veritas processes without requiring a full installation or configuration.

Configuration files

This appendix includes the following topics:

- [About I/O fencing configuration files](#)

About I/O fencing configuration files

[Table C-1](#) lists the I/O fencing configuration files.

Table C-1 I/O fencing configuration files

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

Table C-1 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfenmode	<p>This file contains the following parameters:</p> <ul style="list-style-type: none"> ■ <code>vxfen_mode</code> <ul style="list-style-type: none"> ■ <code>scsi3</code>—For disk-based fencing ■ <code>customized</code>—For server-based fencing ■ <code>disabled</code>—To run the I/O fencing driver but not do any fencing operations. ■ <code>vxfen_mechanism</code> This parameter is applicable only for server-based fencing. Set the value as <code>cps</code>. ■ <code>scsi3_disk_policy</code> <ul style="list-style-type: none"> ■ <code>dmp</code>—Configure the vxfen module to use DMP devices The disk policy is <code>dmp</code> by default. If you use iSCSI devices, you must set the disk policy as <code>dmp</code>. ■ <code>raw</code>—Configure the vxfen module to use the underlying raw character devices <p>Note: You must use the same SCSI-3 disk policy on all the nodes.</p> <ul style="list-style-type: none"> ■ <code>security</code> This parameter is applicable only for server-based fencing. 1—Indicates that communication with the CP server is in secure mode. This setting is the default. 0—Indicates that communication with the CP server is in non-secure mode. ■ List of coordination points This list is required only for server-based fencing configuration. Coordination points in 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 <code>/etc/vxfen.d/vxfenmode_cps</code> for more information on how to specify the coordination points and multiple IP addresses for each CP server. ■ <code>single_cp</code> 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. ■ <code>autoseed_gab_timeout</code> This parameter enables GAB automatic seeding of the cluster even when some cluster nodes are unavailable. This feature requires that I/O fencing is enabled. 0—Turns the GAB auto-seed feature on. Any value greater than 0 indicates the number of seconds that GAB must delay before it automatically seeds the cluster. -1—Turns the GAB auto-seed feature off. This setting is the default.

Table C-1 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	<p>When I/O fencing starts, the <code>vxfen</code> startup script creates this <code>/etc/vxfentab</code> file on each node. The startup script uses the contents of the <code>/etc/vxfendg</code> and <code>/etc/vxfenmode</code> files. Any time a system is rebooted, the fencing driver reinitializes the <code>vxfentab</code> file with the current list of all the coordinator points.</p> <p>Note: The <code>/etc/vxfentab</code> file is a generated file; do not modify this file.</p> <p>For disk-based I/O fencing, the <code>/etc/vxfentab</code> file on each node contains a list of all paths to each coordinator disk along with its unique disk identifier. A space separates the path and the unique disk identifier. An example of the <code>/etc/vxfentab</code> file in a disk-based fencing configuration on one node resembles as follows:</p> <ul style="list-style-type: none"> ■ Raw disk: <pre style="margin-left: 20px;">/dev/sdx HITACHI%5F1724-100%20%20FASTT%5FDISKS%5F6 00A0B8000215A5D000006804E795D075 /dev/sdy HITACHI%5F1724-100%20%20FASTT%5FDISKS%5F6 00A0B8000215A5D000006814E795D076 /dev/sdz HITACHI%5F1724-100%20%20FASTT%5FDISKS%5F6 00A0B8000215A5D000006824E795D077</pre> ■ DMP disk: <pre style="margin-left: 20px;">/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</pre> <p>For server-based fencing, the <code>/etc/vxfentab</code> file also includes the security settings information.</p> <p>For server-based fencing with single CP server, the <code>/etc/vxfentab</code> file also includes the <code>single_cp</code> settings information.</p>

Response files

This appendix includes the following topics:

- [About response files](#)
- [Upgrading SVS using response files](#)
- [Uninstalling SVS using response files](#)
- [Syntax in the response file](#)
- [Response file variables to install, upgrade, or uninstall Symantec VirtualStore](#)
- [Sample response file for Symantec VirtualStore installation](#)
- [Sample response file for Symantec VirtualStore configuration](#)

About response files

The installer or product installation script generates a response file during any installation, configuration, upgrade (except rolling 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.

Upgrading SVS using response files

Typically, you can use the response file that the installer generates after you perform SVS upgrade on one system to upgrade SVS on other systems. You can also create a response file using the `makeresponsefile` option of the installer.

To perform automated SVS upgrade

- 1 Make sure the systems where you want to upgrade SVS meet the upgrade requirements.
- 2 Make sure the pre-upgrade tasks are completed.
- 3 Copy the response file to one of the systems where you want to upgrade SVS.
- 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  
# ./installsvs -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

Uninstalling SVS using response files

Typically, you can use the response file that the installer generates after you perform SVS uninstallation on one cluster to uninstall SVS on other clusters.

To perform an automated uninstallation

- 1 Make sure that you meet the prerequisites to uninstall SVS.
- 2 Copy the response file to the system where you want to uninstall SVS.
- 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/uninstallsvs<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 Veritas installer”](#) on page 37.

Syntax in the response file

The syntax of the Perl statements that are included in the response file variables varies. It can depend on whether the variables require scalar or list values.

For example, in the case of a string value:

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value", "value", "value"];
```

Response file variables to install, upgrade, or uninstall Symantec VirtualStore

[Table D-1](#) lists the response file variables that you can define to configure SVS.

Table D-1

Variable	Description
CFG{opt}{install}	Installs SVS RPMs. Configuration can be performed at a later time using the <code>-configure</code> option. List or scalar: scalar Optional or required: optional

Table D-1 (continued)

Variable	Description
CFG{accepteula}	<p>Specifies whether you agree with the EULA.pdf file on the media.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>
\$CFG{opt}{vxkeyless}	<p>Installs the product with keyless license.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{systems}	<p>List of systems on which the product is to be installed or uninstalled.</p> <p>List or scalar: list</p> <p>Optional or required: required</p>
CFG{prod}	<p>Defines the product to be installed or uninstalled.</p> <p>List or scalar: scalar</p> <p>Optional or required: required</p>
CFG{opt}{keyfile}	<p>Defines the location of an ssh keyfile that is used to communicate with all remote systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{pkgpath}	<p>Defines a location, typically an NFS mount, from which all remote systems can install product RPMs. The location must be accessible from all target systems.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>
CFG{opt}{tmppath}	<p>Defines the location where a working directory is created to store temporary files and the RPMs that are needed during the install. The default location is /var/tmp.</p> <p>List or scalar: scalar</p> <p>Optional or required: optional</p>

Table D-1 (continued)

Variable	Description
CFG{opt}{rsh}	Defines that <i>rsh</i> must be used instead of <i>ssh</i> as the communication method between systems. List or scalar: scalar Optional or required: optional
CFG{opt}{logpath}	Mentions the location where the log files are to be copied. The default location is <i>/opt/VRTS/install/logs</i> . List or scalar: scalar Optional or required: optional
\$CFG{opt}{prodmode}	List of modes for product List or scalar: list Optional or required: optional
CFG{opt}{uninstall}	Uninstalls SVS RPMs. List or scalar: scalar Optional or required: optional
CFG{mirrordgname}{system}	Splits the target disk group name for a system. List or scalar: scalar Optional or required: optional
CFG{splitmirror}{system}	Indicates the system where you want a split mirror backup disk group created. List or scalar: scalar Optional or required: optional

Sample response file for Symantec VirtualStore installation

The following example shows a response file for installing Symantec VirtualStore.

```
#####
#Auto generated svcs responsefile #
#####
```

```
our %CFG;
$CFG{accepteula}=1;
$CFG{opt}{rsh}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vr}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installallpkgs}=1;
$CFG{prod}="SVS60";
$CFG{systems}=[ qw( sys1 sys2 ) ];

1;
```

Sample response file for Symantec VirtualStore configuration

The following example shows a response file for configuring Symantec VirtualStore.

```
#####
#Auto generated svcs responsefile #
#####

our %CFG;
$CFG{accepteula}=1;
$CFG{fencingenabled}="0";
$CFG{lltoverudp}="0";
$CFG{opt}{configure}=1;
$CFG{opt}{gco}=1;
$CFG{opt}{vr}=1;
$CFG{opt}{installrecpkgs}=1;
$CFG{opt}{vxkeyless}=1;
$CFG{prod}="SVS60";
$CFG{reconfigurevcs}=1;
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=127;
$CFG{vcs_clustername}="uxrt_lin";
$CFG{vcs_lltlink1}{ "sys1" }="eth1";
$CFG{vcs_lltlink1}{ "sys2" }="eth1";
$CFG{vcs_lltlink2}{ "sys1" }="eth2";
$CFG{vcs_lltlink2}{ "sys2" }="eth2";
```

```
$CFG{vcs_userenpw}=[ qw(JlmElgLimHmKumGlj) ];  
$CFG{vcs_username}=[ qw(admin) ];  
$CFG{vcs_userpriv}=[ qw(Administrators) ];  
$CFG{opt}{logpath}="/tmp/log";
```

```
1;
```

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](#)

About configuring secure shell or remote shell communication modes before installing products

You can install products to remote systems using either secure shell (ssh) or remote shell (rsh). Symantec recommends that you use ssh as it is more secure than rsh.

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

Note: The script- and Web-based installers support establishing passwordless communication for you.

SVS components

This appendix includes the following topics:

- [Symantec VirtualStore installation RPMs](#)
- [Veritas Cluster Server installation RPMs](#)
- [Veritas Cluster File System installation RPMs](#)
- [Veritas Storage Foundation obsolete and reorganized installation RPMs](#)

Symantec VirtualStore installation RPMs

[Table F-1](#) shows the RPM name and contents for each English language RPM for Symantec VirtualStore. The table also gives you guidelines for which RPMs to install based on whether you want the minimum, recommended, or advanced configuration.

When you install all Symantec VirtualStore and Veritas Cluster Server (VCS) RPMs, the combined functionality is called Symantec VirtualStore and High Availability.

See [“Veritas Cluster Server installation RPMs”](#) on page 343.

Table F-1 Symantec VirtualStore RPMs

RPMs	Contents	Configuration
VRTSaslapm	Veritas Array Support Library (ASL) and Array Policy Module (APM) binaries Required for the support and compatibility of various storage arrays.	Minimum
VRTSperl	Perl 5.14.2 for Veritas	Minimum
VRTSsvs	Symantec VirtualStore	Minimum

Table F-1 Symantec VirtualStore RPMs (*continued*)

RPMs	Contents	Configuration
VRTSvlic	Veritas License Utilities Installs the license key layout files required to decode the Storage Foundation license keys. Provides the standard license key utilities vxlicrep, vxlicinst, and vxlictest.	Minimum
VRTSvxf	Veritas File System binaries Required for VxFS file system support.	Minimum
VRTSvxvm	Veritas Volume Manager binaries, scripts, and utilities. Required for VxVM volume manager support.	Minimum
VRTSdbed	Veritas Storage Foundation for Databases	Recommended
VRTSob	Veritas Enterprise Administrator	Recommended
VRTSodm	Veritas ODM Driver for VxFS Veritas Extension for Oracle Disk Manager is a custom storage interface designed specifically for Oracle9i and 10g. Oracle Disk Manager allows Oracle 9i and 10g to improve performance and manage system bandwidth.	Recommended

Table F-1 Symantec VirtualStore RPMs (*continued*)

RPMs	Contents	Configuration
VRTSsfpci601	<p>Veritas Storage Foundation Common Product Installer</p> <p>The Storage Foundation Common Product installer RPM contains the installer libraries and product scripts that perform the following:</p> <ul style="list-style-type: none"> ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ removing nodes ■ etc. <p>You can use these script to simplify the native operating system installations, configurations, and upgrades.</p>	Minimum
VRTSsfmh	<p>Veritas Storage Foundation Managed Host</p> <p>Veritas Storage Foundation Managed Host is now called Veritas Operations Manager (VOM).</p> <p>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 package on a server, and add this managed host to the Central Server. The VRTSsfmcs package is not part of this release. You can download it separately from:</p> <p>http://www.symantec.com/veritas-operations-manager</p>	Minimum
VRTSspt	Veritas Software Support Tools	Recommended
VRTSvcsdr	Contains the binaries for Veritas Cluster Server disk reservation.	Recommended

Table F-1 Symantec VirtualStore 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
VRTSilmconv	Symantec Logical Volume Manager (LVM) volume converter Converts offline Linux LVM managed volumes to VxVM volumes by rearranging media contents.	All

Veritas Cluster Server installation RPMs

[Table F-2](#) shows the RPM name and contents for each English language RPM for Veritas Cluster Server (VCS). The table also gives you guidelines for which RPMs to install based on 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 [“Symantec VirtualStore installation RPMs”](#) on page 340.

Table F-2 VCS installation RPMs

RPM	Contents	Configuration
VRTSgab	Veritas Cluster Server group membership and atomic broadcast services	Minimum
VRTSilt	Veritas Cluster Server low-latency transport	Minimum
VRTSamf	Veritas Cluster Server Asynchronous Monitoring Framework	Minimum
VRTSvcS	Veritas Cluster Server	Minimum
VRTSvcSag	Veritas Cluster Server Bundled Agents	Minimum

Table F-2 VCS installation RPMs (*continued*)

RPM	Contents	Configuration
VRTSvxfen	Veritas I/O Fencing	Minimum
VRTSvcsea	Consolidated database and enterprise agent RPMs	Recommended
VRTScps	Veritas Coordination Point Server The Coordination Point Server is an alternate mechanism for I/O fencing. It implements I/O fencing through a client/server architecture and can provide I/O fencing for multiple VCS clusters.	All

Veritas Cluster File System installation RPMs

[Table F-3](#) shows the RPM name and contents for each English language RPM for Veritas Cluster File System (CFS). 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 CFS RPMs and all the RPMs that comprise Storage Foundation and Veritas Cluster Server, the resulting functionality is called Storage Foundation Cluster File System.

See [“Symantec VirtualStore installation RPMs”](#) on page 340.

See [“Veritas Cluster Server installation RPMs”](#) on page 343.

Table F-3 CFS installation RPMs

RPM	Contents	Configuration
VRTScavf	Veritas Cluster Server Agents for Storage Foundation Cluster File System	Minimum
VRTSglm	Veritas Group Lock Manager for Storage Foundation Cluster File System	Minimum
VRTSgms	Veritas Group Messaging Services for Storage Foundation Cluster File System	Recommended

Veritas Storage Foundation obsolete and reorganized installation RPMs

Table F-4 lists the RPMs that are obsolete or reorganized for Symantec VirtualStore.

Table F-4 Veritas Storage Foundation obsolete and reorganized RPMs

RPM	Description
Obsolete and reorganized for 6.0.4	
VRTSat	Obsolete
VRTSatClient	Obsolete
VRTSatServer	Obsolete
Obsolete and reorganized for 5.1	
Infrastructure	
SYMCIma	Obsolete
VRTSaa	Included in VRTSsfmh
VRTSccg	Included in VRTSsfmh
VRTSdbms3	Obsolete
VRTSicsco	Obsolete
VRTSjre	Obsolete
VRTSjre15	Obsolete
VRTSmh	Included in VRTSsfmh
VRTSobc33	Obsolete
VRTSobweb	Obsolete
VRTSobgui	Obsolete
VRTSpbx	Obsolete
VRTSsfm	Obsolete
VRTSweb	Obsolete
Product RPMs	

Table F-4 Veritas Storage Foundation obsolete and reorganized RPMs
(continued)

RPM	Description
VRTSaclib	<p>Obsolete</p> <p>The following information is for installations, upgrades, and uninstalls using the script- or Web-based installer.</p> <ul style="list-style-type: none"> ■ For fresh installations VRTSaclib is not installed. ■ For upgrades, VRTSaclib is not uninstalled. ■ For uninstalls, VRTSaclib 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
VRTSdb2ed-common	Included in VRTSdbed
VRTSdbcom-common	Included in VRTSdbed
VRTSdbed-common	Included in VRTSdbed
VRTSdcli	Obsolete
VRTSddlpr	Obsolete
VRTSdsa	Obsolete
VRTSfsman	Included in the product's main RPM.
VRTSfsmnd	Included in the product's main RPM.

Table F-4 Veritas Storage Foundation obsolete and reorganized RPMs
(continued)

RPM	Description
VRTSfspro	Included in VRTSsfmh
VRTSmapro-common	Included in VRTSsfmh
VRTSodm-common	Included in VRTSodm
VRTSodm-platform	Included in VRTSodm
VRTSorgui-common	Obsolete
VRTSvcldb	Included in VRTSvcsea
VRTSvcsmn	Included in VRTSvc
VRTSvcsor	Included in VRTSvcsea
VRTSvcsvr	Included in VRTSvc
VRTSvdid	Obsolete
VRTSvmman	Included in the product's main RPM.
VRTSvmpro	Included in VRTSsfmh
VRTSvrpro	Included in VRTSob
VRTSvrw	Obsolete
VRTSvxfs-common	Included in VRTSvxfs
VRTSvxfs-platform	Included in VRTSvxfs
VRTSvxmsa	Obsolete
VRTSvxvm-common	Included in VRTSvxvm
VRTSvxvm-platform	Included in VRTSvxvm
Documentation	All Documentation RPMs obsolete

Sample SVS cluster setup diagrams for CP server-based I/O fencing

This appendix includes the following topics:

- [Configuration diagrams for setting up server-based I/O fencing](#)

Configuration diagrams for setting up server-based I/O fencing

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:

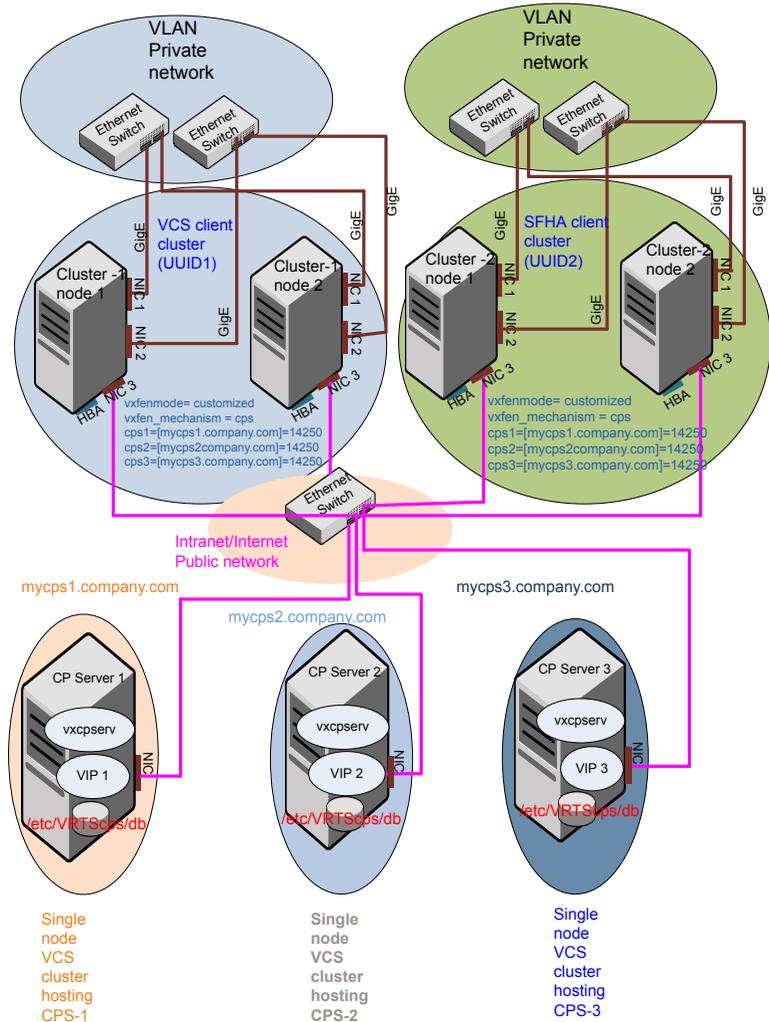
- Two unique client clusters that are served by 3 CP servers:
See [Figure G-1](#) on page 349.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
See [Figure G-2](#) on page 351.
- Two node campus cluster that is served by remote CP server and 2 SCSI-3 disks:
See [Figure G-4](#) on page 354.
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:
See [Figure G-6](#) on page 358.

Two unique client clusters served by 3 CP servers

Figure G-1 displays a configuration where two unique client clusters are being served by 3 CP servers (coordination points). Each client cluster has its own unique user ID (UUID1 and UUID2).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

Figure G-1 Two unique client clusters served by 3 CP servers



Client cluster served by highly available CPS and 2 SCSI-3 disks

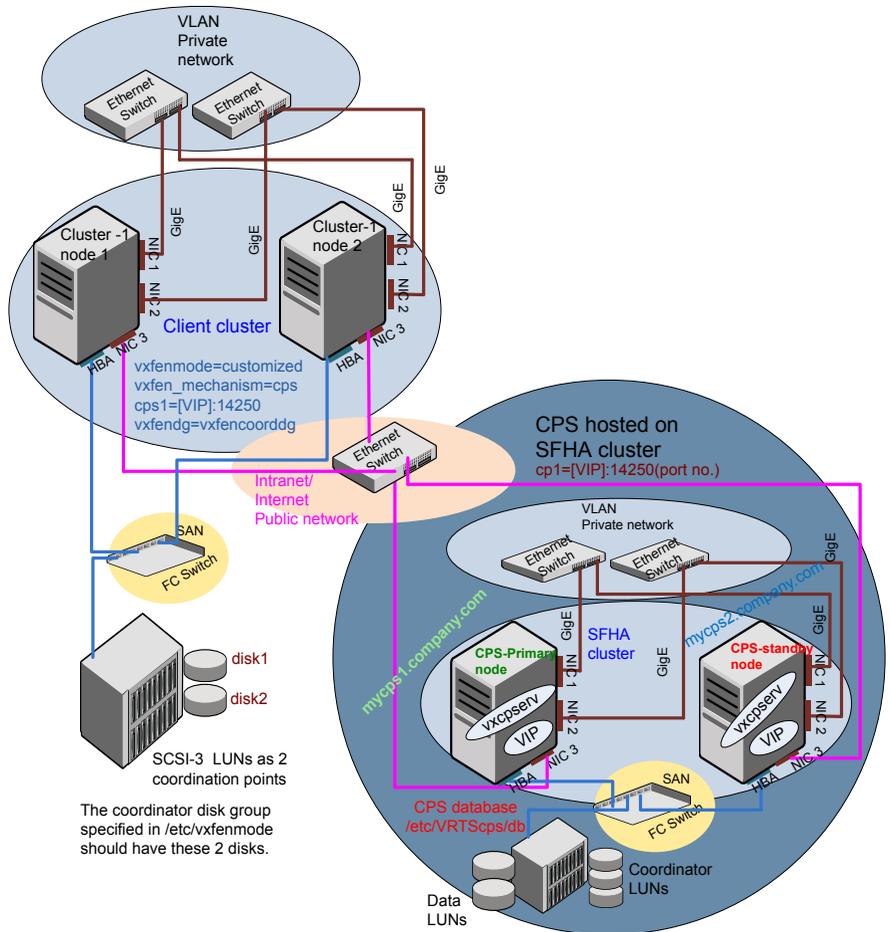
Figure G-2 displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to customized with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

The two SCSI-3 disks are part of the disk group vxfcntlcoordg. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure G-3 Client cluster served by highly available CP server and 2 SCSI-3 disks



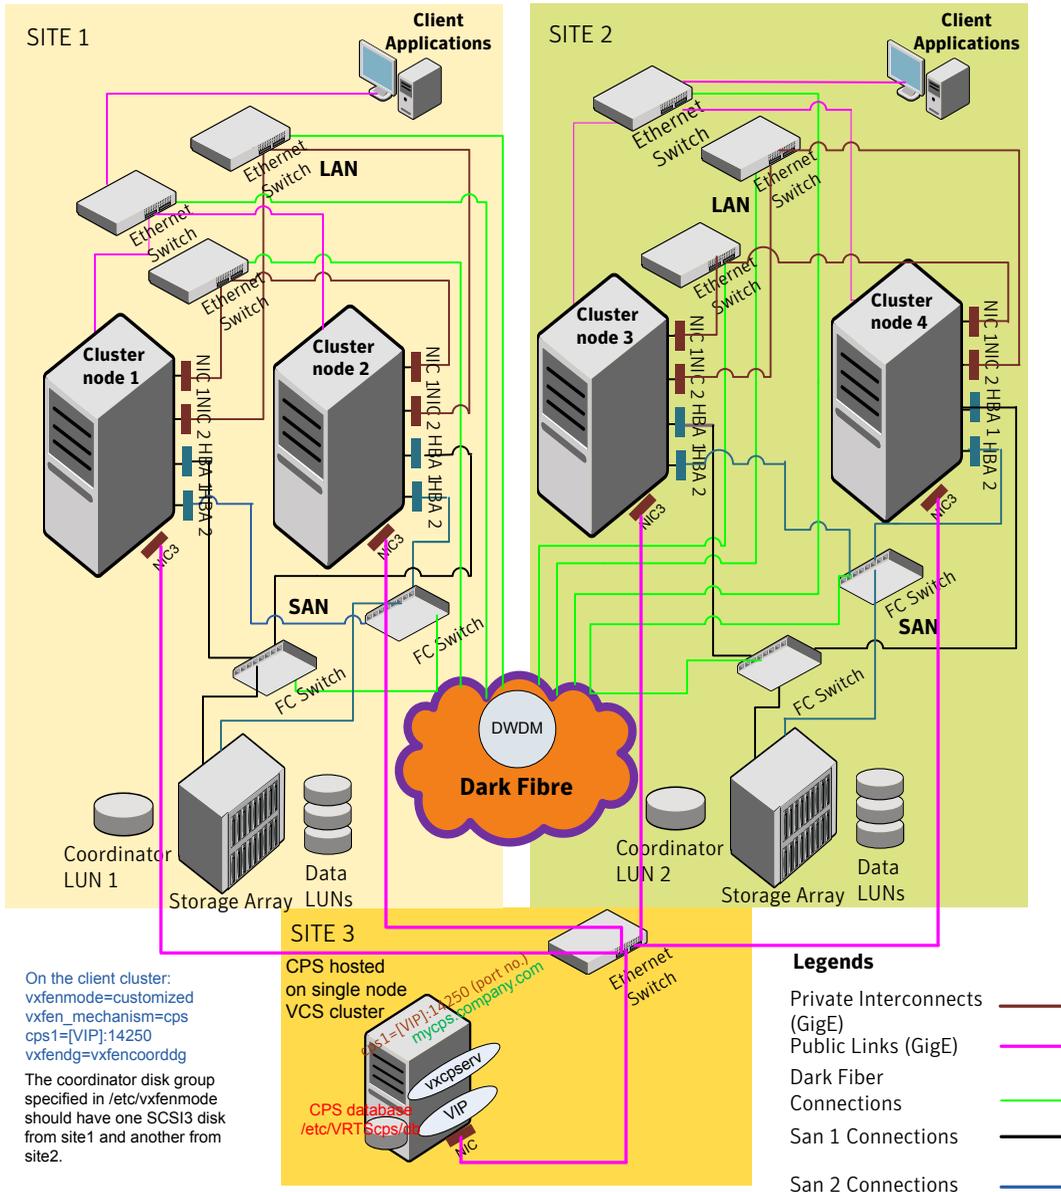
Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure G-4 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 G-4 Two node campus cluster served by remote CP server and 2 SCSI-3



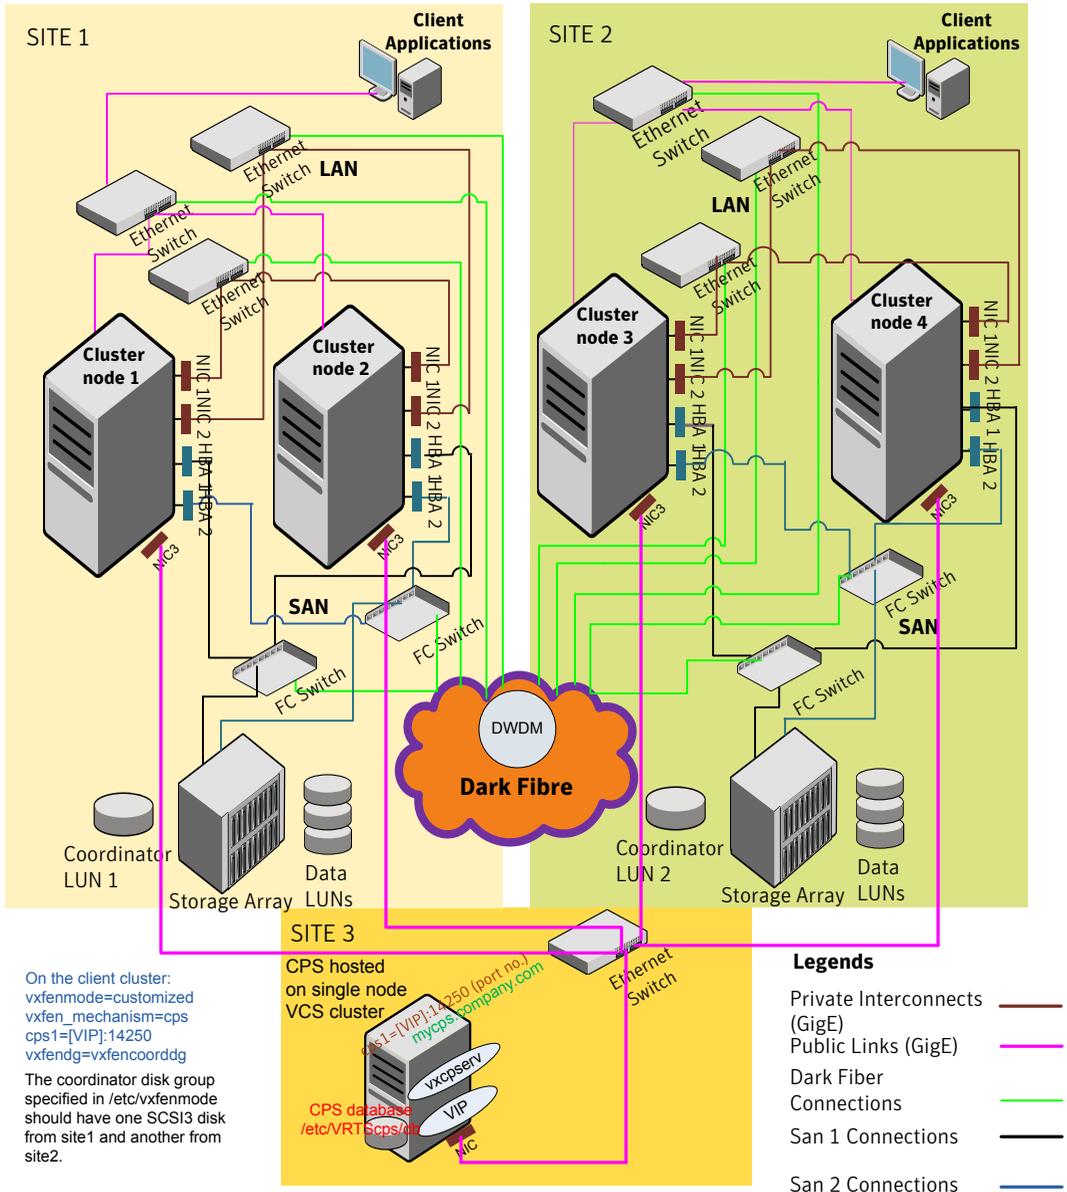
Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure G-5 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 G-5 Two node campus cluster served by remote CP server and 2 SCSI-3



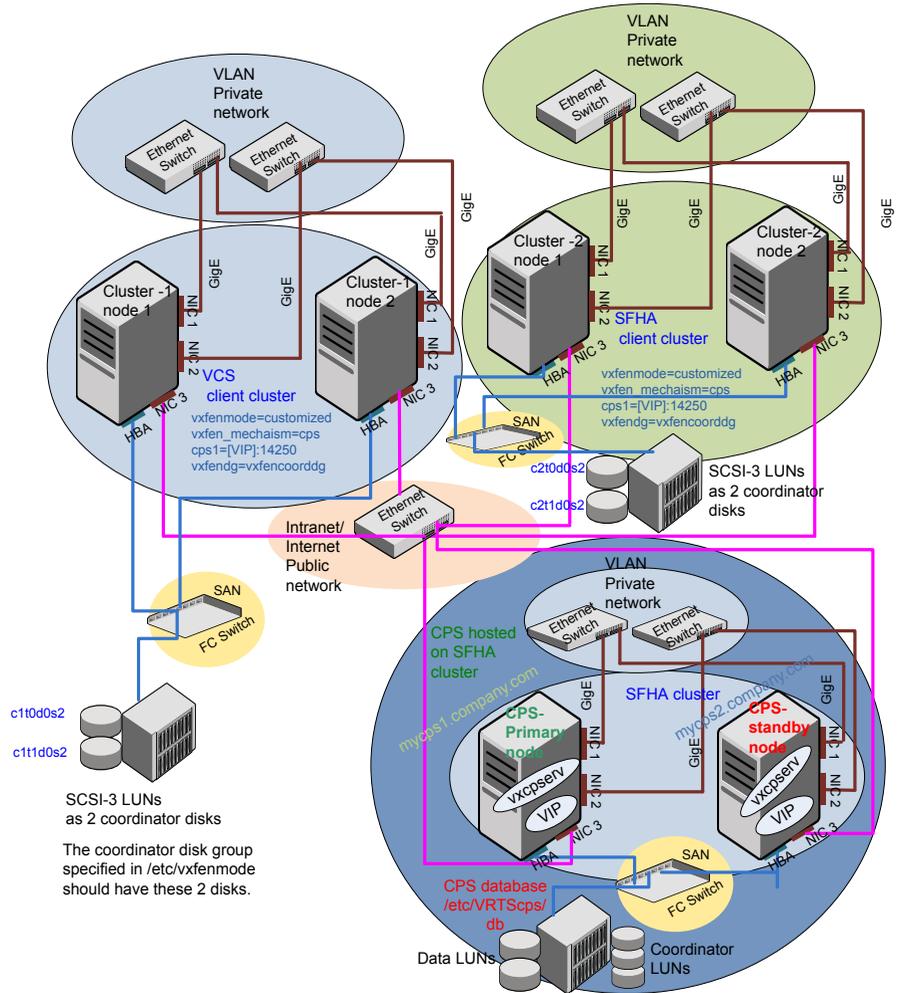
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

Figure G-6 displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNS (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen` mechanism set to `cps`.

The 2 SCSI-3 disks are: `c1t0d0s2` and `c1t1d0s2` which are part of disk group `vxfencoordg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure G-6 Multiple client clusters served by highly available CP server and 2 SCSI-3 disks



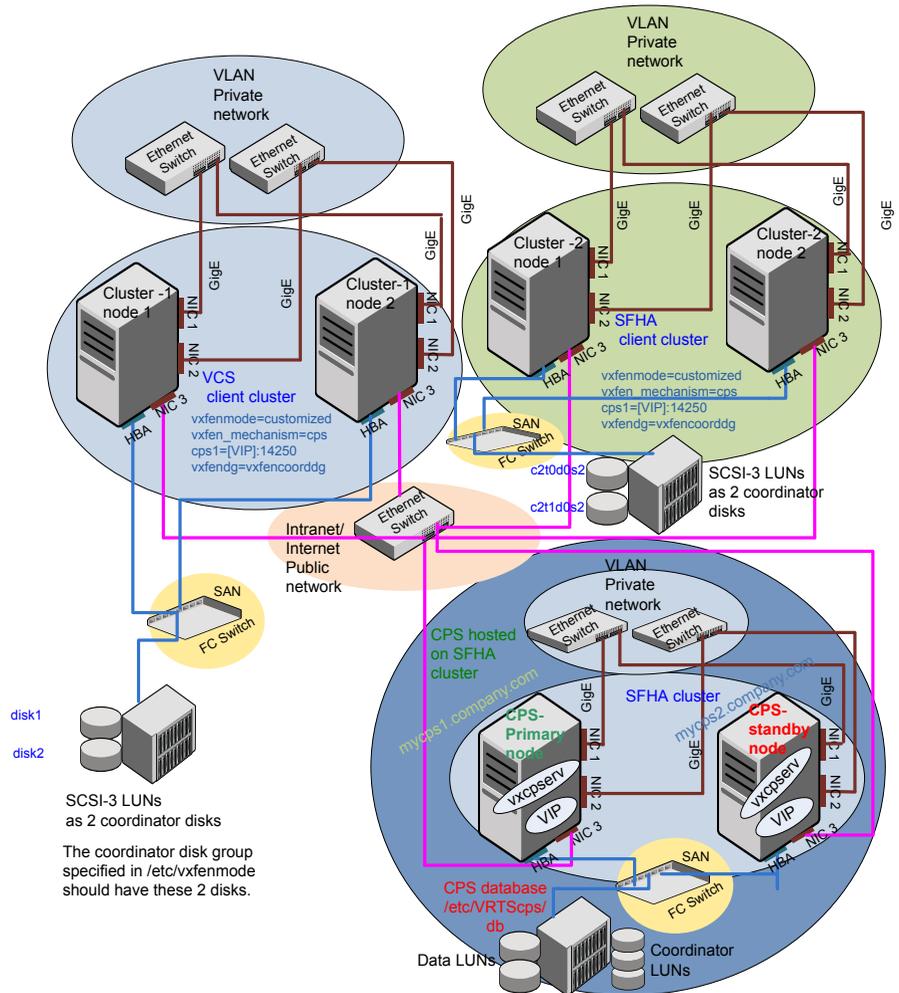
Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

Figure G-7 displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to `customized` with `vxfen_mechanism` set to `cps`.

The two SCSI-3 disks are part of the disk group `vxfencoorddg`. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure G-7 Multiple client clusters served by highly available CP server and 2 SCSI-3 disks



Configuring LLT over UDP using IPv6

This appendix includes the following topics:

- [Using the UDP layer of IPv6 for LLT](#)
- [Manually configuring LLT over UDP using IPv6](#)

Using the UDP layer of IPv6 for LLT

6.0.4 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.
 See [“Selecting UDP ports”](#) on page 362.
- 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 364.

The link command in the `/etc/llttab` file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See [“Sample configuration: direct-attached links”](#) on page 363.
- See [“Sample configuration: links crossing IP routers”](#) on page 364.

Note that some of the fields in [Table H-1](#) differ from the command for standard LLT links.

[Table H-1](#) describes the fields of the link command that are shown in the `/etc/llttab` file examples.

Table H-1 Field description for link command in `/etc/llttab`

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device name of the UDP protocol; for example udp6.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp6" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “Selecting UDP ports” on page 362.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IPv6 address</i>	IPv6 address of the link on the local node.
<i>mcast-address</i>	"-" is the default for clusters spanning routers.

The set-addr command in the /etc/llttab file

The `set-addr` command in the `/etc/llttab` file is required when the multicast feature of LLT is disabled, such as when LLT must cross IP routers.

See “[Sample configuration: links crossing IP routers](#)” on page 364.

[Table H-2](#) describes the fields of the `set-addr` command.

Table H-2 Field description for `set-addr` command in `/etc/llttab`

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.
<i>link tag-name</i>	The string that LLT uses to identify the link; for example link1, link2,....
<i>address</i>	IPv6 address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file `/etc/services`. You should also use the `netstat` command to list the UDP ports currently in use. For example:

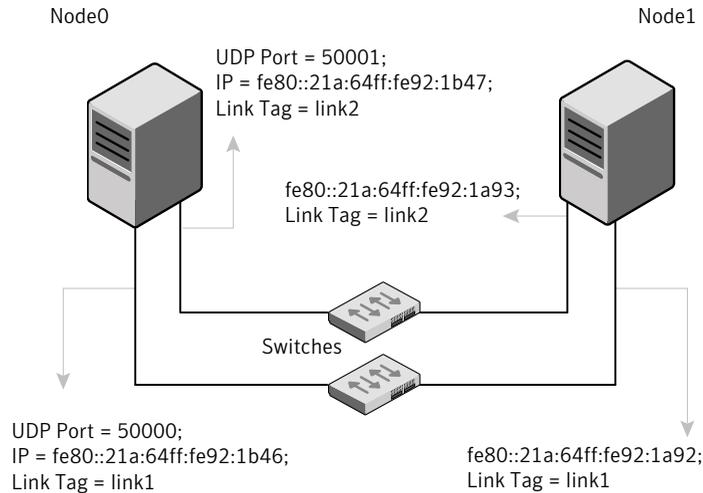
```
# netstat -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address      State
udp      0      0 *:32768          *:*
udp      0      0 *:956            *:*
udp      0      0 *:tftp           *:*
udp      0      0 *:sunrpc         *:*
udp      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.

Sample configuration: direct-attached links

Figure H-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure H-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```

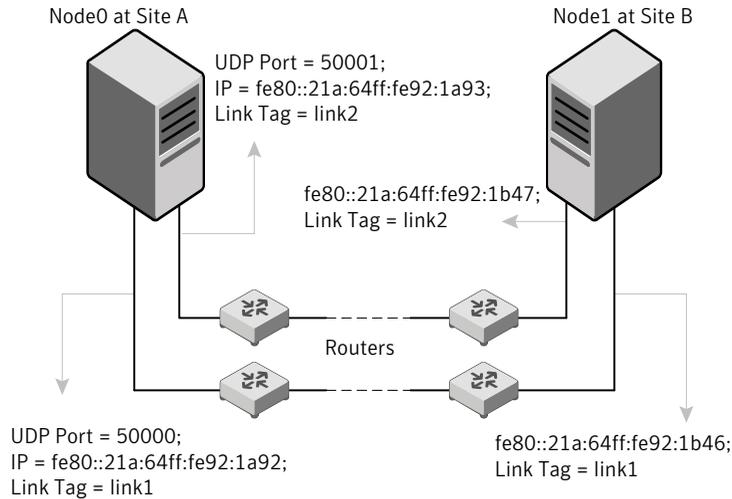
The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

Sample configuration: links crossing IP routers

Figure H-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 H-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
```

```
set-addr 0 link1 fe80::21a:64ff:fe92:1b46
set-addr 0 link2 fe80::21a:64ff:fe92:1b47
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

The `/etc/lltab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 1 link1 fe80::21a:64ff:fe92:1a92
set-addr 1 link2 fe80::21a:64ff:fe92:1a93
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb 0
set-arp 0
```

Configuring LLT over UDP using IPv4

This appendix includes the following topics:

- [Using the UDP layer for LLT](#)
- [Manually configuring LLT over UDP using IPv4](#)

Using the UDP layer for LLT

SVS 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/l1ttab` explicitly depending on the subnet for each link. See [“Broadcast address in the /etc/l1ttab file”](#) on page 367.

- 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 369.
- Set the broadcast address correctly for direct-attached (non-routed) links. See [“Sample configuration: direct-attached links”](#) on page 370.
- 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 372.

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/llttab` file

Review the link command information in this section for the `/etc/llttab` file. See the following information for sample configurations:

- See [“Sample configuration: direct-attached links”](#) on page 370.
- See [“Sample configuration: links crossing IP routers”](#) on page 372.

[Table I-1](#) describes the fields of the link command that are shown in the `/etc/llttab` file examples. Note that some of the fields differ from the command for standard LLT links.

Table I-1 Field description for link command in `/etc/llttab`

Field	Description
<i>tag-name</i>	A unique string that is used as a tag by LLT; for example link1, link2,....
<i>device</i>	The device path of the UDP protocol; for example udp. A place holder string. On other unix platforms like Solaris or HP, this entry points to a device file (for example, <code>/dev/udp</code>). Linux does not have devices for protocols. So this field is ignored.
<i>node-range</i>	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
<i>link-type</i>	Type of link; must be "udp" for LLT over UDP.
<i>udp-port</i>	Unique UDP port in the range of 49152-65535 for the link. See “Selecting UDP ports” on page 369.
<i>MTU</i>	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the <code>lltstat -l</code> command to display the current value.
<i>IP address</i>	IP address of the link on the local node.
<i>bcast-address</i>	<ul style="list-style-type: none"> ■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the `/etc/llttab` file

The `set-addr` command in the `/etc/llttab` file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See [“Sample configuration: links crossing IP routers”](#) on page 372.

[Table I-2](#) describes the fields of the `set-addr` command.

Table I-2 Field description for set-addr command in /etc/llttab

Field	Description
<i>node-id</i>	The ID of the cluster node; for example, 0.
<i>link tag-name</i>	The string that LLT uses to identify the link; for example link1, link2,....
<i>address</i>	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
 - Ports from the range of well-known ports, 0 to 1023
 - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -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/llttab` depending on the subnet that the links are on.

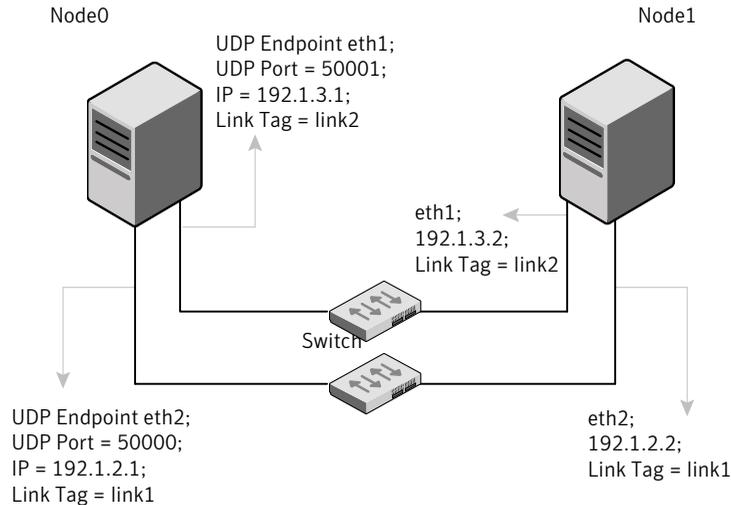
An example of a typical `/etc/llttab` file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

```
# cat /etc/llttab  
set-node nodexyz  
set-cluster 100  
  
link link1 udp - udp 50000 - 192.168.30.1 192.168.30.255  
link link2 udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: direct-attached links

[Figure I-1](#) depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure I-1 A typical configuration of direct-attached links that use LLT over UDP



The configuration that the `/etc/llttab` file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the `/etc/llttab` file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the `/etc/llttab` file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address 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:

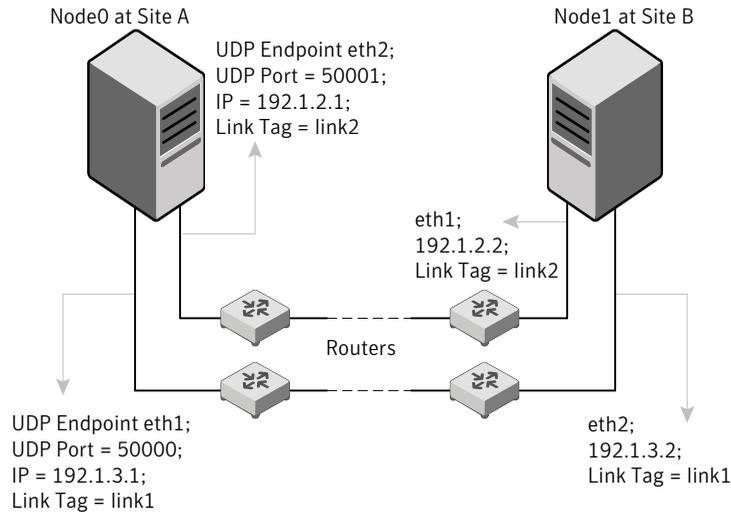
```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
```

```
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 I-2 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure I-2 A typical configuration of links crossing an IP router



The configuration that the following `/etc/llttab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/llttab` file.

```
set-node Node1
set-cluster 1

link link1 udp - udp 50000 - 192.1.3.1 -
link link2 udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 0 link1 192.1.1.1
set-addr 0 link2 192.1.2.1
set-addr 2 link1 192.1.5.2
set-addr 2 link2 192.1.6.2
```

```
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3
```

```
#disable LLT broadcasts
set-bcasthb   0
set-arp       0
```

The `/etc/lltab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 udp - udp 50000 - 192.1.1.1 -
link link2 udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr      1 link1 192.1.3.1
set-addr      1 link2 192.1.4.1
set-addr      2 link1 192.1.5.2
set-addr      2 link2 192.1.6.2
set-addr      3 link1 192.1.7.3
set-addr      3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb   0
set-arp       0
```

Compatibility issues when installing Symantec VirtualStore with other products

This appendix includes the following topics:

- [Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present](#)
- [Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present](#)

Installing, uninstalling, or upgrading Storage Foundation products when other Veritas products are present

Installing Storage Foundation when other Veritas products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where SFM or VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the SFM Central Server and Managed Host RPMs as is.
- When uninstalling Storage Foundation products where SFM or VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where SFM or VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when ApplicationHA is already present

If you plan to install or upgrade Storage Foundation on systems where ApplicationHA has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not consider VCS as an installed product even though it uses the bundled VRTSvcS RPM.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer does not allow the installation or upgrade for products that use VCS. The following products cannot be installed or upgraded: VCS, SFHA, SFCFS, SFCFSHA, SFRAC, SFCFSRAC or SFSYBASECE.
- When you install or upgrade Storage Foundation products where ApplicationHA is present, the installer allows the installation or upgrade of VM, FS, SF, or DMP.
- When you uninstall Storage Foundation products where ApplicationHA is present, the installer does not uninstall VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpx, VRTSicsco, and VRTSat.

Index

A

- about
 - Symantec VirtualStore 19
- Adding
 - connection to Citrix Desktop Studio 260
- adding
 - users 101
- agents
 - disabling 305
- applications, stopping 208
- attributes
 - UseFence 121

B

- backup boot disk group 234–235
 - rejoining 234

C

- cables
 - cross-over Ethernet 275
- Citrix Desktop Studio
 - adding a connection to 260
- Citrix XenDesktop
 - about adding virtual machine clones 258
 - connecting to a vCenter Server 259
- Citrix XenDesktop extension
 - creating virtual machine clones 263
- cluster
 - removing a node from 292
- Clustered NFS
 - setting up 238
- commands
 - vxdisksetup (initializing disks) 113
 - vxlicinst 108
 - vxlicrep 108
- configuring
 - rsh 43
 - ssh 43
 - VirtualStore Web Server for cluster failover 241
 - VMware View 255

- configuring Symantec VirtualStore
 - script-based installer 88
- configuring VCS
 - adding users 101
 - event notification 102–103
 - global clusters 105
 - starting 89
- coordinator disks
 - DMP devices 23–24
 - for I/O fencing 23–24
 - setting up 119
- creating virtual machine clones
 - specifying where to create 251
 - using Citrix XenDesktop 263
 - using Symantec FileSnap 249

D

- data disks
 - for I/O fencing 23
- deployment
 - testing desktop group 268
- desktop group
 - testing the deployment 268
- disabling the agents 305
- disk space
 - requirements for disk space 33, 49
- disks
 - adding and initializing 113
 - coordinator 119
 - testing with vxfcntlshdw 114
 - verifying node access 115

E

- ESX servers
 - adding storage to 248
- Ethernet controllers 275

F

- Fibre Channel fabric 47
- freezing service groups 208

G

- global clusters
 - configuration 105
- guest operating system
 - specifying if customizing for Linux or Windows 252

H

- hubs
 - independent 275

I

- I/O fencing
 - checking disks 114
 - setting up 118
 - shared storage 114
- I/O fencing requirements
 - non-SCSI-3 31
- installer
 - Symantec VirtualStore 55
- Installing
 - SVS with the Web-based installer 159
 - Symantec VirtualStore using the installer 55
- installing
 - post 106
- installing VEA
 - planning 203

L

- license keys
 - adding with vxlicinst 108
 - replacing demo key 108
- licenses
 - information about 108
- licensing 33
 - add-on 32
 - CDS 33
 - full 32
- LLT
 - interconnects 42

M

- media speed 42
 - optimizing 42

N

- network customization parameters
 - specifying for guest operating systems if using DHCP or Static IP 254
- nodes
 - adding application nodes
 - configuring GAB 282
 - configuring LLT 282
 - configuring VXFEN 282
 - starting Volume Manager 281
- non-SCSI-3 fencing
 - manual configuration 146
 - setting up 146
- non-SCSI-3 I/O fencing
 - requirements 31
- non-SCSI3 fencing
 - setting up 132
 - using installsvs 132
- number of virtual machines to be created
 - specifying 251

O

- optimizing
 - media speed 42

P

- persistent reservations
 - SCSI-3 43
- planning to upgrade VVR 202
- pre-requisites
 - virtual machine customization 243
 - VMware View Integration 244
- preinstallation 202
- preparing to upgrade VVR 208

R

- registered
 - VMware vCenter Servers 245
- Registering
 - VirtualStore cluster 242
- rejoining
 - backup boot disk group 234
- removing
 - the Replicated Data Set 306
- removing a system from a cluster 292
- Replicated Data Set
 - removing the 306

- requirements for disk space
 - disk space requirements 33, 49
- rolling upgrade 225, 228
 - versions 225
- rsh 90
 - configuration 43

S

- SAN
 - see Storage Area Network 47
- script-based installer
 - Symantec VirtualStore configuration overview 88
- SCSI-3
 - persistent reservations 43
- SCSI-3 persistent reservations
 - verifying 118
- service groups
 - freezing 208
- setting up
 - Clustered NFS 238
 - VirtualStore vSphere Service Console 241
- SMTP email notification 102
- SNMP trap notification 103
- Solution
 - Symantec VirtualStore 19
- specifying
 - number of virtual machine clones to be created 251
- specifying guest operating system
 - if customizing for Linux or Windows 252
- specifying network customization parameters for guest operating systems
 - if customizing for DHCP or Static IP 254
- ssh 90
 - configuration 43
- starting configuration
 - installvcs program 90
 - Veritas product installer 89
- stopping
 - applications 208
- storage
 - adding to ESX servers 248
- Storage Area Network 47
- SVS
 - coordinator disks 119
- Symantec FileSnap
 - creating virtual machine clones with 249
- Symantec VirtualStore
 - about 19

- Symantec VirtualStore *(continued)*
 - configuring 88
 - installer 55
 - solution 19

T

- tunables file
 - about setting parameters 313
 - parameter definitions 318
 - preparing 317
 - setting for configuration 314
 - setting for installation 314
 - setting for upgrade 314
 - setting parameters 317
 - setting with no other operations 315
 - setting with un-integrated response file 316

U

- unregistering
 - VirtualStore cluster 244
- unsuccessful upgrade 235
- upgrading
 - rolling 225
- upgrading VVR
 - from 4.1 204
 - planning 202
 - preparing 208
- Useful links
 - VMware 245
- Using
 - VirtualStore Plug-in for VMware vSphere 247

V

- vCenter Server
 - connecting to Citrix XenDesktop 259
- VEA installation
 - planning 203
- verifying
 - NIC configuration 107
 - virtual machine clones 257
 - VirtualStore cluster 243
- virtual machine clones
 - adding to Citrix XenDesktop 258
 - specifying where to create 251
 - verifying 257
- virtual machine customization
 - pre-requisites 243

- VirtualStore
 - WMware Vsphere environment 19
- VirtualStore cluster
 - registering 242
 - unregistering 244
 - verifying 243
- VirtualStore Plug-in for VMware vSphere
 - using 247
- VirtualStore vSphere Service Console
 - configure 241
- VirtualStore Web Server for cluster failover
 - configure 241
- VMotion 20
- VMware 19
 - useful links 245
- VMware vCenter Servers
 - registered 245
- VMware View
 - configuring 255
- VMware View Integration
 - pre-requisites 244
- vradmin
 - delpri 307
 - stoprep 306
- VVR 4.1
 - planning an upgrade from 204
- vxdisksetup command 113
- vxlicinst command 108
- vxlicrep command 108
- vxplex
 - used to remove mirrors of root disk volumes 303

W

- Web-based installer 159