Symantec™ Cluster Server 6.2 Installation Guide - AIX

January 2015

Symantec

Symantec™ Cluster Server Installation Guide

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

Product version: 6.2

Document version: 6.2 Rev 2

Legal Notice

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

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

Support agreement resources

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

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

About Symantec Connect

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

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

Contents

Section 1

Installation overview and planning

Chapter 1

Introducing Symantec Cluster Server

This chapter includes the following topics:

About Symantec™ Cluster Server

Symantec™ Cluster Server by Symantec is a high-availability solution for applications and services configured in a cluster. Symantec Cluster Server (VCS) monitors systems and application services, and restarts services when hardware or software fails.

About VCS basics

A single VCS cluster consists of multiple systems that are connected in various combinations to storage devices. When a system is part of a VCS cluster, it is called a node. VCS monitors and controls applications running in the cluster on nodes, and restarts applications in response to a variety of hardware or software faults.

Applications can continue to operate with little or no downtime. In some cases, such as NFS, this continuation is transparent to high-level applications and users. In

other cases, a user might have to retry an operation, such as a Web server reloading a page.

Figure 1-1 illustrates a typical VCS configuration of four nodes that are connected to shared storage.

Figure 1-1 Example of a four-node VCS cluster

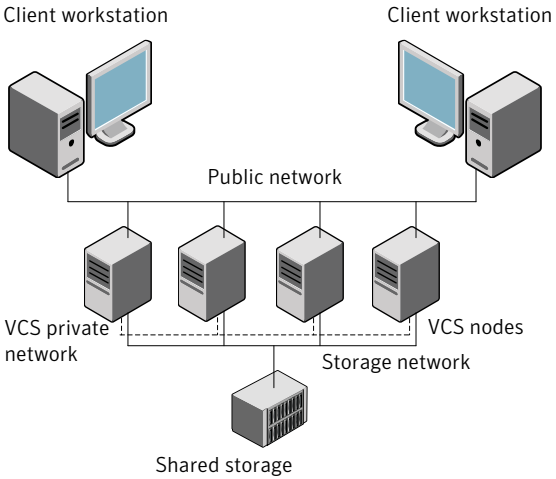

Client workstations receive service over the public network from applications running on VCS nodes. VCS monitors the nodes and their services. VCS nodes in the cluster communicate over a private network.

About multiple nodes

VCS runs in a replicated state on each node in the cluster. A private network enables the nodes to share identical state information about all resources. The private network also recognizes active nodes, nodes that join or leave the cluster, and failed nodes. The private network requires two communication channels to guard against network partitions.

About shared storage

A VCS hardware configuration typically consists of multiple nodes that are connected to shared storage through I/O channels. Shared storage provides multiple systems with an access path to the same data. It also enables VCS to restart applications on alternate nodes when a node fails, which ensures high availability.

VCS nodes can only access physically-attached storage.

Figure 1-2 illustrates the flexibility of VCS shared storage configurations.

Figure 1-2 Two examples of shared storage configurations

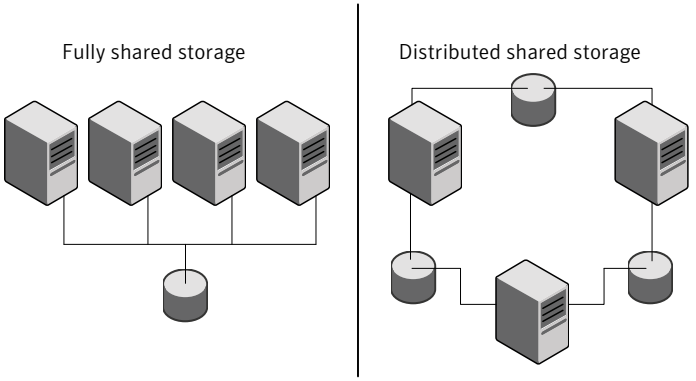

About LLT and GAB

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

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

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

About network channels for heartbeating

For the VCS private network, two network channels must be available to carry heartbeat information. These network connections also transmit other VCS-related information.

Each cluster configuration requires at least two network channels between the systems. The requirement for two channels protects your cluster against network partitioning. For more information on network partitioning, refer to the *Symantec Cluster Server Administrator's Guide*.

Figure 1-3 illustrates a two-node VCS cluster where the nodes sys1 and sys2 have two private network connections.

Figure 1-3 Two Ethernet connections connecting two nodes

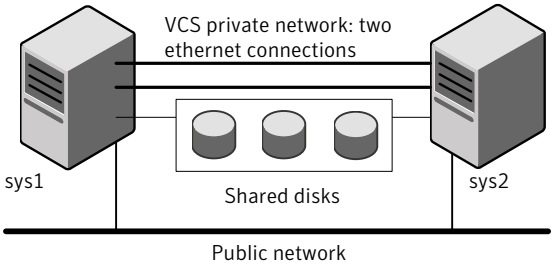

About preexisting network partitions

A preexisting network partition refers to failure in the communication channels that occurs while the systems are down and VCS cannot respond. When the systems start, VCS seeding reduces vulnerability to network partitioning, regardless of the cause of the failure.

About VCS seeding

To protect your cluster from a preexisting network partition, VCS uses the concept of seeding. Seeding is a function of GAB that determines whether or not all nodes have joined a cluster. For this determination, GAB requires that you declare the number of nodes in the cluster. Note that only seeded nodes can run VCS.

GAB automatically seeds nodes under the following conditions:

- An unseeded node communicates with a seeded node
- All nodes in the cluster are unseeded but can communicate with each other

When the last system starts and joins the cluster, the cluster seeds and starts VCS on all nodes. You can then bring down and restart nodes in any combination. Seeding remains in effect as long as at least one instance of VCS is running somewhere in the cluster.

Perform a manual seed to run VCS from a cold start when one or more systems of the cluster are unavailable. VCS does not start service groups on a system until it has a seed. However, if you have I/O fencing enabled in your cluster, you can still configure GAB to automatically seed the cluster even when some cluster nodes are unavailable.

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

About VCS features

VCS offers the following features that you can configure during VCS configuration:

VCS notifications	See "About VCS notifications" on page 27.
VCS global clusters	See "About global clusters" on page 27.
I/O fencing	See "About I/O fencing" on page 27.

About VCS notifications

You can configure both Simple Network Management Protocol (SNMP) and Simple Mail Transfer Protocol (SMTP) notifications for VCS. Symantec recommends you to configure at least one of these notifications. You have the following options:

- Configure SNMP trap notification of VCS events using the VCS Notifier component.
- Configure SMTP email notification of VCS events using the VCS Notifier component.

See the *Symantec Cluster Server Administrator's Guide* for details on configuring these notifications.

About global clusters

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

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

About I/O fencing

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

The fencing operation determines the following:

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

This decision prevents possible data corruption. The installer installs the I/O fencing driver, part of VRTSvxfen fileset, when you install VCS. To protect data on shared disks, you must configure I/O fencing after you install and configure VCS.

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

Disk-based I/O fencing	I/O fencing that uses coordinator disks is referred to as disk-based I/O fencing. Disk-based I/O fencing ensures data integrity in a single cluster.
Server-based I/O fencing	I/O fencing that uses at least one CP server system is referred to as server-based I/O fencing. Server-based fencing can include only CP servers, or a mix of CP servers and coordinator disks. Server-based I/O fencing ensures data integrity in clusters. In virtualized environments that do not support SCSI-3 PR, VCS supports non-SCSI-3 I/O fencing. See "About I/O fencing for VCS in virtual machines that do not support SCSI-3 PR" on page 33.
Majority-based I/O fencing	Majority-based I/O fencing mode does not need coordination points to provide protection against data corruption and data consistency in a clustered environment. Symantec designed majority-based I/O fencing mode to be used in stand-alone appliances. You can configure I/O fencing in majority-based mode, but as a best practice that where possible, utilize Coordination Point servers and or shared SCSI-3 disks to be used as coordination points.

See " About planning to configure I/O fencing" on page 90.

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

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

About VCS optional components

You can add the following optional components to VCS:

Veritas Operations Manager	See "About Veritas Operations Manager" on page 29.
Cluster Manager (Java console)	See "About Cluster Manager (Java Console)" on page 29.
VCS Simulator	See "About VCS Simulator" on page 30.

About Veritas Operations Manager

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

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

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

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

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

About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers administration capabilities for your cluster. Use the different views in the Java Console to monitor and manage clusters and Symantec Cluster Server (VCS) objects, including service groups, systems, resources, and resource types. You cannot manage the new features of releases 6.0 and later using the Java Console.

See *Symantec Cluster Server Administrator's Guide*.

You can download the console from http://www.symantec.com/operations-manager/support.

About VCS Simulator

VCS Simulator enables you to simulate and test cluster configurations. Use VCS Simulator to view and modify service group and resource configurations and test failover behavior. VCS Simulator can be run on a stand-alone system and does not require any additional hardware. You can install VCS Simulator only on a Windows operating system.

VCS Simulator runs an identical version of the VCS High Availability Daemon (HAD) as in a cluster, ensuring that failover decisions are identical to those in an actual cluster.

You can test configurations from different operating systems using VCS Simulator. For example, you can run VCS Simulator to test configurations for VCS clusters on Windows, AIX, HP-UX, Linux, and Solaris operating systems. VCS Simulator also enables creating and testing global clusters.

You can administer VCS Simulator from the Java Console or from the command line.

To download VCS Simulator, go to http://www.symantec.com/operations-manager/support.

About Symantec Operations Readiness Tools

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

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

Table 1-1 Datacenter tasks and the SORT tools

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

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

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

SORT is available at no additional charge.

To access SORT, go to:

https://sort.symantec.com

About configuring VCS clusters for data integrity

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

See " About planning to configure I/O fencing" on page 90.

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

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

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

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

See "Setting up non-SCSI-3 I/O fencing in virtual environments using installvcs" on page 177.

See "Setting up non-SCSI-3 fencing in virtual environments manually" on page 291.

About I/O fencing components

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

- Data disks—Store shared data
  See "About data disks" on page 34.
- Coordination points—Act as a global lock during membership changes
  See "About coordination points" on page 34.

About data disks

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

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

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

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. VCS 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 VCS configuration.
  You can configure coordinator disks to use Veritas Volume Manager's Dynamic Multi-pathing (DMP) feature. Dynamic Multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. So, you can configure I/O fencing to use DMP devices. I/O

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

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

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

- Coordination point servers

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

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

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

About preferred fencing

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

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

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

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

See "Enabling or disabling the preferred fencing policy" on page 181.

Chapter 2

System requirements

This chapter includes the following topics:

Release notes

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

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

https://sort.symantec.com/documents

Important preinstallation information for VCS

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

- Preinstallation checklist for your configuration. Go to the SORT installation checklist tool. From the drop-down lists, select the information for the Symantec product you want to install, and click **Generate Checklist**.
- Hardware compatibility list for information about supported hardware: http://www.symantec.com/docs/TECH211575
- For important updates regarding this release, review the Late-Breaking News Technote on the Symantec Technical Support website: http://www.symantec.com/docs/TECH211540
- VCS supports clusters with up to 64 nodes. Symantec has tested and qualified VCS configurations of up to 32 nodes at the time of the release. For more updates on this support, see the Late-Breaking News Technote.
  Every system where you want to install VCS must meet the hardware and the software requirements.

Hardware requirements for VCS

Table 2-1 lists the hardware requirements for a VCS cluster.

Table 2-1 Hardware requirements for a VCS cluster

Item	Description
VCS nodes	From 1 to 32 systems running a supported AIX operating system. **Note:** VCS is capable of supporting clusters with up to 64 nodes. Symantec has tested and qualified VCS configurations of up to 32 nodes at the time of the release. For more updates on this support, see the Late-Breaking News TechNote.
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.

Table 2-1 Hardware requirements for a VCS cluster *(continued)*

Item	Description
Disks	Typical VCS configurations require that the applications are configured to use shared disks/storage to enable migration of applications between systems in the cluster. The VCS I/O fencing feature requires that all data and coordinator disks support SCSI-3 Persistent Reservations (PR). **Note:** VCS 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 90.
Disk space	See "Disk space requirements" on page 39. **Note:** VCS may require more temporary disk space during installation than the specified disk space.
Ethernet controllers	In addition to the built-in public Ethernet controller, VCS requires at least one more Ethernet interface per system. Symantec recommends two additional interfaces. You can also configure aggregated interfaces. Symantec recommends that you turn off the spanning tree on the LLT switches, and set port-fast on.
Fibre Channel or SCSI host bus adapters	Typical VCS configuration requires at least one SCSI or Fibre Channel Host Bus Adapter per system for shared data disks.
RAM	Each VCS node requires at least 1024 megabytes.

Disk space requirements

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

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

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

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

See "About the script-based installer" on page 50.

Supported operating systems

For information on supported operating systems for various components of VCS, see the *Symantec Cluster Server Release Notes*.

Supported software for VCS

VCS supports the following volume managers and file systems:

- Logical Volume Manager (LVM)
- Journaled File System (JFS) and Enhanced Journaled File System (JFS2) on LVM

VCS supports the following versions of Symantec Storage Foundation:

Symantec Storage Foundation: Veritas Volume Manager (VxVM) with Veritas File System (VxFS)

- Storage Foundation 6.2
  - VxVM 6.2 with VxFS 6.2
- Storage Foundation 6.1
  - VxVM 6.1 with VxFS 6.1

Note: VCS supports the previous and the next versions of Storage Foundation to facilitate product upgrades.

For supported database versions of enterprise agents, refer the support matrix at http://www.symantec.com/business/support/index?page=content&id=DOC4039.

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 41.
- CP servers
  See "CP server requirements" on page 41.

To configure disk-based fencing or to configure server-based fencing with at least one coordinator disk, make sure a version of Veritas Volume Manager (VxVM) that supports SCSI-3 persistent reservations (SCSI-3 PR) is installed on the VCS cluster.

See the *Symantec Storage Foundation and High Availability Installation Guide*.

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

See "Non-SCSI-3 I/O fencing requirements" on page 45.

Coordinator disk requirements for I/O fencing

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

- For disk-based I/O fencing, you must have at least three coordinator disks or there must be odd number of coordinator disks.
- The coordinator disks must be DMP devices.
- Each of the coordinator disks must use a physically separate disk or LUN. Symantec recommends using the smallest possible LUNs for coordinator disks.
- Each of the coordinator disks should exist on a different disk array, if possible.
- The coordinator disks must support SCSI-3 persistent reservations.
- Coordinator devices can be attached over iSCSI protocol but they must be DMP devices and must support SCSI-3 persistent reservations.
- Symantec recommends using hardware-based mirroring for coordinator disks.
- Coordinator disks must not be used to store data or must not be included in disk groups that store user data.
- Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.
- The coordinator disk size must be at least 128 MB.

CP server requirements

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

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

Upgrade considerations for CP servers

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

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

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

See the *Symantec Storage Foundation High Availability Installation Guide*.

See "Hardware requirements for VCS" on page 38.

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

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

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

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

Table 2-2 CP server hardware requirements

Hardware required	Description
Disk space	To host the CP server on a VCS cluster or SFHA cluster, each host requires the following file system space: ■ 550 MB in the /opt directory (additionally, the language pack requires another 15 MB) ■ 300 MB in /usr ■ 20 MB in /var ■ 10 MB in /etc (for the CP server database) See "Disk space requirements" on page 39.

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

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

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

Table 2-3 CP server supported operating systems and versions

CP server	Operating system and version

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

Following are the CP server networking requirements and recommendations:

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

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

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

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

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

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

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

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

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

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

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

Non-SCSI-3 I/O fencing requirements

Supported virtual environment for non-SCSI-3 fencing:

- IBM P Server LPARs with VIOS running
  Guest operating system: AIX 6.1 or 7.1

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

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

Number of nodes supported

VCS supports cluster configurations with up to 64 nodes.

Checking installed product versions and downloading maintenance releases and patches

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

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

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

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

- Available maintenance releases
- Available patch releases

To check your systems and download maintenance releases and patches

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

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

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

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

   You can obtain installer patches automatically or manually.

   See "Obtaining installer patches" on page 46.

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

See "Disabling external network connection attempts" on page 47.

Obtaining installer patches

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

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

You can access installer patches automatically or manually.

To download installer patches automatically

- Starting with Symantec Cluster Server version 6.1, installer patches are downloaded automatically. No action is needed on your part.

  If you are running Symantec Cluster Server version 6.1 or later, and your system has Internet access, the installer automatically imports any needed installer patch, and begins using it.

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

See "Disabling external network connection attempts" on page 47.

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

To download installer patches manually

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

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

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

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

5. Navigate to the installation media or to the installation directory.
6. To start using the patch, run the `installer` command with the `-require` option. For example, enter the following:

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

Disabling external network connection attempts

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

To disable external network connection attempts

- Disable inter-process communication (IPC).

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

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

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

Chapter 3

Planning to install VCS

This chapter includes the following topics:

VCS installation methods

Table 3-1 lists the different methods you can choose to install and configure VCS:

Table 3-1 VCS installation methods

Method	Description
Interactive installation using the script-based installer	You can use one of the following script-based installers: ■ Product installer Use to install and configure multiple Symantec products. ■ `installvcs` program Use to install and configure just VCS. The script-based installer asks you a series of questions and installs and configures VCS based on the information you provide.
Interactive installation using the web-based installer	You can use a web-interface to install and configure VCS.

Table 3-1 VCS installation methods *(continued)*

Method	Description
Automated installation using the VCS response files	Use response files to perform unattended installations. You can generate a response file in one of the following ways: ■ Use the automatically generated response file after a successful installation. ■ Use the `-makeresponsefile` option to create a response file.
Manual installation using the AIX commands and utilities	You can install VCS using the operating system commands like `installp` and then manually configure VCS as described in the section on Manual installation. You can also install VCS using the NIM utility.

About the script-based installer

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

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

- The general product installer (`installer`). The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
  See "Installing VCS using the installer" on page 85.
- Product-specific installation scripts (`installvcs`). The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installvcs` script is identical to running the general product installer and specifying VCS from the list of products to install. Use the product-specific installation scripts to install or configure individual products you download electronically. You can find these scripts at the root of the product media. These scripts are also installed with the product.

Table 3-2 Product installation scripts

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

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

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

For example, for the 6.2 version:

```
# /opt/VRTS/install/installvcs62 -configure
```

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

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

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

See "Installation script options" on page 500.

About the VCS installation program

You can access the `installvcs` program from the command line or through the product installer.

The VCS installation program is interactive and manages the following tasks:

- Licensing VCS
- Installing VCS filesets on multiple cluster systems
- Configuring VCS, by creating several detailed configuration files on each system
- Starting VCS processes

You can choose to configure different optional features, such as the following:

- SNMP and SMTP notification
- VCS configuration in secure mode
- The wide area Global Cluster Option feature
- Cluster Virtual IP address

Review the highlights of the information for which installvcs prompts you as you proceed to configure.

See "About preparing to install VCS" on page 66.

The uninstallvcs, a companion to installvcs, uninstalls VCS filesets.

See "Preparing to uninstall VCS" on page 478.

Features of the script-based installer

The script-based installer supports installing, configuring, upgrading, and uninstalling VCS. In addition, the script-based installer also provides command options to perform the following tasks:

- Check the systems for VCS installation requirements.
  See “Performing automated preinstallation check” on page 75.
- Upgrade VCS if a previous version of VCS currently runs on a cluster.
  See “Upgrading VCS using the script-based installer” on page 344.
- Start or stop VCS processes
  See “Starting and stopping processes for the Symantec products ” on page 555.
- Enable or disable a cluster to run in secure mode
  See the *Symantec Cluster Server Administrator's Guide*.
- Configure I/O fencing for the clusters to prevent data corruption
  See “Setting up disk-based I/O fencing using installvcs” on page 154.
  See “Setting up server-based I/O fencing using installvcs” on page 164.
  See “Setting up non-SCSI-3 I/O fencing in virtual environments using installvcs” on page 177.
- Create a single-node cluster
  See “Creating a single-node cluster using the installer program” on page 526.
- Add a node to an existing cluster
  See “Adding nodes using the VCS installer” on page 449.
- Create an installp_bundle to install VCS using the NIM utility.
  See “Installing VCS using NIM and the installer” on page 247.
- Perform automated installations using the values that are stored in a configuration file.
  See “Installing VCS using response files” on page 210.
  See “Configuring VCS using response files” on page 215.
  See “Upgrading VCS using response files” on page 375.

Interacting with the installvcs

As you run the program, you are prompted to answer yes or no questions. A set of responses that resemble **[y, n, q, ?] (y)** typically follow these questions. The response within parentheses is the default, which you can select by pressing the Enter key. Enter the **?** character to get help to answer the prompt. Enter **q** to quit the installation.

Installation of VCS filesets takes place only after you have confirmed the information. However, you must remove the partially installed VCS files before you run the installvcs again.

See “Preparing to uninstall VCS” on page 478.

During the installation, the installer prompts you to type information. The installer expects your responses to be within a certain range or in a specific format. The

installer provides examples. If you are prompted to enter an item from a list, enter your selection exactly as it is shown in the list.

The installer also prompts you to answer a series of questions that are related to a configuration activity. For such questions, you can enter the **b** character to return to the first prompt in the series. When the installer displays a set of information items you have entered, you are prompted to confirm it. If you answer **n**, the program lets you reenter all of the information for the set.

You can install the VCS Java Console on a single system, which is not required to be part of the cluster. Note that the installvcs does not install the VCS Java Console.

See "Installing the Java Console" on page 422.

About the web-based installer

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

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

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

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

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

See "Before using the web-based installer" on page 185.

See "Starting the web-based installer" on page 186.

About response files

The installer generates a "response file" after performing an installer task such as installation, configuration, uninstallation, or upgrade. These response files contain the details that you provided to the installer questions in the form of values for the response file variables. The response file also contains descriptions and explanations of the variables and their values.

You can also create a response file using the `-makeresponsefile` option of the installer.

The installer displays the location of the response file at the end of each successful installer task. The installer saves the response file in the default location for the install-related log files: /opt/VRTS/install/logs. If you provided a different log path using the `-logpath` option, the installer saves the response file in the path that you specified.

The format of the response file name is:
/opt/VRTS/install/logs/*installscript-YYYYMMDDHHSSxxx*
/*installscript-YYYYMMDDHHSSxxx*.response, where:

- *installscript* may be, for example: installer, webinstaller, installvcs, or uninstallvcs
- *YYYYMMDDHHSS* is the current date when the *installscript* is run and *xxx* are three random letters that the script generates for an installation instance

For example:
/opt/VRTS/install/logs/installer-200910101010ldS/installer-200910101010ldS.response

You can customize the response file as required to perform unattended installations using the `-responsefile` option of the installer. This method of automated installations is useful in the following cases:

- To perform multiple installations to set up a large VCS cluster.
  See "Installing VCS using response files" on page 210.
- To upgrade VCS on multiple systems in a large VCS cluster.
  See "Upgrading VCS using response files" on page 375.
- To uninstall VCS from multiple systems in a large VCS cluster.
  See "Uninstalling VCS using response files" on page 484.

Syntax in the response file

The syntax of the Perl statements that is included in the response file variables varies. It can depend on whether the variables require scalar or list values.

For example, in the case of a string value:

```
$CFG{Scalar_variable}="value";
```

or, in the case of an integer value:

```
$CFG{Scalar_variable}=123;
```

or, in the case of a list:

```
$CFG{List_variable}=["value 1 ", "value 2 ", "value 3 "];
```

About installation and configuration methods

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

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

Table 3-3 Installation and configuration methods

Method	Description

The script-based installer	Using the script-based installer, you can install Symantec products from a driver system running a supported platform to target computers running any supported platform.
	
	To install your Symantec product using the installer, choose one of the following:
	
	■ The general product installer:
	`installer`
	The general product installer script provides a menu that simplifies the selection of installation and configuration options. Use the general product installer if you want to install multiple products from a disc.
	
	■ Product-specific installation scripts:
	`installvcs<version>`
	The product-specific installation scripts provide command-line interface options. Installing and configuring with the `installvcs` script is identical to running the general product installer and specifying VCS from the list of products to install.
	Use the product-specific installation scripts to install or configure individual products you download electronically.
	
	See "About the script-based installer" on page 50.

The web-based installer	Using the web-based installer, you can install Symantec products from a driver system running a supported platform to target computers running any supported platform The web-based installer provides an interface to manage the installation and configuration from a remote site using a standard web browser. `webinstaller` See "About the web-based installer" on page 54.

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

Method	Description
Deployment Server	Using the Deployment Server, you can store multiple release images in one central location and deploy them to systems of any supported platform. See "About the Deployment Server" on page 303.
Silent installation using response files	Response files automate installation and configuration by using the information that is stored in a specified file instead of prompting you for information. You can use any of the above options to generate a response file. You can then customize the response file for another system. Run the product installation script with the response file option to install silently on one or more systems. See "Installing VCS using response files" on page 210.
Install Bundles	Beginning with version 6.1, you can easily install or upgrade your systems directly to a base, maintenance, or patch level in one step using Install Bundles. The installer installs both releases as if they were combined in the same release image. The various scripts, filesets, and patch components are merged, and multiple releases are installed together as if they are one combined release. See "Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches" on page 341.
Network Installation Manager (NIM)	You can perform many advanced NIM installation tasks using the NIM command interface and the System Management Interface Tool (SMIT). Use the product installer or the product-specific installation script to generate a NIM `installp` bundle. Use the generated `installp` bundle to install Symantec filesets from your NIM server.
mksysb utility	You can use the mksysb utility to back up the system image. This image can be installed on another host.

Typical VCS cluster setup models

VCS clusters support different failover configurations, storage configurations, and cluster topologies.

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

Some of the typical VCS setup models are as follows:

- Basic VCS cluster with two nodes
  See "Typical configuration of two-node VCS cluster" on page 58.
- VCS clusters in secure mode
  See "Typical configuration of VCS clusters in secure mode" on page 58.
- VCS clusters centrally managed using Veritas Operations Manager (VOM)
  See "Typical configuration of VOM-managed VCS clusters" on page 59.
- VCS clusters with I/O fencing for data protection
  See "Typical VCS cluster configuration with disk-based I/O fencing" on page 94.
  See "Typical VCS cluster configuration with server-based I/O fencing" on page 95.
- VCS clusters such as global clusters, replicated data clusters, or campus clusters for disaster recovery
  See the *Symantec Cluster Server Administrator's Guide* for disaster recovery cluster configuration models.

Typical configuration of two-node VCS cluster

Figure 3-1 illustrates a simple VCS cluster setup with two nodes.

Figure 3-1 Typical two-node VCS cluster

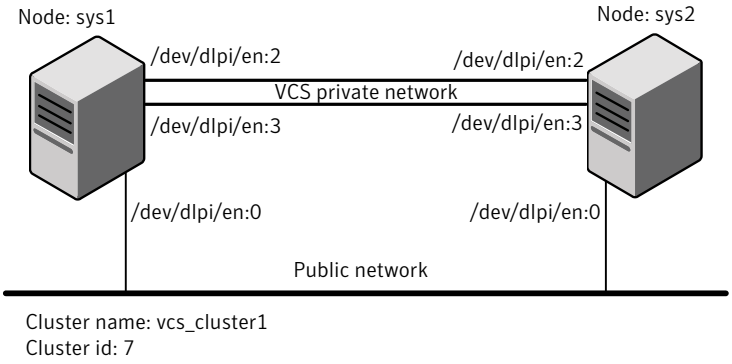

Cluster name: vcs_cluster1
Cluster id: 7

Typical configuration of VCS clusters in secure mode

Enabling secure mode for VCS guarantees that all inter-system communication is encrypted and that security credentials of users are verified.

Figure 3-2 illustrates typical configuration of VCS clusters in secure mode.

Figure 3-2 Typical configuration of VCS clusters in secure mode

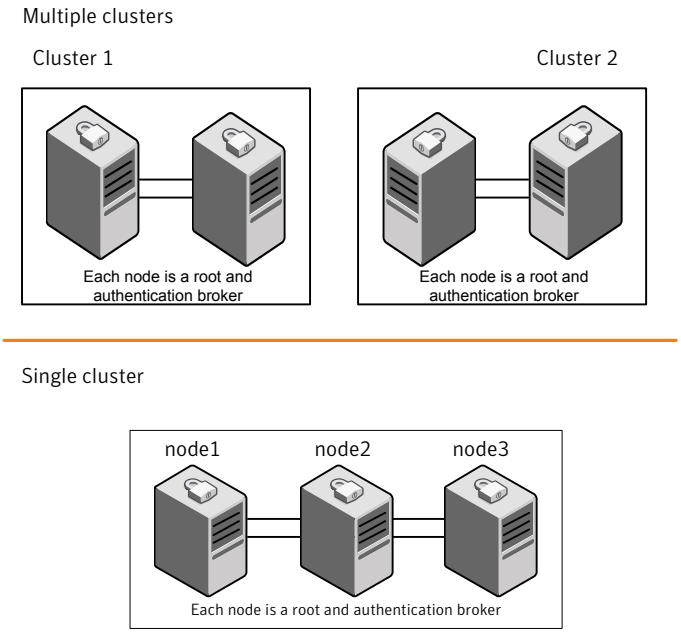

Typical configuration of VOM-managed VCS clusters

Veritas Operations Manager (VOM) provides a centralized management console for Symantec Storage Foundation and High Availability products.

See "About Veritas Operations Manager" on page 29.

Figure 3-3 illustrates a typical setup of VCS clusters that are centrally managed using Veritas Operations Manager.

Figure 3-3 Typical configuration of VOM-managed clusters

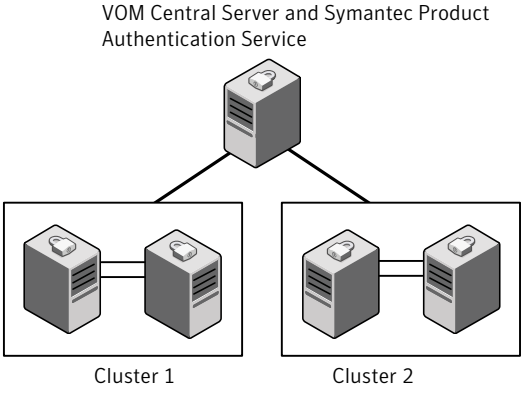

Chapter 4

Licensing VCS

This chapter includes the following topics:

About Symantec product licensing

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

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

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

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

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

Within 60 days of choosing this option, you must install a valid license key corresponding to the license level entitled, or continue with keyless licensing by managing the systems with a management server. If you do not comply with the above terms, continuing to use the Symantec product is a violation of your End User License Agreement, and results in warning messages

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

http://go.symantec.com/sfhakeyless

If you upgrade to this release from a previous release of the Symantec software, the installer asks whether you want to upgrade the key to the new version. The existing license keys may not activate new features in this release.

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

- Run the `vxkeyless` command to set the product level for the products you have purchased. This option also requires that you manage the server or cluster with a management server.
  See "Setting or changing the product level for keyless licensing" on page 244.
  See the `vxkeyless(1m)` manual page.
- Use the `vxlicinst` command to install a valid product license key for the products you have purchased.
  See "Installing Symantec product license keys" on page 63.
  See the `vxlicinst(1m)` manual page.

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

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

Obtaining VCS license keys

This product includes a License Key certificate. The certificate specifies the product keys and the number of product licenses purchased. A single key lets you install the product on the number and type of systems for which you purchased the license. A key may enable the operation of more products than are specified on the certificate. However, you are legally limited to the number of product licenses purchased. The product installation procedure describes how to activate the key.

To register and receive a software license key, go to the Symantec Licensing Portal at the following location:

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

Make sure you have your Software Product License document. You need information in this document to retrieve and manage license keys for your Symantec product. After you receive the license key, you can install the product.

Click the **Get Help** link at this site for contact information and for useful links.

The VRTSvlic fileset enables product licensing. For information about the commands that you can use after the installing VRTSvlic:

See "Installing Symantec product license keys" on page 63.

You can only install the Symantec software products for which you have purchased a license. The enclosed software discs might include other products for which you have not purchased a license.

Installing Symantec product license keys

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

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

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

To install or change a license

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

   ```
   # cd /opt/VRTS/bin

   # ./vxlicinst -k license key
   ```

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

   ```
   # vxdctl license init
   ```

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

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

See "Setting or changing the product level for keyless licensing" on page 244.

Section 2

Preinstallation tasks

Chapter 5

Preparing to install VCS

This chapter includes the following topics:

About preparing to install VCS

Before you perform the preinstallation tasks, make sure you reviewed the installation requirements, set up the basic hardware, and planned your VCS setup.

Performing preinstallation tasks

Table 5-1 lists the tasks you must perform before proceeding to install VCS.

Table 5-1 Preinstallation tasks

Task	Reference
Obtain license keys if you do not want to use keyless licensing.	See "Obtaining VCS license keys" on page 62.
Set up the private network.	See "Setting up the private network" on page 67.
Enable communication between systems.	See "About configuring secure shell or remote shell communication modes before installing products" on page 543.
Set up ssh on cluster systems.	See "Manually configuring passwordless ssh" on page 544.

Table 5-1 Preinstallation tasks *(continued)*

Task	Reference
Set up shared storage for I/O fencing (optional)	See "Setting up shared storage" on page 70.
Creating root user	
Set the PATH and the MANPATH variables.	See "Setting the PATH variable" on page 72. See "Setting the MANPATH variable" on page 73.
Review basic instructions to optimize LLT media speeds.	See "Optimizing LLT media speed settings on private NICs" on page 73.
Review guidelines to help you set the LLT interconnects.	See "Guidelines for setting the media speed of the LLT interconnects" on page 73.
Mount the product disc	See "Mounting the product disc" on page 74.
Verify the systems before installation	See "Performing automated preinstallation check" on page 75.

Setting up the private network

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

You can use network switches instead of hubs.

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

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

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

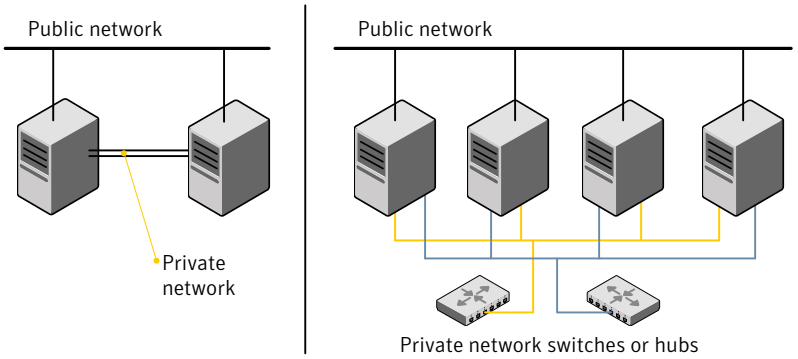

You need to configure at least two independent networks between the cluster nodes with a network switch for each network. You can also interconnect multiple layer 2 switches for advanced failure protection. Such connections for LLT are called cross-links.

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

Figure 5-2 Private network setup with crossed links

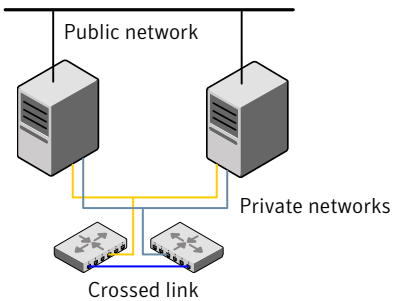

Symantec recommends one of the following two configurations:

- Use at least two private interconnect links and one public link. The public link can be a low priority link for LLT. The private interconnect link is used to share cluster status across all the systems, which is important for membership arbitration and high availability. The public low priority link is used only for heartbeat communication between the systems.
- If your hardware environment allows use of only two links, use one private interconnect link and one public low priority link. If you decide to set up only two links (one private and one low priority link), then the cluster must be configured

to use I/O fencing, either disk-based or server-based fencing configuration. With only two links, if one system goes down, I/O fencing ensures that other system can take over the service groups and shared file systems from the failed node.

To set up the private network

1 Install the required network interface cards (NICs).

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

2 Connect the VCS private Ethernet controllers on each system.

3 Use crossover Ethernet cables, switches, or independent hubs for each VCS communication network. Note that the crossover Ethernet cables are supported only on two systems.

Ensure that you meet the following requirements:

- The power to the switches or hubs must come from separate sources.
- On each system, you must use two independent network cards to provide redundancy.
- If a network interface is part of an aggregated interface, you must not configure the network interface under LLT. However, you can configure the aggregated interface under LLT.
- When you configure Ethernet switches for LLT private interconnect, disable the spanning tree algorithm on the ports used for the interconnect.

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

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

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

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

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

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

You can also manually configure LLT.

See “Configuring LLT manually” on page 251.

About using ssh or rsh with the installer

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

See “Installation script options” on page 500.

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

- When you perform installer sessions using a response file.

See “About configuring secure shell or remote shell communication modes before installing products” on page 543.

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: VCS 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 90.

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

Setting the SCSI identifier value

SCSI adapters are typically set with a default identifier value of 7. Each device on a SCSI bus must have a unique SCSI identifier value. When more than one system

is connected to a SCSI bus, you must change the SCSI identifier to a unique number. You must make this change to one or more systems, usually the unique number is 5 or 6.

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

Figure 5-3 Cabling the shared storage

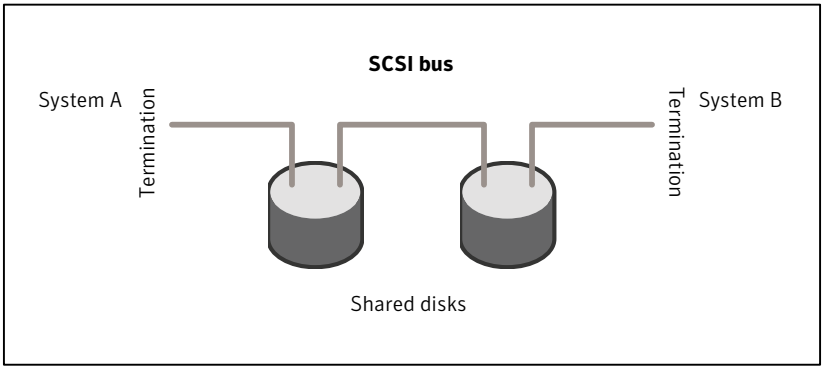

To set the SCSI identifier value

1 Determine the SCSI adapters on each system:

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

2 Verify the SCSI ID of each adapter:

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

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

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

4 Shut down all systems in the cluster.

5 Cable the shared storage as illustrated in Figure 5-3.

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

Setting up Fibre Channel

Perform the following steps to set up Fibre Channel.

To set up Fibre Channel

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

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

2 Reboot each system:

```
# shutdown -Fr
```

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

Setting the PATH variable

To set the PATH variable

- Do one of the following:
  - For the Bourne Shell (sh), Bourne-again Shell (bash), or Korn shell (ksh), type:

    ```
    # PATH=/opt/VRTS/bin:$PATH; export $PATH
    ```

  - For the C Shell (csh) or enhanced C Shell (tcsh), type:

    ```
    # setenv PATH :/opt/VRTS/bin:$PATH
    ```

Setting the MANPATH variable

Set the MANPATH variable to view the manual pages.

To set the MANPATH variable

- Do one of the following:
  - For the Bourne Shell (sh), Bourne-again Shell (bash), or Korn shell (ksh), type:

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

  - For the C Shell (csh) or enhanced C Shell (tcsh), type:

    ```
    % setenv MANPATH /usr/share/man:/opt/VRTS/man
    ```

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.

VCS considerations for Blade server environments

Typically, a server in the Blade environment has only two NICs. The following considerations need to be observed while configuring VCS in Blade environment in private networks:

- If your heartbeat links do not use TCP/IP and are not routable, you must use separate and dedicated physical networks. This will guard against inadvertent split brains due to inappropriate routing configurations.
- Out of the two heartbeat links, one must be dedicated and the other can be a low-priority heartbeat shared on the public IP NIC. It is assumed that the two nodes in the cluster have public IPs on the same subnet and wire.
- The size each packet of traffic on the public NIC must be 64 bytes/second and must not interfere with the public traffic.

Mounting the product disc

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

To mount the product disc

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

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

2. Determine the device access name of the disc drive. For example, enter:

   ```
   # lsdev -C -c cdrom
   ```

   The output resembles:

   ```
   cd0 Available 1G-19-00 IDE DVD-ROM Drive
   ```

   In this example, `cd0` is the disc's device access name.

3 Make sure that the `/cdrom` file system is created:

```
# cat /etc/filesystems
```

If the `/cdrom` file system exists, the output contains a listing that resembles:

```
.
.
/cdrom:
dev = /dev/cd0
vfs = cdrfs
mount = false
options = ro
account = false
.
.
```

4 If the `/cdrom` file system does not exist, create it:

```
# crfs -v cdrfs -p ro -d cd0 -m /cdrom
```

5 Insert the product disc with the VCS software into a drive that is connected to the system.

6 Mount the disc:

```
# mount /cdrom
# cd /cdrom
```

Performing automated preinstallation check

Before you begin the installation of VCS software, you can check the readiness of the systems where you plan to install VCS. The command to start the preinstallation check is:

```
installvcs -precheck system1 system2 ...
```

You can also run the `installer -precheck` command.

See "About Symantec Operations Readiness Tools" on page 30.

To check the systems

1 Navigate to the folder that contains the installvcs.

```
# cd /cdrom/cluster_server
```

2 Start the preinstallation check:

```
# ./installvcs -precheck sys1 sys2
```

The program proceeds in a noninteractive mode to examine the systems for licenses, filesets, disk space, and system-to-system communications.

3 Review the output as the program displays the results of the check and saves the results of the check in a log file.

Reformatting VCS configuration files on a stopped cluster

When you manually edit VCS configuration files (for example, the main.cf or types.cf file) you can potentially create formatting issues that may cause the installer to interpret the cluster configuration information incorrectly.

If you have manually edited any of the configuration files, you need to perform one of the following before you run the installation program:

- On a running cluster, perform an `haconf -dump` command. This command saves the configuration files and ensures that they do not have formatting errors before you run the installer.
- On cluster that is not running, perform the `hacf -cftocmd` and then the `hacf -cmdtocf` commands to format the configuration files.

Note: Remember to make back up copies of the configuration files before you edit them.

You also need to use this procedure if you have manually changed the configuration files before you perform the following actions using the installer:

- Upgrade VCS
- Uninstall VCS

For more information about the main.cf and types.cf files, refer to the *Symantec Cluster Server Administrator's Guide*.

To display the configuration files in the correct format on a running cluster

- Run the following commands to display the configuration files in the correct format:

```
# haconf -dump
```

To display the configuration files in the correct format on a stopped cluster

- Run the following commands to display the configuration files in the correct format:

```
# hacf -cftocmd config

# hacf -cmdtocf config
```

Getting your VCS installation and configuration information ready

The VCS installer prompts you for some information during the installation and configuration process. Review the following information and make sure you have made the necessary decisions and you have the required information ready before you perform the installation and configuration.

Table 5-2 lists the information you need to install the VCS filesets.

Table 5-2 Information to install the VCS filesets

Information	Description and sample value	Your value
System names	The system names where you plan to install VCS Example: **sys1**, **sys2**	
The required license keys	If you decide to use keyless licensing, you do not need to obtain license keys. However, you require to set up management server within 60 days to manage the cluster. See "About Symantec product licensing" on page 61. Depending on the type of installation, keys can include: ■ A valid site license key ■ A valid demo license key ■ A valid license key for VCS global clusters See "Obtaining VCS license keys" on page 62.	

Table 5-2 Information to install the VCS filesets *(continued)*

Information	Description and sample value	Your value
Decide which filesets to install	■ Minimum filesets—provides basic VCS functionality. ■ Recommended filesets—provides full functionality of VCS without advanced features. ■ All filesets—provides advanced feature functionality of VCS. The default option is to install the recommended filesets. See "Viewing the list of VCS filesets" on page 242.	

Table 5-3 lists the information you need to configure VCS cluster name and ID.

Table 5-3 Information you need to configure VCS cluster name and ID

Information	Description and sample value	Your value
A name for the cluster	The cluster name must begin with a letter of the alphabet. The cluster name can contain only the characters "a" through "z", "A" through "Z", the numbers "0" through "9", the hyphen "-", and the underscore "_". Example: **my_cluster**	
A unique ID number for the cluster	A number in the range of 0-65535. If multiple distinct and separate clusters share the same network, then each cluster must have a unique cluster ID. Example: **12133**	

Table 5-4 lists the information you need to configure VCS private heartbeat links.

Table 5-4 Information you need to configure VCS private heartbeat links

Information	Description and sample value	Your value
Decide how you want to configure LLT	You can configure LLT over Ethernet or LLT over UDP. The LLT over Ethernet is the typical configuration for LLT. If the cluster nodes are across routers, then use the LLT over UDP configuration after ensuring that you meet all the prerequisites. See "Using the UDP layer for LLT" on page 529.	

Table 5-4 Information you need to configure VCS private heartbeat links *(continued)*

Information	Description and sample value	Your value
Decide which configuration mode you want to choose	Installer provides you with four options: **1** Configure heartbeat links using LLT over Ethernet **2** Configure heartbeat links using LLT over UDP **3** Automatically detect configuration for LLT over Ethernet You must manually enter details for options 1, 2 whereas the installer detects the details for option 3.	
For option 1: LLT over Ethernet	■ The device names of the NICs that the private networks use among systems A network interface card or an aggregated interface. Do not use the network interface card that is used for the public network, which is typically en0 en1. Example: **en2**, **en3** ■ Choose whether to use the same NICs on all systems. If you want to use different NICs, enter the details for each system.	
For option 2: LLT over UDP	For each system, you must have the following details: ■ The device names of the NICs that the private networks use among systems ■ IP address for each NIC ■ UDP port details for each NIC	

Table 5-5 lists the information you need to configure virtual IP address of the cluster (optional).

Table 5-5 Information you need to configure virtual IP address

Information	Description and sample value	Your value
The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. A network interface card or an aggregated interface. Example: en0 en1	

Table 5-5 Information you need to configure virtual IP address *(continued)*

Information	Description and sample value	Your value
A virtual IP address of the NIC	You can enter either an IPv4 or an IPv6 address. This virtual IP address becomes a resource for use by the ClusterService group. The "Cluster Virtual IP address" can fail over to another cluster system. Example IPv4 address: 192.168.1.16 Example IPv6 address: 2001:454e:205a:110:203:baff:feee:10	
The netmask for the virtual IPv4 address	The subnet that you use with the virtual IPv4 address. Example: 255.255.240.0	
The prefix for the virtual IPv6 address	The prefix length for the virtual IPv6 address. Example: 64	
The NetworkHosts IP addresses	IP addresses that are used to check the adapter connections. Example: 192.168.1.17	

Table 5-6 lists the information you need to add VCS users.

Table 5-6 Information you need to add VCS users

Information	Description and sample value	Your value
User names	VCS usernames are restricted to 1024 characters. Example: **smith**	
User passwords	VCS passwords are restricted to 255 characters. Enter the password at the prompt. **Note:** VCS leverages native authentication in secure mode. Therefore, user passwords are not needed in secure mode.	
To decide user privileges	Users have three levels of privileges: Administrator, Operator, or Guest. Example: Administrator	

Table 5-7 lists the information you need to configure SMTP email notification (optional).

Table 5-7 Information you need to configure SMTP email notification (optional)

Information	Description and sample value	Your value
The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. A network interface card or an aggregated interface. Examples: **en0 en1**	
The domain-based address of the SMTP server	The SMTP server sends notification emails about the events within the cluster. Example: **smtp.symantecexample.com**	
The email address of each SMTP recipient to be notified	Example: **john@symantecexample.com**	
To decide the minimum severity of events for SMTP email notification	Events have four levels of severity, and the severity levels are cumulative: ■ Information VCS sends notifications for important events that exhibit normal behavior. ■ Warning VCS sends notifications for events that exhibit any deviation from normal behavior. Notifications include both Warning and Information type of events. ■ Error VCS sends notifications for faulty behavior. Notifications include both Error, Warning, and Information type of events. ■ Critical VCS sends notifications for a critical error that can lead to data loss or corruption. Notifications include both Severe Error, Error, Warning, and Information type of events. Example: **Error**	

Table 5-8 lists the information you need to configure SNMP trap notification (optional).

Table 5-8 Information you need to configure SNMP trap notification (optional)

Information	Description and sample value	Your value
The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. A network interface card or an aggregated interface. Examples: **en0 en1**	

Table 5-8 Information you need to configure SNMP trap notification (optional) *(continued)*

Information	Description and sample value	Your value
The port number for the SNMP trap daemon	The default port number is 162.	
The system name for each SNMP console	Example: **sys5**	
To decide the minimum severity of events for SNMP trap notification	Events have four levels of severity, and the severity levels are cumulative: ■ Information VCS sends notifications for important events that exhibit normal behavior. ■ Warning VCS sends notifications for events that exhibit any deviation from normal behavior. Notifications include both Warning and Information type of events. ■ Error VCS sends notifications for faulty behavior. Notifications include both Error, Warning, and Information type of events. ■ Critical VCS sends notifications for a critical error that can lead to data loss or corruption. Notifications include both Severe Error, Error, Warning, and Information type of events. Example: **Error**	

Table 5-9 lists the information you need to configure global clusters (optional).

Table 5-9 Information you need to configure global clusters (optional)

Information	Description and sample value	Your value
The name of the public NIC	You can use the same NIC that you used to configure the virtual IP of the cluster. Otherwise, specify appropriate values for the NIC. A network interface card or an aggregated interface. Examples: **en0 en1**	

Table 5-9 Information you need to configure global clusters (optional) *(continued)*

Information	Description and sample value	Your value
The virtual IP address of the NIC	You can enter either an IPv4 or an IPv6 address. You can use the same virtual IP address that you configured earlier for the cluster. Otherwise, specify appropriate values for the virtual IP address. Example IPv4 address: **192.168.1.16** Example IPv6 address: **2001:454e:205a:110:203:baff:feee:10**	
The netmask for the virtual IPv4 address	You can use the same netmask that you used to configure the virtual IP of the cluster. Otherwise, specify appropriate values for the netmask. Example: **255.255.240.0**	
The prefix for the virtual IPv6 address	The prefix length for the virtual IPv6 address. Example: **64**	
The NetworkHosts IP addresses	You can use the same NetworkHosts IP address that you used to configure the virtual IP of the cluster. Otherwise, specify appropriate values for the NetworkHosts IP address when you are prompted. Example: **192.168.1.15**	

Review the information you need to configure I/O fencing.

See " About planning to configure I/O fencing" on page 90.

Section 3

Installation using the script-based installer

Chapter 6

Installing VCS

This chapter includes the following topics:

Installing VCS using the installer

Perform the following steps to install VCS.

To install VCS

1. Confirm that you are logged in as the superuser and you mounted the product disc.

   See "Mounting the product disc" on page 74.

2. Start the installation program. If you obtained VCS from an electronic download site, which does not include the product installer, use the installvcs.

   Product installer — Perform the following steps to start the product installer:

   1. Start the installer.

      ```
      # ./installer
      ```

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

   2. From the opening Selection Menu, choose I for "Install a Product."

   3. From the displayed list of products to install, choose: Symantec Cluster Server.

installvcs program

Perform the following steps to start the product installer:

1. Navigate to the folder that contains the installvcs.

   ```
   # cd dvd_mount/cluster_server
   ```

2. Start the installvcs.

   ```
   # ./installvcs
   ```

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

3. Enter **y** to agree to the End User License Agreement (EULA).

   ```
   Do you agree with the terms of the End User License Agreement
   as specified in the cluster_server/EULA/<lang>/EULA_VCS_Ux_6.2.pdf
   file present on media? [y,n,q,?] y
   ```

4. Choose the VCS filesets that you want to install.

   See "Symantec Cluster Server installation filesets" on page 495.

   Based on what filesets you want to install, enter one of the following:

   1 Installs only the minimal required VCS filesets that provides basic functionality of the product.

   2 Installs the recommended VCS filesets that provides complete functionality of the product. This option does not install the optional VCS filesets.

   Note that this option is the default.

   3 Installs all the VCS filesets.

   You must choose this option to configure any optional VCS feature.

   4 Displays the VCS filesets for each option.

   ```
   Select the filesets to be installed on all systems? [1-4,q,?]
   (2) 3
   ```

5. Enter the names of the systems where you want to install VCS.

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

For a single-node VCS installation, enter one name for the system.

See “Creating a single-node cluster using the installer program” on page 526.

The installer does the following for the systems:

- Checks that the local system that runs the installer can communicate with remote systems.
  If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases. If password-less communication is not present, you can provide the location of `ssh_key` file which isused in every communication.
  If the default communication method ssh fails, the installer attempts to use rsh.
- Makes sure the systems use one of the supported operating systems.
- Makes sure that the systems have the required operating system patches.
  If the installer reports that any of the patches are not available, install the patches on the system before proceeding with the VCS installation.
- Checks for product licenses.
- Checks whether a previous version of VCS is installed.
  If a previous version of VCS is installed , the installer provides an option to upgrade to VCS 6.2.
  See “About upgrading to VCS 6.2” on page 335.
- Checks for the required file system space and makes sure that any processes that are running do not conflict with the installation.
  If requirements for installation are not met, the installer stops and indicates the actions that you must perform to proceed with the process.
- Checks whether any of the filesets already exists on a system.
  If the current version of any fileset exists, the installer removes the fileset from the installation list for the system. If a previous version of any fileset exists, the installer replaces the fileset with the current version.

6 Review the list of filesets and patches that the installer would install on each node.

The installer installs the VCS filesets and patches on the systems sys1 and sys2.

7 Select the license type.

```
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)
```

Based on what license type you want to use, enter one of the following:

1 You must have a valid license key. Enter the license key at the prompt:

```
Enter a VCS license key: [b,q,?]
XXXX-XXXX-XXXX-XXXX-XXXX
```

If you plan to configure global clusters, enter the corresponding license keys when the installer prompts for additional licenses.

```
Do you wish to enter additional licenses? [y,n,q,b] (n) y
```

2 The keyless license option enables you to install VCS without entering a key. However, to ensure compliance, keyless licensing requires that you manage the systems with a management server.

For more information, go to the following website:

http://go.symantec.com/sfhakeyless

Note that this option is the default.

The installer registers the license and completes the installation process.

8 To install the Global Cluster Option, enter y at the prompt.

9 To configure VCS, enter y at the prompt. You can also configure VCS later.

```
Would you like to configure VCS on sys1 sys2 [y,n,q] (n) n
```

See "Overview of tasks to configure VCS using the script-based installer" on page 130.

10 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
```

The installer provides an option to collect data about the installation process each time you complete an installation, upgrade, configuration, or uninstall of the product. The installer transfers the contents of the install log files to an internal Symantec site. The information is used only to gather metrics about how you use the installer. No personal customer data is collected, and no information will be shared by any other parties. Information gathered may include the product and the version installed or upgraded, how many systems were installed, and the time spent in any section of the install process.

11 The installer checks for online updates and provides an installation summary.

12 After the installation, note the location of the installation log files, the summary file, and the response file for future reference.

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

summary file	Lists the filesets that are installed on each system.
log file	Details the entire installation.
response file	Contains the installation information that can be used to perform unattended or automated installations on other systems. See "Installing VCS using response files" on page 210.

Chapter 7

Preparing to configure VCS clusters for data integrity

This chapter includes the following topics:

About planning to configure I/O fencing

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

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

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

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

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

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

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

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

See Figure 7-2 on page 93.

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

Figure 7-1 Workflow to configure I/O fencing

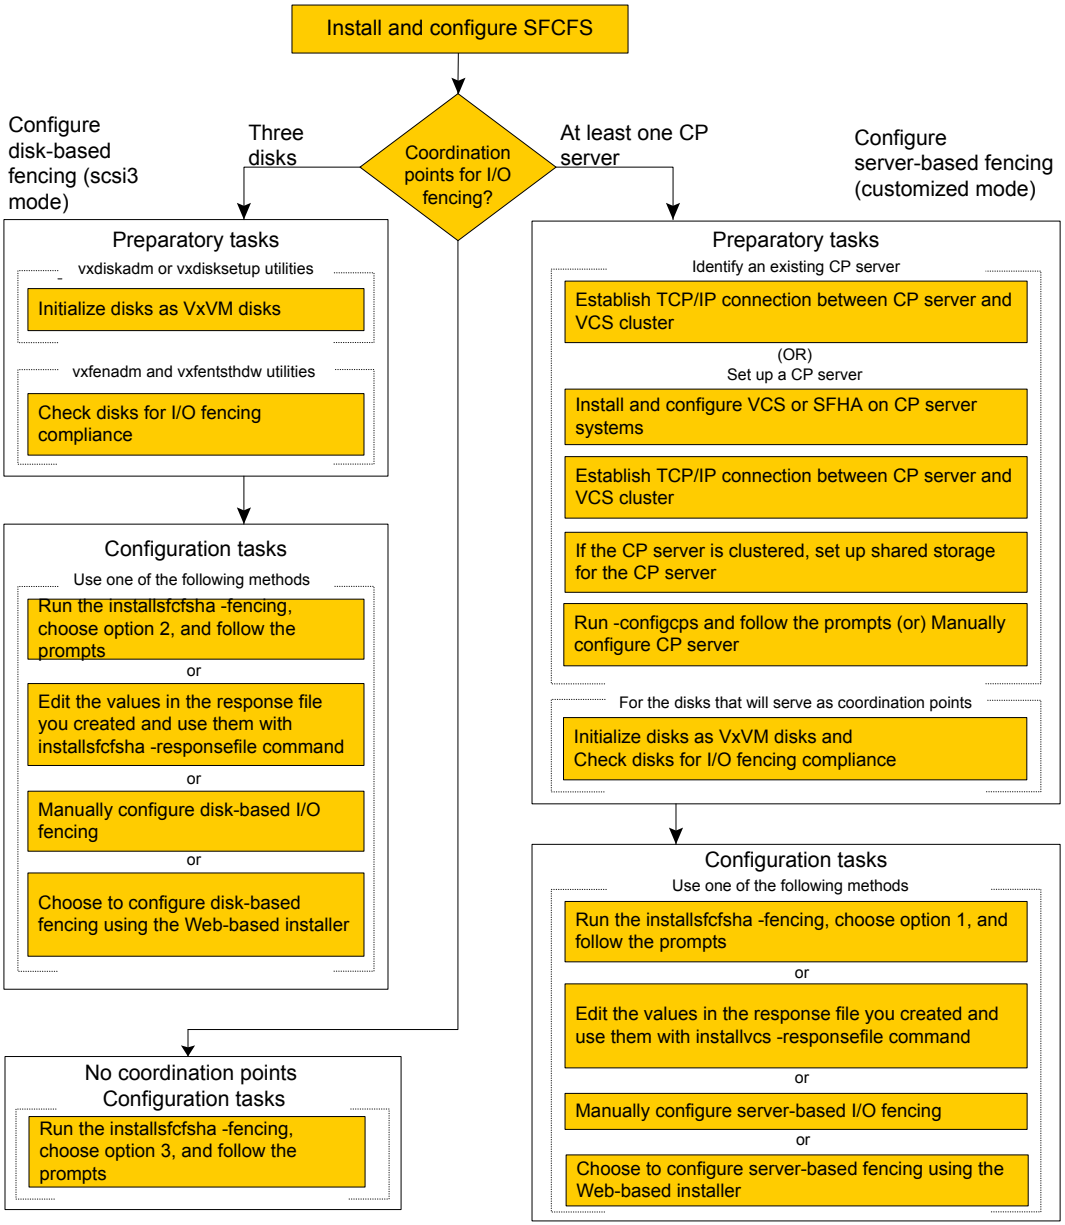

Figure 7-2 illustrates a high-level flowchart to configure non-SCSI-3 I/O fencing for the VCS cluster in virtual environments that do not support SCSI-3 PR.

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

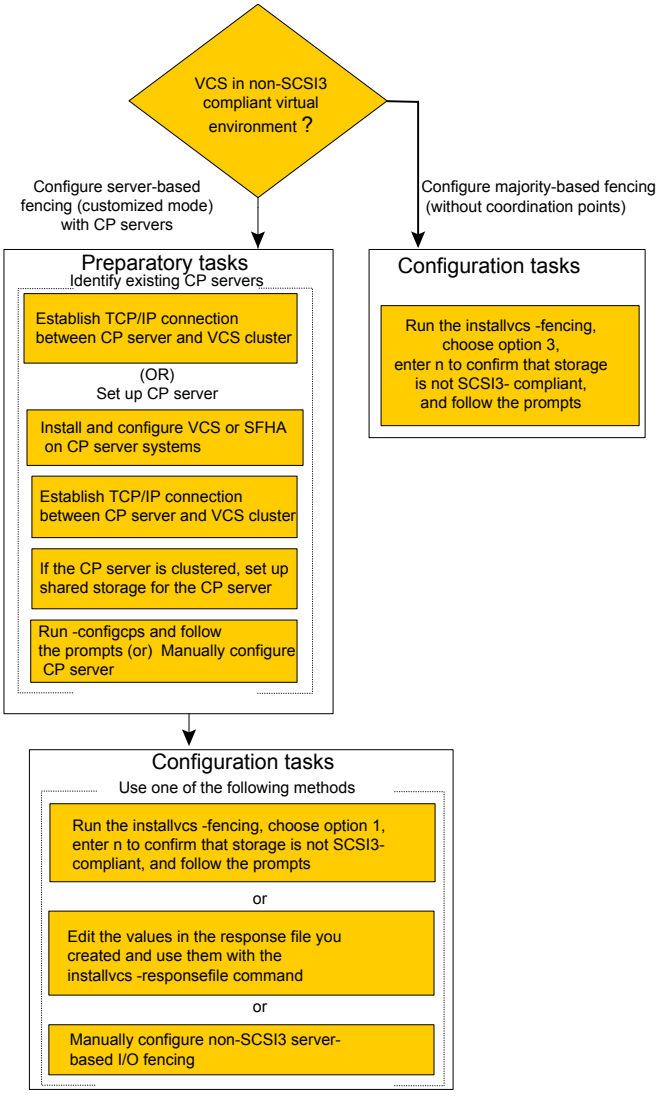

After you perform the preparatory tasks, you can use any of the following methods to configure I/O fencing:

Using the installvcs	See "Setting up disk-based I/O fencing using installvcs" on page 154. See "Setting up server-based I/O fencing using installvcs" on page 164. See "Setting up non-SCSI-3 I/O fencing in virtual environments using installvcs" on page 177. See "Setting up majority-based I/O fencing using installvcs" on page 179.
Using the web-based installer	See "Configuring VCS for data integrity using the web-based installer" on page 196.
Using response files	See "Response file variables to configure disk-based I/O fencing" on page 228. See "Response file variables to configure server-based I/O fencing" on page 231. See "Response file variables to configure non-SCSI-3 I/O fencing" on page 234. See "Response file variables to configure majority-based I/O fencing" on page 236. See "Configuring I/O fencing using response files" on page 227.
Manually editing configuration files	See "Setting up disk-based I/O fencing manually" on page 271. See "Setting up server-based I/O fencing manually" on page 276. See "Setting up non-SCSI-3 fencing in virtual environments manually" on page 291. See "Setting up majority-based I/O fencing manually " on page 297.

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

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

Typical VCS cluster configuration with disk-based I/O fencing

Figure 7-3 displays a typical VCS configuration with two nodes and shared storage. The configuration uses three coordinator disks for I/O fencing.

Figure 7-3 Typical VCS cluster configuration with disk-based I/O fencing

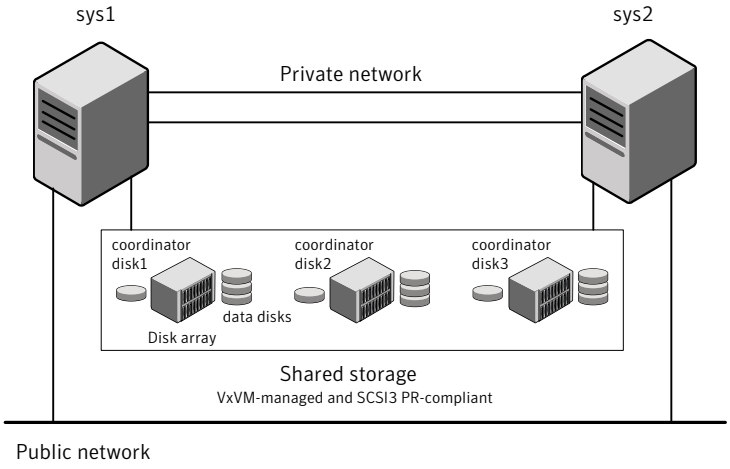

Typical VCS cluster configuration with server-based I/O fencing

Figure 7-4 displays a configuration using a VCS cluster (with two nodes), a single CP server, and two coordinator disks. The nodes within the VCS cluster are connected to and communicate with each other using LLT links.

Figure 7-4 CP server, VCS cluster, and coordinator disks

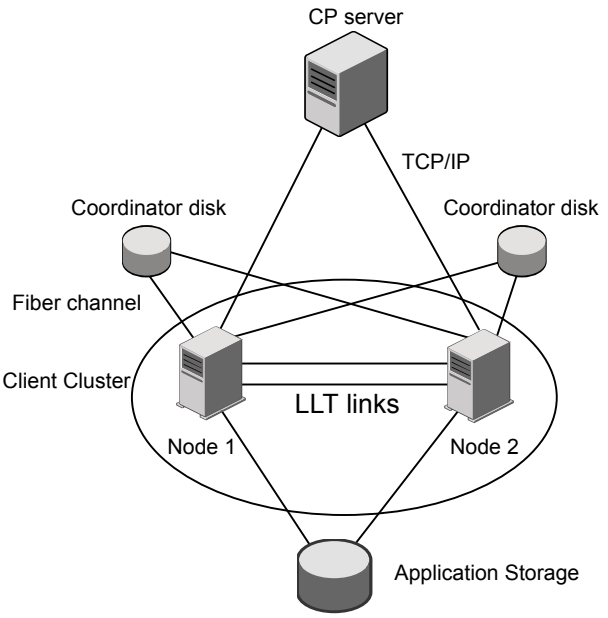

Recommended CP server configurations

Following are the recommended CP server configurations:

- Multiple application clusters use three CP servers as their coordination points
  See Figure 7-5 on page 97.
- Multiple application clusters use a single CP server and single or multiple pairs of coordinator disks (two) as their coordination points
  See Figure 7-6 on page 98.
- Multiple application clusters use a single CP server as their coordination point
  This single coordination point fencing configuration must use a highly available CP server that is configured on an SFHA cluster as its coordination point.
  See Figure 7-7 on page 98.

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

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

Figure 7-5 displays a configuration using three CP servers that are connected to multiple application clusters.

Figure 7-5 Three CP servers connecting to multiple application clusters

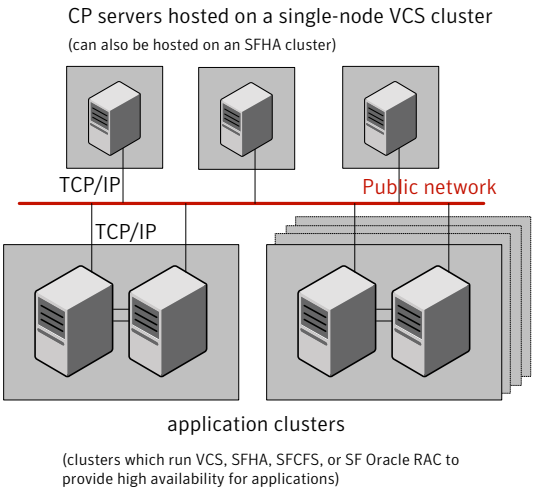

Figure 7-6 displays a configuration using a single CP server that is connected to multiple application clusters with each application cluster also using two coordinator disks.

Figure 7-6 Single CP server with two coordinator disks for each application cluster

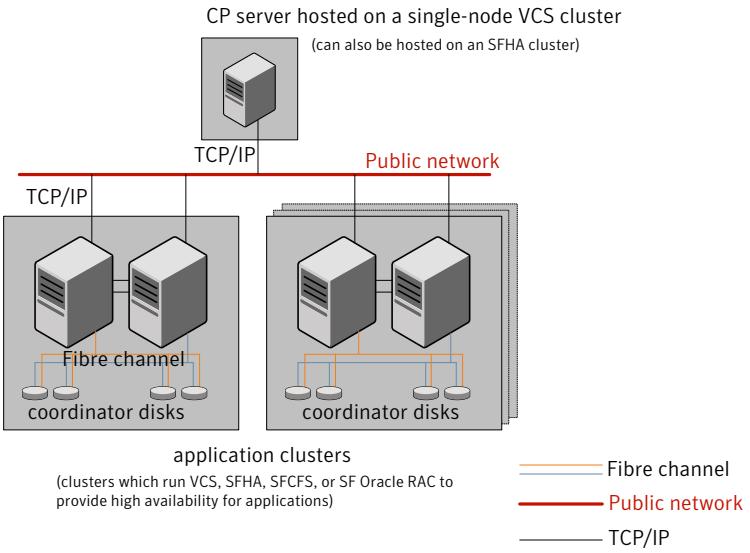

Figure 7-7 displays a configuration using a single CP server that is connected to multiple application clusters.

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

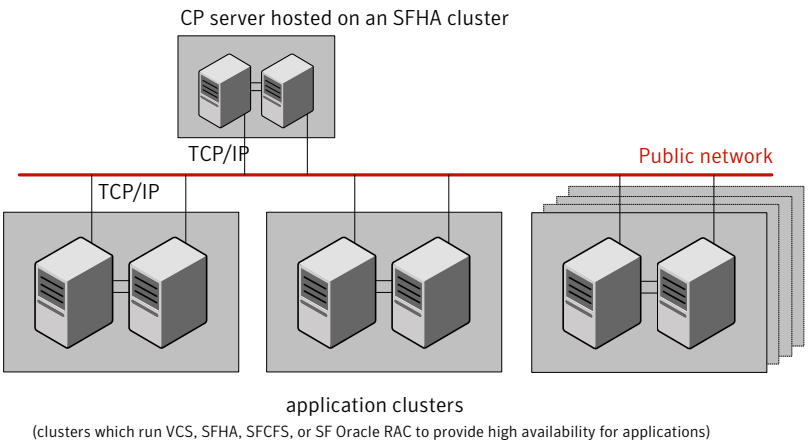

See "Configuration diagrams for setting up server-based I/O fencing" on page 560.

Setting up the CP server

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

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

Task	Reference
Plan your CP server setup	See "Planning your CP server setup" on page 99.
Install the CP server	See "Installing the CP server using the installer" on page 101.
Configure the CP server cluster in secure mode	See "Configuring the CP server cluster in secure mode" on page 101.
Set up shared storage for the CP server database	See "Setting up shared storage for the CP server database" on page 102.
Configure the CP server	See " Configuring the CP server using the installer program" on page 103. See "Configuring the CP server using the web-based installer" on page 115. See "Configuring the CP server manually" on page 116. See "Configuring CP server using response files" on page 123.
Verify the CP server configuration	See "Verifying the CP server configuration" on page 127.

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 VCS 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 for IPM-based communication or HTTPS communication or both.

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

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

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

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

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

  See "CP server requirements" on page 41.

6 Have the following information handy for CP server configuration:

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

Installing the CP server using the installer

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

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

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

CP server setup uses a single system

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

During installation of VCS 6.2, VRTScps will come under recommended set of filesets.

Proceed to configure the CP server.

See " Configuring the CP server using the installer program" on page 103.

See "Configuring the CP server manually" on page 116.

CP server setup uses multiple systems

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

Meet the following requirements for CP server:

- During installation of SFHA 6.2, VRTScps will come under recommended set of filesets.

See the *Symantec 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 cluster in secure mode

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

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

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

To configure the CP server cluster in secure mode

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

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

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

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

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

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

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

  See "About the script-based installer" on page 50.

Setting up shared storage for the CP server database

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

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

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

To set up shared storage for the CP server database

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

For example:

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

2 Create a mirrored volume over the disk group.

For example:

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

3 Create a file system over the volume.

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

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

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

Configuring the CP server using the installer program

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

Perform one of the following procedures:

For CP servers on single-node VCS cluster:	See "To configure the CP server on a single-node VCS cluster" on page 104.
For CP servers on an SFHA cluster:	See "To configure the CP server on an SFHA cluster" on page 109.

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

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

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

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

   See "About the script-based installer" on page 50.
3. Installer checks the cluster information and prompts if you want to configure CP Server on the cluster.

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

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

   The installer then runs the following preconfiguration checks:

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

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

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

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

Enter the name of the CP Server.

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

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

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

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

---

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

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

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

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

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

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

Note that both IPv4 and IPv6 addresses are supported.

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

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

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

```
Symantec recommends secure communication between
the CP server  and application clusters. Enabling security
requires Symantec Product Authentication Service to be installed
and configured on the cluster. Do you want to enable Security for
the communications? [y,n,q,b] (y) n
```

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

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

15 Verify and confirm the CP server configuration information.

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

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

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

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

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

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

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

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

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

```
Enter a valid network interface on sys1 for NIC resource - 1:  en0
Enter a valid network interface on sys1 for NIC resource - 2:  en1
```

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

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

20 Enter the networkhosts information for each NIC resource.

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

Do you want to add NetworkHosts attribute for the NIC device en0
on system sys1? [y,n,q] y
Enter a valid IP address to configure NetworkHosts for NIC en0
on system sys1: 10.200.56.22

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

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

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

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

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

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

The Symantec coordination point server is ONLINE

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

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

```
For example:
# hagrp -state CPSSG
#Group Attribute System Value
CPSSG State.... |ONLINE|
```

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

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

To configure the CP server on an SFHA cluster

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

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

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

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

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

See "About the script-based installer" on page 50.

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

Enter **y** to confirm.

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

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

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

The installer then runs the following preconfiguration checks:

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

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

Enter the name of the CP server.

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

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

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

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

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

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

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

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

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

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

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

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

```
Symantec recommends secure communication between the CP server and application clusters.
Enabling security requires Symantec Product Authentication Service to be
installed and configured on the cluster.
Do you want to enable Security for the communications? [y,n,q,b] (y)
```

14 Enter absolute path of the database.

```
CP Server uses an internal database to store the client information.
As the CP Server is being configured on SFHA cluster, the database should reside
on shared storage with vxfs file system. Please refer to documentation for
information on setting up of shared storage for CP server database.
Enter absolute path of the database: [b] /cpsdb
```

15 Verify and confirm the CP server configuration information.

```
CP Server configuration verification:

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

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

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

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

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

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

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

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

```
Enter a valid network interface on sys1 for NIC resource - 1: en0
Enter a valid network interface on sys1 for NIC resource - 2: en1
```

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

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

20 Enter the networkhosts information for each NIC resource.

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


Do you want to add NetworkHosts attribute for the NIC device en0
on system sys1? [y,n,q] y
Enter a valid IP address to configure NetworkHosts for NIC en0
on system sys1: 10.200.56.22


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

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

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

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

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

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

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

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

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

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

24 Select the CP Server database volume.

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

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

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

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

```
1) newcpsvol
```

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

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

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

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

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

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

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

```
For example:
# hagrp -state CPSSG
#Group Attribute System Value
CPSSG State cps1 |ONLINE|
CPSSG State cps2 |OFFLINE|
```

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

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

Configuring the CP server using the web-based installer

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

To configure VCS on a cluster

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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

Task	Configure CP server
Product	Symantec Cluster Server

Click **Next**.

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

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

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

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

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

Click **Next**.

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

Click **Next**.

8 Configure network hosts for the CP server.

Click **Next**.

9 Configure disk group for the CP server.

Click **Next**.

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

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

Click **Next**.

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

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

Configuring the CP server manually

Perform the following steps to manually configure the CP server.

The CP server supports both IPM-based secure communication and HTTPS-based secure communication. CP servers that are configured for IPM-based secure communication support client nodes that are running prior to 6.1 versions of the product. However, CP servers that are configured for HTTP-based communication

only support client nodes that are running the 6.1 or later version of the product. Client nodes with product versions prior to 6.1 are not supported for HTTPS-based communication.

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

Table 7-2 Tasks to configure the CP server manually

Task	Reference
Configure CP server manually for IPM-based secure communication	See "Configuring the CP server manually for IPM-based secure communication" on page 117.
Configure CP server manually for HTTPS-communication	See "Configuring the CP server manually for HTTPS-based communication" on page 118. See "Generating the key and certificates manually for the CP server" on page 119. See "Completing the CP server configuration" on page 123.

Configuring the CP server manually for IPM-based secure communication

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

To manually configure the CP server

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

```
# hastop -local
```

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

See "Sample configuration files for CP server" on page 518.

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

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

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

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

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

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

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

5 Start VCS on all the cluster nodes.

```
# hastart
```

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

```
# hagrp -state CPSSG
```

Output similar to the following appears:

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

Configuring the CP server manually for HTTPS-based communication

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

To manually configure the CP server

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

```
# hastop -local
```

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

See "Sample configuration files for CP server" on page 518.

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

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

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

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

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

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

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

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

5 Manually generate keys and certificates for the CP server.

See "Generating the key and certificates manually for the CP server" on page 119.

Generating the key and certificates manually for the CP server

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

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

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

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

To manually generate keys and certificates for the CP server:

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

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

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

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

The sample configuration uses the following values:

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

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

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

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

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

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

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

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

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

4 Generate a self-signed CA certificate.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

10 Complete the CP server configuration.

See "Completing the CP server configuration" on page 123.

Completing the CP server configuration

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

1 Start VCS on all the cluster nodes.

```
# hastart
```

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

```
# hagrp -state CPSSG
```

Output similar to the following appears:

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

Configuring CP server using response files

You can configure a CP server using a generated responsefile.

On a single node VCS cluster:

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

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

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

  See "About the script-based installer" on page 50.

On a SFHA cluster:

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

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

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

  See "About the script-based installer" on page 50.

Response file variables to configure CP server

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

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

Variable	List or Scalar	Description
CFG{opt}{configcps}	Scalar	This variable performs CP server configuration task
CFG{cps_singlenode_config}	Scalar	This variable describes if the CP server will be configured on a singlenode VCS cluster
CFG{cps_sfha_config}	Scalar	This variable describes if the CP server will be configured on a SFHA cluster
CFG{cps_unconfig}	Scalar	This variable describes if the CP server will be unconfigured
CFG{cpsname}	Scalar	This variable describes the name of the CP server
CFG{cps_db_dir}	Scalar	This variable describes the absolute path of CP server database
CFG{cps_security}	Scalar	This variable describes if security is configured for the CP server
CFG{cps_reuse_cred}	Scalar	This variable describes if reusing the existing credentials for the CP server
CFG{cps_https_vips}	List	This variable describes the virtual IP addresses for the CP server configured for HTTPS-based communication
CFG{cps_ipm_vips}	List	This variable describes the virtual IP addresses for the CP server configured for IPM-based communication
CFG{cps_https_ports}	List	This variable describes the port number for the virtual IP addresses for the CP server configured for HTTPS-based communication
CFG{cps_ipm_ports}	List	This variable describes the port number for the virtual IP addresses for the CP server configured for IPM-based communication
CFG{cps_nic_list}{cpsvip<n>}	List	This variable describes the NICs of the systems for the virtual IP address

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

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

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

Review the response file variables and their definitions.

See Table 7-3 on page 124.

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

$CFG{cps_db_dir}="/etc/VRTScps/db";
$CFG{cps_https_ports}=[ qw(443) ];
$CFG{cps_https_vips}=[ qw(192.169.0.220) ];
$CFG{cps_ipm_ports}=[ qw(14250) ];
$CFG{cps_ipm_vips}=[ qw(192.169.0.221) ];
$CFG{cps_netmasks}=[ qw(255.255.252.0 255.255.252.0) ];
$CFG{cps_nic_list}{cpsvip1}=[ qw(en0) ];
$CFG{cps_nic_list}{cpsvip2}=[ qw(en0) ];
$CFG{cps_security}="0";
$CFG{cps_singlenode_config}=1;
$CFG{cps_vip2nicres_map}{"192.169.0.220"}=1;
$CFG{cps_vip2nicres_map}{"192.169.0.221"}=1;
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(cps1) ];
$CFG{vcs_clusterid}=64505;
$CFG{vcs_clustername}="single";

1;
```

Sample response file for configuring the CP server on SFHA cluster

Review the response file variables and their definitions.

See Table 7-3 on page 124.

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

$CFG{cps_db_dir}="/cpsdb";
$CFG{cps_diskgroup}="cps_dg1";
```

```
$CFG{cps_https_ports}=[ qw(50006 50007) ];
$CFG{cps_https_vips}=[ qw(10.198.90.6 10.198.90.7) ];
$CFG{cps_ipm_ports}=[ qw(14250) ];
$CFG{cps_ipm_vips}=[ qw(10.198.90.8) ];
$CFG{cps_netmasks}=[ qw(255.255.248.0 255.255.248.0 255.255.248.0) ];
$CFG{cps_network_hosts}{cpsnic1}=[ qw(10.198.88.18) ];
$CFG{cps_network_hosts}{cpsnic2}=[ qw(10.198.88.18) ];
$CFG{cps_newdg_disks}=[ qw(emc_clariion0_249) ];
$CFG{cps_newvol_volsize}=10;
$CFG{cps_nic_list}{cpsvip1}=[ qw(en0 en0) ];
$CFG{cps_nic_list}{cpsvip2}=[ qw(en0 en0) ];
$CFG{cps_nic_list}{cpsvip3}=[ qw(en0 en0) ];
$CFG{cps_security}="0";
$CFG{cps_sfha_config}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.6"}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.7"}=1;
$CFG{cps_vip2nicres_map}{"10.198.90.8"}=1;
$CFG{cps_volume}="volcps";
$CFG{cpsname}="cps1";
$CFG{opt}{configcps}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{noipc}=1;
$CFG{prod}="SFHA62";
$CFG{systems}=[ qw(cps1 cps2) ];
$CFG{vcs_clusterid}=49604;
$CFG{vcs_clustername}="sfha2233";

1;
```

Verifying the CP server configuration

Perform the following steps to verify the CP server configuration.

To verify the CP server configuration

1. Verify that the following configuration files are updated with the information you provided during the CP server configuration process:
   - /etc/vxcps.conf (CP server configuration file)
   - /etc/VRTSvcs/conf/config/main.cf (VCS configuration file)
   - /etc/VRTScps/db (default location for CP server database for a single-node cluster)

- /cps_db (default location for CP server database for a multi-node cluster)

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

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

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

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

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

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

Chapter 8

Configuring VCS

This chapter includes the following topics:

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

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

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

Task	Reference
Start the software configuration	See "Starting the software configuration" on page 130.
Specify the systems where you want to configure VCS	See "Specifying systems for configuration" on page 131.
Configure the basic cluster	See "Configuring the cluster name" on page 132. See "Configuring private heartbeat links" on page 132.
Configure virtual IP address of the cluster (optional)	See "Configuring the virtual IP of the cluster" on page 136.
Configure the cluster in secure mode (optional)	See "Configuring Symantec Cluster Server in secure mode" on page 138.
Add VCS users (required if you did not configure the cluster in secure mode)	See "Adding VCS users" on page 146.
Configure SMTP email notification (optional)	See "Configuring SMTP email notification" on page 147.
Configure SNMP email notification (optional)	See "Configuring SNMP trap notification" on page 148.
Configure global clusters (optional) **Note:** You must have enabled global clustering when you installed VCS.	See "Configuring global clusters" on page 150.
Complete the software configuration	See "Completing the VCS configuration" on page 151.

Starting the software configuration

You can configure VCS using the product installer or the installvcs command.

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

To configure VCS 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 Cluster Server

To configure VCS using the installvcs program

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

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

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

   See "About the script-based installer" on page 50.

   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 VCS. 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 VCS.

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

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

The installer does the following tasks:

- Checks that the local node running the installer can communicate with remote nodes
  If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases. If ssh binaries cannot communicate with remote nodes, the installer tries rsh binaries. And if both ssh and rsh binaries fail, the installer prompts to help the user to setup ssh or rsh binaries.
- Makes sure that the systems are running with the supported operating system
- Checks whether VCS is installed
- Exits if VCS 6.2 is not installed

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

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

See " About planning to configure I/O fencing" on page 90.

Configuring the cluster name

Enter the cluster information when the installer prompts you.

To configure the cluster

1 Review the configuration instructions that the installer presents.

2 Enter a unique cluster name.

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

Configuring private heartbeat links

You now configure the private heartbeat links that LLT uses.

See “Setting up the private network” on page 67.

VCS provides the option to use LLT over Ethernet or LLT 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.

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

See “Using the UDP layer for LLT” on page 529.

The following procedure helps you configure LLT heartbeat links.

To configure private heartbeat links

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

- Option 1: Configure the heartbeat links using LLT over Ethernet (answer installer questions)
  Enter the heartbeat link details at the installer prompt to configure LLT over Ethernet.
  Skip to step 2.
- Option 2: Configure the heartbeat links using LLT over UDP (answer installer questions)
  Make sure that each NIC you want to use as heartbeat link has an IP address configured. Enter the heartbeat link details at the installer prompt to configure LLT over UDP. If you had not already configured IP addresses to the NICs, the installer provides you an option to detect the IP address for a given NIC.
  Skip to step 3.
- Option 3: 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.
  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 en0 en1.)

```
Enter the NIC for the first private heartbeat link on sys1:
[b,q,?] en2
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat link on sys1:
[b,q,?] en3
Would you like to configure a third private heartbeat link?
[y,n,q,b,?](n)

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

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 or step 5 for option 3.

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

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

To configure the virtual IP of the cluster

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

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

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

Do one of the following:

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

```
Active NIC devices discovered on sys1: en0 en1
Enter the NIC for Virtual IP of the Cluster to use on sys1:
[b,q,?](en0 en1)
```

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 en0 en1 to be the public NIC used by all systems
[y,n,q,b,?] (y)
```

5 Enter the virtual IP address for the cluster.

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

For IPv4:

- Enter the virtual IP address.

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

- Confirm the default netmask or enter another one:

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

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

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

- Verify and confirm the Cluster Virtual IP information.

  ```
  Cluster Virtual IP verification:

        NIC: en0 en1
        IP: 192.168.1.16
        Netmask: 255.255.240.0

        NetworkHosts: 192.168.1.17

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

For IPv6

- Enter the virtual IP address.

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

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

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

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

  ```
  Enter the NetworkHosts IP addresses, separated
  by spaces: [b,q,?] 2001:db8::1 2001:db8::2
  ```

- Verify and confirm the Cluster Virtual IP information.

  ```
  Cluster Virtual IP verification:

        NIC: en0 en1
        IP: 2001:454e:205a:110:203:baff:feee:10
        Prefix: 64

        NetworkHosts: 2001:db8::1 2001:db8::2

  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 "Setting up trust relationships for your VCS cluster" on page 139.

See "Configuring a secure cluster node by node" on page 141.

Configuring Symantec Cluster Server in secure mode

Configuring VCS in secure mode ensures that all the communication between the systems is encrypted and users are verified against security credentials. VCS user names and passwords are not used when a cluster is running in secure mode.

To configure VCS in secure mode

1 To install VCS in secure mode, run the command:

```
# installvcs<version> -security
```

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

See "About the script-based installer" on page 50.

2 The installer displays the following question before the install stops the product processes:

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

3 To verify the cluster is in secure mode after configuration, run the command:

```
# haclus -value SecureClus
```

The command returns 1 if cluster is in secure mode, else returns 0.

Setting up trust relationships for your VCS cluster

If you need to use an external authentication broker for authenticating VCS users, you must set up a trust relationship between VCS and the broker. For example, if Veritas Operations Manager (VOM) is your external authentication broker, the trust relationship ensures that VCS accepts the credentials that VOM issues.

Perform the following steps to set up a trust relationship between your VCS cluster and a broker.

To set up a trust relationship

1. Ensure that you are logged in as superuser on one of the nodes in the cluster.
2. Enter the following command:

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

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

   See "About the script-based installer" on page 50.

   The installer specifies the location of the log files. It then lists the cluster information such as cluster name, cluster ID, node names, and service groups.
3. When the installer prompts you for the broker information, specify the IP address, port number, and the data directory for which you want to establish trust relationship with the broker.

   ```
   Input the broker name of IP address: 15.193.97.204

   Input the broker port: (14545)
   ```

   Specify a port number on which broker is running or press Enter to accept the default port.

   ```
   Input the data directory to setup trust with: (/var/VRTSvcs/
   vcsauth/data/HAD)
   ```

   Specify a valid data directory or press Enter to accept the default directory.
4. The installer performs one of the following actions:
   - If you specified a valid directory, the installer prompts for a confirmation.

     ```
     Are you sure that you want to setup trust for the VCS cluster
     with the broker 15.193.97.204 and port 14545? [y,n,q] y
     ```

     The installer sets up trust relationship with the broker for all nodes in the cluster and displays a confirmation.

     ```
     Setup trust with broker 15.193.97.204 on cluster node1
     ........Done

     Setup trust with broker 15.193.97.204 on cluster node2
     ........Done
     ```

     The installer specifies the location of the log files, summary file, and response file and exits.

- If you entered incorrect details for broker IP address, port number, or directory name, the installer displays an error. It specifies the location of the log files, summary file, and response file and exits.

Configuring a secure cluster node by node

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

Table 8-2 lists the tasks that you must perform to configure a secure cluster.

Table 8-2 Configuring a secure cluster node by node

Task	Reference
Configure security on one node	See "Configuring the first node" on page 141.
Configure security on the remaining nodes	See "Configuring the remaining nodes" on page 142.
Complete the manual configuration steps	See "Completing the secure cluster configuration" on page 143.

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/installvcs<version> -securityonenode
   ```

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

   See "About the script-based installer" on page 50.

   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/installvcs<version> -securityonenode
   ```

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

   See "About the script-based installer" on page 50.

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

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

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

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

   Select the option you would like to perform [1-2,q.?]  2
   Enter the security conf file directory: [b]
   ```

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

Completing the secure cluster configuration

Perform the following manual steps to complete the configuration.

To complete the secure cluster configuration

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

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

   # /opt/VRTSvcs/bin/hagrp -list Frozen=0

   # /opt/VRTSvcs/bin/hagrp -freeze groupname -persistent

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

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

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

3. On all nodes, stop the CmdServer.

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

4 To grant access to all users, add or modify `SecureClus=1` and `DefaultGuestAccess=1` in the cluster definition.

For example:

To grant read access to everyone:

```
Cluster clus1 (
SecureClus=1
DefaultGuestAccess=1
)
```

Or

To grant access to only root:

```
Cluster clus1 (
SecureClus=1
)
```

Or

To grant read access to specific user groups, add or modify SecureClus=1 and GuestGroups={} to the cluster definition.

For example:

```
cluster clus1 (
SecureClus=1
GuestGroups={staff, guest}
```

5 Modify `/etc/VRTSvcs/conf/config/main.cf` file on the first node, and add `-secure` to the WAC application definition if GCO is configured.

For example:

```
Application wac (
                StartProgram = "/opt/VRTSvcs/bin/wacstart -secure"
                StopProgram = "/opt/VRTSvcs/bin/wacstop"
                MonitorProcesses = {"/opt/VRTSvcs/bin/wac -secure"}
                RestartLimit = 3
                )
```

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

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

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

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

8 On all nodes, start CmdServer.

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

9 On the first node, unfreeze the service groups.

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

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

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

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

Adding VCS users

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

To add VCS users

1 Review the required information to add VCS users.

2 Reset the password for the Admin user, if necessary.

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

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

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

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

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

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

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

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

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

Configuring SMTP email notification

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

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

To configure SMTP email notification

1 Review the required information to configure the SMTP email notification.

2 Specify whether you want to configure the SMTP notification.

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

See "Configuring SNMP trap notification" on page 148.

3 Provide information to configure SMTP notification.

Provide the following information:

- Enter the SMTP server's host name.

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

- Enter the email address of each recipient.

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

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

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

4 Add more SMTP recipients, if necessary.

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

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

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

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

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

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

5 Verify and confirm the SMTP notification information.

```
SMTP Address: smtp.example.com
Recipient: ozzie@example.com receives email for Warning or
higher events
Recipient: harriet@example.com receives email for Error or
higher events

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

Configuring SNMP trap notification

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

Note that you can also configure the notification after installation.

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

To configure the SNMP trap notification

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

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

   See "Configuring global clusters" on page 150.

3. Provide information to configure SNMP trap notification.

   Provide the following information:

   - Enter the SNMP trap daemon port.

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

   - Enter the SNMP console system name.

     ```
     Enter the SNMP console system name: [b,q,?] sys5
     ```

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

     ```
     Enter the minimum severity of events for which SNMP traps
     should be sent to sys5 [I=Information, W=Warning, E=Error,
     S=SevereError]: [b,q,?] E
     ```

4. Add more SNMP consoles, if necessary.

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

     ```
     Would you like to add another SNMP console? [y,n,q,b] (n) y
     Enter the SNMP console system name: [b,q,?] sys4
     Enter the minimum severity of events for which SNMP traps
     should be sent to sys4 [I=Information, W=Warning,
     E=Error, S=SevereError]: [b,q,?] S
     ```

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

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

5 Verify and confirm the SNMP notification information.

```
SNMP Port: 162
Console: sys5 receives SNMP traps for Error or
higher events
Console: sys4 receives SNMP traps for SevereError or
higher events

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

Configuring global clusters

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster. If not, the installer proceeds to configure VCS based on the configuration details you provided. You can also run the gcoconfig utility in each cluster later to update the VCS configuration file for global cluster.

You can configure global clusters to link clusters at separate locations and enable wide-area failover and disaster recovery. The installer adds basic global cluster information to the VCS configuration file. You must perform additional configuration tasks to set up a global cluster.

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

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

To configure the global cluster option

1 Review the required information to configure the global cluster option.

2 Specify whether you want to configure the global cluster option.

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

3 Provide information to configure this cluster as global cluster.

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

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

Completing the VCS configuration

After you enter the VCS configuration information, the installer prompts to stop the VCS processes to complete the configuration process. The installer continues to create configuration files and copies them to each system. The installer also configures a cluster UUID value for the cluster at the end of the configuration. After the installer successfully configures VCS, it restarts VCS and its related processes.

To complete the VCS configuration

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

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

2. Review the output as the installer stops various processes and performs the configuration. The installer then restarts VCS 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 VCS 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 VCS using response files" on page 215.

Verifying and updating licenses on the system

After you install VCS, 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 152.

See "Updating product licenses" on page 152.

Checking licensing information on the system

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

To check licensing information

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

```
# vxlicrep
```

2 Review the following output to determine the following information:

- The license key
- The type of license
- The product for which it applies
- Its expiration date, if any. Demo keys have expiration dates. Permanent keys and site keys do not have expiration dates.

```
License Key                    = xxx-xxx-xxx-xxx-xxx
Product Name                   = Veritas Cluster Server
Serial Number                  = xxxxx
License Type                   = PERMANENT
OEM ID                         = xxxxx

Features :=
Platform                       = AIX
Version                        = 6.2
Tier                           = 0
Reserved                       = 0
Mode                           = VCS
CPU_Tier                       = 0
```

Updating product licenses

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

See "Replacing a VCS demo license with a permanent license" on page 153.

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 VCS demo license with a permanent license

When a VCS 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 VCS on all nodes in the cluster:

```
# hastop -all -force
```

This command does not shut down any running applications.

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

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

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

```
# vxlicrep
```

5 Start VCS on each node:

```
# hastart
```

Chapter 9

Configuring VCS clusters for data integrity

This chapter includes the following topics:

Setting up disk-based I/O fencing using installvcs

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

Initializing disks as VxVM disks

Perform the following procedure to initialize disks as VxVM disks.

To initialize disks as VxVM disks

1 Scan for the new hdisk devices.

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

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

```
# lsdev -Cc disk
```

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

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

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

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

```
# vxdisk list EMC0_17
```

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

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

  ```
  # vxdisksetup -i device_name
  ```

  The example specifies the CDS format:

  ```
  # vxdisksetup -i EMC0_17
  ```

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

Configuring disk-based I/O fencing using installvcs

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

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

1 Start the installvcs with `-fencing` option.

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

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

See "About the script-based installer" on page 50.

The installvcs 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 VCS 6.2 is configured properly.

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

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,b,q] 2
```

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

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

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

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

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

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

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

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

6 Verify that the coordinator disks you chose meet the I/O fencing requirements.

You must verify that the disks are SCSI-3 PR compatible using the vxfentsthdw utility and then return to this configuration program.

See "Checking shared disks for I/O fencing" on page 160.

7 After you confirm the requirements, the program creates the coordinator disk group with the information you provided.

8 Verify and confirm the I/O fencing configuration information that the installer summarizes.

9 Review the output as the configuration program does the following:

- Stops VCS and I/O fencing on each node.
- Configures disk-based I/O fencing and starts the I/O fencing process.
- Updates the VCS configuration file main.cf if necessary.
- Copies the /etc/vxfenmode file to a date and time suffixed file /etc/vxfenmode-*date-time*. This backup file is useful if any future fencing configuration fails.
- Updates the I/O fencing configuration file /etc/vxfenmode.
- Starts VCS on each node to make sure that the VCS is cleanly configured to use the I/O fencing feature.

10 Review the output as the configuration program displays the location of the log files, the summary files, and the response files.

11 Configure the Coordination Point Agent.

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

12 Enter a name for the service group for the Coordination Point Agent.

```
Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen) vxfen
```

13 Set the level two monitor frequency.

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

14 Decide the value of the level two monitor frequency.

```
Enter the value of the LevelTwoMonitorFreq attribute: [b,q,?] (5)
```

Installer adds Coordination Point Agent and updates the main configuration file.

See "Configuring CoordPoint agent to monitor coordination points" on page 289.

Refreshing keys or registrations on the existing coordination points for disk-based fencing using the installvcs

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for disk-based I/O fencing using the installvcs

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

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

where, `<version>` is the specific release version.

See "About the script-based installer" on page 50.

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

Note down the location of log files that you can access if there is a problem with the configuration process.

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

The program checks that the local node running the script can communicate with the remote nodes and checks whether VCS 6.2 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type the number corresponding to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,q]
```

4 Ensure that the disk group constitution that is used by the fencing module contains the same disks that are currently used as coordination disks.

5 Verify the coordination points.

```
For example,
Disk Group: fendg
Fencing disk policy: dmp
Fencing disks:
 emc_clariion0_62
 emc_clariion0_65
 emc_clariion0_66
```

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

```
Successfully completed the vxfenswap operation
```

The keys on the coordination disks are refreshed.

6 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

7 Do you want to view the summary file? [y,n,q] **(n)**.

Checking shared disks for I/O fencing

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

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

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

- If you test disks in DMP format, use the VxVM command `vxdisk list` to get the DMP path name.
- If you test disks in raw format for Active/Passive disk arrays, you must use an active enabled path with the vxfentsthdw command. Run the `vxdmpadm getsubpaths dmpnodename=`*`enclosure-based_name`* command to list the active enabled paths.
  DMP opens the secondary (passive) paths with an exclusive flag in Active/Passive arrays. So, if you test the secondary (passive) raw paths of the disk, the vxfentsthdw command may fail due to DMP's exclusive flag.

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

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

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

- Verifying the Array Support Library (ASL)
  See "Verifying Array Support Library (ASL)" on page 161.
- Verifying that nodes have access to the same disk
  See "Verifying that the nodes have access to the same disk" on page 161.
- Testing the shared disks for SCSI-3
  See "Testing the disks using vxfentsthdw utility" on page 162.

Verifying Array Support Library (ASL)

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

To verify Array Support Library (ASL)

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

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

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

The following output is a sample:

```
# vxddladm listsupport all


LIBNAME             VID                 PID
============================================================
libvx3par.so        3PARdata            VV
libvxCLARiiON.so    DGC                 All
libvxFJTSYe6k.so    FUJITSU             E6000
libvxFJTSYe8k.so    FUJITSU             All
libvxcompellent.so  COMPELNT            Compellent Vol
libvxcopan.so       COPANSYS            8814, 8818
libvxddns2a.so      DDN                 S2A 9550, S2A 9900,
                                        S2A 9700
```

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

```
# vxdisk scandisks
```

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

Verifying that the nodes have access to the same disk

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

To verify that the nodes have access to the same disk

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

   ```
   # vxfenadm -i diskpath
   ```

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

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

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

   From node A, enter:

   ```
   # vxfenadm -i /dev/rhdisk75

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

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

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

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

Testing the disks using vxfentsthdw utility

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

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

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

```
The disk /dev/rhdisk75 is ready to be configured for I/O Fencing on
node sys1
```

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

To test the disks using vxfentsthdw utility

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

See "About configuring secure shell or remote shell communication modes before installing products" on page 543.

2 From one node, start the utility.

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

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

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

Do you still want to continue : [y/n] (default: n) y
Enter the first node of the cluster: sys1
Enter the second node of the cluster: sys2
```

4 Review the output as the utility performs the checks and reports its activities.

5 If a disk is ready for I/O fencing on each node, the utility reports success for each node. For example, the utility displays the following message for the node sys1.

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

ALL tests on the disk /dev/rhdisk75 have PASSED
The disk is now ready to be configured for I/O fencing on node
sys1
```

6 Run the vxfentsthdw utility for each disk you intend to verify.

Note: Only dmp disk devices can be used as coordinator disks.

Setting up server-based I/O fencing using installvcs

You can configure server-based I/O fencing for the VCS cluster using the installvcs.

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 90.

See "Recommended CP server configurations" on page 96.

This section covers the following example procedures:

Mix of CP servers and coordinator disks	See "To configure server-based fencing for the VCS cluster (one CP server and two coordinator disks)" on page 164.
Single CP server	See "To configure server-based fencing for the VCS cluster (single CP server)" on page 168.

To configure server-based fencing for the VCS 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 VCS cluster. The VCS cluster is also referred to as the application cluster or the client cluster.
     See "Setting up the CP server" on page 99.
   - The coordination disks are verified for SCSI3-PR compliance.
     See "Checking shared disks for I/O fencing" on page 160.
2. Start the installvcs with the `-fencing` option.

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

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

   See "About the script-based installer" on page 50.

   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 VCS 6.2 is configured properly.

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

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,b,q] 1
```

5 Make sure that the storage supports SCSI3-PR, and answer y at the following prompt.

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

6 Provide the following details about the coordination points at the installer prompt:

- Enter the total number of coordination points including both servers and disks. This number should be at least 3.

```
Enter the total number of co-ordination points including both
Coordination Point servers and disks: [b] (3)
```

- Enter the total number of coordinator disks among the coordination points.

```
Enter the total number of disks among these:
[b] (0) 2
```

7 Provide the following CP server details at the installer prompt:

- Enter the total number of virtual IP addresses or the total number of fully qualified host names for each of the CP servers.

```
How many IP addresses would you like to use to communicate
 to Coordination Point Server #1?: [b,q,?] (1) 1
```

- Enter the virtual IP addresses or the fully qualified host name for each of the CP servers. The installer assumes these values to be identical as viewed from all the application cluster nodes.

```
Enter the Virtual IP address or fully qualified host name #1
for the HTTPS Coordination Point Server #1:
[b] 10.209.80.197
```

The installer prompts for this information for the number of virtual IP addresses you want to configure for each CP server.

- Enter the port that the CP server would be listening on.

```
Enter the port that the coordination point server 10.198.90.178
would be listening on or accept the default port
suggested: [b] (443)
```

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

- Choose the coordinator disks from the list of available disks that the installer displays. Ensure that the disk you choose is available from all the VCS (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) rhdisk75
2) rhdisk76
3) rhdisk77

Please enter a valid disk which is available from all the
cluster nodes for co-ordination point [1-3,q] 1
```

- If you have not already checked the disks for SCSI-3 PR compliance in step 1, check the disks now.
  The installer displays a message that recommends you to verify the disks in another window and then return to this configuration procedure.
  Press Enter to continue, and confirm your disk selection at the installer prompt.
- Enter a disk group name for the coordinator disks or accept the default.

```
Enter the disk group name for coordinating disk(s):
[b] (vxfencoorddg)
```

9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 3
Coordination Point Server ([VIP or FQHN]:Port):
    1. 10.209.80.197 ([10.209.80.197]:443)
SCSI-3 disks:
    1. rhdisk75
    2. rhdisk76
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk policy used for customized fencing: dmp
```

The installer initializes the disks and the disk group and deports the disk group on the VCS (application cluster) node.

10 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

11 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

```
Updating client cluster information on Coordination Point Server 10.209.80.197

Adding the client cluster to the Coordination Point Server 10.209.80.197 .......... Done

Registering client node sys1 with Coordination Point Server 10.209.80.197...... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.209.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 ..Done

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

See "About I/O fencing configuration files" on page 515.

12 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

13 Configure the CP agent on the VCS (application cluster). The Coordination Point Agent monitors the registrations on the coordination points.

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

Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen)
```

14 Additionally the coordination point agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidently deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles.

Note that for the LevelTwoMonitorFreq attribute to be applicable there must be disks as part of the Coordinator Disk Group.

```
Enter the value of the LevelTwoMonitorFreq attribute: (5)

Adding Coordination Point Agent via sys1 .... Done
```

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

16 Verify the fencing configuration using:

```
# vxfenadm -d
```

17 Verify the list of coordination points.

```
# vxfenconfig -l
```

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

1 Make sure that the CP server is configured and is reachable from the VCS cluster. The VCS cluster is also referred to as the application cluster or the client cluster.

2 See “Setting up the CP server” on page 99.

3 Start the installvcs with `-fencing` option.

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

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

See "About the script-based installer" on page 50.

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 VCS 6.2 is configured properly.

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

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,q] 1
```

6 Make sure that the storage supports SCSI3-PR, and answer y at the following prompt.

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

7 Enter the total number of coordination points as **1**.

```
Enter the total number of co-ordination points including both
Coordination Point servers and disks: [b] (3) 1
```

Read the installer warning carefully before you proceed with the configuration.

8 Provide the following CP server details at the installer prompt:

- Enter the total number of virtual IP addresses or the total number of fully qualified host names for each of the CP servers.

  ```
  How many IP addresses would you like to use to communicate
  to Coordination Point Server #1? [b,q,?] (1) 1
  ```

- Enter the virtual IP address or the fully qualified host name for the CP server. The installer assumes these values to be identical as viewed from all the application cluster nodes.

```
Enter the Virtual IP address or fully qualified host name
#1 for the Coordination Point Server #1:
[b] 10.209.80.197
```

The installer prompts for this information for the number of virtual IP addresses you want to configure for each CP server.

- Enter the port that the CP server would be listening on.

```
Enter the port in the range [49152, 65535] which the
Coordination Point Server 10.209.80.197
would be listening on or simply accept the default
port suggested: [b] (443)
```

9 Verify and confirm the coordination points information for the fencing configuration.

For example:

```
Total number of coordination points being used: 1
Coordination Point Server ([VIP or FQHN]:Port):
    1. 10.209.80.197 ([10.209.80.197]:443)
```

10 If the CP server is configured for security, the installer sets up secure communication between the CP server and the VCS (application cluster).

After the installer establishes trust between the authentication brokers of the CP servers and the application cluster nodes, press Enter to continue.

11 Verify and confirm the I/O fencing configuration information.

```
CPS Admin utility location: /opt/VRTScps/bin/cpsadm
Cluster ID: 2122
Cluster Name: clus1
UUID for the above cluster: {ae5e589a-1dd1-11b2-dd44-00144f79240c}
```

12 Review the output as the installer updates the application cluster information on each of the CP servers to ensure connectivity between them. The installer then populates the `/etc/vxfenmode` file with the appropriate details in each of the application cluster nodes.

The installer also populates the `/etc/vxfenmode` file with the entry `single_cp=1` for such single CP server fencing configuration.

```
Updating client cluster information on Coordination Point Server 10.209.80.197

Adding the client cluster to the Coordination Point Server 10.209.80.197 .......... Done

Registering client node sys1 with Coordination Point Server 10.209.80.197...... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

Registering client node sys2 with Coordination Point Server 10.209.80.197 ..... Done
Adding CPClient user for communicating to Coordination Point Server 10.209.80.197 .... Done
Adding cluster clus1 to the CPClient user on Coordination Point Server 10.209.80.197 .. Done

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

See "About I/O fencing configuration files" on page 515.

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 VCS (application cluster).

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

Enter a non-existing name for the service group for
Coordination Point Agent: [b] (vxfen)

Adding Coordination Point Agent via sys1 ... Done
```

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

Refreshing keys or registrations on the existing coordination points for server-based fencing using the installvcs

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss might occur because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose registrations of the cluster nodes, the cluster might panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points for server-based I/O fencing using the installvcs

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

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

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

See "About the script-based installer" on page 50.

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

Note the location of log files that you can access if there is a problem with the configuration process.

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

The program checks that the local node running the script can communicate with the remote nodes and checks whether VCS 6.2 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type the number corresponding to the option that suggests to refresh registrations or keys on the existing coordination points.

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q] 5
```

4 Ensure that the `/etc/vxfentab` file contains the same coordination point servers that are currently used by the fencing module.

Also, ensure that the disk group mentioned in the `/etc/vxfendg` file contains the same disks that are currently used by the fencing module as coordination disks.

5 Verify the coordination points.

```
For example,
Total number of coordination points being used: 3
Coordination Point Server ([VIP or FQHN]:Port):
      1. 10.198.94.146 ([10.198.94.146]:443)
      2. 10.198.94.144 ([10.198.94.144]:443)
SCSI-3 disks:
      1. emc_clariion0_61
Disk Group name for the disks in customized fencing: vxfencoorddg
Disk policy used for customized fencing: dmp
```

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

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

Successfully completed the vxfenswap operation
```

The keys on the coordination disks are refreshed.

7 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

8 Do you want to view the summary file? [y,n,q] **(n)**.

Setting the order of existing coordination points for server-based fencing using the installvcs

This section describes the reasons, benefits, considerations, and the procedure to set the order of the existing coordination points for server-based fencing.

About deciding the order of existing coordination points

You can decide the order in which coordination points can participate in a race during a network partition. In a network partition scenario, I/O fencing attempts to

contact coordination points for membership arbitration based on the order that is set in the `vxfentab` file.

When I/O fencing is not able to connect to the first coordination point in the sequence it goes to the second coordination point and so on. To avoid a cluster panic, the surviving subcluster must win majority of the coordination points. So, the order must begin with the coordination point that has the best chance to win the race and must end with the coordination point that has the least chance to win the race.

For fencing configurations that use a mix of coordination point servers and coordination disks, you can specify either coordination point servers before coordination disks or disks before servers.

Note: Disk-based fencing does not support setting the order of existing coordination points.

Considerations to decide the order of coordination points

- Choose the coordination points based on their chances to gain membership on the cluster during the race and hence gain control over a network partition. In effect, you have the ability to save a partition.
- First in the order must be the coordination point that has the best chance to win the race. The next coordination point you list in the order must have relatively lesser chance to win the race. Complete the order such that the last coordination point has the least chance to win the race.

Setting the order of existing coordination points using the installvcs

To set the order of existing coordination points

1 Start the installvcs with `-fencing` option.

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

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

See "About the script-based installer" on page 50.

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

Note the location of log files that you can access if there is a problem with the configuration process.

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

The program checks that the local node running the script can communicate with remote nodes and checks whether VCS 6.2 is configured properly.

3 Review the I/O fencing configuration options that the program presents. Type the number corresponding to the option that suggests to set the order of existing coordination points.

For example:

```
Select the fencing mechanism to be configured in this
Application Cluster [1-6,q] 6


Installer will ask the new order of existing coordination points.
Then it will call vxfenswap utility to commit the
coordination points change.
```

Warning: The cluster might panic if a node leaves membership before the coordination points change is complete.

4 Review the current order of coordination points.

```
Current coordination points order:
(Coordination disks/Coordination Point Server)
Example,
1)  /dev/vx/rdmp/emc_clariion0_65,/dev/vx/rdmp/emc_clariion0_66,
/dev/vx/rdmp/emc_clariion0_62
2)  [10.198.94.144]:443
3)  [10.198.94.146]:443
b)  Back to previous menu
```

5 Enter the new order of the coordination points by the numbers and separate the order by space [1-3,b,q] **3 1 2**.

```
New coordination points order:
(Coordination disks/Coordination Point Server)
Example,
1) [10.198.94.146]:443
2) /dev/vx/rdmp/emc_clariion0_65,/dev/vx/rdmp/emc_clariion0_66,
/dev/vx/rdmp/emc_clariion0_62
3) [10.198.94.144]:443
```

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

```
Preparing vxfenmode.test file on all systems...
Running vxfenswap...
Successfully completed the vxfenswap operation
```

7 Do you want to send the information about this installation to Symantec to help improve installation in the future? [y,n,q,?] **(y)**.

8 Do you want to view the summary file? [y,n,q] **(n)**.

9 Verify that the value of `vxfen_honor_cp_order` specified in the `/etc/vxfenmode` file is set to **1**.

```
For example,
vxfen_mode=customized
vxfen_mechanism=cps
port=443
scsi3_disk_policy=dmp
cps1=[10.198.94.146]
vxfendg=vxfencoorddg
cps2=[10.198.94.144]
vxfen_honor_cp_order=1
```

10 Verify that the coordination point order is updated in the output of the `vxfenconfig -l` command.

```
For example,
I/O Fencing Configuration Information:
======================================

single_cp=0
[10.198.94.146]:443 {e7823b24-1dd1-11b2-8814-2299557f1dc0}
/dev/vx/rdmp/emc_clariion0_65  60060160A38B1600386FD87CA8FDDD11
/dev/vx/rdmp/emc_clariion0_66  60060160A38B1600396FD87CA8FDDD11
/dev/vx/rdmp/emc_clariion0_62  60060160A38B16005AA00372A8FDDD11
[10.198.94.144]:443 {01f18460-1dd2-11b2-b818-659cbc6eb360}
```

Setting up non-SCSI-3 I/O fencing in virtual environments using installvcs

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

1 Start the installvcs with `-fencing` option.

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

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

See "About the script-based installer" on page 50.

The installvcs 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 VCS 6.2 is configured properly.

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

```
Select the fencing mechanism to be configured in this
Application Cluster
[1-7,q] 1
```

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

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

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

6 For server-based fencing, enter the number of CP server coordination points you want to use in your setup.

7 For server-based fencing, enter the following details for each CP server:

- Enter the virtual IP address or the fully qualified host name.
- Enter the port address on which the CP server listens for connections. The default value is 443. You can enter a different port address. Valid values are between 49152 and 65535.

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

8 For server-based fencing, verify and confirm the CP server information that you provided.

9 Verify and confirm the VCS 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 for only server-based fencing, :
  - Registers each node of the VCS cluster with the CP server.
  - Adds CP server user to the CP server.
  - Adds VCS cluster to the CP server user.
- Updates the following configuration files on each node of the VCS cluster
  - `/etc/vxfenmode` file
  - `/etc/default/vxfen` file
  - `/etc/vxenviron` file
  - `/etc/llttab` file
  - `/etc/vxfentab` (only for server-based fencing)

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

For server-based fencing, confirm to configure the CP agent on the VCS 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 majority-based I/O fencing using installvcs

You can configure majority-based fencing for the cluster using the installvcs .

Perform the following steps to confgure majority-based I/O fencing

1 Start the installvcs with the -fencing option.

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

Where *version* is the specific release version. The installvcs starts with a copyright message and verifies the cluster information.

See "About the script-based installer" on page 50.

Note: Make a note of the log file location which you can access in the event of any issues with the configuration process.

2 Confirm that you want to proceed with the I/O fencing configuration at the prompt. The program checks that the local node running the script can communicate with remote nodes and checks whether VCS is configured properly.

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

```
Select the fencing mechanism to be configured in this
Application Cluster [1-7,b,q] 3
```

Note: The installer will ask the following question. Does your storage environment support SCSI3 PR? [y,n,q,?] Input 'y' if your storage environment supports SCSI3 PR. Other alternative will result in installer configuring non-SCSI3 fencing(NSF).

4 The installer then populates the /etc/vxfenmode file with the appropriate details in each of the application cluster nodes.

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

5 Review the output as the installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

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

7 Verify the fencing configuration.

```
# vxfenadm -d
```

Enabling or disabling the preferred fencing policy

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

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

Preferred fencing is not applicable to majority-based I/O fencing.

See "About preferred fencing" on page 35.

To enable preferred fencing for the I/O fencing configuration

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

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

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

- Make the VCS configuration writable.

```
# haconf -makerw
```

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

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

- Set the value of the system-level attribute FencingWeight for each node in the cluster.
  For example, in a two-node cluster, where you want to assign sys1 five times more weight compared to sys2, run the following commands:

```
# hasys -modify sys1 FencingWeight 50
# hasys -modify sys2 FencingWeight 10
```

- Save the VCS configuration.

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

- Verify fencing node weights using:

  ```
  # vxfenconfig -a
  ```

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

- Make the VCS configuration writable.

  ```
  # haconf -makerw
  ```

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

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

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

  ```
  # hagrp -modify service_group Priority 1
  ```

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

- Save the VCS configuration.

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

5 To enable site-based race policy, perform the following steps:

- Make the VCS configuration writable.

  ```
  # haconf -makerw
  ```

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

  ```
  # haclus -modify PreferredFencingPolicy Site
  ```

- Set the value of the site-level attribute Preference for each site.

  ```
  For example,
  # hasite -modify Pune Preference 2
  ```

- Save the VCS configuration.

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

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

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

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

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

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

```
# haconf -makerw
# haclus -modify PreferredFencingPolicy Disabled
# haconf -dump -makero
```

Section 4

Installation using the Web-based installer

Chapter 10

Installing VCS

This chapter includes the following topics:

Before using the web-based installer

The web-based installer requires the following configuration.

Table 10-1 Web-based installer requirements

System	Function	Requirements
Target system	The systems where you plan to install the Symantec products.	Must be a supported platform for VCS 6.2.
Installation server	The server where you start the installation. The installation media is accessible from the installation server.	Must be at one of the supported operating system update levels.
Administrative system	The system where you run the web browser to perform the installation.	Must have a web browser. Supported browsers: ■ Internet Explorer 6, 7, and 8 ■ Firefox 3.x and later

Starting the web-based installer

This section describes starting the web-based installer.

To start the web-based installer

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

```
# ./webinstaller start
```

The `webinstaller` script displays a URL. Note this URL.

Note: If you do not see the URL, please check your firewall and iptables settings. If you have configured a firewall, ensure that the firewall settings allow access to the port 14172. You can alternatively use the `-port` option to use a free port instead.

You can use the following command to display the details about ports used by webinstaller and its status:

```
# ./webinstaller status
```

2 On the administrative server, start the web browser.

3 Navigate to the URL that the script displayed.

4 Certain browsers may display the following message:

```
Secure Connection Failed
```

Obtain a security exception for your browser.

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

5 Log in as superuser.

Obtaining a security exception on Mozilla Firefox

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

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

To obtain a security exception

1 Click **Or you can add an exception** link.

2 Click **I Understand the Risks**, or **You can add an exception**.

3 Click **Get Certificate** button.

4 Uncheck **Permanently Store this exception checkbox (recommended)**.

5 Click **Confirm Security Exception** button.

6 Enter root in *User Name* field and root password of the web server in the *Password* field.

Performing a preinstallation check with the web-based installer

This section describes performing a preinstallation check with the web-based installer.

To perform a preinstallation check

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

2 On the Select a task and a product page, select **Perform a Pre-installation Check** from the **Task** drop-down list.

3 Select **Symantec Cluster Server** 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.

Installing VCS with the web-based installer

This section describes installing VCS with the Symantec web-based installer.

To install VCS using the web-based installer

1 Perform preliminary steps.

See "Performing a preinstallation check with the web-based installer" on page 187.

2 Start the web-based installer.

See "Starting the web-based installer" on page 186.

3 Select **Install a Product** from the **Task** drop-down list.

4 Select **Symantec Cluster Server** 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 Choose minimal, recommended, or all filesets. Click **Next**.

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 VCS on the selected system.

10 After the installation completes, you must choose your licensing method.

On the license page, select one of the following radio buttons:

- Enable keyless licensing and complete system licensing later

---

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

For more information, go to the following website:

http://go.symantec.com/sfhakeyless

Click **Next**

Complete the following information:

- Choose whether you want to enable Global Cluster option.
- Click **Next**.

- Enter a valid license key
  If you have a valid license key, input the license key and click **Next**.

11 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 VCS.

   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.

12 If you are prompted, enter the option to specify whether you want to send your installation information to Symantec.

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

   Click **Finish**.

Chapter 11

Configuring VCS

This chapter includes the following topics:

Configuring VCS using the web-based installer

Before you begin to configure VCS using the web-based installer, review the configuration requirements.

See "Getting your VCS installation and configuration information ready" on page 77.

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 VCS on a cluster

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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

Task	Configure a Product
Product	Symantec Cluster Server

Click **Next**.

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

Example: **sys1 sys2**

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

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

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

Would you like to configure I/O fencing on the cluster?, click **Yes**.

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

See "Configuring VCS for data integrity using the web-based installer" on page 196.

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

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

See "About the script-based installer" on page 50.

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

Cluster Name	Enter a unique cluster name.
Cluster ID	Enter a unique cluster ID. Note that you can have the installer check to see if the cluster ID is unique. Symantec recommends that you use the installer to check for duplicate cluster IDs in multi-cluster environments.
Check duplicate cluster ID	Select the check box if you want the installer to verify if the given cluster ID is unique in your private network. The verification is performed after you specify the heartbeat details in the following pages. The verification takes some time to complete.
LLT Type	Select an LLT type from the list. You can choose to configure LLT over UDP or LLT over Ethernet.
Number of Heartbeats	Choose the number of heartbeat links you want to configure. See "Setting up the private network" on page 67.
Additional Low Priority Heartbeat NIC	Select the check box if you want to configure a low priority link. The installer configures one heartbeat link as low priority link. See "Setting up the private network" on page 67.
Unique Heartbeat NICs per system	For LLT over Ethernet, select the check box if you do not want to use the same NIC details to configure private heartbeat links on other systems. For LLT over UDP, this check box is selected by default.

Click **Next**.

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

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

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

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

Click **Next**.

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

Security

To configure a secure VCS 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 `installvcs<version>`.

Virtual IP

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

VCS Users

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

SMTP

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

SNMP

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

GCO

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

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

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

Click **Next**.

If virtual NICs exist in your setup, the NetworkHosts Configuration page displays.

8 The installer displays the following question before the install stops the product processes:

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**

- To grant read access only to root users, type **n**. The installer grants read access read access to the root users.

- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

9 Enter the details of the network hosts.

- If each system uses a separate NIC, select the **Configure NetworkHosts for every system separately** check box.
- Select a NIC and enter the network host details.
- If GCO is configured, enter the network host details for GCO.
- Click **Next**.

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

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

If you did not choose to configure I/O fencing in step 4, then skip to step 13. Go to step 12 to configure fencing.

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

Configure Coordination Point client based fencing	Choose this option to configure server-based I/O fencing.
Configure disk based fencing	Choose this option to configure disk-based I/O fencing.
Configure majority based fencing	Choose this option to configure majority based I/O fencing.

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

See “Configuring VCS for data integrity using the web-based installer” on page 196.

13 Click **Next** to complete the process of configuring VCS.

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 VCS for data integrity using the web-based installer

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

See "Configuring VCS using the web-based installer" on page 190.

See " About planning to configure I/O fencing" on page 90.

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 196.
- See "Configuring server-based fencing for data integrity using the web-based installer" on page 199.
- See "Configuring fencing in disabled mode using the web-based installer" on page 201.
- See "Configuring fencing in majority mode using the web-based installer" on page 202.
- See "Replacing, adding, or removing coordination points using the web-based installer" on page 203.
- See "Refreshing keys or registrations on the existing coordination points using web-based installer" on page 205.
- See "Setting the order of existing coordination points using the web-based installer" on page 206.

Configuring disk-based fencing for data integrity using the web-based installer

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

See "Configuring VCS using the web-based installer" on page 190.

See " About planning to configure I/O fencing" on page 90.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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 Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 On the Select Fencing Type page, select the `Configure disk-based fencing` option.

6 In the Confirmation dialog box that appears, confirm whether your storage environment supports SCSI-3 PR.

You can configure non-SCSI-3 fencing in a virtual environment that is not SCSI-3 PR compliant.

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

Click **Next**.

8 On the Configure Fencing page, specify the following information:

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

- If you selected one of the disk groups that is listed, choose the fencing disk policy for the disk group.
- If you selected the **Create a new disk group** option, make sure you have SCSI-3 PR enabled disks, and click **Yes** in the confirmation dialog box.

  Click **Next**.

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

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

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

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

10 Verify and confirm the I/O fencing configuration information.

The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

11 If you want to configure the Coordination Point agent on the client cluster, do the following:

- At the prompt for configuring the Coordination Point agent on the client cluster, click **Yes** and enter the Coordination Point agent service group name.
- If you want to set the LevelTwoMonitorFreq attribute, click Yes at the prompt and enter a value (0 to 65535).
- Follow the rest of the prompts to complete the Coordination Point agent configuration.

12 Click **Next** to complete the process of configuring I/O fencing.

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

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

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

Configuring server-based fencing for data integrity using the web-based installer

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

See "Configuring VCS using the web-based installer" on page 190.

See " About planning to configure I/O fencing" on page 90.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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 Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 On the Select Fencing Type page, select the `Configure Coordination Point client based fencing` option.

6 In the Confirmation dialog box that appears, confirm whether your storage environment supports SCSI-3 PR.

You can configure non-SCSI-3 fencing in a virtual environment that is not SCSI-3 PR compliant.

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

Click **Next**.

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

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

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

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

10 In the Confirmation dialog box that appears, confirm whether the coordination points information you provided is correct, and click **Yes**.

11 Verify and confirm the I/O fencing configuration information.

The installer stops and restarts the VCS and the fencing processes on each application cluster node, and completes the I/O fencing configuration.

12 If you want to configure the Coordination Point agent on the client cluster, do the following:

- At the prompt for configuring the Coordination Point agent on the client cluster, click **Yes** and enter the Coordination Point agent service group name.
- Follow the rest of the prompts to complete the Coordination Point agent configuration.

13 Click **Next** to complete the process of configuring I/O fencing.

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

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

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

Configuring fencing in disabled mode using the web-based installer

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

See "Configuring VCS using the web-based installer" on page 190.

See " About planning to configure I/O fencing" on page 90.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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 Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 Fencing may be enabled, installer may prompt whether you want to reconfigure it.

Click **Yes**.

6 On the Select Fencing Type page, select the `Configure fencing in disabled mode` option.

7 Installer stops VCS before applying the selected fencing mode to the cluster.

Note: Unfreeze any frozen service group and unmount any file system that is mounted in the cluster.

Click **Yes**.

8 Installer restarts VCS on all systems of the cluster. I/O fencing is disabled.

9 Verify and confirm the I/O fencing configuration information.

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

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

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

Configuring fencing in majority mode using the web-based installer

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

See "Configuring VCS using the web-based installer" on page 190.

See " About planning to configure I/O fencing" on page 90.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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 Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 Fencing may be enabled, installer may prompt whether you want to reconfigure it.

Click **Yes**.

6 On the Select Fencing Type page, select the `Configure fencing in majority mode` option.

7 Installer stops VCS before applying the selected fencing mode to the cluster.

Note: Unfreeze any frozen service group and unmount any file system that is mounted in the cluster.

Click **Yes**.

8 Installer restarts VCS on all systems of the cluster. I/O fencing is in majority mode.

9 Verify and confirm the I/O fencing configuration information.

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

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

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

Replacing, adding, or removing coordination points using the web-based installer

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

This procedure does not apply to majority-based I/O fencing.

See "Configuring VCS using the web-based installer" on page 190.

See " About planning to configure I/O fencing" on page 90.

To configure VCS for data integrity

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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 Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O Fencing on the cluster.

4 On the Select Cluster page, click **Next** if the installer completes the cluster verification successfully.

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

5 Fencing may be enabled, installer may prompt whether you want to reconfigure it.

Click **Yes**.

6 On the Select Fencing Type page, select the `Replace/Add/Remove coordination points` option.

7 The installer prompts to select the coordination points you want to remove from the currently configured coordination points.

Click **Next**.

8 Provide the number of Coordination point server and disk coordination points to be added to the configuration.

Click **Next**.

9 Provide the number of virtual IP addresses or Fully Qualified Host Name (FQHN) used for each coordination point server.

Click **Next**.

10 Provide the IP or FQHN and port number for each coordination point server.

Click **Next**.

11 Installer prompts to confirm the online migration coordination point servers.

Click **Yes**.

12 Installer proceeds with migration of the new coordination point servers. VCS is restarted during configuration.

Click **Next**.

13 You can add a Coordination Point agent to the client cluster and also provide name to the agent.

14 Click **Next**.

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

16 Select the check box to specify whether you want to send your installation information to Symantec.

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

Refreshing keys or registrations on the existing coordination points using web-based installer

This procedure does not apply to majority-based I/O fencing.

You must refresh registrations on the coordination points in the following scenarios:

- When the CoordPoint agent notifies VCS about the loss of registration on any of the existing coordination points.
- A planned refresh of registrations on coordination points when the cluster is online without having an application downtime on the cluster.

Registration loss may happen because of an accidental array restart, corruption of keys, or some other reason. If the coordination points lose the registrations of the cluster nodes, the cluster may panic when a network partition occurs.

Warning: Refreshing keys might cause the cluster to panic if a node leaves membership before the coordination points refresh is complete.

To refresh registrations on existing coordination points using web-based installer

1. Start the web-based installer.

   See "Starting the web-based installer" on page 186.

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 Cluster Server

   Click **Next**.

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

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

6. The installer may prompt you to reconfigure fencing if it is already enabled. Click **Yes** to reconfigure fencing.

7 On the **Select Fencing Type** page, select the `Refresh keys/registrations on the existing coordination points` option.

8 Ensure that the `/etc/vxfenmode` file contains the same coordination point servers that are currently used by the fencing module.

9 Ensure that the disk group mentioned in the `/etc/vxfenmode` file contains the same disks that are currently used by the fencing module as coordination disks.

10 Installer lists the reasons for the loss of registrations.

Click **OK**.

11 Verify the coordination points.

Click **Yes** if the information is correct.

12 Installer updates the client cluster information on the coordination point servers.

Click **Next**.

Installer prepares the `vxfenmode` file on all nodes and runs the `vxfenswap` utility to refresh registrations on the coordination points.

13 On the **Completion** page, view the `summary` file, `log` file, or `response` file to confirm the configuration.

14 Select the check box to specify whether you want to send your installation information to Symantec.

Click **Finish**.

Setting the order of existing coordination points using the web-based installer

This section describes the reasons, benefits, considerations, and the procedure to set the order of the existing coordination points using the web-based installer.

It does not apply to majority-based I/O fencing.

About deciding the order of existing coordination points

You can decide the order in which coordination points can participate in a race during a network partition. In a network partition scenario, I/O fencing attempts to contact coordination points for membership arbitration based on the order that is set in the `vxfenmode` file.

When I/O fencing is not able to connect to the first coordination point in the sequence it goes to the second coordination point and so on. To avoid a cluster panic, the surviving subcluster must win majority of the coordination points. So, the order must

begin with the coordination point that has the best chance to win the race and must end with the coordination point that has the least chance to win the race

For fencing configurations that use a mix of coordination point servers and coordination disks, you can either specify coordination point servers before coordination point disks or disks before servers.

Note: Disk-based fencing does not support setting the order of existing coordination points.

Considerations to decide the order of coordination points

- Choose coordination points based on their chances gain membership on the cluster during the race and hence gain control over a network partition. In effect, you have the ability to save a partition.
- First in the order must be the coordination point that has the best chance to win the race. The next coordination point you list in the order must have relatively lesser chance to win the race. Complete the order such that the last coordination point has the least chance to win the race.

Setting the order of existing coordination points using the web-based installer

To set the order of existing coordination points for server-based fencing using the web-based installer

1 Start the web-based installer.

See "Starting the web-based installer" on page 186.

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 Cluster Server

Click **Next**.

3 Verify the cluster information that the installer presents and confirm whether you want to configure I/O fencing on the cluster.

4 On the **Select Cluster** page, enter the system name and click **Yes**.

5 On the **Select Cluster** page, click **Next** if the installer completes the cluster verification successfully.

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

6 The installer may prompt you to reconfigure fencing if it is already enabled. Click **Yes** to reconfigure fencing.

Click **Yes**.

7 On the **Select Fencing Type** page, select the `Set the order of existing coordination points` option.

8 Confirm **OK** at the installer message about the procedure.

9 Decide the new order by moving the existing coordination points to the box on the window in the order you want. If you want to change the current order of coordination points, click **Reset** and start again.

10 Click **Next** if the information is correct.

11 On the **Confirmation** window, click **Yes**.

Installer prepares the `vxfenmode` file on all nodes and runs the `vxfenswap` utility to update the new order of coordination points.

12 On the **Completion** page, view the summary file, log file, or response file to confirm the configuration.

13 Select the check box to specify whether you want to send your installation information to Symantec.

Click **Finish**.

Section 5

Automated installation using response files

Chapter 12

Performing an automated VCS installation

This chapter includes the following topics:

Installing VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS installation on one cluster to install VCS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

```
# ./installer -makeresponsefile
```

See "About the script-based installer" on page 50.

To install VCS using response files

1. Make sure the systems where you want to install VCS meet the installation requirements.
2. Make sure that the preinstallation tasks are completed.

   See "Performing preinstallation tasks" on page 66.
3. Copy the response file to one of the cluster systems where you want to install VCS.

   See "Sample response file for installing VCS" on page 213.

4 Edit the values of the response file variables as necessary.

See "Response file variables to install VCS" on page 211.

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

# ./installvcs -responsefile /tmp/response_file
```

Where `/tmp/response_file` is the response file's full path name.

7 Complete the VCS post-installation tasks.

For instructions, see the chapter *Performing post-installation and configuration tasks* in this document.

Response file variables to install VCS

Table 12-1 lists the response file variables that you can define to install VCS.

Table 12-1 Response file variables specific to installing VCS

Variable	List or Scalar	Description
CFG{opt}{install}	Scalar	Installs VCS filesets. (Required)
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 installed, uninstalled or configured. Required
CFG{prod}	Scalar	Defines the product to be installed. The value is VCS62 for VCS. (Required)

Table 12-1 Response file variables specific to installing VCS *(continued)*

Variable	List or Scalar	Description
CFG{opt}{installallpkgs} or CFG{opt}{installrecpkgs} or CFG{opt}{installminpkgs}	Scalar	Instructs the installer to install VCS filesets based on the variable that has the value set to 1: ■ installallpkgs: Installs all filesets ■ installrecpkgs: Installs recommended filesets ■ installminpkgs: Installs minimum filesets **Note:** The installer requires only one of these variable values to be set to 1. (Required)
CFG{opt}{rsh}	Scalar	Defines that *rsh* must be used instead of ssh as the communication method between systems. (Optional)
CFG{opt}{gco}	Scalar	Defines that the installer must enable the global cluster option. You must set this variable value to 1 if you want to configure global clusters. (Optional)
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{patchpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product patches. The location must be accessible from all target systems. (Optional)

Table 12-1 Response file variables specific to installing VCS *(continued)*

Variable	List or Scalar	Description
CFG{opt}{pkgpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product filesets. The location must be accessible from all target systems. (Optional)
CFG{opt}{tmppath}	Scalar	Defines the location where a working directory is created to store temporary files and the filesets that are needed during the install. The default location is /var/tmp. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. **Note:** The installer copies the response files and summary files also to the specified *logpath* location. (Optional)
CFG{opt}{vxkeyless}	Scalar	Installs the product with keyless license if the value is set to 1. If the value is set to 0, you must define the CFG{keys}{system} variable with the license keys. (Optional)
CFG{keys} {system}	Scalar	List of keys to be registered on the system if the variable $CFG{opt}{vxkeyless} is set to 0. (Optional)

Sample response file for installing VCS

Review the response file variables and their definitions.

See "Response file variables to install VCS" on page 211.

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

$CFG{accepteula}=1;
$CFG{opt}{install}=1;
$CFG{opt}{installrecpkgs}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{uuid} = "16889f4e-1dd2-11b2-a559-afce02598e1b";
1;
```

Chapter 13

Performing an automated VCS configuration

This chapter includes the following topics:

Configuring VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS configuration on one cluster to configure VCS on other clusters. You can also create a response file using the `-makeresponsefile` option of the installer.

```
# ./installer -makeresponsefile -configure
```

```
# ./installvcs -makeresponsefile -configure
```

To configure VCS using response files

1. Make sure the VCS filesets are installed on the systems where you want to configure VCS.
2. Copy the response file to one of the cluster systems where you want to configure VCS.

   See "Sample response file for configuring Symantec Cluster Server" on page 225.

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 Cluster Server" on page 216.

4 Start the configuration from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/installvcs<version>
-responsefile /tmp/response_file
```

Where `<version>` is the specific release version, and `/tmp/response_file` is the response file's full path name.

See "About the script-based installer" on page 50.

Response file variables to configure Symantec Cluster Server

Table 13-1 lists the response file variables that you can define to configure VCS.

Table 13-1 Response file variables specific to configuring Symantec Cluster Server

Variable	List or Scalar	Description
CFG{opt}{configure}	Scalar	Performs the configuration if the filesets are already installed. (Required) Set the value to 1 to configure VCS.
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)
CFG{prod}	Scalar	Defines the product to be configured. The value is VCS62 for VCS. (Required)

Table 13-1 Response file variables specific to configuring Symantec Cluster Server *(continued)*

Variable	List or Scalar	Description
CFG{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{secusrgrps}	List	Defines the user groups which get read access to the cluster. (Optional)
CFG {rootsecusrgrps}	Scalar	Defines the read access to the cluster only for root and other users or usergroups which are granted explicit privileges on VCS objects. (Optional)
CFG{opt}{rsh}	Scalar	Defines that *rsh* must be used instead of ssh as the communication method between systems. (Optional)
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. **Note:** The installer copies the response files and summary files also to the specified *logpath* location. (Optional)
CFG{uploadlogs}	Scalar	Defines a Boolean value 0 or 1. The value 1 indicates that the installation logs are uploaded to the Symantec website. The value 0 indicates that the installation logs are not uploaded to the Symantec website. (Optional)

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service group (csgnic, csgvip, and csgnetmask) must be defined if any are defined. The same is true for the SMTP notification (smtpserver, smtprecp, and smtprsev), the SNMP trap notification (snmpport, snmpcons, and snmpcsev), and the Global Cluster Option (gconic, gcovip, and gconetmask).

Table 13-2 lists the response file variables that specify the required information to configure a basic VCS cluster.

Table 13-2 Response file variables specific to configuring a basic VCS 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 VCS configuration, defines if fencing is enabled. Valid values are 0 or 1. (Required)

Table 13-3 lists the response file variables that specify the required information to configure LLT over Ethernet.

Table 13-3 Response file variables specific to configuring private LLT over Ethernet

Variable	List or Scalar	Description
CFG{vcs_lltlink#} {"system"}	Scalar	Defines the NIC to be used for a private heartbeat link on each system. Atleast two LLT links are required per system (lltlink1 and lltlink2). You can configure up to four LLT links. See "Setting up the private network" on page 67. You must enclose the system name within double quotes. (Required)
CFG{vcs_lltlinklowpri#} {"system"}	Scalar	Defines a low priority heartbeat link. Typically, lltlinklowpri is used on a public network link to provide an additional layer of communication. If you use different media speed for the private NICs, you can configure the NICs with lesser speed as low-priority links to enhance LLT performance. For example, lltlinklowpri1, lltlinklowpri2, and so on. You must enclose the system name within double quotes. (Optional)

Table 13-4 lists the response file variables that specify the required information to configure LLT over UDP.

Table 13-4 Response file variables specific to configuring LLT over UDP

Variable	List or Scalar	Description
CFG{lltoverudp}=1	Scalar	Indicates whether to configure heartbeat link using LLT over UDP. (Required)

Table 13-4 Response file variables specific to configuring LLT over UDP *(continued)*

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_address} {<sys1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG {vcs_udplinklowpri<n>_address} {<sys1>}	Scalar	Stores the IP address (IPv4 or IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)
CFG{vcs_udplink<n>_port} {<sys1>}	Scalar	Stores the UDP port (16-bit integer value) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG{vcs_udplinklowpri<n>_port} {<sys1>}	Scalar	Stores the UDP port (16-bit integer value) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 13-4 Response file variables specific to configuring LLT over UDP *(continued)*

Variable	List or Scalar	Description
CFG{vcs_udplink<n>_netmask} {<sys1>}	Scalar	Stores the netmask (prefix for IPv6) that the heartbeat link uses on node1. You can have four heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective heartbeat links. (Required)
CFG {vcs_udplinklowpri<n>_netmask} {<sys1>}	Scalar	Stores the netmask (prefix for IPv6) that the low priority heartbeat link uses on node1. You can have four low priority heartbeat links and <n> for this response file variable can take values 1 to 4 for the respective low priority heartbeat links. (Required)

Table 13-5 lists the response file variables that specify the required information to configure virtual IP for VCS cluster.

Table 13-5 Response file variables specific to configuring virtual IP for VCS cluster

Variable	List or Scalar	Description
CFG{vcs_csgnic} {system}	Scalar	Defines the NIC device to use on a system. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_csgvip}	Scalar	Defines the virtual IP address for the cluster. (Optional)
CFG{vcs_csgnetmask}	Scalar	Defines the Netmask of the virtual IP address for the cluster. (Optional)

Table 13-6 lists the response file variables that specify the required information to configure the VCS cluster in secure mode.

Table 13-6 Response file variables specific to configuring VCS cluster in secure mode

Variable	List or Scalar	Description
CFG{vcs_eat_security}	Scalar	Specifies if the cluster is in secure enabled mode or not.
CFG{opt}{securityonenode}	Scalar	Specifies that the securityonenode option is being used.
CFG{securityonenode_menu}	Scalar	Specifies the menu option to choose to configure the secure cluster one at a time. ■ 1—Configure the first node ■ 2—Configure the other node
CFG{secusrgrps}	List	Defines the user groups which get read access to the cluster. List or scalar: list Optional or required: optional
CFG{rootsecusrgrps}	Scalar	Defines the read access to the cluster only for root and other users or user groups which are granted explicit privileges in VCS objects. (Optional)
CFG{security_conf_dir}	Scalar	Specifies the directory where the configuration files are placed.
CFG{opt}{security}	Scalar	Specifies that the security option is being used.
CFG{vcs_eat_security_fips}	Scalar	Specifies that the enabled security is FIPS compliant.

Table 13-7 lists the response file variables that specify the required information to configure VCS users.

Table 13-7 Response file variables specific to configuring VCS users

Variable	List or Scalar	Description
CFG{vcs_userenpw}	List	List of encoded passwords for VCS users The value in the list can be "Administrators Operators Guests" **Note:** The order of the values for the vcs_userenpw list must match the order of the values in the vcs_username list. (Optional)
CFG{vcs_username}	List	List of names of VCS users (Optional)
CFG{vcs_userpriv}	List	List of privileges for VCS users **Note:** The order of the values for the vcs_userpriv list must match the order of the values in the vcs_username list. (Optional)

Table 13-8 lists the response file variables that specify the required information to configure VCS notifications using SMTP.

Table 13-8 Response file variables specific to configuring VCS notifications using SMTP

Variable	List or Scalar	Description
CFG{vcs_smtpserver}	Scalar	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for web notification. (Optional)
CFG{vcs_smtprecp}	List	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. (Optional)

Table 13-8 Response file variables specific to configuring VCS notifications using SMTP *(continued)*

Variable	List or Scalar	Description
CFG{vcs_smtprsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SMTP recipients are to receive. Note that the ordering of severity levels must match that of the addresses of SMTP recipients. (Optional)

Table 13-9 lists the response file variables that specify the required information to configure VCS notifications using SNMP.

Table 13-9 Response file variables specific to configuring VCS notifications using SNMP

Variable	List or Scalar	Description
CFG{vcs_snmpport}	Scalar	Defines the SNMP trap daemon port (default=162). (Optional)
CFG{vcs_snmpcons}	List	List of SNMP console system names (Optional)
CFG{vcs_snmpcsev}	List	Defines the minimum severity level of messages (Information, Warning, Error, SevereError) that listed SNMP consoles are to receive. Note that the ordering of severity levels must match that of the SNMP console system names. (Optional)

Table 13-10 lists the response file variables that specify the required information to configure VCS global clusters.

Table 13-10 Response file variables specific to configuring VCS global clusters

Variable	List or Scalar	Description
CFG{vcs_gconic} {system}	Scalar	Defines the NIC for the Virtual IP that the Global Cluster Option uses. You can enter 'all' as a system value if the same NIC is used on all systems. (Optional)
CFG{vcs_gcovip}	Scalar	Defines the virtual IP address to that the Global Cluster Option uses. (Optional)
CFG{vcs_gconetmask}	Scalar	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. (Optional)

Sample response file for configuring Symantec Cluster Server

Review the response file variables and their definitions.

See "Response file variables to configure Symantec Cluster Server" on page 216.

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

$CFG{opt}{configure}=1;
$CFG{opt}{gco}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=13221;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_csgnetmask}="255.255.255.0";
$CFG{vcs_csgnic}{all}="en0";
$CFG{vcs_csgvip}="10.10.12.1";
$CFG{vcs_gconetmask}="255.255.255.0";
```

```
$CFG{vcs_gcovip}="10.10.12.1";
$CFG{vcs_lltlink1}{sys1}="en2";
$CFG{vcs_lltlink1}{sys2}="en2";
$CFG{vcs_lltlink2}{sys1}="en3";
$CFG{vcs_lltlink2}{sys2}="en3";

$CFG{vcs_smtprecp}=[ qw(earnie@symantecexample.com) ];
$CFG{vcs_smtprsev}=[ qw(SevereError) ];
$CFG{vcs_smtpserver}="smtp.symantecexample.com";
$CFG{vcs_snmpcons}=[ qw(neptune) ];
$CFG{vcs_snmpcsev}=[ qw(SevereError) ];
$CFG{vcs_snmpport}=162;
1;
```

Chapter 14

Performing an automated I/O fencing configuration using response files

This chapter includes the following topics:

Configuring I/O fencing using response files

Typically, you can use the response file that the installer generates after you perform I/O fencing configuration to configure I/O fencing for VCS.

To configure I/O fencing using response files

1. Make sure that VCS 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 90.
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 231.

   See "Sample response file for configuring server-based I/O fencing" on page 233.

   See "Sample response file for configuring non-SCSI-3 I/O fencing" on page 235.

   See "Sample response file for configuring majority-based I/O fencing" on page 236.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to configure disk-based I/O fencing" on page 228.

   See "Response file variables to configure server-based I/O fencing" on page 231.

   See "Response file variables to configure non-SCSI-3 I/O fencing" on page 234.

   See "Response file variables to configure majority-based I/O fencing" on page 236.
5. Start the configuration from the system to which you copied the response file. For example:

   ```
   # /opt/VRTS/install/installvcs<version>
   -responsefile /tmp/response_file
   ```

   Where *<version>* is the specific release version, and */tmp/response_file* is the response file's full path name.

   See "About the script-based installer" on page 50.

Response file variables to configure disk-based I/O fencing

Table 14-1 lists the response file variables that specify the required information to configure disk-based I/O fencing for VCS.

Table 14-1 Response file variables specific to configuring disk-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. ■ 1—Configure Coordination Point client-based I/O fencing ■ 2—Configure disk-based I/O fencing ■ 3—Configure majority-based I/O fencing ■ 4—Configure I/O fencing in disabled mode ■ 5—Replace/Add/Remove coordination points ■ 6—Refresh keys/registrations on the existing coordination points ■ 7—Set the order of existing coordination points (Required)
CFG{fencing_dgname}	Scalar	Specifies the disk group for I/O fencing. (Optional) **Note:** You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.
CFG{fencing_newdg_disks}	List	Specifies the disks to use to create a new disk group for I/O fencing. (Optional) **Note:** You must define the fencing_dgname variable to use an existing disk group. If you want to create a new disk group, you must use both the fencing_dgname variable and the fencing_newdg_disks variable.

Table 14-1 Response file variables specific to configuring disk-based I/O fencing *(continued)*

Variable	List or Scalar	Description
CFG{fencing_cpagent_monitor_freq}	Scalar	Specifies the frequency at which the Coordination Point Agent monitors for any changes to the Coordinator Disk Group constitution. **Note:** Coordination Point Agent can also monitor changes to the Coordinator Disk Group constitution such as a disk being accidently deleted from the Coordinator Disk Group. The frequency of this detailed monitoring can be tuned with the LevelTwoMonitorFreq attribute. For example, if you set this attribute to 5, the agent will monitor the Coordinator Disk Group constitution every five monitor cycles. If LevelTwoMonitorFreq attribute is not set, the agent will not monitor any changes to the Coordinator Disk Group. 0 means not to monitor the Coordinator Disk Group constitution.
CFG {fencing_config_cpagent}	Scalar	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.
CFG {fencing_cpagentgrp}	Scalar	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the **fencing_config_cpagent** field is given a value of '0'.

Sample response file for configuring disk-based I/O fencing

Review the disk-based I/O fencing response file variables and their definitions.

See "Response file variables to configure disk-based I/O fencing" on page 228.

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

$CFG{fencing_config_cpagent}=1;
$CFG{fencing_cpagent_monitor_freq}=5;
$CFG{fencing_cpagentgrp}="vxfen";
$CFG{fencing_dgname}="fencingdg1";
$CFG{fencing_newdg_disks}=[ qw(emc_clariion0_155
 emc_clariion0_162 emc_clariion0_163) ];
$CFG{fencing_option}=2;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{fencing_cpagent_monitor_freq}=5;

$CFG{prod}="VCS62";

$CFG{systems}=[ qwsys1sys2 ];
$CFG{vcs_clusterid}=32283;
$CFG{vcs_clustername}="clus1";
1;
```

Response file variables to configure server-based I/O fencing

You can use a coordination point server-based fencing response file to configure server-based customized I/O fencing.

Table 14-2 lists the fields in the response file that are relevant for server-based customized I/O fencing.

Table 14-2 Coordination point server (CP server) based fencing response file definitions

Response file field	Definition
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the `fencing_config_cpagent` field is given a value of '0'.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers.
CFG {fencing_reusedg}	This response file field indicates whether to reuse an existing DG name for the fencing configuration in customized fencing (CP server and coordinator disks). Enter either a "1" or "0". Entering a "1" indicates reuse, and entering a "0" indicates do not reuse. When reusing an existing DG name for the mixed mode fencing configuration. you need to manually add a line of text , such as "$CFG{fencing_reusedg}=0" or "$CFG{fencing_reusedg}=1" before proceeding with a silent installation.
CFG {fencing_dgname}	The name of the disk group to be used in the customized fencing, where at least one disk is being used.
CFG {fencing_disks}	The disks being used as coordination points if any.
CFG {fencing_ncp}	Total number of coordination points being used, including both CP servers and disks.
CFG {fencing_ndisks}	The number of disks being used.

Table 14-2 Coordination point server (CP server) based fencing response file definitions *(continued)*

Response file field	Definition
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server.
CFG {fencing_cps_ports}	The port that the virtual IP address or the fully qualified host name of the CP server listens on.
CFG{fencing_option}	Specifies the I/O fencing configuration mode. ■ 1—Configure Coordination Point client-based I/O fencing ■ 2—Configure disk-based I/O fencing ■ 3—Configure majority-based I/O fencing ■ 4—Configure I/O fencing in disabled mode ■ 5—Replace/Add/Remove coordination points ■ 6—Refresh keys/registrations on the existing coordination points ■ 7—Set the order of existing coordination points

Sample response file for configuring server-based I/O fencing

The following is a sample response file used for server-based I/O fencing:

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.200.117.145) ];
$CFG{fencing_cps_vips}{"10.200.117.145"}=[ qw(10.200.117.145) ];
$CFG{fencing_dgname}="vxfencoorddg";
$CFG{fencing_disks}=[ qw(emc_clariion0_37 emc_clariion0_13
emc_clariion0_12) ];
$CFG{fencing_scsi3_disk_policy}="dmp";
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=2;
$CFG{fencing_cps_ports}{"10.200.117.145"}=443;
$CFG{fencing_reusedg}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=1256;
```

```
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Response file variables to configure non-SCSI-3 I/O fencing

Table 14-3 lists the fields in the response file that are relevant for non-SCSI-3 I/O fencing.

See "About I/O fencing for VCS in virtual machines that do not support SCSI-3 PR" on page 33.

Table 14-3 Non-SCSI-3 I/O fencing response file definitions

Response file field	Definition
CFG{non_scsi3_fencing}	Defines whether to configure non-SCSI-3 I/O fencing. Valid values are 1 or 0. Enter 1 to configure non-SCSI-3 I/O fencing.
CFG {fencing_config_cpagent}	Enter '1' or '0' depending upon whether you want to configure the Coordination Point agent using the installer or not. Enter "0" if you do not want to configure the Coordination Point agent using the installer. Enter "1" if you want to use the installer to configure the Coordination Point agent. **Note:** This variable does not apply to majority-based fencing.
CFG {fencing_cpagentgrp}	Name of the service group which will have the Coordination Point agent resource as part of it. **Note:** This field is obsolete if the `fencing_config_cpagent` field is given a value of '0'. This variable does not apply to majority-based fencing.
CFG {fencing_cps}	Virtual IP address or Virtual hostname of the CP servers. **Note:** This variable does not apply to majority-based fencing.

Table 14-3 Non-SCSI-3 I/O fencing response file definitions *(continued)*

Response file field	Definition
CFG {fencing_cps_vips}	The virtual IP addresses or the fully qualified host names of the CP server. **Note:** This variable does not apply to majority-based fencing.
CFG {fencing_ncp}	Total number of coordination points (CP servers only) being used. **Note:** This variable does not apply to majority-based fencing.
CFG {fencing_cps_ports}	The port of the CP server that is denoted by *cps* . **Note:** This variable does not apply to majority-based fencing.

Sample response file for configuring non-SCSI-3 I/O fencing

The following is a sample response file used for non-SCSI-3 I/O fencing :

```
$CFG{fencing_config_cpagent}=0;
$CFG{fencing_cps}=[ qw(10.198.89.251 10.198.89.252 10.198.89.253) ];
$CFG{fencing_cps_vips}{"10.198.89.251"}=[ qw(10.198.89.251) ];
$CFG{fencing_cps_vips}{"10.198.89.252"}=[ qw(10.198.89.252) ];
$CFG{fencing_cps_vips}{"10.198.89.253"}=[ qw(10.198.89.253) ];
$CFG{fencing_ncp}=3;
$CFG{fencing_ndisks}=0;
$CFG{fencing_cps_ports}{"10.198.89.251"}=443;
$CFG{fencing_cps_ports}{"10.198.89.252"}=443;
$CFG{fencing_cps_ports}{"10.198.89.253"}=443;
$CFG{non_scsi3_fencing}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(sys1 sys2) ];
$CFG{vcs_clusterid}=1256;
$CFG{vcs_clustername}="clus1";
$CFG{fencing_option}=1;
```

Response file variables to configure majority-based I/O fencing

Table 14-4 lists the response file variables that specify the required information to configure disk-based I/O fencing for VCS.

Table 14-4 Response file variables specific to configuring majority-based I/O fencing

Variable	List or Scalar	Description
CFG{opt}{fencing}	Scalar	Performs the I/O fencing configuration. (Required)
CFG{fencing_option}	Scalar	Specifies the I/O fencing configuration mode. ■ 1—Coordination Point Server-based I/O fencing ■ 2—Coordinator disk-based I/O fencing ■ 3—Disabled-based fencing ■ 4—Online fencing migration ■ 5—Refresh keys/registrations on the existing coordination points ■ 6—Change the order of existing coordination points ■ 7—Majority-based fencing (Required)

Sample response file for configuring majority-based I/O fencing

```
$CFG{fencing_option}=7;
$CFG{config_majority_based_fencing}=1;
$CFG{opt}{configure}=1;
$CFG{opt}{fencing}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(sys1  sys2) ];
```

```
$CFG{vcs_clusterid}=59082;
$CFG{vcs_clustername}="clus1";
```

Section 6

Manual installation

Chapter 15

Performing preinstallation tasks

This chapter includes the following topics:

Modifying /etc/pse.conf to enable the Ethernet driver

Before you install VCS, examine the /etc/pse.conf file on each system to see if the Ethernet driver is configured. If the driver is not configured, you must modify the file and restart the system.

To enable the Ethernet driver

1. Check to see if the Ethernet driver is configured in the /etc/pse.conf file:

   ```
   # egrep 'ethernet driver' /etc/pse.conf
   ```

2. In the output, examine the line containing the "ethernet driver" expression:

   ```
   #d+ dlpi en /dev/dlpi/en # streams dlpi ethernet driver
   ```

3. If the comment symbol ("#") precedes the line, the Ethernet driver is not configured. Using vi or another text editor, edit the file:

   ```
   # vi /etc/pse.conf
   ```

4. Find the section in the file labeled "#PSE drivers" and look for the line in step 2.

   Uncomment the line by removing the initial "#" symbol.

5 Save and close the file.

6 To configure the driver, restart the system or execute the following command.

```
# strload -f /etc/pse.conf
```

7 Repeat step 1 through step 6 on each system in the cluster.

Requirements for installing VCS

Review requirements before you install.

See “Important preinstallation information for VCS” on page 38.

Chapter 16

Manually installing VCS

This chapter includes the following topics:

About VCS manual installation

You can manually install and configure VCS instead of using the installvcs.

A manual installation takes a lot of time, patience, and care. Symantec recommends that you use the installvcs instead of the manual installation when possible.

Installing VCS software manually

If you manually install VCS software to upgrade your cluster, make sure to back up the previous VCS configuration files before you start the installation. The configuration files that you must back up are as follows:

- All the files from `/etc/VRTSvcs/conf/config` directory.
- /etc/llttab
- /etc/gabtab
- /etc/llthosts
- /etc/amftab
- /etc/default/vcs
- /etc/default/amf
- /etc/default/vxfen

- /etc/default/gab
- /etc/default/llt

Table 16-1 lists the tasks that you must perform when you manually install and configure VCS 6.2.

Table 16-1 Manual installation tasks for VCS 6.2

Task	Reference
Modify /etc/pse.conf to enable the ethernet driver.	See "Modifying /etc/pse.conf to enable the Ethernet driver" on page 239.
Install VCS software manually on each node in the cluster.	See "Installing VCS filesets for a manual installation" on page 243.
Add a license key.	See "Adding a license key for a manual installation" on page 244.
Copy the installation guide to each node.	See "Copying the installation guide to each node" on page 247.
Configure LLT and GAB.	■ See "Configuring LLT manually" on page 251. ■ See "Configuring GAB manually" on page 254.
Configure VCS.	See "Configuring VCS manually" on page 254.
Start LLT, GAB, and VCS services.	See "Starting LLT, GAB, and VCS after manual configuration" on page 261.
Modify the VCS configuration.	See "Modifying the VCS configuration" on page 270.
Replace demo license with a permanent license.	See "Replacing a VCS demo license with a permanent license for manual installations" on page 246.

Viewing the list of VCS filesets

During the VCS installation, the installer prompts you with an option to choose the VCS filesets to install. You can view the list of filesets that each of these options would install using the installer command-line option.

Manual installation or upgrade of the product requires you to install the filesets in a specified order. For example, you must install some filesets before other filesets because of various product dependencies. The following installer command options list the filesets in the order in which you must install these filesets.

Table 16-2 describes the VCS fileset installation options and the corresponding command to view the list of filesets.

Table 16-2 Installer command options to view VCS filesets

Option	Description	Command option to view the list of filesets
1	Installs only the minimal required VCS filesets that provide basic functionality of the product.	`installvcs -minpkgs`
2	Installs the recommended VCS filesets that provide complete functionality of the product. This option does not install the optional VCS filesets.	`installvcs -recpkgs`
3	Installs all the VCS filesets. You must choose this option to configure any optional VCS feature.	`installvcs -allpkgs`

To view the list of VCS filesets

1 Navigate to the directory from where you can start the installvcs.

```
# cd cluster_server
```

2 Run the following command to view the list of filesets. Based on what filesets you want to install, enter the appropriate command option:

```
# ./installvcs -minpkgs
```

Or

```
# ./installvcs -recpkgs
```

Or

```
# ./installvcs -allpkgs
```

Installing VCS filesets for a manual installation

All filesets are installed into the `/opt` directory and a few files are installed into the `/etc` and `/var` directories.

You can create lists of the filesets to install.

See "Viewing the list of VCS filesets" on page 242.

If you copied the Symantec filesets to `/tmp/install`, navigate to the directory and perform the following on each system:

To install VCS filesets on a node

- Install the required filesets in the order shown:

```
# installp -a -d VRTSperl.bff VRTSperl
# installp -a -d VRTSvlic.bff VRTSvlic
# installp -a -d VRTSspt.bff VRTSspt
# installp -a -d VRTSveki.bff VRTSveki
# installp -a -d VRTSllt.bff VRTSllt
# installp -a -d VRTSgab.bff VRTSgab
# installp -a -d VRTSvxfen.bff VRTSvxfen
# installp -a -d VRTSamf.bff VRTSamf
# installp -a -d VRTSvcs.bff VRTSvcs
# installp -a -d VRTScps.bff VRTScps
# installp -a -d VRTSvcsag.bff VRTSvcsag
# installp -a -d VRTSvcsea.bff VRTSvcsea
# installp -a -d VRTSacclib.bff VRTSacclib
# installp -a -d VRTSsfmh.bff VRTSsfmh
# installp -a -d VRTSsvbs.bff VRTSvbs
# installp -a -d VRTSvcswiz.bff VRTSvcswiz
# installp -a -d VRTSsfcpi62.bff
```

See "Symantec Cluster Server installation filesets" on page 495.

Adding a license key for a manual installation

After you have installed all filesets on each cluster node, use the `vxlicinst` command to add the VCS license key on each system:

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

Setting or changing the product level for keyless licensing

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

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

http://go.symantec.com/vom

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

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

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

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

To set or change the product level

1. Change your current working directory:

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

2. View the current setting for the product level.

   ```
   # vxkeyless -v display
   ```

3. View the possible settings for the product level.

   ```
   # vxkeyless displayall
   ```

4. Set the desired product level.

   ```
   # vxkeyless set prod_levels
   ```

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

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

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

See "Installing Symantec product license keys" on page 63.

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.

Checking licensing information on the system for a manual installation

Use the `vxlicrep` utility to display information about all Symantec licenses on a system. For example, enter:

```
# vxlicrep
```

From the output, you can determine the following:

- The license key
- The type of license
- The product for which it applies
- Its expiration date, if one exists
  Demo keys have expiration dates, while permanent keys and site keys do not.

Replacing a VCS demo license with a permanent license for manual installations

When a VCS demo key license expires, you can replace it with a permanent license using the `vxlicinst` program.

See "Checking licensing information on the system" on page 152.

Copying the installation guide to each node

After you install VCS, Symantec recommends that you copy the PDF version of this guide from the installation disc to the /opt/VRTS/docs directory on each node to make it available for reference. The PDF is located at `docs/cluster_server/vcs_install_version_platform.pdf` where *version* is the release version and *platform* is the name of the operating system.

Installing VCS using NIM and the installer

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

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

In the following samples, the LPP resource uses LPP-7100-up2date and its relevant SPOT resource is spot-7100-up2date.

Preparing the installation bundle on the NIM server

You need to prepare the installation bundle on the NIM server before you use NIM to install VCS filesets. The following actions are executed on the NIM server.

Note: Make sure that the appropriate NIM LPP_SOURCE and SPOT resources are present on the NIM server.

To prepare the installation bundle

1. Insert and mount the installation media.
2. Choose an LPP source:

```
# lsnim |grep -i lpp_source
LPP-7100-up2date resources lpp_source
```

3. Navigate to the product directory on the installation media and run the `installvcs` command to prepare the bundle resource:

```
# ./installvcs -nim LPP-7100-up2date
```

The installation program copies the necessary filesets and patches to the LPP resource directory.

4 Enter a name for the bundle, for example *VCS62*.

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

```
# lsnim -l VCS62
VCS62:
class = resources
type = installp_bundle
Rstate = ready for use
prev_state = unavailable for use
location = /opt/VRTS/nim/VCS62.bundle
alloc_count = 0
server = master
```

Installing VCS on the NIM client using SMIT on the NIM server

You can install VCS on the NIM client using the SMIT tool on the NIM server.

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

To install VCS

1 On the NIM server, start SMIT.

```
# smitty nim
```

2 In the menu, select **Perform NIM Software Installation and Maintenance Tasks**.

3 In the menu, select **Install and Update Software**.

4 In the menu, select **Install Software Bundle**.

5 Select the systems from the list on which to install the software bundle.

6 In the menu, select the LPP_SOURCE. In this example, specify **LPP-7100-up2date**.

7 In the menu, select the bundle, for example, **VCS62**.

8 For the `installp` flags, specify that the ACCEPT new license agreements flag has a **yes** value.

9 Press the Enter key to start the installation. Note that it may take some time to finish.

10 After the installation completes, configure VCS.

For instructions, see the chapter *Configuring VCS* in this document.

Installing VCS and the operating system on the NIM client using SMIT

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

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

To install VCS and the operating system

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

   ```
   # smitty nim_bosinst
   ```

2. In the menu, select the standalone target.
3. In the menu, select **spot - Install a copy of a SPOT resource**.
4. In the menu, select the spot resource **spot-7100-up2date**.
5. In the menu, select the LPP_SOURCE. In this example, select **LPP-7100-up2date**.
6. In the menu, select the following options:
   - For the ACCEPT new license agreements option, specify **yes**.
   - For the Additional Bundles to Install option, specify **VCS62**.
7. For the `installp` flags, specify that the ACCEPT new license agreements flag has a **yes** value.
8. After the installation completes, configure VCS.

   For instructions, see the chapter *Configuring VCS* in this document.

Chapter 17

Manually configuring VCS

This chapter includes the following topics:

About configuring VCS manually

This section describes the procedures to manually configure VCS.

Note: For manually configuring VCS in single node mode, you can skip steps about configuring LLT manually and configuring GAB manually.

Configuring LLT manually

VCS uses the Low Latency Transport (LLT) protocol for all cluster communications as a high-performance, low-latency replacement for the IP stack. LLT has two major functions.

It handles the following tasks:

- Traffic distribution
- Heartbeat traffic

To configure LLT over Ethernet, perform the following steps on each node in the cluster:

- Set up the file /etc/llthosts.
  See "Setting up /etc/llthosts for a manual installation" on page 251.
- Set up the file /etc/lltab.
  See "Setting up /etc/lltab for a manual installation" on page 251.
- Edit the following file on each node in the cluster to change the values of the LLT_START and the LLT_STOP environment variables to 1:
  /etc/default/llt

You can also configure LLT over UDP.

See "Using the UDP layer for LLT" on page 529.

Setting up /etc/llthosts for a manual installation

The file llthosts(4) is a database. It contains one entry per system that links the LLT system ID (in the first column) with the LLT host name. You must ensure that contents of this file are identical on all the nodes in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.

Use vi or another editor, to create the file /etc/llthosts that contains the entries that resemble:

```
0 sys1
1 sys2
```

Setting up /etc/lltab for a manual installation

The /etc/lltab file must specify the system's ID number (or its node name), its cluster ID, and the network links that correspond to the system. In addition, the file can contain other directives. Refer also to the sample lltab file in /opt/VRTSllt.

See "About LLT directives in /etc/lltab file" on page 252.

Use vi or another editor to create the file /etc/llttab that contains the entries that resemble:

```
set-node sys1
set-cluster 2
link en1 /dev/dlpi/en:1 - ether - -
link en2 /dev/dlpi/en:2 - ether - -
```

The first line must identify the system where the file exists. In the example, the value for `set-node` can be: sys1 or 0. The next line, beginning with the `set-cluster` command, identifies the cluster number, which must be a unique number when more than one cluster is configured on the same physical network connection. The next two lines, beginning with the `link` command, identify the two private network cards that the LLT protocol uses. The order of directives must be the same as in the sample `llttab` file in `/opt/VRTSllt`.

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. For example:

Use vi or another editor to create the file /etc/llttab that contains the entries that resemble:

```
set-node sys1
set-cluster 2
link en1 /dev/dlpi/en:1 - ether - -
link en2 /dev/dlpi/en:2 - ether - -
link-lowpri en3 /dev/dlpi/en:3 - ether - -
```

See “Setting up the private network” on page 67.

About LLT directives in /etc/llttab file

Table 17-1 lists the LLT directives in /etc/llttab file for LLT over Ethernet.

Table 17-1 LLT directives

Directive	Description
`set-node`	Assigns the system ID or symbolic name. The system ID number must be unique for each system in the cluster, and must be in the range 0-63. The symbolic name corresponds to the system ID, which is in /etc/llthosts file. Note that LLT fails to operate if any systems share the same ID.

Table 17-1 LLT directives *(continued)*

Directive	Description
`set-cluster`	Assigns a unique cluster number. Use this directive when more than one cluster is configured on the same physical network connection. LLT uses a default cluster number of zero.
`link`	Attaches LLT to a network interface. At least one link is required, and up to eight are supported. LLT distributes network traffic evenly across all available network connections unless you mark the link as low-priority using the link-lowpri directive or you configured LLT to use destination-based load balancing. The first argument to link is a user-defined tag shown in the `lltstat(1M)` output to identify the link. It may also be used in `llttab` to set optional static MAC addresses. The second argument to link is the device name of the network interface. Its format is device_name:device_instance_number. The remaining four arguments to link are defaults; these arguments should be modified only in advanced configurations. There should be one link directive for each network interface. LLT uses an unregistered Ethernet SAP of 0xcafe. If the SAP is unacceptable, refer to the `llttab(4)` manual page for information on how to customize SAP. Note that IP addresses do not need to be assigned to the network device; LLT does not use IP addresses in LLT over Ethernet mode.
`link-lowpri`	Use this directive in place of `link` for public network interfaces. This directive prevents VCS communication on the public network until the network is the last link, and reduces the rate of heartbeat broadcasts. If you use private NICs with different speed, use "link-lowpri" directive in place of "link" for all links with lower speed. Use the "link" directive only for the private NIC with higher speed to enhance LLT performance. LLT uses low-priority network links for VCS communication only when other links fail.

For more information about the LLT directives, refer to the `llttab`(4) manual page.

Additional considerations for LLT for a manual installation

You must attach each network interface that is configured for LLT to a separate and distinct physical network.

Configuring GAB manually

VCS uses the Group Membership Services/Atomic Broadcast (GAB) protocol for cluster membership and reliable cluster communications. GAB has two major functions.

It handles the following tasks:

- Cluster membership
- Cluster communications

To configure GAB

1 Set up an /etc/gabtab configuration file on each node in the cluster using vi or another editor. The following example shows an /etc/gabtab file:

```
/sbin/gabconfig -c -nN
```

Where the `-c` option configures the driver for use. The `-n`N option specifies that the cluster is not formed until at least N systems are ready to form the cluster. Symantec recommends that you set N to be the total number of systems in the cluster.

Warning: Symantec does not recommend the use of the `-c -x` option or `-x` option for `/sbin/gabconfig`. Using `-c -x` or `-x` can lead to a split-brain condition.

---

2 Edit the following file on each node in the cluster to change the values of the GAB_START and the GAB_STOP environment variables to 1:

/etc/default/gab

Configuring VCS manually

VCS configuration requires the types.cf and main.cf files on each system in the cluster. Both of the files are in the /etc/VRTSvcs/conf/config directory.

main.cf file — The main.cf configuration file requires the following minimum essential elements:

- An "include" statement that specifies the file, types.cf, which defines the VCS bundled agent resource type definitions.
- The name of the cluster.
- The name of the systems that make up the cluster.

types.cf file	Note that the "include" statement in main.cf refers to the types.cf file. This text file describes the VCS bundled agent resource type definitions. During new installations, the types.cf file is automatically copied in to the /etc/VRTSvcs/conf/config directory.

When you manually install VCS, the file /etc/VRTSvcs/conf/config/main.cf contains only the line:

```
include "types.cf"
```

For a full description of the main.cf file, and how to edit and verify it, refer to the *Symantec Cluster Server Administrator's Guide*.

To configure VCS manually

1. Log on as superuser, and move to the directory that contains the configuration file:

   ```
   # cd /etc/VRTSvcs/conf/config
   ```

2. Use vi or another text editor to edit the main.cf file, defining your cluster name and system names. Refer to the following example.

   An example main.cf for a two-node cluster:

   ```
   include "types.cf"
   cluster VCSCluster2 ( )
   system sys1 ( )
   system sys2 ( )
   ```

   An example main.cf for a single-node cluster:

   ```
   include "types.cf"
   cluster VCSCluster1 ( )
   system sn1 ( )
   ```

3. Save and close the main.cf file.

Configuring the cluster UUID when creating a cluster manually

You need to configure the cluster UUID when you manually create a cluster.

To configure the cluster UUID when you create a cluster manually

- On one node in the cluster, perform the following command to populate the cluster UUID on each node in the cluster.

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -configure nodeA
        nodeB ... nodeN
```

Where *nodeA*, *nodeB*, through *nodeN* are the names of the cluster nodes.

Configuring VCS in single node mode

In addition to the steps mentioned in the manual configuration section, complete the following steps to configure VCS in single node mode.

See "Configuring VCS manually" on page 254.

To configure VCS in single node mode

1. Edit the following file to change the value of the ONENODE environment variable to **yes**.

   ```
   /etc/default/vcs
   ```

2. If the single node is intended only to manage applications, you can disable LLT, GAB, I/O fencing kernel modules.

   Note: Disabling VCS kernel modules means that you cannot make the applications highly available across multiple nodes.

See "Disabling LLT, GAB, and I/O fencing on a single node cluster" on page 256.

See "Enabling LLT, GAB, and I/O fencing" on page 259.

Disabling LLT, GAB, and I/O fencing on a single node cluster

This section discusses how to disable kernel modules on a single node VCS cluster.

Typically, LLT, GAB, and I/O fencing kernel modules are loaded on a node when you install VCS. However, you can disable LLT, GAB, and I/O fencing modules if you do not require high availability for the applications. You can continue to manage applications on the single node and use the application restart capabilities of VCS.

If you later decide to extend the cluster to multiple nodes, you can enable these modules and make the applications highly available across multiple nodes.

Note: If VCS engine hangs on the single node cluster with GAB disabled, GAB cannot detect the hang state and cannot take action to restart VCS. For such a condition, you need to detect that VCS engine has hung and take corrective action. For more information, refer to the 'About GAB client process failure' section in the *Symantec Cluster Server Administrator's Guide*.

See "Disabling LLT, GAB, and I/O fencing" on page 257.

Disabling LLT, GAB, and I/O fencing

Complete the following procedures to disable the kernel modules.

To disable I/O fencing

1 Edit the following the file to set the value of `VXFEN_START` to **0**.

```
/etc/default/vxfen
```

2 Stop the `vxfen` module.

```
# /etc/methods/vxfenext -stop -dvxfen > /dev/null 2>&1
```

3 Remove the following Object Disk Manager (ODM) entries.

```
# odmdelete -q"name=vxfend" -o CuAt >/dev/null 2>&1
# odmdelete -q"name=vxfend" -o CuDv >/dev/null 2>&1
# odmdelete -q"uniquetype='vxfend/node/vxfend'" -o PdAt \
>/dev/null 2>&1
# odmdelete -q"uniquetype='vxfend/node/vxfend'" -o PdDv \
>/dev/null 2>&1
# odmdelete -q"rule='/etc/methods/vxfenext -define -dvxfend'" \
-o Config_Rules >/dev/null 2>&1
# odmdelete -q"name=vxfen" -o CuAt >/dev/null 2>&1
# odmdelete -q"name=vxfen" -o CuDv >/dev/null 2>&1
# odmdelete -q"uniquetype='vxfen/node/vxfen'" -o PdAt \
>/dev/null 2>&1
# odmdelete -q"uniquetype='vxfen/node/vxfen'" -o PdDv \
>/dev/null 2>&1
# odmdelete -q"rule='/etc/methods/vxfenext -define -dvxfen'"\
-o Config_Rules >/dev/null 2>&1
```

To disable GAB

1 Edit the following the file to set the value of `GAB_START` to **0**.

```
/etc/default/gab
```

2 Unload the GAB driver.

```
# /usr/sbin/slibclean

# /etc/methods/gabkext -stop
```

3 Remove the following Object Disk Manager (ODM) entries.

```
# odmdelete -q"name=gab" -o CuAt > /dev/null 2>&1

# odmdelete -q"name=gab" -o CuDv > /dev/null 2>&1

# odmdelete -q"uniquetype='gab/node/gab'" -o PdDv \

> /dev/null 2>&1

# odmdelete -q"uniquetype='gab/node/gab'" -o PdAt \

> /dev/null 2>&1

# odmdelete -q"rule='/etc/methods/gabkext -define'" \

-o Config_Rules > /dev/null 2>&1
```

To disable LLT

1 Unload the LLT driver. You can ignore any errors that are displayed as a result of the command.

```
# /usr/sbin/strload -u -d /usr/lib/drivers/pse/llt
```

2 Remove the following Object Disk Manager (ODM) entry.

```
# odmdelete -q"rule='/etc/methods/loadllt'" -o Config_Rules
```

3 Edit the following file to set the value of `LLT_START` to **0**.

```
/etc/default/llt
```

Enabling LLT, GAB, and I/O fencing

Complete the following procedures to enable the kernel modules.

To enable LLT:

1 In the `/etc/default/llt` file, ensure LLT_START=1.

2 Create `llt.odm` in any directory with the following text:

```
Config_Rules:
     phase = 2
     seq = 21
     rule = "/etc/methods/loadllt"
```

3 Run the following commands with the file path.

```
# odmadd $DIR/llt.odm
# /etc/methods/loadllt
```

To enable GAB:

1 In the `/etc/default/gab` file, ensure GAB_START=1.

2 Run the following commands:

```
# /usr/sbin/slibclean
# /usr/bin/odmadd /usr/lib/methods/gab.odm
# /etc/methods/gabkext -start
```

To enable I/O fencing:

1 In the `/etc/default/vxfen` file, ensure VXFEN_START=1.

2 Run the following commands:

```
# /usr/sbin/slibclean
# /usr/bin/odmadd /usr/lib/methods/vxfen.odm
# /usr/bin/odmadd /usr/lib/methods/vxfend.odm
# /etc/methods/vxfenext -start -dvxfen
# /etc/methods/vxfenext -start -dvxfend
```

3 Reboot the nodes.

After enabling kernel modules, you can run the `init` scripts for each kernel module to start LLT, GAB, and I/O fencing

Starting LLT, GAB, and VCS after manual configuration

After you have configured LLT, GAB, and VCS, use the following procedures to start LLT, GAB, and VCS.

To start LLT

1 On each node, run the following command to start LLT:

```
# /etc/init.d/llt.rc start
```

If LLT is configured correctly on each node, the console output resembles:

```
Loading LLT Driver...
Starting LLT:
LLT: loading module...
Loaded   kernel_version on kernel kernel_version
LLT: configuring module
where, kernel_version is the kernel version
of the Linux operating system
```

2 On each node, run the following command to verify that LLT is running:

```
# /sbin/lltconfig
LLT is running
```

To start GAB

1 On each node, run the following command to start GAB:

```
# /etc/init.d/gab.rc start
```

If GAB is configured correctly on each node, the console output resembles:

```
GAB: Starting
GAB: Starting Done
```

2 On each node, run the following command to verify that GAB is running:

```
# /sbin/gabconfig -a
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
```

To start VCS

- On each node, type:

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

See "Verifying the cluster" on page 433.

About configuring cluster using VCS Cluster Configuration wizard

Consider the following before configuring a cluster using VCS Cluster Configuration wizard

- The VCS Cluster Configuration wizard allows you to configure a VCS cluster and add a node to the cluster.
  See "Configuring a cluster by using the VCS cluster configuration wizard" on page 265.
- Symantec recommends that you first configure application monitoring using the wizard before using VCS commands to add additional components or modify the existing configuration. Apart from configuring application availability, the wizard also sets up the other components required for successful application monitoring.

Before configuring a VCS cluster using the VCS Cluster Configuration wizard

Ensure that you complete the following tasks before launching the VCS Cluster Configuration wizard to configure a VCS cluster:

- Install Symantec Cluster Server (VCS) on the system on which you want to configure the VCS cluster.
- You must have the following user privileges when you attempt to configure the VCS cluster:
  - Configure Application Monitoring (Admin) privileges when you launch the wizard from the vSphere client.
  - Admin role privileges if you launch the wizard through VOM
- Install the application and the associated components that you want to monitor on the system.
- If you have configured a firewall, ensure that your firewall settings allow access to ports used by Symantec Cluster Server installer, wizards, and services. Verify that the following ports are not blocked by the firewall:

VMware environment	443, 5634, 14152, and 14153
Physical environment	5634, 14161, 14162, 14163, and 14164 At least one port from 14161, 14162, 14163, and 14164 must be open.

- You must not select bonded interfaces for cluster communication. A bonded interface is a logical NIC, formed by grouping several physical NICs together. All NICs in a bond have an identical MAC address, due to which you may experience the following issues:
  - Single Sign On (SSO) configuration failure.
  - The wizard may fail to discover the specified network adapters.
  - The wizard may fail to discover or validate the specified system name.
- The host name of the system must be resolvable through the DNS server or locally, using /etc/hosts file entries.

Launching the VCS Cluster Configuration wizard

You must launch the VCS Cluster Configuration wizard from the system where the disk residing on the shared datastore is attached.

You can launch the VCS Cluster Configuration wizard from:

- VMware vSphere Client
  See Launching the VCS Cluster Configuration wizard from VMware vSphere Client.
- A browser window
  See Launching the VCS Cluster Configuration wizard from a browser window.

Launching the VCS Cluster Configuration wizard from VMware vSphere Client

To launch the wizard from the VMware vSphere Client:

1. Launch the VMware vSphere Client and connect to the VMware vCenter Server that hosts the virtual machine.
2. From the vSphere Client's Inventory view in the left pane, select the virtual machine where you want to configure the VCS cluster.
3. Select the **Symantec High Availability** tab.

   The tab displays various menus based on the what is configured on the system. The menu options launch the appropriate wizard panel based on the tasks that you choose to perform.

Launching the VCS Cluster Configuration wizard from a browser window

You can launch the VCS Cluster Configuration wizard from the Symantec High Availability view.

1. Open a browser window and enter the following URL:

   https://<IP_or_HostName>:5634/vcs/admin/application_health.html

   where <IP_or_HostName> is the IP address or host name of the system on which you want to configure the cluster.
2. Click the **Configure cluster** link on the Symantec High Availability view page to launch the wizard.

Note: At various stages of cluster configuration, the Symantec High Availability view offers different configuration options. These options launch appropriate wizard panels based on the tasks that you choose to perform.

See "Configuring a cluster by using the VCS cluster configuration wizard" on page 265.

See "Adding a system to a VCS cluster" on page 268.

Refer to the *Administering application monitoring from the Symantec High Availability view* section in *Symantec Cluster Server Administrator's Guide* for more information on the configurations possible from the Symantec High Availability view.

Configuring a cluster by using the VCS cluster configuration wizard

Perform the following steps to configure a Symantec Cluster Server (VCS) cluster by using the VCS Cluster Configuration wizard.

To configure a VCS cluster

1. Access the Symantec High Availability view (for any system belonging the required cluster).

   See "Launching the VCS Cluster Configuration wizard" on page 263.

2. Review the information on the Welcome panel and click **Next**.

   The Configuration Inputs panel appears.

   The local system is by default selected as a cluster system.

3. If you do not want to add more systems to the cluster, skip this step. You can add systems later using the same wizard.

   To add a system to the cluster, click **Add System**.

   In the Add System dialog box, specify the following details for the system that you want to add to the VCS cluster and click **OK**.

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
User name	Specify the user account for the system. Typically, this is the root user. The root user should have the necessary privileges.
Password	Specify the password for the user account you specified.
Use the specified user account on all systems	Select this check box to use the specified user account on all the cluster systems that have the same user name and password.

4. On the Configuration Inputs panel, do one of the following actions:
   - To add another system to the cluster, click Add System and repeat step 3.
   - To modify the specified User name or Password for a cluster system, use the edit icon.
   - Click **Next**

5 If you do not want to modify the security settings for the cluster, click **Next**, and proceed to step 7.

By default, the wizard configures single sign-on for secure cluster communication. If you want to modify the security settings for the cluster, click **Advanced Settings**.

6 In the Advanced settings dialog box, specify the following details and click **OK**.

Use Single Sign-on	Select to configure single sign-on using VCS Authentication Service for cluster communication. This option is enabled by default.
Use VCS user privileges	Select to configure a user with administrative privileges to the cluster. Specify the username and password and click OK.

7 On the Network Details panel, select the type of network protocol to configure the VCS cluster network links (Low Latency Transport or LLT module), and then specify the adapters for network communication.

The wizard configures the VCS cluster communication links using these adapters. You must select a minimum of two adapters per cluster system.

Note: By default, the LLT links are configured over Ethernet.

Select **Use MAC address for cluster communication (LLT over Ethernet)** or select **Use IP address for cluster communication (LLT over UDP)**, and specify the following details for each cluster system.

- To configure LLT over Ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- To configure LLT over UDP, select the type of IP protocol (IPv4 or IPv6), and then specify the required details for each communication link.

Network Adapter	Select a network adapter for the communication links. You must select a different network adapter for each communication link.
IP Address	Displays the IP address.

Port	Specify a unique port number for each link. For IPv4 and IPv6, the port range is from 49152 to 65535. A specified port for a link is used for all the cluster systems on that link.
Subnet mask (IPv4)	Displays the subnet mask details.
Prefix (IPv6)	Displays the prefix details.

By default, one of the links is configured as a low-priority link on a public network interface. The second link is configured as a high-priority link.

To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

Note: Symantec recommends that you configure one of the links on a public network interface. You can assign the link on the public network interface as a low-priority link for minimal VCS cluster communication over the link.

---

8 On the Configuration Summary panel, specify a cluster name and unique cluster ID and then click **Validate**.

Note: If multiple clusters exist in your network, the wizard validates if the specified cluster ID is a unique cluster ID among all clusters accessible from the current system. Among clusters that are not accessible from the current system, you must ensure that the cluster ID you specified is unique

---

9 Review the VCS Cluster Configuration Details and then click **Next** to proceed with the configuration

10 On the Implementation panel, the wizard creates the VCS cluster.

The wizard displays the status of the configuration task. After the configuration is complete, click **Next**.

If the configuration task fails, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure and run the wizard again to configure the VCS cluster.

11 On the Finish panel, click **Finish** to complete the wizard workflow.

This completes the VCS cluster configuration.

Adding a system to a VCS cluster

Perform the following steps to add a system to a Symantec Cluster Server (VCS) cluster by using the VCS Cluster Configuration wizard.

The system from where you launch the wizard must be part of the cluster to which you want to add a new system.

To add a system to a VCS cluster

1. Access the Symantec High Availability view (for any system belonging to the required cluster).

   See "Launching the VCS Cluster Configuration wizard" on page 263.

2. Click **Actions** > **Add System to VCS Cluster**.

   The VCS Cluster Configuration Wizard is launched.

3. Review the information on the Welcome panel and click **Next**.

   The Configuration Inputs panel appears, along with the cluster name, and a table of existing cluster systems.

4. To add a system to the cluster, click **Add System**.

5. In the Add System dialog box, specify the following details for the system that you want to add to the VCS cluster and click **OK**.

System Name or IP address	Specify the name or IP address of the system that you want to add to the VCS cluster.
User name	Specify the user account for the system. Typically, this is the root user. The root user should have the necessary privileges.
Password	Specify the password for the user account you specified.
Use the specified user account on all systems	Select this check box to use the specified user account on all the cluster systems that have the same user name and password.

6. On the Configuration Inputs panel, do one of the following actions:

   - To add another system to the cluster, click **Add System** and repeat step 4.
   - To modify the User name or Password for a cluster system, use the edit icon.

- Click **Next**

7 On the Network Details panel, specify the adapters for network communication (Low Latency Transport or LLT module of VCS) for the system. The wizard configures the VCS cluster communication links using these adapters. You must select a minimum of two adapters.

Note: You cannot modify the existing type of cluster communication (LLT over Ethernet or LLT over UDP).

- If the existing cluster uses LLT over Ethernet, select the adapter for each network communication link. You must select a different network adapter for each communication link.
- If the existing cluster uses LLT over UDP, select the type of IP protocol (IPv4 or IPv6), and then specify the required details for each communication link.

Network Adapter	Select a network adapter for the communication links. You must select a different network adapter for each communication link.
IP Address	Displays the IP address.
Port	Specify a unique port number for each link. For IPv4 and IPv6, the port range is from 49152 to 65535. A specified port for a link is used for all the cluster systems on that link.
Subnet mask (IPv4)	Displays the subnet mask details.
Prefix (IPv6)	Displays the prefix details.

By default, one of the links is configured as a low-priority link on a public network interface. The other link is configured as a high-priority link.

To change a high-priority link to a low-priority link, click **Modify**. In the Modify low-priority link dialog box, select the link and click **OK**.

Note: Symantec recommends that you configure one of the links on a public network interface. You can assign the link on the public network interface as a low-priority link for minimal VCS cluster communication over the link.

8 On the Configuration Summary panel, review the VCS Cluster Configuration Details.

9 On the Implementation panel, the wizard creates the VCS cluster.

The wizard displays the status of the configuration task. After the configuration is complete, click **Next**.

If the configuration task fails, click **Diagnostic information** to check the details of the failure. Rectify the cause of the failure and run the wizard again to add the required system to the VCS cluster.

10 On the Finish panel, click **Finish** to complete the wizard workflow.

Modifying the VCS configuration

After the successful installation of VCS, you can modify the configuration of VCS using several methods. You can dynamically modify the configuration from the command line, Veritas Operations Manager, or the Cluster Manager (Java Console). For information on management tools, refer to the *Symantec Cluster Server Administrator's Guide*.

You can also edit the main.cf file directly. For information on the structure of the main.cf file, refer to the *Symantec Cluster Server Administrator's Guide*.

Configuring the ClusterService group

When you have installed VCS, and verified that LLT, GAB, and VCS work, you can create a service group to include the optional features. These features include the VCS notification components and the Global Cluster option. If you manually added VCS to your cluster systems, you must manually create the ClusterService group. You can refer to the configuration examples of a system with a ClusterService group. See the *Symantec Cluster Server Administrator's Guide* for more information.

See “Sample main.cf file for VCS clusters” on page 512.

Chapter 18

Manually configuring the clusters for data integrity

This chapter includes the following topics:

- Setting up disk-based I/O fencing manually
- Setting up server-based I/O fencing manually
- Setting up non-SCSI-3 fencing in virtual environments manually
- Setting up majority-based I/O fencing manually

Setting up disk-based I/O fencing manually

Table 18-1 lists the tasks that are involved in setting up I/O fencing.

Table 18-1

Task	Reference
Initializing disks as VxVM disks	See "Initializing disks as VxVM disks" on page 154.
Identifying disks to use as coordinator disks	See "Identifying disks to use as coordinator disks" on page 272.
Checking shared disks for I/O fencing	See "Checking shared disks for I/O fencing" on page 160.
Setting up coordinator disk groups	See "Setting up coordinator disk groups" on page 272.
Creating I/O fencing configuration files	See "Creating I/O fencing configuration files" on page 273.

Table 18-1 *(continued)*

Task	Reference
Modifying VCS configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 274.
Configuring CoordPoint agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 289.
Verifying I/O fencing configuration	See "Verifying I/O fencing configuration" on page 276.

Identifying disks to use as coordinator disks

Make sure you initialized disks as VxVM disks.

See "Initializing disks as VxVM disks" on page 154.

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 160.

Setting up coordinator disk groups

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

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

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

The following example procedure assumes that the disks have the device names hdisk10, hdisk11, and hdisk12.

To create the vxfencoorddg disk group

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

```
# vxdg init vxfencoorddg hdisk10 hdisk11 hdisk12
```

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

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

3 Deport the coordinator disk group:

```
# vxdg deport vxfencoorddg
```

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

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

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

```
# vxdg deport vxfencoorddg
```

Creating I/O fencing configuration files

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

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

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

1 On each nodes, type:

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

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

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

2 On all cluster nodes specify the use of DMP disk policy in the `/etc/vxfenmode` file.

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

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

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

4 Ensure that you edit the following file on each node in the cluster to change the values of the VXFEN_START and the VXFEN_STOP environment variables to 1:

/etc/default/vxfen

Modifying VCS configuration to use I/O fencing

After you add coordination points and configure I/O fencing, add the UseFence = SCSI3 cluster attribute to the VCS configuration file /etc/VRTSvcs/conf/config/main.cf.

If you reset this attribute to UseFence = None, VCS does not make use of I/O fencing abilities while failing over service groups. However, I/O fencing needs to be disabled separately.

To modify VCS configuration to enable I/O fencing

1 Save the existing configuration:

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

2 Stop VCS on all nodes:

```
# hastop -all
```

3 To ensure High Availability has stopped cleanly, run `gabconfig -a`.

In the output of the commans, check that Port h is not present.

4 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

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

5 Make a backup of the main.cf file on all the nodes:

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

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

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

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

7 Save and close the file.

8 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

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

9 Using rcp or another utility, copy the VCS configuration file from a node (for example, sys1) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp sys1:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

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

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

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

- Start VCS on the node where main.cf is modified.

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

- Start VCS on all other nodes once VCS on first node reaches RUNNING state.

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

Verifying I/O fencing configuration

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

To verify I/O fencing configuration

1 On one of the nodes, type:

```
# vxfenadm -d
```

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

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

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

   * 0 (sys1)
   1 (sys2)

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

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

```
# vxfenconfig -l
```

Setting up server-based I/O fencing manually

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

Table 18-2 Tasks to set up server-based I/O fencing manually

Task	Reference
Preparing the CP servers for use by the VCS cluster	See "Preparing the CP servers manually for use by the VCS cluster" on page 277.

Table 18-2 Tasks to set up server-based I/O fencing manually *(continued)*

Task	Reference
Generating the client key and certificates on the client nodes manually	See "Generating the client key and certificates manually on the client nodes " on page 280.
Modifying I/O fencing configuration files to configure server-based I/O fencing	See "Configuring server-based fencing on the VCS cluster manually" on page 282.
Modifying VCS configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 274.
Configuring Coordination Point agent to monitor coordination points	See "Configuring CoordPoint agent to monitor coordination points" on page 289.
Verifying the server-based I/O fencing configuration	See "Verifying server-based I/O fencing configuration" on page 290.

Preparing the CP servers manually for use by the VCS cluster

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

Table 18-3 displays the sample values used in this procedure.

Table 18-3 Sample values in procedure

CP server configuration component	Sample name
CP server	cps1
Node #1 - VCS cluster	sys1
Node #2 - VCS cluster	sys2
Cluster name	clus1
Cluster UUID	{f0735332-1dd1-11b2}

To manually configure CP servers for use by the VCS cluster

1 Determine the cluster name and uuid on the VCS cluster.

For example, issue the following commands on one of the VCS 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 VCS 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 Cluster Server Administrator's Guide*.

3 Add the VCS 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 disabled, then add the user name "cpsclient@hostname" to the server.

5 Add the users to the CP server.

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

```
# cpsadm -s cps1.symantecexample.com -a add_user -e\
cpsclient@hostname\
 -f cps_operator -g vx

User cpsclient@hostname
successfully added
```

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

For example, issue the following command on the CP server (cps1.symantecexample.com) for VCS 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 cpsclient@hostname\
 -f cps_operator -g vx

Cluster successfully added to user
 cpsclient@hostname  privileges.
```

See “Generating the client key and certificates manually on the client nodes ” on page 280.

Generating the client key and certificates manually on the client nodes

The client node that wants to connect to a CP server using HTTPS must have a private key and certificates signed by the Certificate Authority (CA) on the CP server

The client uses its private key and certificates to establish connection with the CP server. The key and the certificate must be present on the node at a predefined location. Each client has one client certificate and one CA certificate for every CP server, so, the certificate files must follow a specific naming convention. Distinct certificate names help the `cpsadm` command to identify which certificates have to be used when a client node connects to a specific CP server.

The certificate names must be as follows: ca_*cps-vip*.crt and client _*cps-vip*.crt

Where, *cps-vip* is the VIP or FQHN of the CP server listed in the `/etc/vxfenmode` file. For example, for a sample VIP, *192.168.1.201*, the corresponding certificate name is ca_*192.168.1.201*.

To manually set up certificates on the client node

1 Create the directory to store certificates.

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

Note: Since the openssl utility might not be available on client nodes, Symantec recommends that you access the CP server using SSH to generate the client keys or certificates on the CP server and copy the certificates to each of the nodes.

2 Generate the private key for the client node.

```
# /usr/bin/openssl genrsa -out client_private.key 2048
```

3 Generate the client CSR for the cluster. CN is the UUID of the client's cluster.

```
# /usr/bin/openssl req -new -key client_private.key\
-subj '/C=countryname/L=localityname/OU=COMPANY/CN=CLUS_UUID'\
-out client_192.168.1.201.csr
```

Where, *countryname* is the country code, *localityname* is the city, *COMPANY* is the name of the company, and *CLUS_UUID* is the certificate name.

4 Generate the client certificate by using the CA key and the CA certificate. Run this command from the CP server.

```
# /usr/bin/openssl x509 -req -days days -in
client_192.168.1.201.csr\
-CA /var/VRTScps/security/certs/ca.crt -CAkey\
/var/VRTScps/security/keys/ca.key -set_serial 01 -out
client_192.168.10.1.crt
```

Where, *days* is the days you want the certificate to remain valid, *192.168.1.201* is the VIP or FQHN of the CP server.

5 Copy the client key, client certificate, and CA certificate to each of the client nodes at the following location.

Copy the client key at `/var/VRTSvxfen/security/keys/client_private.key`. The client is common for all the client nodes and hence you need to generate it only once.

Copy the client certificate at `/var/VRTSvxfen/security/certs/client_`*`192.168.1.201`*`.crt`.

Copy the CA certificate at `/var/VRTSvxfen/security/certs/ca_`*`192.168.1.201`*`.crt`

Note: Copy the certificates and the key to all the nodes at the locations that are listed in this step.

6 If the client nodes need to access the CP server using the FQHN and or the host name, make a copy of the certificates you generated and replace the VIP with the FQHN or host name. Make sure that you copy these certificates to all the nodes.

7 Repeat the procedure for every CP server.

8 After you copy the key and certificates to each client node, delete the client keys and client certificates on the CP server.

Configuring server-based fencing on the VCS cluster manually

The configuration process for the client or VCS 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)
- CP server or CP servers
- Coordinator disk group (if applicable to your I/O fencing configuration)
- Set the order of coordination points

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

See "Setting up coordinator disk groups" on page 272.

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

To configure server-based fencing on the VCS cluster manually

1. Use a text editor to edit the following file on each node in the cluster:

   `/etc/default/vxfen`

   You must change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1.

2. Use a text editor to edit the `/etc/vxfenmode` file values to meet your configuration specifications.

   - If your server-based fencing configuration uses a single highly available CP server as its only coordination point, make sure to add the `single_cp=1` entry in the `/etc/vxfenmode` file.
   - If you want the `vxfen` module to use a specific order of coordination points during a network partition scenario, set the `vxfen_honor_cp_order` value to be **1**. By default, the parameter is disabled.

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

   See "Sample vxfenmode file output for server-based fencing" on page 283.

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

   For example:

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

Sample vxfenmode file output for server-based fencing

The following is a sample vxfenmode file for server-based fencing:

```
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
```

```
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

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

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

#
# security parameter is deprecated release 6.1 onwards
# since communication with CP server will always happen
# over HTTPS which is inherently secure. In pre-6.1 releases,
# it was used to configure secure communication to the
# cp server using VxAT (Veritas Authentication Service)
# available options:
# 0 - don't use Veritas Authentication Service for cp server
#   communication
# 1  - use Veritas Authentication Service for cp server
#   communication
security=1

#
# vxfen_honor_cp_order determines the order in which vxfen
# should use the coordination points specified in this file.
#
# available options:
```

```
# 0 - vxfen uses a sorted list of coordination points specified
# in this file,
# the order in which coordination points are specified does not matter.
#   (default)
# 1 - vxfen uses the coordination points in the same order they are
#    specified in this file

# Specify 3 or more odd number of coordination points in this file,
# each one in its own line. They can be all-CP servers,
# all-SCSI-3 compliant coordinator disks, or a combination of
# CP servers and SCSI-3 compliant coordinator disks.
# Please ensure that the CP server coordination points
# are numbered sequentially and in the same order
# on all the cluster nodes.
#
# Coordination Point Server(CPS) is specified as follows:
#
#  cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames
# over different subnets, all of the IPs/names can be specified
# in a comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,
...,[<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#  is the serial number of the CPS as a coordination point; must
#  start with 1.
# <vip>
#  is the virtual IP address of the CPS, must be specified in
#  square brackets ("[]").
# <vhn>
#  is the virtual hostname of the CPS, must be specified in square
#  brackets ("[]").
# <port>
#  is the port number bound to a particular <vip/vhn> of the CPS.
#  It is optional to specify a <port>. However, if specified, it
#  must follow a colon (":") after <vip/vhn>. If not specified, the
#  colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>,
```

```
# a default port can be specified as follows:
#
# port=<default_port>
#
#  Where <default_port> is applicable to all the <vip/vhn>s for
#  which a <port> is not specified. In other words, specifying
#  <port> with a <vip/vhn> overrides the <default_port> for that
#  <vip/vhn>. If the <default_port> is not specified, and there
#  are <vip/vhn>s for which <port> is not specified, then port
#  number 14250 will be used for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[cps1.company.com]
#  cps2=[192.168.0.25]
# cps3=[cps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [cps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
# - if default port 57777 were not specified, port 14250
# would be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
#  vxfendg=<coordinator disk group name>
# Example:
#  vxfendg=vxfencoorddg
#
# Examples of different configurations:
#  1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
```

```
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
# vxfendg=
# Note: The disk group specified in case should have three disks
# cps1=[cps1.company.com]
# cps2=[cps2.company.com]
# cps3=[cps3.company.com]
# port=443
```

Table 18-4 defines the vxfenmode parameters that must be edited.

Table 18-4 vxfenmode file parameters

vxfenmode File Parameter	Description
vxfen_mode	Fencing mode of operation. This parameter must be set to “customized”.
vxfen_mechanism	Fencing mechanism. This parameter defines the mechanism that is used for fencing. If one of the three coordination points is a CP server, then this parameter must be set to “cps”.
scsi3_disk_policy	Configure the vxfen module to use DMP devices, "dmp". **Note:** The configured disk policy is applied on all the nodes.
security	Deprecated from release 6.1 onwards. Security parameter is deprecated release 6.1 onwards as communication between CP servers and application clusters happens over the HTTPS protocol which is inherently secure. In releases prior to 6.1, the security parameter was used to configure secure communication to the CP server using the VxAT (Veritas Authentication Service) options. The options are: ■ 0 - Do not use Veritas Authentication Service for CP server communication ■ 1 - Use Veritas Authentication Service for CP server communication

Table 18-4 vxfenmode file parameters *(continued)*

vxfenmode File Parameter	Description

cps1, cps2, or vxfendg	Coordination point parameters.
	Enter either the virtual IP address or the FQHN (whichever is accessible) of the CP server.
	`cps<number>=[virtual_ip_address/virtual_host_name]:port`
	Where *port* is optional. The default port value is 443.
	If you have configured multiple virtual IP addresses or host names over different subnets, you can specify these as comma-separated values. For example:
	`cps1=[192.168.0.23],[192.168.0.24]:58888, [cps1.company.com]`
	Note: Whenever coordinator disks are used in an I/O fencing configuration, a disk group has to be created (vxfencoorddg) and specified in the /etc/vxfenmode file. Additionally, the customized fencing framework also generates the /etc/vxfentab file which specifies the security setting and the coordination points (all the CP servers and the disks from disk group specified in /etc/vxfenmode file).

port	Default port for the CP server to listen on. If you have not specified port numbers for individual virtual IP addresses or host names, the default port number value that the CP server uses for those individual virtual IP addresses or host names is 443. You can change this default port value using the port parameter.
single_cp	Value 1 for single_cp parameter indicates that the server-based fencing uses a single highly available CP server as its only coordination point. Value 0 for single_cp parameter indicates that the server-based fencing uses at least three coordination points.
vxfen_honor_cp_order	Set the value to 1 for vxfen module to use a specific order of coordination points during a network partition scenario. By default the parameter is disabled. The default value is 0.

Configuring CoordPoint agent to monitor coordination points

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

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

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

To configure CoordPoint agent to monitor coordination points

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

3. Configure the Phantom resource for the vxfen disk group.

```
# haconf -makerw
# hares -add RES_phantom_vxfen Phantom vxfen
# hares -modify RES_phantom_vxfen Enabled 1
# haconf -dump -makero
```

4 Verify the status of the agent on the VCS 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
```

5 Access the engine log to view the agent log. The agent log is written to the engine log.

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

To view the debug logs in the engine log, change the dbg level for that node using the following commands:

```
# haconf -makerw

# hatype -modify Coordpoint LogDbg 10

# haconf -dump -makero
```

The agent log can now be viewed at the following location:

/var/VRTSvcs/log/engine_A.log

Note: The Coordpoint agent is always in the online state when the I/O fencing is configured in the majority or the disabled mode. For both these modes the I/O fencing does not have any coordination points to monitor. Thereby, the Coordpoint agent is always in the online state.

Verifying server-based I/O fencing configuration

Follow the procedure described below to verify your server-based I/O fencing configuration.

To verify the server-based I/O fencing configuration

1 Verify that the I/O fencing configuration was successful by running the `vxfenadm` command. For example, run the following command:

```
# vxfenadm -d
```

Note: For troubleshooting any server-based I/O fencing configuration issues, refer to the *Symantec Cluster Server Administrator's Guide*.

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

```
# vxfenconfig -l
```

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

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

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

1 Configure I/O fencing either in majority-based fencing mode with no coordination points or in server-based fencing mode only with CP servers as coordination points.

See "Setting up server-based I/O fencing manually" on page 276.

See "Setting up majority-based I/O fencing manually " on page 297.

2 Make sure that the VCS cluster is online and check that the fencing mode is customized mode or majority mode.

```
# vxfenadm -d
```

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

```
# haclus -value UseFence
```

4 On each node, edit the /etc/vxenviron file as follows:

```
data_disk_fencing=off
```

5 Enter the following command to change the vxfen_min_delay parameter value:

```
# chdev -l vxfen -P -a vxfen_vxfnd_tmt=25
```

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

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

Refer to the sample /etc/vxfenmode file.

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

- Run the following command:

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

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

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

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

- Make the VCS configuration file writable:

```
# haconf -makerw
```

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

```
# hares -modify <dg_resource> MonitorReservation 0

# hares -modify <dg_resource> Reservation "NONE"
```

- Run the following command to verify the value:

```
# hares -list Type=DiskGroup MonitorReservation!=0

# hares -list Type=DiskGroup Reservation!="NONE"
```

The command should not list any resources.

- Modify the default value of the Reservation attribute at type-level.

```
# haattr -default DiskGroup Reservation "NONE"
```

- Make the VCS configuration file read-only

```
# haconf -dump -makero
```

9 Make sure that the UseFence attribute in the VCS configuration file main.cf is set to SCSI-3.

10 To make these VxFEN changes take effect, stop and restart VxFEN and the dependent modules

- On each node, run the following command to stop VCS:

```
# /etc/init.d/vcs.rc stop
```

- After VCS takes all services offline, run the following command to stop VxFEN:

```
# /etc/init.d/vxfen.rc stop
```

- On each node, run the following commands to restart VxFEN and VCS:

```
# /etc/init.d/vxfen.rc start
# /etc/init.d/vcs.rc start
```

Sample /etc/vxfenmode file for non-SCSI-3 fencing

```
#
# vxfen_mode determines in what mode VCS I/O Fencing should work.
#
# available options:
# scsi3      - use scsi3 persistent reservation disks
# customized - use script based customized fencing
# disabled   - run the driver but don't do any actual fencing
#
vxfen_mode=customized

# vxfen_mechanism determines the mechanism for customized I/O
# fencing that should be used.
#
# available options:
# cps      - use a coordination point server with optional script
#            controlled scsi3 disks
#
```

```
vxfen_mechanism=cps

#
# scsi3_disk_policy determines the way in which I/O fencing
# communicates with the coordination disks. This field is
# required only if customized coordinator disks are being used.
#
# available options:
# dmp - use dynamic multipathing
#
scsi3_disk_policy=dmp

#
# Seconds for which the winning sub cluster waits to allow for the
# losing subcluster to panic & drain I/Os. Useful in the absence of
# SCSI3 based data disk fencing loser_exit_delay=55
#
# Seconds for which vxfend process wait for a customized fencing
# script to complete. Only used with vxfen_mode=customized
# vxfen_script_timeout=25

# security parameter is deprecated release 6.1 onwards since
# communication with CP server will always happen over HTTPS
# which is inherently secure. In pre-6.1 releases, it was used
# to configure secure communication to the cp server using
# VxAT (Veritas Authentication Service) available options:
# 0 - don't use Veritas Authentication Service for cp server
#   communication
# 1  - use Veritas Authentication Service for cp server
#   communication
security=1

#
# vxfen_honor_cp_order determines the order in which vxfen
# should use the coordination points specified in this file.
#
# available options:
# 0 - vxfen uses a sorted list of coordination points specified
# in this file, the order in which coordination points are specified
# does not matter.
#   (default)
# 1 - vxfen uses the coordination points in the same order they are
#    specified in this file
```

```
# Specify 3 or more odd number of coordination points in this file,
# each one in its own line. They can be all-CP servers, all-SCSI-3
# compliant coordinator disks, or a combination of CP servers and
# SCSI-3 compliant coordinator disks.
# Please ensure that the CP server coordination points are
# numbered sequentially and in the same order on all the cluster
# nodes.
#
# Coordination Point Server(CPS) is specified as follows:
#
#  cps<number>=[<vip/vhn>]:<port>
#
# If a CPS supports multiple virtual IPs or virtual hostnames
# over different subnets, all of the IPs/names can be specified
# in a comma separated list as follows:
#
# cps<number>=[<vip_1/vhn_1>]:<port_1>,[<vip_2/vhn_2>]:<port_2>,
# ...,[<vip_n/vhn_n>]:<port_n>
#
# Where,
# <number>
#  is the serial number of the CPS as a coordination point; must
#  start with 1.
# <vip>
#  is the virtual IP address of the CPS, must be specified in
#  square brackets ("[]").
# <vhn>
#  is the virtual hostname of the CPS, must be specified in square
#  brackets ("[]").
# <port>
#  is the port number bound to a particular <vip/vhn> of the CPS.
#  It is optional to specify a <port>. However, if specified, it
#  must follow a colon (":") after <vip/vhn>. If not specified, the
#  colon (":") must not exist after <vip/vhn>.
#
# For all the <vip/vhn>s which do not have a specified <port>,
# a default port can be specified as follows:
#
# port=<default_port>
#
#  Where <default_port> is applicable to all the <vip/vhn>s for which a
#  <port> is not specified. In other words, specifying <port> with a
```

```
#  <vip/vhn> overrides the <default_port> for that <vip/vhn>.
# If the <default_port> is not specified, and there are <vip/vhn>s for
# which <port> is not specified, then port number 14250 will be used
# for such <vip/vhn>s.
#
# Example of specifying CP Servers to be used as coordination points:
# port=57777
# cps1=[192.168.0.23],[192.168.0.24]:58888,[cps1.company.com]
#  cps2=[192.168.0.25]
# cps3=[cps2.company.com]:59999
#
# In the above example,
# - port 58888 will be used for vip [192.168.0.24]
# - port 59999 will be used for vhn [cps2.company.com], and
# - default port 57777 will be used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
# - if default port 57777 were not specified, port 14250 would be
#  used for all remaining <vip/vhn>s:
#    [192.168.0.23]
#    [cps1.company.com]
#    [192.168.0.25]
#
# SCSI-3 compliant coordinator disks are specified as:
#
#  vxfendg=<coordinator disk group name>
# Example:
#  vxfendg=vxfencoorddg
#
# Examples of different configurations:
#  1. All CP server coordination points
# cps1=
# cps2=
# cps3=
#
# 2. A combination of CP server and a disk group having two SCSI-3
# coordinator disks
# cps1=
# vxfendg=
# Note: The disk group specified in this case should have two disks
#
# 3. All SCSI-3 coordinator disks
```

```
# vxfendg=
# Note: The disk group specified in case should have three disks
# cps1=[cps1.company.com]
# cps2=[cps2.company.com]
# cps3=[cps3.company.com]
# port=443
```

Setting up majority-based I/O fencing manually

Table 18-5 lists the tasks that are involved in setting up I/O fencing.

Task	Reference
Creating I/O fencing configuration files	Creating I/O fencing configuration files
Modifying VCS configuration to use I/O fencing	Modifying VCS configuration to use I/O fencing
Verifying I/O fencing configuration	Verifying I/O fencing configuration

Creating I/O fencing configuration files

To update the I/O fencing files and start I/O fencing

1 On all cluster nodes, run the following command

```
# cp /etc/vxfen.d/vxfenmode_majority /etc/vxfenmode
```

2 To check the updated /etc/vxfenmode configuration, enter the following command on one of the nodes.

```
# cat /etc/vxfenmode
```

3 Ensure that you edit the following file on each node in the cluster to change the values of the `VXFEN_START` and the `VXFEN_STOP` environment variables to 1.

```
/etc/default/vxfen
```

Modifying VCS configuration to use I/O fencing

After you configure I/O fencing, add the UseFence = SCSI3 cluster attribute to the VCS configuration file /etc/VRTSvcs/conf/config/main.cf.

If you reset this attribute to UseFence = None, VCS does not make use of I/O fencing abilities while failing over service groups. However, I/O fencing needs to be disabled separately.

To modify VCS configuration to enable I/O fencing

1 Save the existing configuration:

```
# haconf -dump -makero
```

2 Stop VCS on all nodes:

```
# hastop -all
```

3 To ensure High Availability has stopped cleanly, run `gabconfig -a`.

In the output of the commans, check that Port h is not present.

4 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

```
# /etc/init.d/vxfen.rc stop
```

5 Make a backup of the main.cf file on all the nodes:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

6 On one node, use vi or another text editor to edit the main.cf file. To modify the list of cluster attributes, add the UseFence attribute and assign its value as SCSI3.

```
cluster clus1(
UserNames = { admin = "cDRpdxPmHpzS." }
Administrators = { admin }
HacliUserLevel = COMMANDROOT
CounterInterval = 5
UseFence = SCSI3
)
```

For fencing configuration in any mode except the disabled mode, the value of the cluster-level attribute UseFence is set to SCSI3.

7 Save and close the file.

8 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

```
# hacf -verify /etc/VRTSvcs/conf/config
```

9 Using rcp or another utility, copy the VCS configuration file from a node (for example, sys1) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp sys1:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

10 Start the I/O fencing driver and VCS. Perform the following steps on each node:

- Start the I/O fencing driver.
  The vxfen startup script also invokes the `vxfenconfig` command, which configures the vxfen driver.

  ```
  # /etc/init.d/vxfen.rc start
  ```

- Start VCS on the node where main.cf is modified.

  ```
  # /opt/VRTS/bin/hastart
  ```

- Start VCS on all other nodes once VCS on first node reaches RUNNING state.

  ```
  # /opt/VRTS/bin/hastart
  ```

Verifying I/O fencing configuration

Verify from the vxfenadm output that the fencing mode reflects the configuration in the /etc/vxfenmode file.

To verify I/O fencing configuration

- On one of the nodes, type:

```
# vxfenadm -d
```

Output similar to the following appears if the fencing mode is majority:

```
I/O Fencing Cluster Information:
================================

 Fencing Protocol Version: 201
 Fencing Mode: MAJORITY
 Cluster Members:

        * 0 (sys1)
          1 (sys2)

 RFSM State Information:
        node   0 in state  8 (running)
        node   1 in state  8 (running)
```

Sample /etc/vxfenmode file for majority-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
# majority   - use majority based fencing
# disabled   - run the driver but don't do any actual fencing
#
# vxfen_mode=majority
```

Section 7

Managing your Symantec deployments

Chapter 19

Performing centralized installations using the Deployment Server

This chapter includes the following topics:

- Deploying Symantec releases
- Connecting the Deployment Server to SORT using a proxy server

About the Deployment Server

The Deployment Server makes it easier to install or upgrade SFHA releases from a central location. The Deployment Server lets you store multiple release images and patches in one central location and deploy them to systems of any supported UNIX or Linux operating system (6.1 or later).

Note: The script-based installer for version 6.1 and higher supports installations from one operating system node onto a different operating system. Therefore, heterogeneous push installations are supported for 6.1 and higher releases only.

Push installations for product versions 5.1, 6.0, or 6.0.1 releases must be executed from a system that is running the same operating system as the target systems. In order to perform push installations for product versions 5.1, 6.0, or 6.0.1 releases on multiple platforms, you must have a separate Deployment Server for each operating system.

The Deployment Server lets you do the following as described in Table 19-1.

Table 19-1 Deployment Server functionality

Feature	Description
Manage repository images	■ View available SFHA releases. ■ Download maintenance and patch release images from the Symantec Operations Readiness Tools (SORT) website into a repository. ■ Load the downloaded release image files from FileConnect and SORT into the repository. ■ View and remove the release image files that are stored in the repository.
Version check systems	■ Discover filesets and patches installed on your systems and informs you of the product and version installed ■ Identify base, maintenance, and patch level upgrades to your system and to download maintenance and patch releases. ■ Query SORT for the most recent updates.

Table 19-1 Deployment Server functionality *(continued)*

Feature	Description
Install or upgrade systems	■ Install base, maintenance, or patch level releases. ■ Install SFHA from any supported UNIX or Linux operating system to any other supported UNIX or Linux operating system. ■ Automatically load the script-based installer patches that apply to that release. ■ Install or upgrade an Install Bundle that is created from the **Define/Modify Install Bundles** menu. ■ Install an Install Template that is created from the **Create Install Templates** menu.
Define or modify Install Bundles	Define or modify Install Bundles and save them using the Deployment Server.
Create Install Templates	Discover installed components on a running system that you want to replicate on to new systems.
Update metadata	Download, load the release matrix updates, and product installer updates for systems behind a firewall. This process happens automatically when you connect the Deployment Server to the Internet, or it can be initiated manually. If the Deployment Server is not connected to the Internet, then the **Update Metadata** option is used to upload current metadata.
Set preferences	Define or reset program settings.
Connecting the Deployment Server to SORT using a proxy server	Use a proxy server, a server that acts as an intermediary for requests from clients, for connecting the Deployment Server to the Symantec Operations Readiness Tools (SORT) website.

Note: The Deployment Server is available only from the command line. The Deployment Server is not available for the web-based installer.

Note: Many of the example outputs used in this chapter are based on Red Hat Enterprise Linux.

Deployment Server overview

After obtaining and installing the Deployment Server and defining a central repository, you can begin managing your deployments from that repository. You

can load and store product images for Symantec products back to version 5.1 in your Deployment Server. The Deployment Server is a central installation server for storing and managing your product updates.

Setting up and managing your repository involves the following tasks:

- Installing the Deployment Server.
  See "Installing the Deployment Server" on page 305.
- Setting up a Deployment Server.
  See "Setting up a Deployment Server" on page 306.
- Finding out which products you have installed, and which upgrades or updates you may need.
  See "Viewing or downloading available release images" on page 313.
- Adding release images to your Deployment Server.
  See "Viewing or downloading available release images" on page 313.
- Removing release images from your Deployment Server.
  See "Viewing or removing repository images stored in your repository" on page 318.
- Defining or modifying Install Bundles to manually install or upgrade a bundle of two or more releases.
  See "Defining Install Bundles" on page 322.
- Creating Install Templates to discover installed components on a system that you want to replicate to another system.
  See "Creating Install Templates" on page 328.

Later, when your repository is set up, you can use it to deploy Symantec products to other systems in your environment.

See "Deploying Symantec product updates to your environment" on page 320.

See "Deploying Symantec releases" on page 330.

Installing the Deployment Server

You can obtain the Deployment Server by either:

- Installing the Deployment Server manually.
- Running the Deployment Server after installing at least one Symantec 6.2 product.

Note: The `VRTSperl` and the `VRTSsfcpi<version>`filesets are included in all Storage Foundation (SF) products, so installing any Symantec 6.2 product lets you access the Deployment Server.

To install the Deployment Server manually without installing a Symantec 6.2 product

1. Log in as superuser.
2. Mount the installation media.

   See "Mounting the product disc" on page 74.
3. Move to the top-level directory on the disc.

   ```
   # cd /mnt/cdrom
   ```
4. Navigate to the following directory:

   ```
   # cd pkgs
   ```
5. Run the following commands to install the `VRTSperl` and the `VRTSsfcpi<version>` filesets:

   ```
   # installp -C
   # installp -aXd ./VRTSperl.bff VRTSperl
   # installp -aXd ./VRTSsfcpi<version>.bff VRTSsfcpi<version>
   ```

To run the Deployment Server

1. Log in as superuser.
2. Navigate to the following directory:

   ```
   # cd /opt/VRTS/install
   ```
3. Run the Deployment Server.

   ```
   # ./deploy_sfha
   ```

Setting up a Deployment Server

Symantec recommends that you create a dedicated Deployment Server to manage your product updates.

A Deployment Server is useful for doing the following tasks:

- Storing release images for the latest upgrades and updates from Symantec in a central repository directory.
- Installing and updating systems directly by accessing the release images that are stored within a central repository.
- Defining or modifying Install Bundles for deploying a bundle of two or more releases.
- Discovering installed components on a system that you want to replicate to another system.
- Installing Symantec products from the Deployment Server to systems running any supported platform.
- Creating a file share on the repository directory provides a convenient, central location from which systems running any supported platform can install the latest Symantec products and updates.

Create a central repository on the Deployment Server to store and manage the following types of Symantec releases:

- Base releases. These major releases and minor releases are available for all Symantec products. They contain new features, and you can download them from FileConnect.
- Maintenance releases. These releases are available for all Symantec products. They contain bug fixes and a limited number of new features, and you can download them from the Symantec Operations Readiness Tools (SORT) website.
- Patches. These releases contain fixes for specific products, and you can download them from the SORT website.

Note: All base releases and maintenance releases can be deployed using the install scripts that are included in the release. Before version 6.0.1, patches were installed manually. From the 6.0.1 release and onwards, install scripts are included with patch releases.

You can set up a Deployment Server with or without Internet access.

- If you set up a Deployment Server that has Internet access, you can download maintenance releases and patches from Symantec directly. Then, you can deploy them to your systems.
  Setting up a Deployment Server that has Internet access
- If you set up a Deployment Server that does not have Internet access, you can download maintenance releases and patches from Symantec on another system that has Internet access. Then, you can load the images onto the Deployment Server separately.

Setting up a Deployment Server that does not have Internet access

Setting up a Deployment Server that has Internet access

Figure 19-1 shows a Deployment Server that can download product images directly from Symantec using the Deployment Server.

Figure 19-1 Example Deployment Server that has Internet access

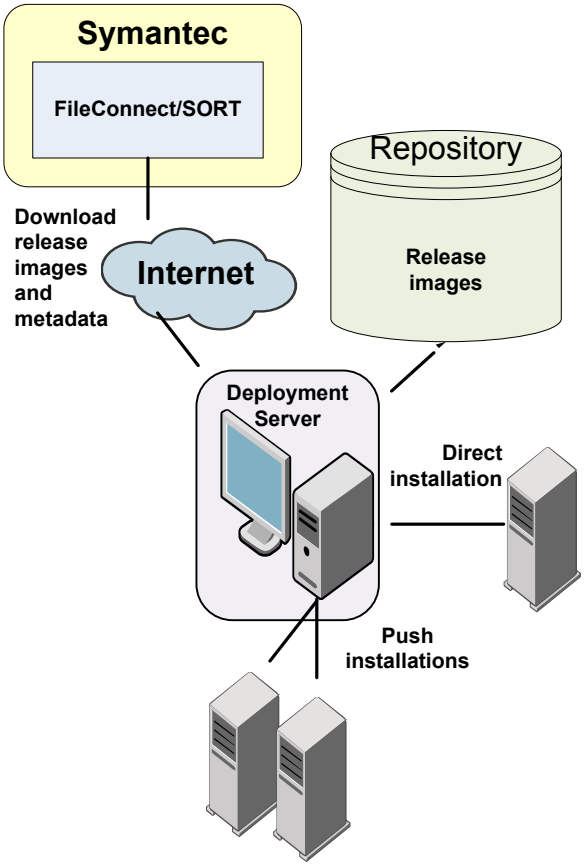

Setting up a Deployment Server that does not have Internet access

Figure 19-2 shows a Deployment Server that does not have Internet access. In this scenario, release images and metadata updates are downloaded from another system. Then, they are copied to a file location available to the Deployment Server, and loaded.

Figure 19-2 Example Deployment Server that does not have Internet access

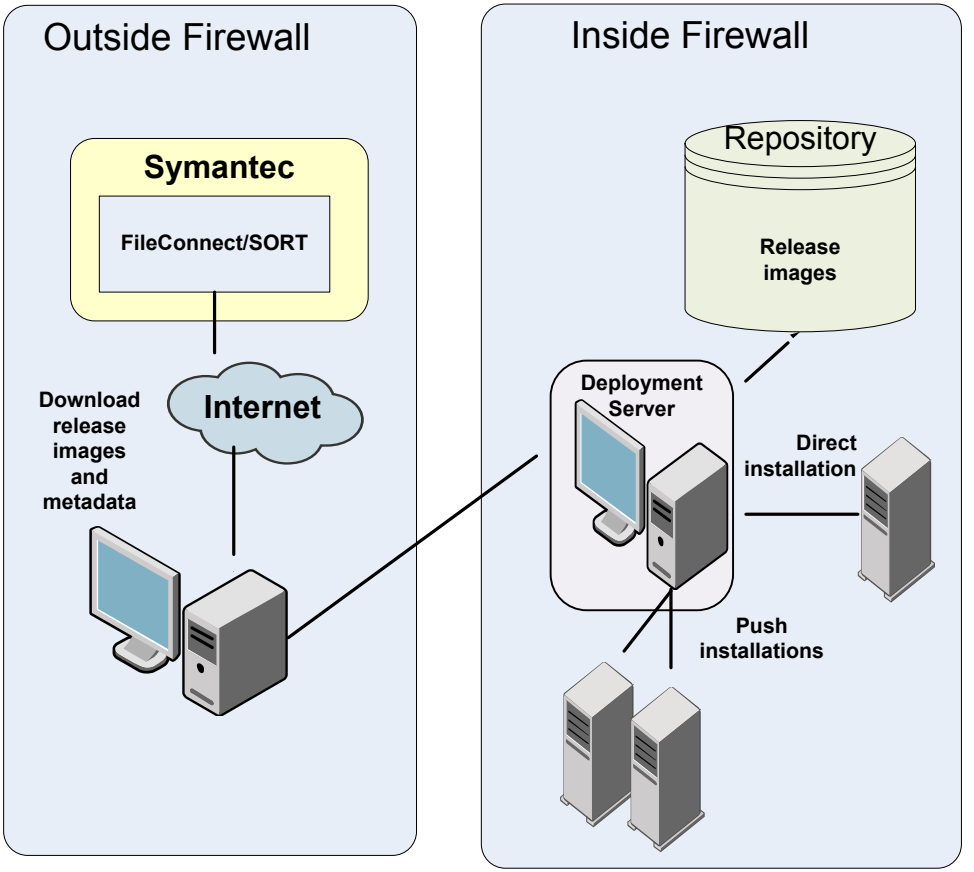

Release image files for base releases must be manually downloaded from FileConnect and loaded in a similar manner.

Setting deployment preferences

You can set preferences for managing the deployment of products dating back to version 5.1.

Note: You can select option **U (Terminology and Usage)** to obtain more information about Deployment Server terminology and usage.

To set deployment preferences

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images             M) Update Metadata
V) Version Check Systems                S) Set Preferences
I) Install/Upgrade Systems              U) Terminology and Usage
B) Define/Modify Install Bundles        ?) Help
T) Create Install Templates             Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **S**, **Set Preferences**.

You see the following output:

```
Current Preferences:

Repository            /opt/VRTS/repository
Selected Platforms    N/A
Save Tar Files        N/A

Preference List:

     1)  Repository
     2)  Selected Platforms
     3)  Save Tar Files
     b)  Back to previous menu

Select a preference to set: [1-3,b,q,?]
```

3 Do one of the following:

- To set the default repository, enter **1**. Then enter the name of the repository in which you want to store your downloads. For example, enter the following:

  ```
  /opt/VRTS/install/ProductDownloads
  ```

  If the specified repository replaces a previous repository, the installer asks if you want to move all your files into the new repository. To move your files to the new repository, enter **y**.

- To add or remove a platform, enter **2**. You are provided selections for adding or deleting a platform. When a single platform is removed, it becomes **N/A**, which means that it is not defined. By default, all platforms are chosen. Once you select to add or remove a platform, the platform is added or removed in the preferences file and the preference platforms are updated.
  If only one platform is defined, no platform, architecture, distribution, and version selection menu is displayed.
- To set the option for saving or removing tar files, enter **3**. At the prompt, if you want to save the tar files after untarring them, enter **y**. Or, if you want to remove tar files after untarring them, enter **n**.
  By default, the installer does not remove tar files after the releases have been untarred.

Specifying a non-default repository location

You can specify a repository location other than the default that has been set within the system preferences by using the command line option. The command line option is mainly used to install a release image from a different repository location. When you use the command line option, the designated repository folder is used instead of the default for the execution time of the script. Using the command line option does not override the repository preference set by the **Set Preference** menu item.

Note: When you specify a non-default repository, you are allowed only to view the repository (**View/Remove Repository**), and use the repository to install or upgrade (**Install/Upgrade Systems**) on other systems.

To use the command line option to specify a non-default repository location

- At the command line, to specify a non-default repository location, enter the following:

  ```
  # ./deploy_sfha -repository repository_path
  ```

  where *repository_path* is the location of the repository.

Downloading the most recent release information

Use one of the following methods to obtain a `.tar` file with the most recent release information:

- Download a copy from the SORT website.
- Run the Deployment Server from a system that has Internet access.

To obtain a data file by downloading a copy from the SORT website

1 Download the `.tar` file from the SORT site at:

https://sort.symantec.com/support/related_links/offline-release-updates

2 Click on **deploy_sfha.tar [Download]**, and save the file to your desktop.

To obtain a data file by running the Deployment Server from a system with Internet access

1 Run the Deployment Server. Enter the following:

```
# /opt/VRTS/install/deploy_sfha
```

2 Select option **M**, **Update Metadata**.

You see the following output:

```
The Update Metadata option is used to load release matrix updates on to systems that do
not have an Internet connection with SORT (https://sort.symantec.com).  Your system has a
connection with SORT and is able to receive updates.  No action is necessary unless you
would like to create a file to update another Deployment Server system.

     1)  Download release matrix updates and installer patches
     2)  Load an update tar file
     b)  Back to previous menu

Select the option: [1-2,b,q,?]
```

3 Select option **1**, **Download release matrix updates and installer patches**.

Loading release information and patches on to your Deployment Server

In this procedure, the Internet-enabled system is the system to which you downloaded the `deploy_sfha.tar` file.

See “Downloading the most recent release information” on page 311.

To load release information and patches on to your Deployment Server

1. On the Internet-enabled system, copy the `deploy_sfha.tar` file you downloaded to a location accessible by the Deployment Server.
2. On the Deployment Server, change to the installation directory. For example, enter the following:

   ```
   # cd /opt/VRTS/install/
   ```

3. Run the Deployment Server. Enter the following:

   ```
   # ./deploy_sfha
   ```

4. Select option **M**, **Update Metadata**, and select option **2**, **Load an update tar file**. Enter the location of the `deploy_sfha.tar` file (the installer calls it a "meta-data tar file").

   ```
   Enter the location of the meta-data tar file: [b]
   (/opt/VRTS/install/deploy_sfha.tar)
   ```

   For example, enter the location of the meta-data tar file:

   ```
   /tmp/deploy_sfha.tar
   ```

Viewing or downloading available release images

You can use the Deployment Server to conveniently view or download available release images to be deployed on other systems in your environment.

Note: If you have Internet access, communication with the Symantec Operations Readiness Tools (SORT) provides the latest release information. If you do not have Internet access, static release matrix files are referenced, and the most recent updates may not be included.

See "Loading release information and patches on to your Deployment Server" on page 312.

To view or download available release images

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

Select the option you would like to perform [1-3,b,q,?]
```

3 Select option **1**, **View/Download Available Releases**, to view or download what is currently installed on your system.

You see a list of platforms and release levels.

```
To view or download available releases, the platform type and release
level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  RHEL7 x86_64              8)  SLES10 x86_64
9)  SLES11 x86_64             10) Solaris 9 Sparc
11) Solaris 10 Sparc          12) Solaris 10 x64
13) Solaris 11 Sparc          14) Solaris 11 x64
b)  Back to previous menu

Select the platform of the release to view/download [1-14,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels.

```
1) Base
2) Maintenance
3) Patch
b) Back to previous menu

Select the level of the <platform> releases to view/download
[1-3,b,q,?]
```

5 Select the number corresponding to the type of release you want to view (Base, Maintenance, or Patch).

You see a list of releases available for download.

```
Available Maintenance releases for aix71:

release_version SORT_release_name       DL OBS AI rel_date   size_KB
=====================================================================
5.1SP1PR1RP2    sfha-aix-5.1SP1PR1RP2   -  Y   Y  2011-09-28 288760
5.1SP1PR1RP3    sfha-aix71-5.1SP1PR1RP3 Y  Y   Y  2012-10-02 290321
5.1SP1PR1RP4    sfha-aix71-5.1SP1PR1RP4 -  -   Y  2013-08-21 304300
6.0RP1          sfha-aix-6.0RP1         -  -   Y  2012-03-22 293980
6.0.3           sfha-aix-6.0.3          -  -   Y  2013-01-31 294041
```

```
Enter the release_version to view details about a release or press 'Enter' to continue [b,q,?]
```

The following are the descriptions for the column headers:

- release_version: The version of the release.
- SORT_release_name: The name of the release, used when accessing SORT (https://sort.symantec.com).
- DL: An indicator that the release is present in your repository.
- OBS: An indicator that the release is obsolete by another higher release.
- AI: An indicator that the release has scripted install capabilities. All base and maintenance releases have auto-install capabilities. Patch releases with auto-install capabilities are available beginning with version 6.1. Otherwise the patch requires a manual installation.
- rel_date: The date the release is available.
- size_KB: The file size of the release in kilobytes.

6 If you are interested in viewing more details about any release, type the release version. For example, enter the following:

```
6.0.3
```

You see the following output:

```
release_version: 6.0.3
release_name: sfha-aix-6.0.3
release_type: MR
release_date: 2013-01-31
install_path: aix/installmr
upload_location: ftp://ftp.veritas.com/pub/support/patchcentral
/AIX/6.0.3/sfha/sfha-aix-6.0.3-patches.tar.gz
obsoletes: 6.0.1.200-fs,6.0.1.200-vm
obsoleted_by: None

Would you like to download this Maintenance Release? [y,n,q] (y) n

Enter the release_version to view the details about a release or press
'Enter' to continue [b,q,?]
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to download a aix Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all releases that are not currently in the repository.

```
1)  5.1SP1PR2RP2
     2)  5.1SP1PR2RP3
     3)  5.1SP1PR2RP4
     4)  5.1SP1PR3RP2
     5)  5.1SP1PR3RP3
     6)  5.1SP1PR3RP4
     7)  6.0RP1
     8)  6.0.3
     9)  6.0.5
    10)  6.1.1
    11)  All non-obsoleted releases
    12)  All releases
     b)  Back to previous menu

Select the patch release to download, 'All non-obsoleted releases' to
download all non-obsoleted releases, or 'All releases' to download
all releases [1-5,b,q] 3
```

8 Select the number corresponding to the release that you want to download. You can download a single release, all non-obsoleted releases, or all releases.

The selected release images are downloaded to the Deployment Server.

```
Downloading sfha-aix-6.0RP1 from SORT - https://sort.symantec.com
Downloading 215118373 bytes (Total 215118373 bytes [205.15 MB]): 100%
Untarring sfha-aix-6.0RP1 ..................................................... Done

sfha-aix-6.0RP1 has been downloaded successfully.
```

9 From the menu, select option **2**, **View/Remove Repository Images**, and follow the prompts to check that the release images are loaded.

See "Viewing or downloading available release images" on page 313.

Viewing or removing repository images stored in your repository

You can use the Deployment Server to conveniently view or remove the release images that are stored in your repository.

To view or remove release images stored in your repository

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images             M) Update Metadata
V) Version Check Systems                S) Set Preferences
I) Install/Upgrade Systems              U) Terminology and Usage
B) Define/Modify Install Bundles        ?) Help
T) Create Install Templates             Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **R**, **Manage Repository Images**.

You see the following output:

```
1)  View/Download Available Releases
2)  View/Remove Repository Images
3)  Load a Release Image
b)  Back to previous menu

Select the option you would like to perform [1-3,b,q,?]
```

3 Select option **2**, **View/Remove Repository Images**, to view or remove the release images currently installed on your system.

You see a list of platforms and release levels if you have downloaded the corresponding Base, Maintenance, or Patch release on that platform.

```
To view or remove repository images, the platform type and
release level type must be selected.

1)  AIX 5.3                   2)  AIX 6.1
3)  AIX 7.1                   4)  HP-UX 11.31
5)  RHEL5 x86_64              6)  RHEL6 x86_64
7)  RHEL7 x86_64              8)  SLES10 x86_64
9)  SLES11 x86_64             10) Solaris 9 Sparc
11) Solaris 10 Sparc          12) Solaris 10 x64
13) Solaris 11 Sparc          14) Solaris 11 x64
b)  Back to previous menu

Select the platform of the release to view/remove [1-14,b,q]
```

4 Select the release level for which you want to get release image information. Enter the platform you want.

You see options for the Symantec release levels if you have downloaded the corresponding Base, Maintenance, or Patch release.

```
1) Base
2) Maintenance
3) Patch
b) Back to previous menu

Select the level of the <platform> releases to view/remove
[1-3,b,q]
```

5 Select the number corresponding to the type of release you want to view or remove (Base, Maintenance, or Patch).

You see a list of releases that are stored in your repository.

```
Stored Repository Releases:

release_version SORT_release_name        OBS AI
================================================
6.0RP1          sfha-aix-6.0RP1           -   Y
6.0.3           sfha-aix-6.0.3            -   Y
```

6 If you are interested in viewing more details about a release image that is stored in your repository, type the release version. For example, enter the following:

```
6.0.3
```

7 If you do not need to check detail information, you can press **Enter**.

You see the following question:

```
Would you like to remove a aix61 Maintenance Release Image?
[y,n,q] (n) y
```

If you select a **y**, you see a menu of all the releases that are stored in your repository that match the selected platform and release level.

```
1)  6.0RP1
2)  6.0.3
b)  Back to previous menu

Select the patch release to remove [1-2,b,q] 1
```

8 Type the number corresponding to the release version you want to remove.

The release images are removed from the Deployment Server.

```
Removing sfha-aix-6.0RP1-patches ........................... Done
sfha-aix-6.0RP1-patches has been removed successfully.
```

Deploying Symantec product updates to your environment

You can use the Deployment Server to deploy release images to the systems in your environment as follows:

- If you are not sure what to deploy, perform a version check. A version check tells you if there are any Symantec products installed on your systems. It suggests patches and maintenance releases, and gives you the option to install updates.
  See "Finding out which releases you have installed, and which upgrades or updates you may need" on page 321.

- If you know which update you want to deploy on your systems, use the Install/Upgrade Systems script to deploy a specific Symantec release.
  See "Deploying Symantec releases" on page 330.

Finding out which releases you have installed, and which upgrades or updates you may need

Use the Version Check option to determine which Symantec product you need to deploy. The Version Check option is useful if you are not sure which releases you already have installed, or you want to know about available releases.

The Version Check option gives you the following information:

- Installed products and their versions (base, maintenance releases, and patches)
- Installed filesets (required and optional)
- Available releases (base, maintenance releases, and patches) relative to the version which is installed on the system

To determine which Symantec product updates to deploy

1 Launch the Deployment Server. For example, enter the following:

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **V**, **Version Check Systems**.

3 At the prompt, enter the system names for the systems you want to check. For example, enter the following:

sys1

You see output for the installed filesets (required, optional, or missing).

You see a list of releases available for download.

```
Available Base Releases for Veritas Storage Foundation HA 6.0.1:
None

Available Maintenance Releases for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name     DL OBS AI rel_date   size_KB
================================================================
6.0.3           sfha-aix-6.0.3        Y  -   -  2013-02-01 212507

Available Public Patches for Veritas Storage Foundation HA 6.0.1:

release_version SORT_release_name      DL OBS AI rel_date   size_KB
================================================================
6.0.1.200-fs    fs-aix-6.0.1.200       -  Y   -  2012-09-20 14346
6.0.1.200-vm    vm-aix-6.0.1.200       -  Y   -  2012-10-10 47880

Would you like to download the available Maintenance or Public Patch
releases which cannot be found in the repository? [y,n,q] (n) y
```

4 If you want to download any of the available maintenance releases or patches, enter **y**.

5 If you have not set a default repository for releases you download, the installer prompts you for a directory. (You can also set the default repository in **Set Preferences**).

See "Setting deployment preferences" on page 309.

6 Select an option for downloading products.

The installer downloads the releases you specified and stores them in the repository.

Defining Install Bundles

You can use Install Bundles to directly install the latest base, maintenance, and patch releases on your system. Install Bundles are a combination of base,

maintenance, and patch releases that can be bundled and installed or upgraded in one operation.

Note: Install Bundles can be defined only from version 6.1 or later. The exception to this rule is base releases 6.0.1, 6.0.2, or 6.0.4 or later with maintenance release 6.0.5 or later.

To define Install Bundles

1 Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images             M) Update Metadata
V) Version Check Systems                S) Set Preferences
I) Install/Upgrade Systems              U) Terminology and Usage
B) Define/Modify Install Bundles        ?) Help
T) Create Install Templates             Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **B**, **Define/Modify Install Bundles**.

You see the following output the first time you enter:

```
Select a Task:

     1)  Create a new Install Bundle
     b)  Back to previous menu

Select the task you would like to perform [1-1,b,q]
```

3 Select option **1**, **Create a new Install Bundle**.

You see the following output:

```
Enter the name of the Install Bundle you would like to define:
{press [Enter] to go back)
```

For example, if you entered:

```
rhel605
```

You see the following output:

```
To create an Install Bundle, the platform type must be selected:

     1)  AIX 5.3                    2)  AIX 6.1
     3)  AIX 7.1                    4)  HP-UX 11.31
     5)  RHEL5 x86_64               6)  RHEL6 x86_64
     7)  RHEL7 x86_64               8)  SLES10 x86_64
     9)  SLES11 x86_64              10) Solaris 9 Sparc
     11) Solaris 10 Sparc           12) Solaris 10 x64
     13) Solaris 11 Sparc           14) Solaris 11 x64
     b)  Back to previous menu


Select the platform of the release for the Install Bundle rhel605:
[1-14,b,q]
```

4 Select the number corresponding to the platform you want to include in the Install Bundle. For example, select the number for the **RHEL5 x86_64** release, **5**.

You see the following output:

```
Details of the Install Bundle: rhel605

Install Bundle Name    rhel605
Platform               RHEL5 x86_64
Base Release           N/A
Maintenance Release    N/A
Patch Releases         N/A

     1)  Add a Base Release
     2)  Add a Maintenance Release
     3)  Add a Patch Release
     4)  Change Install Bundle Name
     b)  Back to previous menu

Select an action to perform on the Install Bundle rhel605 [1-4,b,q]
```

5 Select option **1**, **Add a Base Release**.

You see the following output:

```
     1)  6.0.1
     2)  6.0.2
     3)  6.1
     b)  Back to previous menu

Select the Base Release version to add to the Install Bundle rhel605
[1-3,b,q]
```

6 Select option **1**, **6.0.1**.

You see the following output:

```
Symantec Storage Foundation and High Availability Solutions 6.2 Deployment Server Program
                    pilotlnx11

Details of the Install Bundle: rhel605

Install Bundle Name    rhel605
Platform               RHEL5 x86_64
Base Release           6.0.1
Maintenance Release    N/A
Patch Releases         N/A

     1)  Remove Base Release 6.0.1
     2)  Add a Maintenance Release
     3)  Add a Patch Release
     4)  Change Install Bundle Name
     b)  Back to previous menu

Select an action to perform on the Install Bundle rhel605 [1-4,b,q]
```

7 Select option **2**, **Add a Maintenance Release**.

You see the following output:

```
     1)  6.0.5
     b)  Back to previous menu

Select the Maintenance Release version to add to the Install Bundle
rhel605 [1-1,b,q]
```

8 Select option **1**, **6.0.5**.

You see the following output:

```
Symantec Storage Foundation and High Availability Solutions 6.2 Deployment Server Program
                                  pilotlnx11

Details of the Install Bundle: rhel605

Install Bundle Name    rhel605
Platform               RHEL5 x86_64
Base Release           6.0.1
Maintenance Release    6.0.5
Patch Releases         N/A

     1)  Remove Base Release 6.0.1
     2)  Remove Maintenance Release 6.0.5
     3)  Add a Patch Release
     4)  Save Install Bundle rhel605
     5)  Change Install Bundle Name
     b)  Back to previous menu

Select an action to perform on the Install Bundle rhel605
[1-5,b,q]
```

9 Select option **4**, **Save Install Bundle**.

You see the following output:

```
Install Bundle rhel605 has been saved successfully

Press [Enter] to continue:
```

If there are no releases for the option you selected, you see a prompt saying that there are no releases at this time. You are prompted to continue.

After selecting the desired base, maintenance, or patch releases, you can choose to save your Install Bundle.

The specified Install Bundle is saved on your system. The specified Install Bundle is available as an installation option when using the **I) Install/Upgrade Systems** option to perform an installation or upgrade.

Creating Install Templates

You can use Install Templates to discover installed components (filesets, patches, products, or versions) on a system that you want to replicate. Use Install Templates to automatically install those same components on to other systems.

To create Install Templates

1. Launch the Deployment Server.

```
# /opt/VRTS/install/deploy_sfha
```

2. You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

3. Select option **T**, **Create Install Templates**.
4. You see the following output:

```
Select a Task:

     1)  Create a new Install Template
     b)  Back to previous menu

Select the task you would like to perform [1-1,b,q]
```

5 Select option **1**, **Create a new Install Template**.

You see the following output:

```
Enter the system names separated by spaces for creating an Install Template:
(press [Enter] to go back)
```

For example, if you entered **`rhel89202`** as the system name, you see the following output:

```
Enter the system names separated by spaces for version checking: rhel89202

Checking communication on rhel89202 ..... Done
Checking installed products on rhel89202 ..... Done

Platform of rhel89202:
   Linux RHEL 6.3 x86_64

Installed product(s) on rhel89202:
   Symantec Storage Foundation Cluster File System HA - 6.1.1 - license vxkeyless

Product:
   Symantec Storage Foundation Cluster File System HA - 6.1.1 - license vxkeyless

Packages:
   Installed Required packages for Symantec Storage Foundation Cluster File System HA 6.1.1:
          #PACKAGE     #VERSION
          VRTSamf      6.1.1.000
          VRTSaslapm   6.1.1.000
          ..........   .........
          ..........   .........
          VRTSvxfs     6.1.1.000
          VRTSvxvm     6.1.1.000

   Installed optional packages for Symantec Storage Foundation Cluster File System HA 6.1.1:
          #PACKAGE     #VERSION
          VRTSdbed     6.1.1.000
          VRTSgms      6.1.0.000
          ..........   .........
          ..........   .........
          VRTSvcsdr    6.1.0.000
          VRTSvcsea    6.1.1.000

   Missing optional packages for Symantec Storage Foundation Cluster File System HA 6.1.1:
          #PACKAGE
```

```
          VRTScps
          VRTSfssdk
          VRTSlvmconv

Summary:

Packages:
 17 of 17 required Symantec Storage Foundation Cluster File System HA 6.1.1 packages installed
 8 of 11 optional Symantec Storage Foundation Cluster File System HA 6.1.1 packages installed

Installed Public and Private Hot Fixes for Symantec Storage Foundation Cluster File
System HA 6.1.1:
 None

Would you like to generate a template file based on the above release information? [y,n,q] (y)

 1)  rhel89202
 b)  Back to previous menu

Select a machine list to generate the template file [1-1,b,q]
```

6 Select option **1**, **rhel89202**.

You see the following output:

```
Enter the name of the Install Template you would like to define:
(press [Enter] to go back)
```

7 Enter the name of your Install Template. For example, if you enter **MyTemplate** as the name for your Install Template, you would see the following:

```
Install Template MyTemplate has been saved successfully

Press [Enter] to continue:
```

All of the necessary information is stored in the Install Template you created.

Deploying Symantec releases

You can use the Deployment Server to deploy your licensed Symantec products dating back to version 5.1. If you know which product version you want to install, follow the steps in this section to install it.

You can use the Deployment Server to install the following:

- A single Symantec release
- Two or more releases using defined Install Bundles
  See "Defining Install Bundles" on page 322.
- Installed components on a system that you want to replicate on another system
  See "Creating Install Templates" on page 328.

To deploy a specific Symantec release

1 From the directory in which you installed your Symantec product (version 6.1 or later), launch the Deployment Server with the upgrade and install systems option. For example, enter the following:

```
# /opt/VRTS/install/deploy_sfha
```

You see the following output:

```
Task Menu:

R) Manage Repository Images              M) Update Metadata
V) Version Check Systems                 S) Set Preferences
I) Install/Upgrade Systems               U) Terminology and Usage
B) Define/Modify Install Bundles         ?) Help
T) Create Install Templates              Q) Quit

Enter a Task: [R,M,V,S,I,U,B,?,T,Q]
```

2 Select option **I**, **Install/Upgrade Systems**.

You see the following output:

```
    1)  AIX 5.3
    2)  AIX 6.1
    3)  AIX 7.1
    4)  RHEL5 x86_64
    b)  Back to previous menu

Select the platform of the available release(s) to be
upgraded/installed [1-4,b,q,?]
```

3 Select the number corresponding to the platform for the release you want to deploy. For example, select the number for the **RHEL5 x86_64** release or the **AIX 6.1** release.

You see the following output:

```
1) Install/Upgrade systems using a single release
2) Install/Upgrade systems using an Install Bundle
3) Install systems using an Install Template
b) Back to previous menu

Select the method by which you want to Install/Upgrade your systems
[1-3,b,q]
```

4 Section option **1**, **Install/Upgrade systems using a single release** if you want to deploy a specific Symantec release.

Select a Symantec product release.

The installation script is executed and the release is deployed on the specified server.

To deploy an Install Bundle

1 Follow Steps 1 - 3.

2 Select option **2**, **Install/Upgrade systems using an Install Bundle**.

You see the following output:

```
1) <NameofInstallBundle1>
2) <NameofInstallBundle2>
b) Back to previous menu

Select the bundle to be installed/upgraded [1-2,b,q]
```

You see the following output:

```
Enter the platform target system name(s) separated by spaces:
[press [Enter] to go back)
```

3 Enter the name of the target system for which you want to install or upgrade the Install Bundle.

The installation script for the selected Install Bundle is executed, and the Install Bundle is deployed on the specified target system.

To deploy an Install Template

1. Follow Steps 1 - 3.
2. Select option **3**, **Install/Upgrade systems using an Install Template**.

   You see the following output:

   ```
   1) <NameofInstallTemplate>
   b) Back to previous menu

   Select the template to be installed [1-1,b,q] 1
   ```

   You see the following output:

   ```
   Enter the platform target system name(s) separated by spaces:
   [press [Enter] to go back)
   ```

   The installation script for the selected Install Template is executed, and the Install Template is deployed on the specified target system.

Connecting the Deployment Server to SORT using a proxy server

You can use a proxy server, a server that acts as an intermediary for requests from clients, for connecting the Deployment Server to the Symantec Operations Readiness Tools (SORT) website.

To enable the proxy access, run the following commands to set the shell environment variables before you launch Deployment Server. The shell environment variables enable Deployment Server to use the proxy server myproxy.mydomain.com which connects to port 3128.

```
http_proxy="http://myproxy.mydomain.com:3128"
export http_proxy

ftp_proxy="http://myproxy.mydomain.com:3128"
export ftp_proxy
```

The lines above can be added to the user's shell profile. For the bash shell, the profile is the `~/.bash_profile` file.

Section 8

Upgrading VCS

Chapter 20

Planning to upgrade VCS

This chapter includes the following topics:

About upgrading to VCS 6.2

When you upgrade to VCS 6.2, you need not reconfigure application monitoring with VCS. All existing monitoring configurations are preserved.

You can upgrade VCS using one of the following methods:

- Typical upgrade using product installer or the installvcs
  See "Supported upgrade paths for VCS 6.2" on page 336.
  See "Upgrading VCS using the script-based installer" on page 344.
- Typical upgrade using Veritas web installer

See “Supported upgrade paths for VCS 6.2” on page 336.
See “Upgrading VCS using the web-based installer” on page 346.

- Performing an online upgrade
  Perform a script-based or web-based online upgrade of your installation to upgrade VCS without stopping your applications. The supported upgrade paths for the online upgrades are same as those documented under the script and web-based upgrades.
  See “Supported upgrade paths for VCS 6.2” on page 336.
  See “Upgrading VCS online using the script-based installer” on page 350.
  See “Upgrading VCS online using the web-based installer” on page 351.
- Phased upgrade to reduce downtime
  See “Performing a phased upgrade using the script-based installer” on page 357.
- Automated upgrade using response files
  See “Supported upgrade paths for VCS 6.2” on page 336.
  See “Upgrading VCS using response files” on page 375.
- Upgrade using supported native operating system utility
  Alternate disk
  See “About upgrading VCS using an alternate disk” on page 402.
- Rolling upgrade to minimize downtime
  See “Performing a rolling upgrade of VCS using the web-based installer” on page 390.

You can upgrade VCS 6.2 to Storage Foundation High Availability 6.2 using the product installer or response files.

See the *Symantec Storage Foundation and High Availability Installation Guide*.

Note: In a VMware virtual environment, you can use the vSphere Client to directly install VCS and supported high availability agents (together called guest components) on the guest virtual machines. For details, see the *Symantec High Availability Solution Guide for VMware*.

Supported upgrade paths for VCS 6.2

The following tables describe upgrading to 6.2.

Table 20-1 AIX upgrades using the script- or web-based installer

Symantec product version	AIX 5.3	AIX 6.1	AIX 7.1
5.1 5.1 RPs 5.1 SP1 5.1 SP1 RP1	Upgrade the operating system to AIX 6.1 TL8 or later - but do not upgrade to AIX 7.1. Then use the installer to upgrade your Symantec product to 6.2.	Upgrade the operating system to AIX 6.1 TL8 or later, or AIX 7.1 TL2 or later. Then use the installer to upgrade your Symantec product to 6.2.	N/A
5.1 SP1 RP2 5.1 SP1 RP3 5.1 SP1 RP4	Upgrade the operating system to AIX 6.1 TL8 or later - but do not upgrade to AIX 7.1.Then use the installer to upgrade your Symantec product to 6.2	Upgrade the operating system to AIX 6.1 TL8 or later, or AIX 7.1 TL2 or later. Then use the installer to upgrade your Symantec product to 6.2.	Upgrade the operating system to AIX 7.1 TL2 or later. Then use the installer to upgrade your Symantec product to 6.2.
5.1 SP1 PR1	N/A	N/A	Upgrade the operating system to AIX 7.1 TL2 or later. Then use the installer to upgrade your Symantec product to 6.2.
6.0 6.0 RPs 6.0.1 6.0.3 6.0.5 6.1 6.1.1	N/A	Upgrade the operating system to AIX 6.1 TL8 or later, or AIX 7.1 TL2 or later. Use the installer to upgrade your Symantec product to 6.2.	Upgrade the operating system to AIX 7.1 TL2 or later. Then use the installer to upgrade your Symantec product to 6.2.

Upgrading VCS in secure enterprise environments

In secure enterprise environments, ssh or rsh communication is not allowed between systems. In such cases, the installvcs program can upgrade VCS only on systems with which it can communicate (most often the local system only).

To upgrade VCS in secure enterprise environments with no rsh or ssh communication

1 Run the installvcs program on each node to upgrade the cluster to VCS 6.2.

On each node, the installvcs program updates the configuration, stops the cluster, and then upgrades VCS on the node. The program also generates a cluster UUID on the node. Each node may have a different cluster UUID at this point.

2 Start VCS on the first node.

```
# hastart
```

VCS generates the cluster UUID on this node. Run the following command to display the cluster UUID on the local node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -display systemname
```

3 On each of the other nodes, perform the following steps:

- Set the value of the VCS_HOST environment variable to the name of the first node.
- Display the value of the CID attribute that stores the cluster UUID value:

```
# haclus -value CID
```

- Copy the output of the CID attribute to the file /etc/vx/.uuids/clusuuid.
- Update the VCS_HOST environment variable to remove the set value.
- Start VCS.
  The node must successfully join the already running nodes in the cluster.
  See "Verifying LLT, GAB, and cluster operation" on page 429.

Considerations for upgrading secure VCS 5.x clusters to VCS 6.2

When you upgrade a secure VCS 5.x cluster to VCS 6.2, the upgrade does not migrate the old broker configuration to the new broker because of the change in architecture. Both the old broker (`/opt/VRTSat/bin/vxatd`) and new broker (`/opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver`) continue to run. In such a scenario, you must consider the following:

- The HA commands that you run in VCS 6.2 are processed by the new broker by default. To ensure that the HA commands are processed by the old broker, set the VCS_REMOTE_BROKER environment variable as follows:

```
# export VCS_REMOTE_BROKER=localhost IP,2821
```

See “About enabling LDAP authentication for clusters that run in secure mode” on page 414.

- VCS 6.2 does not prompt non-root users who run HA commands for passwords. In 5.x, non-root users required a password to run HA commands. If you want non-root users to enter passwords before they run HA commands, set the VCS_DOMAINTYPE environment variable to unixpwd.
- Trust relationships are not migrated during the upgrade. If you had configured secure GCO or secure steward, ensure that trust relationships are recreated between the clusters and the steward.
  See “Setting up trust relationships for your VCS cluster” on page 139.
- For WPARs, the HA commands run within the container and use credentials that were deployed by the old broker. However, you can migrate to the newer credentials from the new broker by running `hawparsetup` again.

When the old broker is not used anymore, you can delete the old VRTSat fileset.

Considerations for upgrading VCS to 6.2 on systems configured with an Oracle resource

If you plan to upgrade VCS running on systems configured with an Oracle resource, set the `MonitorOption` attribute to 0 (zero) before you start the upgrade. If you use the product installer for the rolling upgrade, it sets the `MonitorOption` to 0 through its scripts. In a manual upgrade, the `MonitorOption` value must be set to 0 using the `hares` command. When the upgrade is complete, invoke the `build_oraapi.sh` script, and then set the `MonitorOption` to 1 to enable the Oracle health check.

For more information on enabling the Oracle health check, see the *Symantec Cluster Server Agent for Oracle Installation and Configuration Guide*.

Considerations for upgrading secure VCS clusters to VCS 6.2

1. When you upgrade a secure VCS cluster to VCS 6.2, you need to configure one of the following attributes to enable guest access to the cluster.
   - DefaultGuestAccess: Set the value of this attribute to **1** to enable guest access for any authenticated user.

- GuestGroups: This attribute contains of list of user groups who have read access to the cluster. Configure this attribute to control the guest access to the cluster.

2. The non-root and WPAR users need to authenticate again if you perform an upgrade in secure mode of a cluster from VCS 6.x to VCS 6.2. Run the `halogin` command to regenerate the password for non-root users. Use the `hawparsetup` command to update the respective credentials of WPAR users.

   Refer to*Modifying the service group configuration* section under *Configuring VCS in WPARs* chapter of the *Symantec Storage Foundation and High Availability Solutions Virtualization Guide* for steps on how to update the credentials of WPAR users.

Considerations for upgrading secure CP servers

CP server supports Symantec Product Authentication Services (AT) (IPM-based protocol) and HTTPS communication to securely communicate with clusters. For HTTPS communication, you do not need to consider setting up trust relationships.

When you upgrade the CP Server that supports IPM-based protocol, trust relationships are not migrated.

If you upgrade the CP clients after you upgrade the CP server that supports IPM-based protocol, the installer recreates the trust relationships that are established by the client. You do not need to establish the trust relationships manually. However, the CP server and CP clients cannot communicate with each other till trust relationships are established.

If you do not upgrade the CP clients after you upgrade the CP server that supports IPM-based protocol, you must recreate the trust relationships between the CP server and CP clients.

Considerations for upgrading secure CP clients

Passwordless communication from CP clients to CP server must exist for the installer to reconfigure fencing. If passwordless communication does not exist, you must reconfigure fencing manually.

See “Setting up disk-based I/O fencing manually” on page 271.

See “Setting up server-based I/O fencing manually” on page 276.

Setting up trust relationship between CP server and CP clients manually

You need to set up trust relationship only if you use the Symantec Product Authentication Services (AT) (IPM-based protocol) for communication between CP servers and CP server clients.

For each client cluster on release version 6.0 and later, run the following command on the CP server:

```
EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSERVER \
/opt/VRTSvcs/bin/vcsat setuptrust -b client_ip_addres:14149 -s high
```

For each client cluster on release version prior to 6.0, run the following command on the CP server:

```
EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSERVER \
/opt/VRTSvcs/bin/vcsat setuptrust -b client_ip_addres:2821 -s high
```

For each client node on release version 6.0 and later, run the following command:

```
EAT_DATA_DIR=/var/VRTSvcs/vcsauth/data/CPSADM \
/opt/VRTSvcs/bin/vcsat setuptrust -b cpserver_ip_address:14149 -s high
```

For each client node on release version prior to 6.0, run the following command:

```
/bin/echo y | /opt/VRTSvcs/bin/vcsat setuptrust -b \
ip_addres_of_cp_server:14149 -s high
```

Using Install Bundles to simultaneously install or upgrade full releases (base, maintenance, rolling patch), and individual patches

Beginning with version 6.1, Symantec offers you a method to easily install or upgrade your systems directly to a base, maintenance, patch level or a combination of multiple patches and packages together in one step using Install Bundles. With Install Bundles, the installer has the ability to merge so that customers can install or upgrade directly to maintenance or patch levels in one execution. The various scripts, filesets, and patch components are merged, and multiple releases are installed together as if they are one combined release. You do not have to perform two or more install actions to install or upgrade systems to maintenance levels or patch levels.

Releases are divided into the following categories:

Table 20-2 Release Levels

Level	Content	Form factor	Applies to	Release types	Download location
Base	Features	filesets	All products	Major, minor, Service Pack (SP), Platform Release (PR)	FileConnect
Maintenance	Fixes, new features	filesets	All products	Maintenance Release (MR), Rolling Patch (RP)	Symantec Operations Readiness Tools (SORT)
Patch	Fixes	filesets	Single product	P-Patch, Private Patch, Public patch	SORT, Support site

When you install or upgrade using Install Bundles:

- SFHA products are discovered and assigned as a single version to the maintenance level. Each system can also have one or more patches applied.
- Base releases are accessible from FileConnect that requires customer serial numbers. Maintenance and patch releases can be automatically downloaded from SORT. You can download them from the SORT website manually or use the `deploy_sfha` script.
- Patches can be installed using automated installers from the 6.0.1 version or later.
- Patches can now be detected to prevent upgrade conflict. Patch releases are not offered as a combined release. They are only available from Symantec Technical Support on a need basis.

You can use the `-base_path` and `-patch_path` options to import installation code from multiple releases. You can find filesets and patches from different media paths, and merge fileset and patch definitions for multiple releases. You can use these options to use new task and phase functionality to correctly perform required operations for each release component. You can install the filesets and patches in defined phases using these options, which helps you when you want to perform a single start or stop process and perform pre and post operations for all level in a single operation.

Four possible methods of integration exist. All commands must be executed from the highest base or maintenance level install script.

For example:

1. Base + maintenance:

   This integration method can be used when you install or upgrade from a lower version to 6.2.1.

   Enter the following command:

   ```
   # installmr -base_path <path_to_base>
   ```

2. Base + patch:

   This integration method can be used when you install or upgrade from a lower version to 6.2.0.100.

   Enter the following command:

   ```
   # installer -patch_path <path_to_patch>
   ```

3. Maintenance + patch:

   This integration method can be used when you upgrade from version 6.2 to 6.2.1.100.

   Enter the following command:

   ```
   # installmr -patch_path <path_to_patch>
   ```

4. Base + maintenance + patch:

   This integration method can be used when you install or upgrade from a lower version to 6.2.1.100.

   Enter the following command:

   ```
   # installmr -base_path <path_to_base>
   -patch_path <path_to_patch>
   ```

   Note: From the 6.1 or later release, you can add a maximum of five patches using *-patch_path <path_to_patch> -patch2_path <path_to_patch> ... -patch5_path <path_to_patch>*

Chapter 21

Performing a typical VCS upgrade using the installer

This chapter includes the following topics:

Before upgrading VCS using the script-based or web-based installer

As a result of OS upgrade, if VCS is not in running state before upgrade, the installer does not start VCS after the upgrade is completed. You need to manually start it or restart the cluster nodes. Before you upgrade VCS, you first need to remove deprecated resource types and modify changed values.

To prepare to upgrade to VCS 6.2, make sure that all non-global zones are booted and in the running state before you install or upgrade the VCS packages in the global zone. If the non-global zones are not mounted and running at the time of upgrade, you must upgrade each package in each non-global zone manually.

Upgrading VCS using the script-based installer

You can use the product installer to upgrade VCS.

To upgrade VCS using the product installer

1. Log in as superuser and mount the product disc.
2. Start the installer.

   ```
   # ./installer
   ```

   The installer starts the product installation program with a copyright message. It then specifies where it creates the logs. Note the log's directory and name.
3. From the opening Selection Menu, choose: **G** for "Upgrade a Product."
4. Choose **1** for Full Upgrade.
5. Enter the names of the nodes that you want to upgrade. Use spaces to separate node names. Press the Enter key to proceed.

   The installer runs some verification checks on the nodes and displays the following message:

   ```
   VCS supports application zero downtime for full upgrade.
   ```

6. When the verification checks are complete, the installer asks if you agree with the terms of the End User License Agreement. Press **y** to agree and continue.

   The installer lists the filesets to upgrade.
7. The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.
   - Do you want to grant read access to everyone? [y,n,q,?]
     - To grant read access to all authenticated users, type **y**.
     - To grant usergroup specific permissions, type **n**.
   - Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
     - To specify usergroups and grant them read access, type **y**
     - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
   - Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

8 The installer asks if you want to stop VCS processes. Press the Enter key to continue.

The installer stops VCS processes, uninstalls filesets, installs or upgrades filesets, and configures VCS.

The installer lists the nodes that Symantec recommends you restart.

9 The installer asks if you would like to send the information about this installation to Symantec to help improve installation in the future. Enter your response.

The installer displays the location of log files, summary file, and response file.

10 If you want to upgrade CP server systems that use VCS or SFHA to VCS 6.2, make sure that you first upgrade all application clusters to version VCS 6.2. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA, see the *Symantec Cluster Server Installation Guide* or the *Storage Foundation and High Availability Installation Guide*.

If you are upgrading from 4.x, you may need to create new VCS accounts if you used native OS accounts.

Upgrading VCS using the web-based installer

This section describes upgrading VCS with the web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems.

To upgrade VCS

1 Perform the required steps to save any data that you want to preserve. For example, make configuration file backups.

2 If you want to upgrade a high availability (HA) product, take all service groups offline. List all service groups:

```
# /opt/VRTSvcs/bin/hagrp -list
```

For each service group listed, take it offline:

```
# /opt/VRTSvcs/bin/hagrp -offline service_group -any
```

3 Start the web-based installer.

See “Starting the web-based installer” on page 186.

4 On the Select a task and a product page, select **Upgrade a Product** from the Task drop-down menu.

The product is discovered once you specify the system. Click **Next**.

5 Indicate the systems on which to upgrade. Enter one or more system names, separated by spaces. Click **Next**.

6 Installer detects the product that is installed on the specified system. It shows the cluster information and lets you confirm if you want to perform upgrade on the cluster. Select **Yes** and click **Next**.

7 On the License agreement page, select whether you accept the terms of the End User License Agreement (EULA). To continue, select **Yes I agree** and click **Next**.

8 Click **Next** to complete the upgrade.

After the upgrade completes, the installer displays the location of the log and summary files. If required, view the files to confirm the installation status.

9 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant more usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**.
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

10 If you are prompted to restart the systems, enter the following restart command:

```
# /usr/sbin/shutdown -r now
```

11 After the upgrade, if the product is not configured, the web-based installer asks: "Do you want to configure this product?" If the product is already configured, it does not ask any questions.

12 Click **Finish**. The installer prompts you for another task.

13 If you want to upgrade application clusters that use VCS or SFHA to 6.2, make sure that you upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.2. For instructions to upgrade VCS or SFHA, see the *VCS or SFHA Installation Guide*.

Chapter 22

Performing an online upgrade

This chapter includes the following topics:

Limitations of online upgrade

- Online upgrade is available only for VCS and ApplicationHA. If you have Storage Foundation, SFHA, SFCFSHA, or any other solution with VxVM and VxFS installed, then the online upgrade process will not be supported.
- The non-Symantec applications running on the node have zero down time during the online upgrade.
- VCS does not monitor the applications when online upgrade is in progress.
- For upgrades from VCS versions lower than 6.1, upgrade the CP server before performing the online upgrade.

See "Supported upgrade paths for VCS 6.2" on page 336.

See "Upgrading VCS online using the script-based installer" on page 350.

See "Upgrading VCS online using the web-based installer" on page 351.

Upgrading VCS online using the script-based installer

You can use the product installer to upgrade VCS online. The supported upgrade paths are same as those for the script-based installer.

See "Supported upgrade paths for VCS 6.2" on page 336.

To upgrade VCS online using the product installer

1. Log in as superuser and mount the product disc.
2. Start the installer.

   ```
   # ./installer
   ```

   The installer starts the product installation program with a copyright message. It then specifies where it creates the logs.

   Note the directory name and path where the logs get stored.
3. From the opening Selection Menu, choose: **G** for "Upgrade a Product."

   The system prompts you to select the method by which you want to upgrade the product.
4. Choose **3** for **Online Upgrade** from the upgrade options.
5. After selecting the online upgrade method, enter any one system name from the cluster on which you want to perform the online upgrade.

   Even if you specify a single node from the cluster, the installer asks whether you want to perform online upgrade of VCS on the entire cluster, keeping your applications online. After you enter the system name, the installer performs some verification checks and asks the following question:

   ```
   Online upgrade supports application zero downtime.
   Would you like to perform online upgrade on the
   whole cluster? [y,n,q](y)
   ```
6. Enter **y** to initiate the online upgrade.

   Note: You can either exit the installer with the option **q** or cancel the upgrade using **n** and select any other cluster to upgrade at this step.

   The installer runs some verification checks on the nodes and subsequently asks if you agree with the terms of the End User License Agreement.
7. Enter **y** to agree and continue.

   The installer lists the filesets that will be upgraded.

8 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

9 The installer asks if you want to stop VCS processes. Enter **y** to stop the VCS process.

It stops the VCS processes, uninstalls filesets, reinstalls or upgrade filesets, again configures VCS, and starts the processes.

10 The installer asks if you want to stop VCS processes. Enter **y** to stop the VCS process.

It stops the VCS processes, uninstalls filesets, reinstalls or upgrade filesets, again configures VCS, and starts the processes.

Upgrading VCS online using the web-based installer

This section describes upgrading VCS online with the web-based installer. The installer detects and upgrades the product that is currently installed on the specified system or systems. The web-based installer upgrades VCS without stopping your applications. The supported upgrade paths are same as those for the web-based installer upgrade.

See "Supported upgrade paths for VCS 6.2" on page 336.

To upgrade VCS

1 Perform the required steps to save any data that you want to preserve. For example, make configuration file backups.

2 Start the web-based installer.

See “Starting the web-based installer” on page 186.

3 On the Select a task and a product page, select **Online Upgrade [VCS/ApplicationHA only]** from the Task drop-down menu.

The product is discovered once you specify the system. Click **Next**.

4 After selecting the online upgrade method, enter any one system name from the cluster on which you want perform the online upgrade.

The method performs online upgrade of VCS on the entire cluster, keeping your applications online. After you enter the system name, the installer performs some verification checks and asks the following question:

```
Online upgrade supports application zero downtime.
Would you like to perform online upgrade on the
whole cluster? [y,n,q](y)
```

5 Enter **y** to initiate the online upgrade.

Note: You can either exit the installer with the option **q** or cancel the upgrade using **n** and select any other cluster to upgrade at this step.

The installer runs some verification checks on the nodes and subsequently asks if you agree with the terms of the End User License Agreement.

6 Enter **y** to agree and continue.

The installer lists the filesets that will be upgraded.

7 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**

- To grant read access only to root users, type **n**. The installer grants read access read access to the root users.

- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

8 The installer asks if you want to stop VCS processes. Enter **y** to stop the VCS process.

It stops the VCS processes, uninstalls filesets, reinstalls or upgrade filesets, again configures VCS, and starts the processes.

Chapter 23

Performing a phased upgrade of VCS

This chapter includes the following topics:

About phased upgrade

Perform a phased upgrade to minimize the downtime for the cluster.

Depending on the situation, you can calculate the approximate downtime as follows:

Table 23-1

Fail over condition	Downtime
You can fail over all your service groups to the nodes that are up.	Downtime equals the time that is taken to offline and online the service groups.
You have a service group that you cannot fail over to a node that runs during upgrade.	Downtime for that service group equals the time that is taken to perform an upgrade and restart the node.

Prerequisites for a phased upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade.

Planning for a phased upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group.

Some rough guidelines follow:

- Split the cluster into two sub-clusters of equal or near equal size.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.
- Before you start the upgrade, back up the VCS configuration files `main.cf` and `types.cf` which are in the `/etc/VRTSvcs/conf/config/` directory.
- Before you start the upgrade make sure that all the disk groups have the latest backup of configuration files in the `/etc/vx/cbr/bk` directory. If not, then run the following command to take the latest backup.

  `# /etc/vx/bin/vxconfigbackup -| [dir] [dgname|dgid]`

Phased upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the phased upgrade:

- While you perform the upgrades, do not start any modules.
- When you start the installer, only select VCS.
- While you perform the upgrades, do not add or remove service groups to any of the nodes.
- After you upgrade the first half of your cluster (the first subcluster), you need to set up password-less SSH or RSH. Create the connection between an upgraded node in the first subcluster and a node from the other subcluster. The node from the other subcluster is where you plan to run the installer and also plan to upgrade.
- Depending on your configuration, you may find that you cannot upgrade multiple nodes at the same time. You may only be able to upgrade one node at a time.
- For very large clusters, you might have to repeat these steps multiple times to upgrade your cluster.

Phased upgrade example

In this example, you have a secure cluster that you have configured to run on four nodes: node01, node02, node03, and node04. You also have four service groups:

sg1, sg2, sg3, and sg4. For the purposes of this example, the cluster is split into two subclusters. The nodes node01 and node02 are in the first subcluster, which you first upgrade. The nodes node03 and node04 are in the second subcluster, which you upgrade last.

Figure 23-1 Example of phased upgrade set up

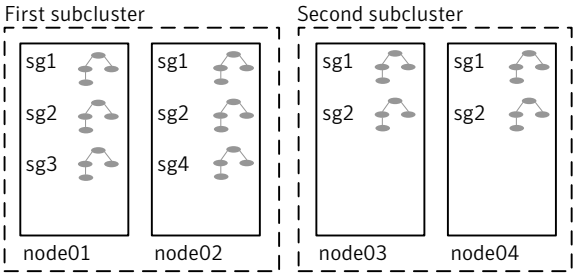

Each service group is running on the nodes as follows:

- sg1 and sg2 are parallel service groups and run on all the nodes.
- sg3 and sg4 are failover service groups. sg3 runs on node01 and sg4 runs on node02.

In your system list, you have each service group that fails over to other nodes as follows:

- sg1 and sg2 are running on all the nodes.
- sg3 and sg4 can fail over to any of the nodes in the cluster.

Phased upgrade example overview

This example's upgrade path follows:

- Move all the failover service groups from the first subcluster to the second subcluster.
- Take all the parallel service groups offline on the first subcluster.
- Upgrade the operating system on the first subcluster's nodes, if required.
- On the first subcluster, start the upgrade using the installation program.
- Get the second subcluster ready.
- Activate the first subcluster. After activating the first cluster, switch the service groups online on the second subcluster to the first subcluster.
- Upgrade the operating system on the second subcluster's nodes, if required.
- On the second subcluster, start the upgrade using the installation program.

- Activate the second subcluster.

See “Performing a phased upgrade using the script-based installer” on page 357.

Performing a phased upgrade using the script-based installer

This section explains how to perform a phased upgrade of VCS on four nodes with four service groups. Note that in this scenario, VCS and the service groups cannot stay online on the second subcluster during the upgrade of the second subcluster. Do not add, remove, or change resources or service groups on any nodes during the upgrade. These changes are likely to get lost after the upgrade.

An example of a phased upgrade follows. It illustrates the steps to perform a phased upgrade. The example makes use of a secure VCS cluster.

You can perform a phased upgrade from VCS 5.1 or other supported previous versions to VCS 6.2.

See “About phased upgrade” on page 354.

See “Phased upgrade example” on page 355.

Moving the service groups to the second subcluster

Perform the following steps to establish the service group's status and to switch the service groups.

To move service groups to the second subcluster

1 On the first subcluster, determine where the service groups are online.

```
# hagrp -state
```

The output resembles:

```
#Group   Attribute System Value
sg1      State     node01 |ONLINE|
sg1      State     node02 |ONLINE|
sg1      State     node03 |ONLINE|
sg1      State     node04 |ONLINE|
sg2      State     node01 |ONLINE|
sg2      State     node02 |ONLINE|
sg2      State     node03 |ONLINE|
sg2      State     node04 |ONLINE|
sg3      State     node01 |ONLINE|
sg3      State     node02 |OFFLINE|
sg3      State     node03 |OFFLINE|
sg3      State     node04 |OFFLINE|
sg4      State     node01 |OFFLINE|
sg4      State     node02 |ONLINE|
sg4      State     node03 |OFFLINE|
sg4      State     node04 |OFFLINE|
```

2 Offline the parallel service groups (sg1 and sg2) from the first subcluster. Switch the failover service groups (sg3 and sg4) from the first subcluster (node01 and node02) to the nodes on the second subcluster (node03 and node04). For SFHA, vxfen sg is the parallel service group.

```
# hagrp -offline sg1 -sys node01
# hagrp -offline sg2 -sys node01
# hagrp -offline sg1 -sys node02
# hagrp -offline sg2 -sys node02
# hagrp -switch sg3 -to node03
# hagrp -switch sg4 -to node04
```

3 On the nodes in the first subcluster, unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -k
```

```
Filesystem            1024-blocks      Free     %Used Iused %Iused Mounted on
/dev/hd4                 20971520     8570080   60%  35736   2%   /
/dev/hd2                  5242880     2284528   57%  55673   9%   /usr
/dev/hd9var               4194304     3562332   16%   5877   1%   /var
/dev/hd3                  6291456     6283832    1%    146   1%   /tmp
/dev/hd1                   262144      261408    1%     62   1%   /home
/dev/hd11admin             262144      184408   30%      6   1%   /admin
/proc                         -           -       -    -     -    /proc
/dev/hd10opt             20971520    5799208      73%  65760  5%   /opt
/dev/vx/dsk/dg2/dg2vol1     10240       7600      26%      4  1%   /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2     10240       7600      26%      4  1%   /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3     10240       7600      26%      4  1%   /mnt/dg2/dg2vol3
```

```
# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

4 On the nodes in the first subcluster, stop all VxVM volumes (for each disk group) that VCS does not manage.

5 Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

6 Freeze the nodes in the first subcluster.

```
# hasys -freeze -persistent node01
# hasys -freeze -persistent node02
```

7 Dump the configuration and make it read-only.

```
# haconf -dump -makero
```

8 Verify that the service groups are offline on the first subcluster that you want to upgrade.

```
# hagrp -state
```

Output resembles:

```
#Group Attribute System Value
sg1  State node01 |OFFLINE|
sg1  State node02 |OFFLINE|
sg1  State node03 |ONLINE|
sg1  State node04 |ONLINE|
sg2  State node01 |OFFLINE|
sg2  State node02 |OFFLINE|
sg2  State node03 |ONLINE|
sg2  State node04 |ONLINE|
sg3  State node01 |OFFLINE|
sg3  State node02 |OFFLINE|
sg3  State node03 |ONLINE|
sg3  State node04 |OFFLINE|
sg4  State node01 |OFFLINE|
sg4  State node02 |OFFLINE|
sg4  State node03 |OFFLINE|
sg4  State node04 |ONLINE|
```

Upgrading the operating system on the first subcluster

You can perform the operating system upgrade on the first subcluster, if required.

Before performing operating system upgrade, it is better to prevent LLT from starting automatically when the node starts. For example, you can do the following:

```
# mv /etc/llttab /etc/llttab.save
```

or you can change the `/etc/default/llt` file by setting LLT_START = **0**.

After you finish upgrading the OS, remember to change the LLT configuration to its original configuration.

Refer to the operating system's documentation for more information.

Upgrading the first subcluster

You now navigate to the installer program and start it.

To start the installer for the phased upgrade

1. Confirm that you are logged on as the superuser and you mounted the product disc.
2. Make sure that you can ssh or rsh from the node where you launched the installer to the nodes in the second subcluster without requests for a password.
3. Navigate to the folder that contains installvcs.

   ```
   # cd cluster_server
   ```

4. Start the installvcs program, specify the nodes in the first subcluster (node1 and node2).

   ```
   # ./installvcs node1 node2
   ```

   The program starts with a copyright message and specifies the directory where it creates the logs.

5. Enter **y** to agree to the End User License Agreement (EULA).

   ```
   Do you agree with the terms of the End User License Agreement
   as specified in the
   cluster_server/EULA/<lang>/EULA_SFHA_Ux_<version>.pdf
   file present on media? [y,n,q,?] y
   ```

6 Review the available installation options.

See “Symantec Cluster Server installation filesets” on page 495.

1 Installs only the minimal required VCS filesets that provides basic functionality of the product.

2 Installs the recommended VCS filesets that provide complete functionality of the product. This option does not install the optional VCS filesets.

Note that this option is the default.

3 Installs all the VCS filesets.

You must choose this option to configure any optional VCS feature.

4 Displays the VCS filesets for each option.

For this example, select 3 for all filesets.

Select the filesets to be installed on all systems? [1-4,q,?] (2) 3

7 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility.

8 When you are prompted, reply **y** to continue with the upgrade.

```
Do you want to continue? [y,n,q] (y)
```

9 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]
  - To grant read access to all authenticated users, type **y**.
  - To grant usergroup specific permissions, type **n**.
- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some

usergroups are not created yet, create the usergroups after configuration if needed. [b]

10 When you are prompted, reply **y** to stop appropriate processes.

```
Do you want to stop VCS processes? [y,n,q] (y)
```

11 The installer ends for the first subcluster with the following output:

```
    Configuring VCS: 100%

    Estimated time remaining: 0:00

    Performing VCS upgrade configuration ..................... Done

Symantec Cluster Server Configure completed successfully


You are performing phased upgrade (Phase 1) on the systems.
Follow  the steps in install guide to upgrade the remaining
systems.

Would you like to send the information about this installation to
Symantec to help improve installation in the future? [y,n,q,?] (y)
```

The upgrade is finished on the first subcluster. Do not reboot the nodes in the first subcluster until you complete the Preparing the second subcluster procedure.

12 In the `/etc/default/llt` file, set LLT_START = 0.

Preparing the second subcluster

Perform the following steps on the second subcluster before rebooting nodes in the first subcluster.

To prepare to upgrade the second subcluster

1 Get the summary of the status of your resources.

```
# hastatus -summ
-- SYSTEM STATE
-- System               State                Frozen

A  node01               EXITED                1
A  node02               EXITED                1
A  node03               RUNNING               0
A  node04               RUNNING               0

-- GROUP STATE
-- Group           System  Probed     AutoDisabled    State

B  SG1             node01   Y          N               OFFLINE
B  SG1             node02   Y          N               OFFLINE
B  SG1             node03   Y          N               ONLINE
B  SG1             node04   Y          N               ONLINE
B  SG2             node01   Y          N               OFFLINE
B  SG2             node02   Y          N               OFFLINE
B  SG2             node03   Y          N               ONLINE
B  SG2             node04   Y          N               ONLINE
B  SG3             node01   Y          N               OFFLINE
B  SG3             node02   Y          N               OFFLINE
B  SG3             node03   Y          N               ONLINE
B  SG3             node04   Y          N               OFFLINE
B  SG4             node01   Y          N               OFFLINE
B  SG4             node02   Y          N               OFFLINE
B  SG4             node03   Y          N               OFFLINE
B  SG4             node04   Y          N               ONLINE
```

2 Unmount all the VxFS file systems that VCS does not manage, for example:

```
# df -k
```

```
Filesystem              1024-blocks     Free   %Used    Iused %Iused Mounted on
/dev/hd4                   20971520  8570080    60%    35736     2% /
/dev/hd2                    5242880  2284528    57%    55673     9% /usr
/dev/hd9var                 4194304  3562332    16%     5877     1% /var
/dev/hd3                    6291456  6283832     1%      146     1% /tmp
/dev/hd1                     262144   261408     1%       62     1% /home
/dev/hd11admin               262144   184408    30%        6     1% /admin
/proc                           -        -      -        -      -  /proc
/dev/hd10opt               20971520  5799208    73%    65760     5% /opt
/dev/vx/dsk/dg2/dg2vol1       10240     7600    26%        4     1% /mnt/dg2/dg2vol1
/dev/vx/dsk/dg2/dg2vol2       10240     7600    26%        4     1% /mnt/dg2/dg2vol2
/dev/vx/dsk/dg2/dg2vol3       10240     7600    26%        4     1% /mnt/dg2/dg2vol3
```

```
# umount /mnt/dg2/dg2vol1
# umount /mnt/dg2/dg2vol2
# umount /mnt/dg2/dg2vol3
```

3 Take the service groups offline on node03 and node04.

```
# hagrp -offline sg1 -sys node03
# hagrp -offline sg1 -sys node04
# hagrp -offline sg2 -sys node03
# hagrp -offline sg2 -sys node04
# hagrp -offline sg3 -sys node03
# hagrp -offline sg4 -sys node04
```

4 Verify the state of the service groups.

```
# hagrp -state
#Group       Attribute    System     Value
SG1            State      node01    |OFFLINE|
SG1            State      node02    |OFFLINE|
SG1            State      node03    |OFFLINE|
SG1            State      node04    |OFFLINE|
SG2            State      node01    |OFFLINE|
SG2            State      node02    |OFFLINE|
SG2            State      node03    |OFFLINE|
SG2            State      node04    |OFFLINE|
SG3            State      node01    |OFFLINE|
SG3            State      node02    |OFFLINE|
SG3            State      node03    |OFFLINE|
SG3            State      node04    |OFFLINE|
```

5 Stop all VxVM volumes (for each disk group) that VCS does not manage.

6 Stop VCS, I/O Fencing, GAB, and LLT on node03 and node04.

```
# /opt/VRTSvcs/bin/hastop -local
# /etc/init.d/vxfen.rc stop
# /etc/init.d/gab.rc stop
# /etc/init.d/llt.rc stop
```

7 Make sure that the VXFEN, GAB, and LLT modules on node03 and node04 are not loaded.

```
# /sbin/vxfenconfig -l
VXFEN vxfenconfig ERROR V-11-2-1087 There are 0 active
coordination points for this node

# /sbin/gabconfig -l
GAB Driver Configuration
Driver state         : Unconfigured
Partition arbitration: Disabled
Control port seed    : Disabled
Halt on process death: Disabled
Missed heartbeat halt: Disabled
Halt on rejoin       : Disabled
Keep on killing      : Disabled
Quorum flag          : Disabled
Restart              : Disabled
Node count           : 0
Disk HB interval (ms): 1000
Disk HB miss count   : 4
IOFENCE timeout (ms) : 15000
Stable timeout (ms)  : 5000

# /usr/sbin/strload -q -d /usr/lib/drivers/pse/llt
/usr/lib/drivers/pse/llt: no
```

Activating the first subcluster

Get the first subcluster ready for the service groups.

Note: These steps fulfill part of the installer's output instructions, see Upgrading the first subcluster step 11.

To activate the first subcluster

1. Start LLT and GAB on one node in the first half of the cluster..
2. Seed node01 in the first subcluster.

```
# gabconfig -x
```

3. On the first half of the cluster, start VCS:

```
# cd /opt/VRTS/install
# ./installvcs<version> -start sys1 sys2
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 50.

4. Make the configuration writable on the first subcluster.

```
# haconf -makerw
```

5. Unfreeze the nodes in the first subcluster.

```
# hasys -unfreeze -persistent node01
# hasys -unfreeze -persistent node02
```

6. Dump the configuration and make it read-only.

```
# haconf -dump -makero
```

7. Bring the service groups online on node01 and node02.

```
# hagrp -online sg1 -sys node01
# hagrp -online sg1 -sys node02
# hagrp -online sg2 -sys node01
# hagrp -online sg2 -sys node02
# hagrp -online sg3 -sys node01
# hagrp -online sg4 -sys node02
```

Upgrading the operating system on the second subcluster

You can perform the operating system upgrade on the second subcluster, if required.

Before performing operating system upgrade, it is better to prevent LLT from starting automatically when the node starts. For example, you can do the following:

```
# mv /etc/lltab /etc/lltab.save
```

or you can change the `/etc/default/llt` file by setting LLT_START = **0**.

After you finish upgrading the OS, remember to change the LLT configuration to its original configuration.

Refer to the operating system's documentation for more information.

Upgrading the second subcluster

Perform the following procedure to upgrade the second subcluster (node03 and node04).

To start the installer to upgrade the second subcluster

1. Confirm that you are logged on as the superuser and you mounted the product disc.
2. Navigate to the folder that contains installvcs.

   ```
   # cd cluster_server
   ```

3. Confirm that VCS is stopped on node03 and node04. Start the installvcs program, specify the nodes in the second subcluster (node3 and node4).

   ```
   # ./installvcs node3 node4
   ```

   The program starts with a copyright message and specifies the directory where it creates the logs.
4. Enter **y** to agree to the End User License Agreement (EULA).

   ```
   Do you agree with the terms of the End User License Agreement
   as specified in the
   cluster_server/EULA/<lang>/EULA_VCS_Ux_<version>.pdf
   file present on media? [y,n,q,?] y
   ```

5 Review the available installation options.

See “Symantec Cluster Server installation filesets” on page 495.

1. Installs only the minimal required VCS filesets that provides basic functionality of the product.
2. Installs the recommended VCS filesets that provide complete functionality of the product. This option does not install the optional VCS filesets.

   Note that this option is the default.
3. Installs all the VCS filesets.

   You must choose this option to configure any optional VCS feature.
4. Displays the VCS filesets for each option.

   For this example, select 3 for all filesets.

   Select the filesets to be installed on all systems? [1-4,q,?] (2) 3

6 The installer performs a series of checks and tests to ensure communications, licensing, and compatibility.

7 When you are prompted, reply **y** to continue with the upgrade.

```
Do you want to continue? [y,n,q] (y)
```

8 When you are prompted, reply **y** to stop VCS processes.

```
Do you want to stop VCS processes? [y,n,q] (y)
```

9 Monitor the installer program answering questions as appropriate until the upgrade completes.

Finishing the phased upgrade

Complete the following procedure to complete the upgrade.

To finish the upgrade

1 Verify that the cluster UUID is the same on the nodes in the second subcluster and the first subcluster. Run the following command to display the cluster UUID:

```
# /opt/VRTSvcs/bin/uuidconfig.pl
-clus -display node1 [node2 ...]
```

If the cluster UUID differs, manually copy the cluster UUID from a node in the first subcluster to the nodes in the second subcluster. For example:

```
# /opt/VRTSvcs/bin/uuidconfig.pl [-rsh] -clus
-copy -from_sys node01 -to_sys node03 node04
```

2 Reboot the node03 and node04 in the second subcluster.

```
# /usr/sbin/shutdown -r
```

The nodes in the second subcluster join the nodes in the first subcluster.

3 In the `/etc/default/llt` file, change the value of the LLT_START attribute.

In the `/etc/default/gab` file, change the value of the GAB_START attribute.

In the `/etc/default/vxfen` file, change the value of the VXFEN_START attribute.

In the `/etc/default/vcs` file, change the value of the VCS_START attribute.

```
LLT_START = 1
GAB_START = 1
VXFEN_START =1
VCS_START =1
```

4 Start LLT and GAB.

```
# /etc/init.d/llt.rc start
```

```
# /etc/init.d/gab.rc start
```

5 Seed node03 and node04 in the second subcluster.

```
# gabconfig -x
```

6 On the second half of the cluster, start VCS:

```
# cd /opt/VRTS/install

# ./installvcs<version> -start sys3 sys4
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 50.

7 Check to see if VCS and its components are up.

```
# gabconfig -a
GAB Port Memberships
===============================================================
Port a gen   nxxxnn membership 0123
Port b gen   nxxxnn membership 0123
Port h gen   nxxxnn membership 0123
```

8 Run an `hastatus -sum` command to determine the status of the nodes, service groups, and cluster.

```
# hastatus -sum

-- SYSTEM STATE
-- System          State           Frozen

A  node01          RUNNING         0
A  node02          RUNNING         0
A  node03          RUNNING         0
A  node04          RUNNING         0

-- GROUP STATE
-- Group    System    Probed   AutoDisabled   State
B  sg1      node01    Y        N              ONLINE
B  sg1      node02    Y        N              ONLINE
B  sg1      node03    Y        N              ONLINE
B  sg1      node04    Y        N              ONLINE
B  sg2      node01    Y        N              ONLINE
B  sg2      node02    Y        N              ONLINE
B  sg2      node03    Y        N              ONLINE
B  sg2      node04    Y        N              ONLINE
B  sg3      node01    Y        N              ONLINE
B  sg3      node02    Y        N              OFFLINE
B  sg3      node03    Y        N              OFFLINE
B  sg3      node04    Y        N              OFFLINE
B  sg4      node01    Y        N              OFFLINE
B  sg4      node02    Y        N              ONLINE
B  sg4      node03    Y        N              OFFLINE
B  sg4      node04    Y        N              OFFLINE
```

9 After the upgrade is complete, start the VxVM volumes (for each disk group) and mount the VxFS file systems.

In this example, you have performed a phased upgrade of VCS. The service groups were down when you took them offline on node03 and node04, to the time VCS brought them online on node01 or node02.

Note: If you want to upgrade application clusters that use CP server based fencing to 6.2, make sure that you first upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.2. However, note that the CP server upgraded to 6.2 can support application clusters on 6.2 (HTTPS-based communication) and application clusters prior to 6.2 (IPM-based communication). When you configure the CP server, the installer asks the VIPs for HTTPS-based communication (if the clients are on release version 6.2) or VIPs for IPM-based communication (if the clients are on a release version prior to 6.2).

For instructions to upgrade VCS or SFHA, see the VCS or SFHA Installation Guide.

Chapter 24

Performing an automated VCS upgrade using response files

This chapter includes the following topics:

Upgrading VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS upgrade on one system to upgrade VCS on other systems.

You can also create a response file using the `makeresponsefile` option of the installer.

```
# ./installer -makeresponsefile
```

To perform automated VCS upgrade

1. Make sure the systems where you want to upgrade VCS meet the upgrade requirements.
2. Make sure the pre-upgrade tasks are completed.
3. Copy the response file to the system where you want to upgrade VCS.

   See "Sample response file for upgrading VCS" on page 378.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to upgrade VCS" on page 376.
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

   # ./installvcs -responsefile /tmp/response_file
   ```

   Where /tmp/*response_file* is the response file's full path name.

Response file variables to upgrade VCS

Table 24-1 lists the response file variables that you can define to upgrade VCS.

Table 24-1 Response file variables specific to upgrading VCS

Variable	List or Scalar	Description
CFG{opt}{upgrade}	Scalar	Upgrades VCS filesets. (Required)
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 upgraded. (Required)

Table 24-1 Response file variables specific to upgrading VCS *(continued)*

Variable	List or Scalar	Description
CFG{prod}	Scalar	Defines the product to be upgraded. The value is VCS62 for VCS. (Optional)
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{opt}{keyfile}	Scalar	Defines the location of an ssh keyfile that is used to communicate with all remote systems. (Optional)
CFG{opt}{pkgpath}	Scalar	Defines a location, typically an NFS mount, from which all remote systems can install product filesets. The location must be accessible from all target systems. (Optional)
CFG{opt}{tmppath}	Scalar	Defines the location where a working directory is created to store temporary files and the filesets that are needed during the install. The default location is /var/tmp. (Optional)
CFG{secusrgrps}	List	Defines the user groups which get read access to the cluster. (Optional)

Table 24-1 Response file variables specific to upgrading VCS *(continued)*

Variable	List or Scalar	Description
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. **Note:** The installer copies the response files and summary files also to the specified *logpath* location. (Optional)
CFG{opt}{rsh}	Scalar	Defines that *rsh* must be used instead of ssh as the communication method between systems. (Optional)

Sample response file for upgrading VCS

Review the response file variables and their definitions.

See "Response file variables to upgrade VCS" on page 376.

```
#
# Configuration Values:
#
our %CFG;

$CFG{accepteula}=1;
$CFG{secusrgrps}=qw{staff usergroup@hostname.cdc.symantec.com}
$CFG{vcs_allowcomms}=1;
$CFG{opt}{upgrade}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw( sys1 sys2) ];
1;
```

Performing rolling upgrade of VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS upgrade on one system to upgrade VCS on other systems.

You can also create a response file using the `makeresponsefile` option of the installer.

To perform automated VCS rolling upgrade

1. Make sure the systems where you want to upgrade VCS meet the upgrade requirements.
2. Make sure the pre-upgrade tasks are completed.
3. Copy the response file to the systems where you want to launch the installer.

   See "Sample response file for VCS using rolling upgrade" on page 381.
4. Edit the values of the response file variables as necessary.

   See "Response file variables to upgrade VCS using rolling upgrade" on page 379.
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

   # ./installvcs -responsefile /tmp/response_file
   ```

   Where /tmp/*response_file* is the response file's full path name.

Response file variables to upgrade VCS using rolling upgrade

Table 24-2 lists the response file variables that you can define to upgrade VCS using rolling upgrade.

Table 24-2 Response file variables for upgrading VCS using rolling upgrade

Variable	Description
CFG{phase1}{0}	A series of $CFG{phase1}{N} items define sub-cluster division. The index N indicatse the order to do RU phase1. The index starts from 0. Each item has a list of node(at least 1). List or scalar: list Optional or required: conditional required Required if rolling upgrade phase1 needs to be performed.
CFG{rollingupgrade_phase2}	The `CFG{rollingupgrade_phase2}` option is used to perform rolling upgrade Phase 2. In the phase, VCS and other agent filesets upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version. List or scalar: scalar Optional or required: conditional required Required if rolling upgrade phase 2 needs to be performed.
CFG{rolling_upgrade}	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade Phase 1 or Phase 2 explicitly.
CFG{systems}	List of systems on which the product is to be installed or uninstalled. List or scalar: list Optional or required: required
CFG{opt}{upgrade}	Upgrades all filesets installed. List or scalar: scalar Optional or required: optional

Table 24-2 Response file variables for upgrading VCS using rolling upgrade *(continued)*

Variable	Description
CFG{secusrgrps}	Defines the user groups which get read access to the cluster. List or scalar: list Optional or required: optional
CFG{rootsecusrgrps}	Defines the read access to the cluster from root users, specific users, or usergroups based on your choice. The selected users or usergroups get explicit privileges on VCS objects. List or scalar: scalar Optional or required: Optional
CFG{accepteula}	Specifies whether you agree with the EULA.pdf file on the media. List or scalar: scalar Optional or required: required

Sample response file for VCS using rolling upgrade

The following example shows a response file for VCS using Rolling Upgrade.

```
our %CFG;
$CFG{accepteula}=1;
$CFG{client_vxfen_warning}=1;
$CFG{fencing_cps}=[ qw(10.198.90.6) ];
$CFG{fencing_cps_ports}{"10.198.90.6"}=50006;
$CFG{fencing_cps_vips}{"10.198.90.6"}=[ qw(10.198.90.6) ];
$CFG{opt}{gco}=1;
$CFG{opt}{noipc}=1;
$CFG{opt}{rolling_upgrade}=1;
$CFG{opt}{rollingupgrade_phase2}=1;
$CFG{opt}{updatekeys}=1;
$CFG{opt}{upgrade}=1;
$CFG{secusrgrps}=qw{staff pilotaix218@cdc.veritas.com};
$CFG{opt}{vr}=1;
$CFG{phase1}{"0"}=[ qw(sys3 sys2) ];
$CFG{phase1}{"1"}=[ qw(sys1) ];
```

```
$CFG{systems}=[ qw(sys1 sys2 sys3) ];
$CFG{vcs_allowcomms}=1;
1;
```

Chapter 25

Performing a rolling upgrade

This chapter includes the following topics:

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 filesets in phase 1 and VCS agent related filesets in phase 2.

Note: You need to perform a rolling upgrade on a completely configured cluster.

If the Oracle agent is configured, set the `MonitorFrequency` to 1 to ensure proper functioning of traditional monitoring during the upgrade.

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 Symantec Cluster Server (VCS) engine HAD, but does not include application downtime.

Figure 25-1 illustrates an example of the installer performing a rolling upgrade for three service groups on a two node cluster.

Figure 25-1 Example of the installer performing a rolling upgrade

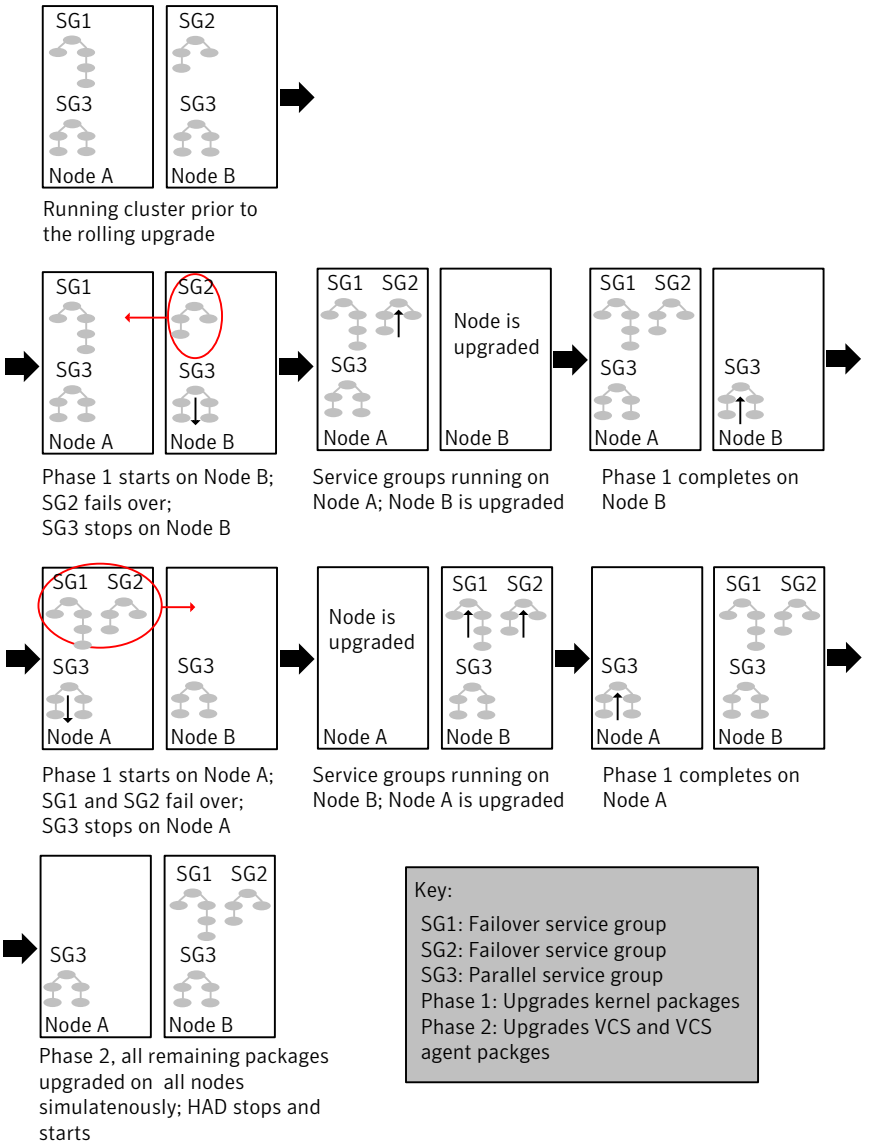

The following limitations apply to rolling upgrades:

- 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 VCS with the script-based installer, the web-based installer, or manually.

The rolling upgrade procedures support only minor operating system upgrades.

Table 25-1 shows the versions of VCS for which you can perform a rolling upgrade to VCS 6.2.

Table 25-1 Supported rolling upgrade paths

Platform	VCS version
AIX 6.1	5.1, 5.1RPs 5.1SP1, 5.1SP1RPs 6.0, 6.0RP1 6.0.1, 6.0.3, 6.0.5 6.1, 6.1.1
AIX 7.1	5.1SP1RPs, 5.1SP1PR1 6.0, 6.0RP1 6.0.1, 6.0.3, 6.0.5 6.1, 6.1.1

Note: Before performing a rolling upgrade from version 5.1SP1RP3 to version 6.2, install patch VRTSvxfen-5.1SP1RP3P2. For downloading the patch, search VRTSvxfen-5.1SP1RP3P2 in Patch Lookup on the SORT website.

Performing a rolling upgrade using the installer

Use a rolling upgrade to upgrade Symantec Cluster Server to the latest release with minimal application downtime.

Performing a rolling upgrade using the script-based installer

Before you start the rolling upgrade, make sure that Symantec Cluster Server (VCS) is running on all the nodes of the cluster.

To perform a rolling upgrade

1 Complete the preparatory steps on the first sub-cluster.

Unmount all VxFS file systems not under VCS control:

```
# umount mount_point
```

2 Log in as superuser and mount the VCS 6.2 installation media.

3 From root, start the installer.

```
# ./installer
```

4 From the menu, select `Upgrade a Product` and from the sub menu, select `Rolling Upgrade`.

5 The installer suggests system names for the upgrade. Press **Enter** to upgrade the suggested systems, or enter the name of any one system in the cluster on which you want to perform a rolling upgrade and then press **Enter**.

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 detects the online service groups, the installer prompts the user to do one of the following:

- Manually switch service groups
- Use the CPI to automatically switch service groups

The downtime is the time that it normally takes for the service group's failover.

Note: It is recommended that you manually switch the service groups. Automatic switching of service groups does not resolve dependency issues.

---

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 filesets, and installs the new filesets. The installer asks if you want to update your licenses to the current version. Select **Yes** or **No**. Symantec recommends that you update your licenses to fully use the new features in the current release.

13 If the cluster has configured Coordination Point Server based fencing, then during upgrade, installer may ask the user to provide the new HTTPS Coordination Point Server.

The installer performs the upgrade configuration and starts the processes. If the boot disk is encapsulated before the upgrade, installer prompts the user to reboot the node after performing the upgrade configuration.

14 Complete the preparatory steps on the nodes that you have not yet upgraded.

Unmount all VxFS file systems not under VCS control on all the nodes.

```
# umount mount_point
```

15 The installer begins phase 1 of the upgrade on the remaining node or nodes. Type **y** to continue the rolling upgrade. If the installer was invoked on the upgraded (rebooted) nodes, you must invoke the installer again.

If the installer prompts to restart nodes, restart 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.

16 When Phase 1 of the rolling upgrade completes, mount all the VxFS file systems that are not under VCS control manually. Begin Phase 2 of the upgrade. Phase 2 of the upgrade includes downtime for the VCS engine (HAD), which does not include application downtime. Type **y** to continue.

17 The installer determines the remaining filesets to upgrade. Press **Enter** to continue.

18 The installer displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

- Do you want to grant read access to everyone? [y,n,q,?]

- To grant read access to all authenticated users, type **y**.
- To grant usergroup specific permissions, type **n**.

- Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
  - To specify usergroups and grant them read access, type **y**
  - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
- Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

19 Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

20 The installer stops Symantec Cluster Server (VCS) processes but the applications continue to run. Type **y** to continue.

The installer performs prestop, uninstalls old filesets, and installs the new filesets. It performs post-installation tasks, and the configuration for the upgrade.

21 If you have network connection to the Internet, the installer checks for updates.

If updates are discovered, you can apply them now.

22 A prompt message appears to ask if the user wants to read the summary file. You can choose **y** if you want to read the install summary file.

23 Upgrade application to the supported version.

24 If you want to upgrade application clusters that use CP server-based fencing to 6.2, make sure that you upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.2. However, note that the CP server upgraded to 6.2 can support application clusters on 6.1 and later (HTTPS-based communication) and application clusters prior to 6.1 (IPM-based communication). When you configure the CP server, the installer asks the VIPs for HTTPS-based communication (if the clients are on release version 6.1 and later) or VIPs for IPM-based communication (if the clients are on a release version prior to 6.1).

For instructions to upgrade VCS or SFHA on the CP server systems, refer to the appropriate installation guide.

Performing a rolling upgrade of VCS using the web-based installer

This section describes using the web-based installer to perform a rolling upgrade. The installer detects and upgrades the product that is currently installed on the specified system or systems. If you want to upgrade to a different product, you may need to perform additional steps.

See "About rolling upgrades" on page 383.

To start the rolling upgrade—phase 1

1 Perform the required steps to save any data that you want to preserve. For example, take backups of configuration files.

2 Start the web-based installer.

See "Starting the web-based installer" on page 186.

3 In the Task pull-down menu, select `Rolling Upgrade`.

The option `Phase-1: Upgrade Kernel packages` is displayed and selected by default.

Click **Next** to proceed.

4 Enter the name of any one system in the cluster on which you want to perform a rolling upgrade.The installer identifies the cluster information of the system and displays the information.

Click **Yes** to confirm the cluster information. The installer now displays the nodes in the cluster that will be upgraded during phase 1 of the upgrade.

5 Review the systems that the installer has chosen for phase 1 of the rolling upgrade. These systems are chosen to minimize downtime during the upgrade.

Click **Yes** to proceed.

The installer validates systems.

6 Review the End User License Agreement (EULA). To continue, select **Yes, I agree** and click **Next**.

7 If you have online failover service groups, the installer prompts you to choose to switch these service groups either manually or automatically. Choose any option and follow the steps to switch all the failover service groups to the other subcluster.

8 The installer stops all processes. Click **Next** to proceed.

The installer removes old software and upgrades the software on the systems that you selected.

9 The installer asks if you want to update your licenses to the current version. Select **Yes** or **No**. Symantec recommends that you update your licenses to fully use the new features in the current release.

10 If the cluster has configured Coordination Point Server-based fencing, then during upgrade, installer asks the user to provide the new HTTPS Coordination Point Server. If you are prompted, restart the product.

The installer starts all the relevant processes and brings all the service groups online if the nodes do not require a restart.

11 Restart the nodes, if required.

Restart the installer.

12 Repeat step 5 through step 11 until the kernel filesets of all the nodes are upgraded. For clusters with larger number of nodes, this process may get repeated several times. Service groups come down and are brought up to accommodate the upgrade.

13 When prompted, perform step 3 through step 11 on the nodes that you have not yet upgraded.

14 When prompted, start phase 2. Click **Yes** to continue with the rolling upgrade.

You may need to restart the web-based installer to perform phase 2.

See “Starting the web-based installer” on page 186.

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 displays the following question before the install stops the product processes. If the cluster was not configured in secure mode before the upgrade, these questions are not displayed.

   - Do you want to grant read access to everyone? [y,n,q,?]
     - To grant read access to all authenticated users, type **y**.
     - To grant usergroup specific permissions, type **n**.
   - Do you want to provide any usergroups that you would like to grant read access?[y,n,q,?]
     - To specify usergroups and grant them read access, type **y**
     - To grant read access only to root users, type **n**. The installer grants read access read access to the root users.
   - Enter the usergroup names separated by spaces that you would like to grant read access. If you would like to grant read access to a usergroup on a specific node, enter like 'usrgrp1@node1', and if you would like to grant read access to usergroup on any cluster node, enter like 'usrgrp1'. If some usergroups are not created yet, create the usergroups after configuration if needed. [b]

5. The installer stops the `HAD` and `CmdServer` processes in phase 2 of the rolling upgrade process but the applications continue to run. Click **Next** to proceed.

6. The installer removes old software and upgrades the software on the systems that you selected. Review the output and click the **Next** button when prompted. Register the software and click **Next** to proceed. The installer starts all the relevant processes and brings all the service groups online.

7. If you have network connection to the Internet, the installer checks for updates.

   If updates are discovered, you can apply them now.

8. A prompt message appears to ask if the user wants to read the summary file. You can choose **y** if you want to read the install summary file.

The upgrade is complete.

Chapter 26

Upgrading VCS using Network Install Manager Alternate Disk Migration

This chapter includes the following topics:

Supported upgrade paths for VCS using NIM ADM

You can perform an upgrade of the product and the operating system using Network Install Manager Alternate Disk Migration (NIM ADM).

The supported upgrade paths are as follows:

AIX version	AIX 5.3
	AIX 6.1
	AIX 7.1
VCS version	5.1 and later

Preparing to upgrade VCS and the operating system using the nimadm utility

Complete the preparatory steps in the following procedure before you upgrade VCS and the operating system.

To prepare to upgrade VCS and the operating system using the nimadm utility

1. Make sure that the VCS installation media is available.
2. Check the status of the physical disks on each node in the cluster.

   Note: The alternate disk must have a physical identifier and must not contain any mounted volume groups.

   ```
   # lspv
   Output similar to the following displays:
   hdisk0          0009710fa9c79877    rootvg    active
   hdisk1          0009710f0b90db93    None
   ```

   If the alternate disk does not have a physical identifier, set the physical identifier for the disk:

   ```
   # chdev -l hdisk1 -a pv=yes
   ```

3. Make sure that the following filesets are installed on the NIM server and the client: bos.alt_disk_install.boot_images, bos.alt_disk_install.rte

Preparing the installation bundle on the NIM server

You need to prepare the installation bundle installp on the NIM server before you use nimadm to upgrade VCS filesets. The following actions are executed on the NIM server.

Note: Make sure that a NIM LPP_SOURCE is present on the NIM server.

To prepare the installation bundle

1. Insert and mount the product installation media.
2. Choose an LPP source:

   ```
   # lsnim |grep -i lpp_source
   LPP-7100-up2date resources lpp_source
   ```

3 Check that the NIM LPP_RESOURCE and corresponding SPOT are in healthy state before you start upgrade:

```
# nim -Fo check LPP-7100-up2date
# nim -Fo check SPOT-7100-up2date
```

4 Navigate to the product directory on the installation media and run the `installvcs` command to prepare the bundle resource:

```
# ./installvcs -nim LPP-7100-up2date
```

The installation program copies the necessary filesets and patches to the LPP resource directory.

5 Enter a name for the bundle, for example *VCS62*.

6 Run the `lsnim -l` command to check that the installp_bundle resource is created successfully.

```
# lsnim -l VCS62
VCS62:
class = resources
type = installp_bundle
Rstate = ready for use
prev_state = unavailable for use
location = /opt/VRTS/nim/VCS62.bundle
alloc_count = 0
server = master
```

Upgrading VCS and the operating system using the `nimadm` utility

This section provides instructions to upgrade VCS and the operating system using the `nimadm` utility. You may perform the steps manually or using the SMIT interface.

In the procedure examples, `hdisk0` indicates the primary or current boot environment and `hdisk1` indicates the alternate or inactive boot environment.

To upgrade VCS and the operating system in a high-availability environment using the `nimadm` utility

Perform the instructions on each node in the cluster from the NIM server.

1 On the primary boot disk, change `/etc/default/llt` to start LLT on the nodes by setting the`LLT_START` attribute to 1:

```
LLT_START=1
```

2 From the NIM server, clone the primary boot disk `rootvg` to an alternate disk.

Manual

Upgrade VCS and the operating system by running the following command on the NIM server:

```
# nimadm -l lpp_source -c nim_client \
-s spot_name -b bundle_name \
-d nimclient_altdisk_name -Y
```

For example:

```
# nimadm -l LPP-7100-up2date -c node1 \
-s spot-7100-up2date -b vcs62 \
-d hdisk1 -Y
```

Using SMIT interface

Start the SMIT menu:

```
# smit nimadm
```

Select the option **Perform NIM Alternate Disk Migration**.

Enter the required information at the prompts:

- Target NIM Client: **sys1**
- NIM LPP_SOURCE resource: **LPP-7100-up2date**
- NIM SPOT resource: **SPOT-7100-up2date**
- Bundle name: **vcs62**
- Target disk(s) to install: **hdisk1**
- Phase to execute: **all**
- Set Client `bootlist` to alternate disk? **yes**
- ACCEPT new license agreements? **yes**

Press **Enter** to start the upgrade on the alternate disk. The upgrade process takes some time.

3 Set the environment variable `FORCE` to `yes` on the alternate boot disk with the upgraded operating system. Perform this step on each node in the cluster.

```
# export FORCE=yes
```

4 Check that the operating system version is correct. Enter the following:

```
# oslevel -s
```

5 Wake up the volume group on the alternate boot disk (hdisk1) that you cloned by running the following command on each node in the cluster:

```
# /usr/sbin/alt_rootvg_op -W -d hdisk1
```

6 Verify that the alternate disk is created:

```
# lspv
Output similar to the following displays:
hdisk0          0009710fa9c79877    rootvg
hdisk1          0009710f0b90db93    altinst_rootvg
```

7 Change directory to /alt_inst/etc/VRTSvcs/conf/config.

```
# cd /alt_inst/etc/VRTSvcs/conf/config
```

8 Back up a copy of the old types.cf file and copy the new one for VCS to use.

```
# mv types.cf types.cf.ORIG
# cp ../types.cf .
```

9 If you did not configure fencing in the existing cluster, but want to use it in your upgraded cluster, perform the instructions in the following section

See "Configuring fencing for an ADI upgrade" on page 409.

10 Move to root and run the `alt_rootvg_op -S` command to put the alternate root to sleep.

```
# cd /
# alt_rootvg_op -S
```

11 Verify that the normal boot list includes the name of the alternate boot disk. By default, the alternate disk installation process changes the boot list to enable the system to boot from the alternate disk.

```
# bootlist -m normal -o
hdisk1
```

12 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

```
# hagrp -offline grp_name -any
```

13 Stop VCS on all nodes:

```
# hastop -all
```

14 Change the `/etc/default/llt` file to prevent LLT from starting automatically after restart by setting the `LLT_START` attribute to 0:

```
LLT_START=0
```

This step ensures that VCS remains operable on the current primary disk in case the alternate disk upgrade fails.

15 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

```
# shutdown -r
```

16 Copy the product installation scripts to the alternate disk:

```
# /opt/VRTS/install/bin/UXRT<version>/add_install_scripts
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 50.

The command copies the installation scripts and uninstallation scripts to the alternate disk.

17 Verify the upgrade.

See "Verifying the upgrade performed using the NIM ADM utility" on page 401.

18 After the systems have booted into their alternate environments, initialize the VxVM disks by running the following command on each node in the cluster:

```
# vxinstall
```

19 On the alternate disk, change `/etc/default/llt` to start LLT on the nodes by setting the LLT_START attribute to 1:

```
LLT_START=1
```

20 Start VCS:

```
# cd /opt/VRTS/install

# installvcs<version> -start
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 50.

21 Verify that all GAB ports are up:

```
# gabconfig -a
```

22 Complete the post-upgrade tasks.

See the chapter "Performing post-upgrade tasks" in this document.

23 If you want to upgrade the CP server systems that use VCS or SFHA to 6.2, make sure that you upgraded all application clusters to version 6.2. Then, upgrade VCS or SFHA on the CP server systems.

For instructions to upgrade VCS or SFHA on the CP server systems, see the *VCS or SFHA installation guide*.

Note: If the operating system version is incorrect, and the `bos.txt.spell` and `bos.txt.tfs` filesets are missed, update these filesets manually through `nim <os_version> lpp_source`.

```
# oslevel -rl 6100-07
Fileset              Actual Level        Recommended ML
-------------------------------------------------------------
bos.txt.spell        5.3.12.0            6.1.6.0
bos.txt.tfs          5.3.12.0            6.1.6.0
```

To update the `bos.txt.spell` fileset manually, do the following:

smitty nim >> Perform NIM Software Installation and Maintenance Tasks >> Install and Update Software >> Install Software >> Select corresponding LPP_SOURCE >> * Software to Install >> Select bos.txt.spell

Follow the same procedure for the `bos.txt.tfs` fileset.

Verifying the upgrade performed using the NIM ADM utility

To ensure that alternate disk installation has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

To verify the upgrade using the NIM ADM utility

1 Verify that the alternate boot environment is active:

```
# lspv | grep rootvg
hdisk0          0009710fa9c79877    old_rootvg
hdisk1          0009710f0b90db93    rootvg          active
```

If there are multiple boot disks for rootvg, run the following command to verify the disk the system has booted from:

```
# bootlist -m normal -o
hdisk1 blv=hd5 pathid=0
```

2 Verify that the version of the upgraded filesets on the alternate boot disk is `6.2.0.0`.

```
# lslpp -h VRTSvcs
```

Note: The `VRTSsfcpi61` fileset still exists on the alternate boot disk. You need to manually uninstall the fileset.

If you upgraded the operating system:

```
# oslevel -s
```

Chapter 27

Upgrading VCS using an alternate disk

This chapter includes the following topics:

About upgrading VCS using an alternate disk

Use the alternate disk installation process to upgrade the operating system and VCS on a production server while the server runs. Perform the upgrade on an alternate or inactive boot environment. After the upgrade, restart the system on the alternate disk to use the updated environment. The instructions in this section assume a working knowledge of the alternate disk installation process. See the operating system documentation for detailed information on alternate disk installations.

Note: Only Technology Level (TL) and Service Pack (SP) releases of the operating system can be upgraded using this procedure.

Upgrading VCS on an alternate disk has the following advantages:

- The server remains active during the time the new boot environment is created and upgraded on the alternate boot device.
- The actual downtime for the upgrade is reduced to the period of time that is required for a single restart.
- The original boot environment is still available for use if the updated environment fails to become active.

Supported upgrade scenarios

The following upgrade scenarios are supported on an alternate disk:

- Upgrading only VCS
  See " Upgrading VCS on an alternate disk" on page 405.
- Upgrading only the operating system (Technology Level (TL) and Service Pack (SP) releases)

  Note: For instructions, see the operating system documentation. No additional steps are required for VCS after the operating system upgrade.

- Upgrading the operating system (Technology Level (TL) and Service Pack (SP) releases) and VCS
  See " Upgrading VCS on an alternate disk" on page 405.

Supported upgrade paths for VCS using alternate disks

You can upgrade the operating system and VCS using an alternate disk from the following versions:

AIX version	Technology Level and Service Pack releases of AIX 6.1/ 7.1
VCS version	5.1 and later

Preparing to upgrade VCS on an alternate disk

Complete the preparatory steps in the following procedure before you upgrade VCS on an alternate disk.

To prepare to upgrade VCS on an alternate disk

1. Make sure that the VCS installation media is available.
2. Check the status of the physical disks on your system.

   Note: The alternate disk must have a physical identifier and must not contain any mounted volume groups.

   ```
   # lspv
   Output similar to the following displays:
   hdisk0          0009710fa9c79877    rootvg    active
   hdisk1          0009710f0b90db93    None
   ```

   If the alternate disk does not have a physical identifier, set the physical identifier for the disk:

   ```
   # chdev -l hdisk1 -a pv=yes
   ```

3. Make sure that the following filesets are installed on the primary disk:
   `bos.alt_disk_install.boot_images, bos.alt_disk_install.rte`
4. Mount the VCS installation media.

   Determine the filesets you want to install on the alternate disk.

   ```
   # ./installvcs -install_option
   ```

   where install_option is one of the following:

   -minpkgs: For installing the minimum set of filesets

   -recpkgs: For installing the recommended filesets

   -allpkgs: For installing all filesets

   Copy the required filesets from the `pkgs` directory on the installation media to a directory on the primary boot disk, for example `/tmp/prod_name`

   If you want to upgrade the operating system along with VCS, copy the necessary operating system filesets and the VCS filesets to a directory on the primary disk, for example `/tmp/prod_name`.

   See the operating system documentation to determine the operating system filesets.

Upgrading VCS on an alternate disk

This section provides instructions to clone the primary boot environment to the alternate disk, upgrade VCS on the alternate disk, and restart the system to start from the alternate disk. You may perform the steps manually or using the SMIT interface.

In the procedure examples, the primary or current boot environment resides on `hdisk0` and the alternate or inactive boot environment resides on `hdisk1`.

To upgrade VCS on an alternate disk in a high-availability environment

Perform the instructions on each node in the cluster.

1 On the primary boot disk, change `/etc/default/llt` to start LLT on the nodes by setting the`LLT_START` attribute to 1:

```
LLT_START=1
```

2 Clone the primary boot disk `rootvg` to an alternate disk.

Manual

Run the following command:

```
# /usr/sbin/alt_disk_copy -I "acNgXY" -P "all" \
-l "/tmp/prod_name" -w "all" -d "hdisk1"
```

Where:

- -d indicates the name of the target disk on which you clone the primary disk.
- -l indicates the full path of the directory that contains the filesets to be upgraded
- -w indicates the list of VCS filesets that you want to upgrade on the alternate boot disk. The option `all` indicates that all the filesets that are contained in the directory you specified (using option -l) must be installed to the alternate boot disk.

Using SMIT interface

Start the SMIT menu and enter the required information at the prompts:

```
# smit alt_clone
```

- Target disk to install: **hdisk1**
- Fileset(s) to install: **all**
- Directory or Device with images (full path of the directory that contains the filesets to be upgraded): **/tmp/prod_name**
- ACCEPT new license agreements? **yes**
- Set `bootlist` to boot from this disk on next restart**yes**

Press **Enter** to start the upgrade on the alternate disk. The upgrade process takes some time.

3 Use the following command to wake up the volume group on the alternate boot disk (hdisk1) that you cloned.

```
# /usr/sbin/alt_rootvg_op -W -d hdisk1
```

4 Verify that the alternate disk is created:

```
# lspv |grep rootvg
Output similar to the following displays:
hdisk0          0009710fa9c79877    rootvg
hdisk1          0009710f0b90db93    altinst_rootvg
```

5 Change directory to /alt_inst/etc/VRTSvcs/conf/config.

```
# cd /alt_inst/etc/VRTSvcs/conf/config
```

6 Back up a copy of the old types.cf file and copy the new one for VCS to use.

```
# mv types.cf types.cf.ORIG
# cp ../types.cf .
```

7 If you did not configure fencing in the existing cluster, but want to use it in your upgraded cluster, perform the instructions in the following section

See "Configuring fencing for an ADI upgrade" on page 409.

8 Set the `LLT_START` attribute to 0 in the `/alt_inst/etc/default/llt` file on the alternate disk to prevent LLT from starting automatically after restart:

```
LLT_START=0
```

9 Move to root and run the `alt_rootvg_op -S` command to put the alternate root to sleep.

```
# cd /
# alt_rootvg_op -S
```

Make sure all the `/alt_*` filesystem gets unmounted successfully.

```
# alt_rootvg_op -S
Putting volume group altinst_rootvg to sleep ...
forced unmount of /alt_inst/var/adm/ras/livedump
forced unmount of /alt_inst/var/adm/ras/livedump
forced unmount of /alt_inst/var
forced unmount of /alt_inst/var
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/usr
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/tmp
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/opt
forced unmount of /alt_inst/home
forced unmount of /alt_inst/home
forced unmount of /alt_inst/admin
forced unmount of /alt_inst/admin
forced unmount of /alt_inst
forced unmount of /alt_inst
Fixing LV control blocks...
Fixing file system superblocks...
#
```

10 Verify that the normal boot list includes the name of the alternate boot disk. By default, the alternate disk installation process changes the boot list to enable the system to boot from the alternate disk.

```
# bootlist -m normal -o
hdisk1
```

11 Stop all applications that use VxFS or VxVM disk groups, whether local or CFS.

```
# hagrp -offline grp_name -any
```

12 Stop VCS on all nodes:

```
# hastop -all
```

13 Restart all the nodes in the cluster. The boot environment on the alternate disk is activated when you restart the nodes.

```
# shutdown -r
```

14 Execute the product installation scripts to replace the old installation scripts with the latest product version.

```
# sh /opt/VRTS/install/bin/UXRT<version>/add_install_scripts
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 50.

Check if the `/opt/VRTS/install` shows the latest installation uninstallation scripts.

```
# ls

.cpi5          installdmp62        installsfha62       showversion
.history       installfs62         installvcs62        uninstalldmp62
bin            installsf62         installvm62         uninstallfs62
deploy_sfha                        logs                uninstallsf62
#
```

15 Verify the upgrade.

See "Verifying the upgrade" on page 411.

16 After the systems have booted into their alternate environments, initialize the VxVM disks by running the following command on each node in the cluster:

```
# vxinstall
```

17 On the alternate disk, change `/etc/default/llt` to start LLT on the nodes by setting the LLT_START attribute to 1:

```
LLT_START=1
```

18 Start VCS:

```
# cd /opt/VRTS/install

# installvcs<version> -start
```

Where <*version*> is the specific release version.

See "About the script-based installer" on page 50.

19 Verify that all GAB ports are up:

```
# gabconfig -a
```

20 If you want to upgrade application clusters that use CP server based fencing to 6.2, make sure that you upgrade VCS or SFHA on the CP server systems. Then, upgrade all application clusters to version 6.2.

For instructions to upgrade VCS or SFHA on the CP server systems, see the *VCS or SFHA installation guide*.

Configuring fencing for an ADI upgrade

If you did not configure fencing in your existing cluster, the configuration files that fencing requires are not present on the primary boot disk. Since these configuration files are missing, ADI cannot install them when it installs the product packages on the alternate disk. The absence of these files means that fencing cannot start after the upgrade. Configure fencing if you plan to use fencing in your upgraded cluster.

Configuring fencing in disabled mode for an ADI upgrade

Perform the following procedure to configure I/O fencing in disabled mode for the ADI upgrade.

To configure I/O fencing in disabled mode for an ADI upgrade

- On all the nodes in the cluster type:

```
# cp /alt_inst/etc/vxfen.d/vxfenmode_disabled \
/alt_inst/etc/vxfenmode
```

Configuring fencing in SCSI-3 mode for an ADI upgrade

Perform the following procedures to configure I/O fencing in SCSI-3 mode for the ADI upgrade.

To update the I/O fencing files and start I/O fencing for an ADI upgrade using SCSI-3

1. Perform the tasks in the following section.

   See "Setting up coordinator disk groups" on page 272.

   When complete return to this procedure.

2. Create the /alt_inst/etc/vxfendg file, which includes the name of the coordinator disk group. On each node, type:

   ```
   # echo "vxfencoorddg" > /alt_inst/etc/vxfendg
   ```

   Where ***vxfencoorddg*** is an example name and changes based on the name that you give to the coordinator disk group. Do not use spaces between the quotes in the "vxfencoorddg" text.

3. On all cluster nodes, type the following command:

   - ```
     # cp /alt_inst/etc/vxfen.d/vxfenmode_scsi3_dmp \
     /alt_inst/etc/vxfenmode
     ```

4. Check the updated /etc/vxfenmode configuration. Enter the following command on one of the nodes. For example:

   ```
   # more /alt_inst/etc/vxfenmode
   ```

Modifying VCS configuration to use I/O fencing

After you add coordinator disks and configure I/O fencing, add the *UseFence* = SCSI3 cluster attribute to the VCS configuration file /alt_inst/etc/VRTSvcs/conf/config/main.cf. If you reset this attribute to *UseFence* = None, VCS does not make use of I/O fencing abilities while it fails over service groups. However, I/O fencing needs to be disabled separately.
   ```

To modify VCS configuration to enable I/O fencing on the alternate disk

1 Make a backup copy of the main.cf file:

```
# cd /alt_inst/etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

2 Use `vi` or another text editor to edit the main.cf file. To modify the list of cluster attributes, add the *UseFence* attribute and assign its value as SCSI3.

```
cluster clus1(
UserNames = { admin = "cDRpdxPmHpzS." }
Administrators = { admin }
HacliUserLevel = COMMANDROOT
CounterInterval = 5
UseFence = SCSI3
)
```

3 Save and close the file.

4 Verify the syntax of the file /alt_inst/etc/VRTSvcs/conf/config/main.cf:

```
# hacf -verify /alt_inst/etc/VRTSvcs/conf/config
```

Verifying the upgrade

To ensure that alternate disk installation has completed successfully, verify that all the nodes have booted from the alternate boot environment and joined the cluster.

To verify the upgrade

1 Verify that the alternate boot environment is active:

```
# lspv |grep rootvg
hdisk0          0009710fa9c79877    old_rootvg
hdisk1          0009710f0b90db93    rootvg          active
```

If there are multiple boot disks for rootvg, run the following command to verify the disk the system has booted from:

```
# bootlist -m normal -o
hdisk1 blv=hd5 pathid=0
```

2 Verify that the version of the upgraded filesets on the alternate boot disk is `6.2.0.0`.

```
# lslpp -h VRTSvcs
```

Note: The `VRTSsfcpi60` fileset still exists on the alternate boot disk. You need to manually uninstall the fileset.

If you upgraded the operating system (TL or SP):

```
# oslevel -s
```

Section 9

Post-installation tasks

Chapter 28

Performing post-installation tasks

This chapter includes the following topics:

About enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the authentication broker. AT supports all common LDAP distributions such as OpenLDAP and Windows Active Directory.

For a cluster that runs in secure mode, you must enable the LDAP authentication plug-in if the VCS users belong to an LDAP domain.

If you have not already added VCS users during installation, you can add the users later.

See the *Symantec Cluster Server Administrator's Guide* for instructions to add VCS users.

Figure 28-1 depicts the VCS cluster communication with the LDAP servers when clusters run in secure mode.

Figure 28-1 Client communication with LDAP servers

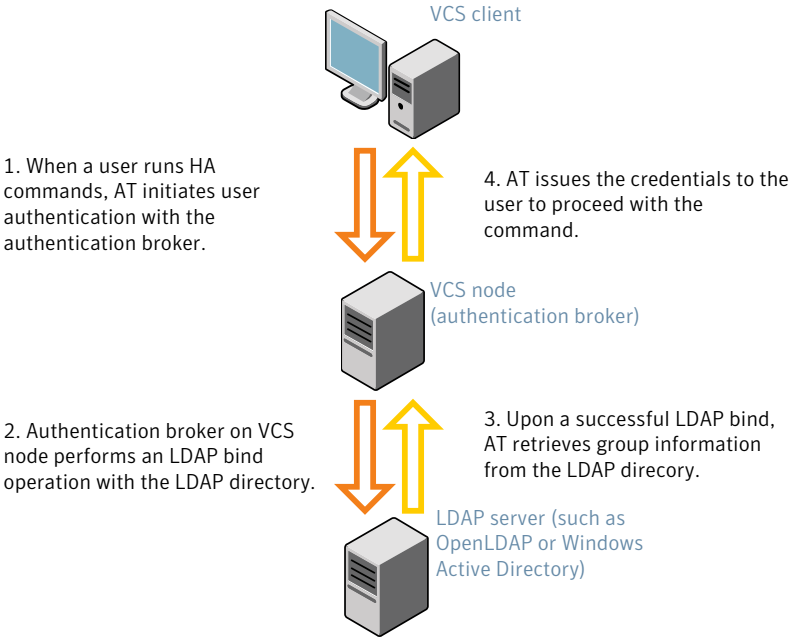

The LDAP schema and syntax for LDAP commands (such as, ldapadd, ldapmodify, and ldapsearch) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
  - UserObjectClass (the default is posixAccount)
  - UserObject Attribute (the default is uid)
  - User Group Attribute (the default is gidNumber)
  - Group Object Class (the default is posixGroup)
  - GroupObject Attribute (the default is cn)
  - Group GID Attribute (the default is gidNumber)
  - Group Membership Attribute (the default is memberUid)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)

- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

Enabling LDAP authentication for clusters that run in secure mode

The following procedure shows how to enable the plug-in module for LDAP authentication. This section provides examples for OpenLDAP and Windows Active Directory LDAP distributions.

Before you enable the LDAP authentication, complete the following steps:

- Make sure that the cluster runs in secure mode.

```
# haclus -value SecureClus
```

The output must return the value as 1.

- Make sure that the AT version is 6.1.6.0 or later.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion
vssat version: 6.1.12.8
```

To enable OpenLDAP authentication for clusters that run in secure mode

1 Run the LDAP configuration tool `atldapconf` using the `-d` option. The `-d` option discovers and retrieves an LDAP properties file which is a prioritized attribute list.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-d -s domain_controller_name_or_ipaddress -u domain_user

Attribute list file name not provided, using AttributeList.txt

Attribute file created.
```

You can use the `catatldapconf` command to view the entries in the attributes file.

2 Run the LDAP configuration tool using the `-c` option. The `-c` option creates a CLI file to add the LDAP domain.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf \
-c -d LDAP_domain_name

Attribute list file not provided, using default AttributeList.txt

CLI file name not provided, using default CLI.txt

CLI for addldapdomain generated.
```

3 Run the LDAP configuration tool `atldapconf` using the `-x` option. The `-x` option reads the CLI file and executes the commands to add a domain to the AT.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/atldapconf -x

Using default broker port 14149

CLI file not provided, using default CLI.txt

Looking for AT installation...

AT found installed at ./vssat

Successfully added LDAP domain.
```

4 Check the AT version and list the LDAP domains to verify that the Windows Active Directory server integration is complete.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showversion

vssat version: 6.1.12.8

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat listldapdomains

Domain Name : mydomain.com

Server URL : ldap://192.168.20.32:389

SSL Enabled : No

User Base DN : CN=people,DC=mydomain,DC=com

User Object Class : account

User Attribute : cn

User GID Attribute : gidNumber

Group Base DN : CN=group,DC=symantecdomain,DC=com

Group Object Class : group

Group Attribute : cn

Group GID Attribute : cn

Auth Type : FLAT

Admin User :

Admin User Password :

Search Scope : SUB
```

5 Check the other domains in the cluster.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat showdomains -p vx
```

The command output lists the number of domains that are found, with the domain names and domain types.

6 Generate credentials for the user.

```
# unset EAT_LOG

# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vssat authenticate \
-d ldap:LDAP_domain_name -p user_name -s user_password -b \
localhost:14149
```

7 Add non-root users as applicable.

```
# useradd user1

# passwd pw1

Changing password for "user1"

user1's New password:

Re-enter user1's new password:

# su user1

# bash

# id

uid=204(user1) gid=1(staff)

# pwd

# mkdir /home/user1

# chown user1 /home/ user1
```

8 Add the non-root user to the VCS configuration.

```
# haconf -makerw
# hauser -add user1
# haconf -dump -makero
```

9 Log in as non-root user and run VCS commands as LDAP user.

```
# cd /home/user1

# ls

# cat .vcspwd

101 localhost mpise LDAP_SERVER ldap

# unset VCS_DOMAINTYPE

# unset VCS_DOMAIN

# /opt/VRTSvcs/bin/hasys -state

   #System        Attribute          Value

cluster1:sysA     SysState           FAULTED

cluster1:sysB     SysState           FAULTED

cluster2:sysC     SysState           RUNNING

cluster2:sysD     SysState           RUNNING
```

Accessing the VCS documentation

The software disc contains the documentation for VCS in Portable Document Format (PDF) in the cluster_server/docs directory. After you install VCS, Symantec recommends that you copy the PDF version of the documents to the /opt/VRTS/docs directory on each node to make it available for reference.

To access the VCS documentation

- Copy the PDF from the software disc (cluster_server/docs/) to the directory /opt/VRTS/docs.

Removing permissions for communication

Make sure you completed the installation of VCS and the verification of disk support for I/O fencing. If you used `rsh`, remove the temporary `rsh` access permissions that you set for the nodes and restore the connections to the public network.

If the nodes use `ssh` for secure communications, and you temporarily removed the connections to the public network, restore the connections.

Chapter 29

Installing or upgrading VCS components

This chapter includes the following topics:

Installing the Java Console

You can administer VCS using the VCS Java-based graphical user interface, Java Console. After VCS has been installed, install the Java Console on a Windows system or AIX system with X-Windows. Review the software requirements for Java Console.

The system from which you run the Java Console can be a system in the cluster or a remote workstation. A remote workstation enables each system in the cluster to be administered remotely.

Review the information about using the Java Console. For more information, refer to the *Symantec Cluster Server Administrator's Guide*.

Software requirements for the Java Console

Cluster Manager (Java Console) is supported on:

- AIX 5.2 and 5.3
- Windows XP and Windows 2003

Note: Make sure that you are using an operating system version that supports JRE 1.6.

Hardware requirements for the Java Console

The minimum hardware requirements for the Java Console are as follows:

- Pentium II 300 megahertz
- 256 megabytes of RAM
- 800x600 display resolution
- 8-bit color depth of the monitor
- A graphics card that is capable of 2D images

Note: Symantec recommends using Pentium III 400MHz or higher, 256MB RAM or higher, and 800x600 display resolution or higher.

The version of the Java™ 2 Runtime Environment (JRE) requires 32 megabytes of RAM.

Symantec recommends using the following hardware:

- 48 megabytes of RAM
- 16-bit color mode
- The KDE and the KWM window managers that are used with displays set to local hosts

Installing the Java Console on AIX

Review the procedure to install the Java console. Before you begin with the procedure, ensure that you have the `gunzip` utility installed on your system.

To install Java console on AIX

1. Log in as superuser on the node where you intend to install the fileset.
2. Create a directory for the installation of the VCS Java Console:

```
# mkdir /tmp/install
```

3. Download the Java GUI utility from http://go.symantec.com/vcsm_download to a temporary directory.

4 Go to the temporary directory and unzip the compressed package file:

```
# cd /tmp/install
# gunzip VRTScscm.rte.gz
```

The following file is now present in the temporary directory:

```
VRTScscm.rte.bff
```

5 Install the fileset using the following command:

```
# installp -a -d VRTScscm.rte.bff VRTScscm.rte
```

6 Answer "y" if prompted.

Installing the Java Console on a Windows system

Review the procedure to install the Java console on a Windows system.

To install the Java Console on a Windows system

1 Download the Java GUI utility from http://go.symantec.com/vcsm_download to a temporary directory.

2 Extract the zipped file to a temporary folder.

3 From this extracted folder, double-click setup.exe.

4 The Symantec Cluster Manager Install Wizard guides you through the installation process.

Upgrading the Java Console

Use one of the following applicable procedures to upgrade Java Console.

To upgrade Java console on AIX

1 Log in as superuser on the node where you intend to install the fileset.

2 Remove the GUI from the previous installation.

```
# installp -u VRTScscm
```

3 Install the VCS Java console.

See “Installing the Java Console on AIX” on page 423.

To upgrade the Java Console on a Windows client

1. Stop Cluster Manager (Java Console) if it is running.
2. Remove Cluster Manager from the system.
   - From the Control Panel, double-click **Add/Remove Programs**
   - Select **Veritas Cluster Manager**.
   - Click **Add/Remove**.
   - Follow the uninstall wizard instructions.
3. Install the new Cluster Manager.

   See "Installing the Java Console on a Windows system" on page 424.

Installing VCS Simulator

You can administer VCS Simulator from the Java Console or from the command line. For more information, see the *Symantec Cluster Server Administrator's Guide*.

Review the software requirements for VCS Simulator.

Software requirements for VCS Simulator

VCS Simulator is supported on:

- Windows XP SP3, Windows 2008, Windows Vista, and Windows 7

Note: Make sure that you are using an operating system version that supports JRE 1.6 or later.

Installing VCS Simulator on Windows systems

This section describes the procedure to install VCS Simulator on Windows systems.

To install VCS Simulator on Windows systems

1. Download VCS Simulator from the following location to a temporary directory.

   http://www.symantec.com/business/cluster-server and click **Utilities**.
2. Extract the compressed files to another directory.
3. Navigate to the path of the Simulator installer file:

   \cluster_server\windows\VCSWindowsInstallers\Simulator
4. Double-click the installer file.

5 Read the information in the Welcome screen and click **Next**.

6 In the Destination Folders dialog box, click **Next** to accepted the suggested installation path or click **Change** to choose a different location.

7 In the Ready to Install the Program dialog box, click **Back** to make changes to your selections or click **Install** to proceed with the installation.

8 In the Installshield Wizard Completed dialog box, click **Finish**.

Reviewing the installation

VCS Simulator installs Cluster Manager (Java Console) and Simulator binaries on the system. The Simulator installation creates the following directories:

Directory	Content
attrpool	Information about attributes associated with VCS objects
bin	VCS Simulator binaries
default_clus	Files for the default cluster configuration
sample_clus	A sample cluster configuration, which serves as a template for each new cluster configuration
templates	Various templates that are used by the Java Console
types	The types.cf files for all supported platforms
conf	Contains another directory called types. This directory contains assorted resource type definitions that are useful for the Simulator. The type definition files are present in platform-specific sub directories.

Additionally, VCS Simulator installs directories for various cluster configurations.

VCS Simulator creates a directory for every new simulated cluster and copies the contents of the sample_clus directory. Simulator also creates a log directory within each cluster directory for logs that are associated with the cluster.

Upgrading VCS Simulator

Use the following procedure to upgrade VCS Simulator.

To upgrade VCS Simulator on a Windows client

1 Stop all instances of VCS Simulator.

2 Stop VCS Simulator, if it is running.

3 Remove VCS Simulator from the system.

- From the Control Panel, double-click **Add/Remove Programs**
- Select **VCS Simulator**.
- Click **Add/Remove**.
- Follow the uninstall wizard instructions.

4 Install the new Simulator.

See "Installing VCS Simulator on Windows systems" on page 425.

Chapter 30

Verifying the VCS installation

This chapter includes the following topics:

About verifying the VCS installation

After you install and configure VCS, you can inspect the contents of the key VCS configuration files that you have installed and modified during the process. These files reflect the configuration that is based on the information you supplied. You can also run VCS commands to verify the status of LLT, GAB, and the cluster.

About the cluster UUID

You can verify the existence of the cluster UUID.

To verify that the cluster UUID exists

- From the prompt, run a cat command.

```
cat /etc/vx/.uuids/clusuuid
```

To display UUID of all the nodes in the cluster

- From the prompt, run the command from any node.

```
/opt/VRTSvcs/bin/uuidconfig.pl -rsh -clus -display -use_llthost
```

Verifying the LLT, GAB, and VCS configuration files

Make sure that the LLT, GAB, and VCS configuration files contain the information you provided during VCS installation and configuration.

To verify the LLT, GAB, and VCS configuration files

1. Navigate to the location of the configuration files:
   - LLT
     /etc/llthosts
     /etc/lltab
   - GAB
     /etc/gabtab
   - VCS
     /etc/VRTSvcs/conf/config/main.cf
2. Verify the content of the configuration files.

   See "About the LLT and GAB configuration files" on page 507.

   See "About the VCS configuration files" on page 511.

Verifying LLT, GAB, and cluster operation

Verify the operation of LLT, GAB, and the cluster using the VCS commands.

To verify LLT, GAB, and cluster operation

1. Log in to any node in the cluster as superuser.
2. Make sure that the PATH environment variable is set to run the VCS commands.

   See "Setting the PATH variable" on page 72.
3. Verify LLT operation.

   See "Verifying LLT" on page 430.
4. Verify GAB operation.

   See "Verifying GAB" on page 432.

5 Verify the cluster operation.

See “Verifying the cluster” on page 433.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

1 Log in as superuser on the node sys1.

2 Run the `lltstat` command on the node sys1 to view the status of LLT.

```
lltstat -n
```

The output on sys1 resembles:

```
LLT node information:
    Node             State           Links
    *0 sys1          OPEN           2
     1 sys2          OPEN           2
```

Each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

If LLT does not operate, the command does not return any LLT links information: If only one network is connected, the command returns the following LLT statistics information:

```
LLT node information:
    Node                State    Links
   * 0 sys1             OPEN       2
     1 sys2             OPEN       2
     2 sys5             OPEN       1
```

3 Log in as superuser on the node sys2.

4 Run the `lltstat` command on the node sys2 to view the status of LLT.

```
lltstat -n
```

The output on sys2 resembles:

```
LLT node information:
    Node             State           Links
     0 sys1          OPEN            2
    *1 sys2          OPEN            2
```

5 To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node sys1 in a two-node cluster:

```
lltstat -nvv active
```

The output on sys1 resembles:

```
Node           State      Link    Status      Address
*0 sys1      OPEN
                          en1 UP     08:00:20:93:0E:34
                          en2 UP     08:00:20:93:0E:38
1 sys2       OPEN
                          en1 UP     08:00:20:8F:D1:F2
                          en2 DOWN
```

The command reports the status on the two active nodes in the cluster, sys1 and sys2.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- An address for each link

However, the output in the example shows different details for the node sys2. The private network connection is possibly broken or the information in the /etc/llttab file may be incorrect.

6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node sys1 in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
  Port  Usage       Cookie
  0     gab          0x0
```

```
      opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects:    0 1
7     gab          0x7
      opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects:    0 1
31    gab          0x1F
      opens:       0 2 3 4 5 6 7 8 9 10 11 ... 60 61 62 63
      connects:    0 1
```

Verifying GAB

Verify the GAB operation using the `gabconfig -a` command. This command returns the GAB port membership information.

The ports indicate the following:

Port a
- Nodes have GAB communication.
- gen a36e0003 is a randomly generated number.
- membership 01 indicates that nodes 0 and 1 are connected.

Port b
- Indicates that the I/O fencing driver is connected to GAB port b.

  Note: Port b appears in the `gabconfig` command output only if you had configured I/O fencing after you configured VCS.
- gen a23da40d is a randomly generated number.
- membership 01 indicates that nodes 0 and 1 are connected.

Port h
- VCS is started.
- gen fd570002 is a randomly generated number
- membership 01 indicates that nodes 0 and 1 are both running VCS

For more information on GAB, refer to the *Symantec Cluster Server Administrator's Guide*.

To verify GAB

1 To verify that GAB operates, type the following command on each node:

```
/sbin/gabconfig -a
```

2 Review the output of the command:

- If GAB operates, the following GAB port membership information is returned: For a cluster where I/O fencing is not configured:

```
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
Port h gen fd570002 membership 01
```

For a cluster where I/O fencing is configured:

```
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
Port b gen a23da40d membership 01
Port h gen fd570002 membership 01
```

Note that port b appears in the `gabconfig` command output only if you had configured I/O fencing. You can also use the `vxfenadm -d` command to verify the I/O fencing configuration.

- If GAB does not operate, the command does not return any GAB port membership information:

```
GAB Port Memberships
===================================
```

- If only one network is connected, the command returns the following GAB port membership information:

```
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
Port a gen a36e0003 jeopardy   ;1
Port h gen fd570002 membership 01
Port h gen fd570002 jeopardy   ;1
```

Verifying the cluster

Verify the status of the cluster using the `hastatus` command. This command returns the system state and the group state.

Refer to the `hastatus(1M)` manual page.

Refer to the *Symantec Cluster Server Administrator's Guide* for a description of system states and the transitions between them.

To verify the cluster

1 To verify the status of the cluster, type the following command:

```
# hastatus -summary
```

The output resembles:

```
-- SYSTEM STATE
-- System               State                Frozen

A  sys1                 RUNNING                 0
A  sys2                 RUNNING                 0

-- GROUP STATE
-- Group           System      Probed  AutoDisabled   State

B  ClusterService  sys1          Y       N            ONLINE
B  ClusterService  sys2          Y       N            OFFLINE
```

2 Review the command output for the following information:

- The system state
  If the value of the system state is RUNNING, the cluster is successfully started.
- The ClusterService group state
  In the sample output, the group state lists the ClusterService group, which is ONLINE on sys1 and OFFLINE on sys2.

Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

Refer to the `hasys(1M)` manual page.

Refer to the *Symantec Cluster Server Administrator's Guide* for information about the system attributes for VCS.

To verify the cluster nodes

- On one of the nodes, type the `hasys -display` command:

```
# hasys -display
```

The example in the following procedure is for SPARC and it shows the output when the command is run on the node sys1. The list continues with similar information for sys2 (not shown) and any other nodes in the cluster.

```
#System    Attribute             Value

sys1       AgentsStopped         0

sys1       AvailableCapacity     100

sys1       CPUThresholdLevel     Critical 90 Warning 80 Note 70
                                 Info 60

sys1       CPUUsage              0

sys1       CPUUsageMonitoring    Enabled 0 ActionThreshold 0
                                 ActionTimeLimit 0 Action NONE
                                 NotifyThreshold 0 NotifyTimeLimit 0

sys1       Capacity              100

sys1       ConfigBlockCount      341

sys1       ConfigCheckSum        57519

sys1       ConfigDiskState       CURRENT

sys1       ConfigFile            /etc/VRTSvcs/conf/config

sys1       ConfigInfoCnt         0

sys1       ConfigModDate         Mon Sep 03 07:14:23 CDT 2012

sys1       ConnectorState        Up

sys1       CurrentLimits

sys1       DiskHbStatus

sys1       DynamicLoad           0

sys1       EngineRestarted       0

sys1       EngineVersion         6.2.00.0

sys1       FencingWeight         0

sys1       Frozen                0

sys1       GUIIPAddr

sys1       HostUtilization       CPU 0 Swap 0

sys1       LLTNodeId             0
```

```
sys1       LicenseType            PERMANENT_SITE

sys1       Limits

sys1       LinkHbStatus           en1 UP en2 UP

sys1       LoadTimeCounter        0

sys1       LoadTimeThreshold      600

sys1       LoadWarningLevel       80

sys1       NoAutoDisable          0

sys1       NodeId                 0

sys1       OnGrpCnt               7

sys1       PhysicalServer

sys1       ShutdownTimeout        600

sys1       SourceFile             ./main.cf

sys1       SwapThresholdLevel     Critical 90 Warning 80 Note 70
                                  Info 60

sys1       SysInfo                Aix:sys1,6,1,00C129B44C00

sys1       SysName                sys1

sys1       SysState               RUNNING

sys1       SystemLocation

sys1       SystemOwner

sys1       SystemRecipients

sys1       TFrozen                0

sys1       TRSE                   0

sys1       UpDownState            Up

sys1       UserInt                0

sys1       UserStr

sys1       VCSFeatures            DR

sys1       VCSMode                VCS
```

Upgrading the disk group version

After you upgrade from previous versions to 6.2, you have to upgrade the disk group version manually.

To upgrade disk group version, you have to first upgrade the cluster protocol version using the `vxdctl upgrade` command.

```
# vxdctl list
Volboot file
version: 3/1
seqno:   0.1
cluster protocol version: 120
hostid:  sys1
hostguid:  {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
#
# vxdctl upgrade
#

# vxdctl list

Volboot file
version: 3/1
seqno:   0.2
cluster protocol version: 140
hostid:  sys1
hostguid:  {fca678ac-e0ef-11e2-b22c-5e26fd3b6f13}
```

Verify if the cluster protocol version shows 140 and disk group version is upgraded to 200.

```
# vxdctl list |grep version

version:   140
#
# vxdg upgrade dg_name
#
# vxdg list dg_name |grep version

version:   200
```

Performing a postcheck on a node

The installer's `postcheck` command can help you to determine installation-related problems and provide troubleshooting information.

See "About using the postcheck option" on page 438.

To run the postcheck command on a node

1 Run the installer with the `-postcheck` option.

```
# ./installer -postcheck system_name
```

2 Review the output for installation-related information.

About using the postcheck option

You can use the installer's post-check to determine installation-related problems and to aid in troubleshooting.

Note: This command option requires downtime for the node.

When you use the `postcheck` option, it can help you troubleshoot the following VCS-related issues:

- The heartbeat link does not exist.
- The heartbeat link cannot communicate.
- The heartbeat link is a part of a bonded or aggregated NIC.
- A duplicated cluster ID exists (if LLT is not running at the check time).
- The VRTSllt pkg version is not consistent on the nodes.
- The llt-linkinstall value is incorrect.
- The /etc/llthosts and /etc/lltab configuration is incorrect.
- the `/etc/gabtab` file is incorrect.
- The incorrect GAB linkinstall value exists.
- The VRTSgab pkg version is not consistent on the nodes.
- The `main.cf` file or the `types.cf` file is invalid.
- The `/etc/VRTSvcs/conf/sysname` file is not consistent with the hostname.
- The cluster UUID does not exist.
- The `uuidconfig.pl` file is missing.

- The VRTSvcs pkg version is not consistent on the nodes.
- The `/etc/vxfenmode` file is missing or incorrect.
- The `/etc/vxfendg file` is invalid.
- The vxfen link-install value is incorrect.
- The VRTSvxfen pkg version is not consistent.

The `postcheck` option can help you troubleshoot the following SFHA or SFCFSHA issues:

- Volume Manager cannot start because the `/etc/vx/reconfig.d/state.d/install-db` file has not been removed.
- Volume Manager cannot start because the `volboot` file is not loaded.
- Volume Manager cannot start because no license exists.
- Cluster Volume Manager cannot start because the CVM configuration is incorrect in the `main.cf` file. For example, the Autostartlist value is missing on the nodes.
- Cluster Volume Manager cannot come online because the node ID in the `/etc/llthosts` file is not consistent.
- Cluster Volume Manager cannot come online because Vxfen is not started.
- Cluster Volume Manager cannot start because gab is not configured.
- Cluster Volume Manager cannot come online because of a CVM protocol mismatch.
- Cluster Volume Manager group name has changed from "cvm", which causes CVM to go offline.

You can use the installer's post-check option to perform the following checks:

General checks for all products:

- All the required filesets are installed.
- The versions of the required filesets are correct.
- There are no verification issues for the required filesets.

Checks for Volume Manager (VM):

- Lists the daemons which are not running (`vxattachd`, `vxconfigbackupd`, `vxesd`, `vxrelocd` ...).
- Lists the disks which are not in 'online' or 'online shared' state (`vxdisk list`).
- Lists the diskgroups which are not in 'enabled' state (`vxdg list`).
- Lists the volumes which are not in 'enabled' state (`vxprint -g <dgname>`).

- Lists the volumes which are in 'Unstartable' state (`vxinfo -g <dgname>`).
- Lists the volumes which are not configured in `/etc/filesystems` .

Checks for File System (FS):

- Lists the VxFS kernel modules which are not loaded (`vxfs/fdd/vxportal`.).
- Whether all VxFS file systems present in `/etc/filesystems` file are mounted.
- Whether all VxFS file systems present in `/etc/filesystems` are in disk layout 6 or higher.
- Whether all mounted VxFS file systems are in disk layout 6 or higher.

Checks for Cluster File System:

- Whether FS and ODM are running at the latest protocol level.
- Whether all mounted CFS file systems are managed by VCS.
- Whether cvm service group is online.

See “Performing a postcheck on a node” on page 438.

Section 10

Adding and removing cluster nodes

Chapter 31

Adding a node to a single-node cluster

This chapter includes the following topics:

Adding a node to a single-node cluster

All nodes in the new cluster must run the same version of VCS. The example procedure refers to the existing single-node VCS node as Node A. The node that is to join Node A to form a multiple-node cluster is Node B.

Table 31-1 specifies the activities that you need to perform to add nodes to a single-node cluster.

Table 31-1 Tasks to add a node to a single-node cluster

Task	Reference
Set up Node B to be compatible with Node A.	See "Setting up a node to join the single-node cluster" on page 443.
■ Add Ethernet cards for private heartbeat network for Node B. ■ If necessary, add Ethernet cards for private heartbeat network for Node A. ■ Make the Ethernet cable connections between the two nodes.	See "Installing and configuring Ethernet cards for private network" on page 444.
Connect both nodes to shared storage.	See "Configuring the shared storage" on page 445.

Table 31-1 Tasks to add a node to a single-node cluster *(continued)*

Task	Reference
■ Bring up VCS on Node A. ■ Edit the configuration file.	See "Bringing up the existing node" on page 445.
If necessary, install VCS on Node B and add a license key. Make sure Node B is running the same version of VCS as the version on Node A.	See "Installing the VCS software manually when adding a node to a single node cluster" on page 445.
Edit the configuration files on Node B.	See "About the VCS configuration files" on page 511.
Start LLT and GAB on Node B.	See "Starting LLT and GAB" on page 446.
■ Start LLT and GAB on Node A. ■ Copy UUID from Node A to Node B. ■ Restart VCS on Node A. ■ Modify service groups for two nodes.	See "Reconfiguring VCS on the existing node" on page 446.
■ Start VCS on Node B. ■ Verify the two-node cluster.	See "Verifying configuration on both nodes" on page 447.

Setting up a node to join the single-node cluster

The new node to join the existing single node that runs VCS must run the same operating system.

To set up a node to join the single-node cluster

1 Do one of the following tasks:

- If VCS is not currently running on Node B, proceed to step 2.
- If the node you plan to add as Node B is currently part of an existing cluster, remove the node from the cluster. After you remove the node from the cluster, remove the VCS filesets and configuration files.
  See "Removing a node from a VCS cluster" on page 468.
- If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS.
- If you renamed the LLT and GAB startup files, remove them.

2 If necessary, install VxVM and VxFS.

See “Installing VxVM or VxFS if necessary” on page 444.

Installing VxVM or VxFS if necessary

If you have either VxVM or VxFS with the cluster option installed on the existing node, install the same version on the new node.

Refer to the appropriate documentation for VxVM and VxFS to verify the versions of the installed products. Make sure the same version runs on all nodes where you want to use shared storage.

Installing and configuring Ethernet cards for private network

Both nodes require Ethernet cards (NICs) that enable the private network. If both Node A and Node B have Ethernet cards installed, you can ignore this step.

For high availability, use two separate NICs on each node. The two NICs provide redundancy for heartbeating.

See “Setting up the private network” on page 67.

To install and configure Ethernet cards for private network

1 Shut down VCS on Node A.

```
# hastop -local
```

2 Shut down the nodes.

```
# shutdown -F
```

3 Install the Ethernet card on Node A.

If you want to use aggregated interface to set up private network, configure aggregated interface.

4 Install the Ethernet card on Node B.

If you want to use aggregated interface to set up private network, configure aggregated interface.

5 Configure the Ethernet card on both nodes.

6 Make the two Ethernet cable connections from Node A to Node B for the private networks.

7 Restart the nodes.

Configuring the shared storage

Make the connection to shared storage from Node B. Configure VxVM on Node B and reboot the node when you are prompted.

See "Setting up shared storage" on page 70.

Bringing up the existing node

Bring up the node.

To bring up the node

1. Log in as superuser.
2. Make the VCS configuration writable.

   ```
   # haconf -makerw
   ```

3. Display the service groups currently configured.

   ```
   # hagrp -list
   ```

4. Freeze the service groups.

   ```
   # hagrp -freeze group -persistent
   ```

   Repeat this command for each service group in step 3.
5. Make the configuration read-only.

   ```
   # haconf -dump -makero
   ```

6. Stop VCS on Node A.

   ```
   # hastop -local -force
   ```

7. Enable the GAB and LLT startup files so they can be used.

   ```
   # mv /etc/rc.d/rc2.d/X92gab /etc/rc.d/rc2.d/S92gab
   # mv /etc/rc.d/rc2.d/X70llt /etc/rc.d/rc2.d/S70llt
   ```

Installing the VCS software manually when adding a node to a single node cluster

Install the VCS 6.2 filesets manually and install the license key.

Refer to the following sections:

- See “Modifying /etc/pse.conf to enable the Ethernet driver” on page 239.
- See “Adding a license key for a manual installation” on page 244.

Creating configuration files

Create the configuration files for your cluster.

To create the configuration files

1. Create the file /etc/llttab for a two-node cluster

   See “Setting up /etc/llttab for a manual installation” on page 251.
2. Create the file /etc/llthosts that list both the nodes.

   See “Setting up /etc/llthosts for a manual installation” on page 251.
3. Create the file /etc/gabtab.

   See “Configuring GAB manually” on page 254.

Starting LLT and GAB

On the new node, start LLT and GAB.

To start LLT and GAB

1. Start LLT on Node B.

   ```
   # /etc/init.d/llt.rc start
   ```
2. Start GAB on Node B.

   ```
   # /etc/init.d/gab.rc start
   ```

Reconfiguring VCS on the existing node

Reconfigure VCS on the existing nodes.

To reconfigure VCS on existing nodes

1. On Node A, create the files /etc/llttab, /etc/llthosts, and /etc/gabtab. Use the files that are created on Node B as a guide, customizing the /etc/llttab for Node A.
2. Start LLT on Node A.

   ```
   # /etc/init.d/llt.rc start
   ```

3 Start GAB on Node A.

```
# /etc/init.d/gab.rc start
```

4 Check the membership of the cluster.

```
# gabconfig -a
```

5 Copy the cluster UUID from the existing node to the new node:

```
# /opt/VRTSvcs/bin/uuidconfig.pl -clus -copy -from_sys \
node_name_in_running_cluster -to_sys new_sys1 ... new_sysn
```

Where you are copying the cluster UUID from a node in the cluster (*node_name_in_running_cluster*) to systems from *new_sys1* through *new_sysn* that you want to join the cluster.

6 Start VCS on Node A.

```
# hastart
```

7 Make the VCS configuration writable.

```
# haconf -makerw
```

8 Add Node B to the cluster.

```
# hasys -add sysB
```

9 Add Node B to the system list of each service group.

- List the service groups.

  ```
  # hagrp -list
  ```

- For each service group that is listed, add the node.

  ```
  # hagrp -modify group SystemList -add sysB 1
  ```

Verifying configuration on both nodes

Verify the configuration for the nodes.

To verify the nodes' configuration

1 On Node B, check the cluster membership.

```
# gabconfig -a
```

2 Start the VCS on Node B.

```
# hastart
```

3 Verify that VCS is up on both nodes.

```
# hastatus
```

4 List the service groups.

```
# hagrp -list
```

5 Unfreeze the service groups.

```
# hagrp -unfreeze group -persistent
```

6 Save the new two-node configuration.

```
# haconf -dump -makero
```

Chapter 32

Adding a node to a multi-node VCS cluster

This chapter includes the following topics:

Adding nodes using the VCS installer

The VCS installer performs the following tasks:

- Verifies that the node and the existing cluster meet communication requirements.
- Verifies the products and filesets installed on the new node.
- Discovers the network interfaces on the new node and checks the interface settings.
- Creates the following files on the new node:

  `/etc/llttab`

  `/etc/VRTSvcs/conf/sysname`
- Updates the following configuration files and copies them on the new node:

  `/etc/llthosts`

  `/etc/gabtab`

  `/etc/VRTSvcs/conf/config/main.cf`
- Copies the following files from the existing cluster to the new node
  /etc/vxfenmode
  /etc/vxfendg

/etc/vx/.uuids/clusuuid
/etc/default/llt
/etc/default/gab
/etc/default/vxfen

- Configures disk-based or server-based fencing depending on the fencing mode in use on the existing cluster.

At the end of the process, the new node joins the VCS 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 VCS cluster using the VCS installer

1 Log in as the root user on one of the nodes of the existing cluster.

2 Run the VCS installer with the -addnode option.

```
# cd /opt/VRTS/install

# ./installvcs<version> -addnode
```

Where <version> is specific to the release version.

See "About the script-based installer" on page 50.

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 VCS cluster. The installer uses the node information to identify the existing cluster.

```
Enter a node name in the VCS cluster to which
you want to add a node: sys1
```

4 Review and confirm the cluster information.

5 Enter the name of the systems that you want to add as new nodes to the cluster.

```
Enter the system names separated by spaces
to add to the cluster: sys5
```

The installer checks the installed products and filesets 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 LLT configuration for the new node must be the same as that of the existing cluster. If your existing cluster uses LLT over UDP, the installer asks questions related to LLT over UDP for the new node.

See "Configuring private heartbeat links" on page 132.

```
Enter the NIC for the first private heartbeat
link on sys5: [b,q,?] en1
```

7 Enter **y** to configure a second private heartbeat link.

Note: At least two private heartbeat links must be configured for high availability of the cluster.

```
Would you like to configure a second private
heartbeat link? [y,n,q,b,?] (y)
```

8 Enter the name of the network interface that you want to configure as the second private heartbeat link.

```
Enter the NIC for the second private heartbeat link
on sys5: [b,q,?] en2
```

9 Depending on the number of LLT links configured in the existing cluster, configure additional private heartbeat links for the new node.

The installer verifies the network interface settings and displays the information.

10 Review and confirm the information.

11 If you have configured SMTP, SNMP, or the global cluster option in the existing cluster, you are prompted for the NIC information for the new node.

```
Enter the NIC for VCS to use on sys5: en3
```

12 If you have enabled security on the cluster, the installer displays the following message:

```
Since the cluster is in secure mode, please check
the main.cf to see whether you would like to modify
the usergroup that you would like to grant read access.
```

To modify the user group to grant read access to secure clusters, use the following commands:

```
haconf -makerw
hauser -addpriv <user group> GuestGroup
haclus -modify <user group> GuestGroup
haconf -dump -makero
```

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.

Manually adding a node to a cluster

The system you add to the cluster must meet the hardware and software requirements.

See "Hardware requirements for VCS" on page 38.

Table 32-1 specifies the tasks that are involved in adding a cluster. The example demonstrates how to add a node saturn to already existing nodes, sys1 and sys2.

Table 32-1 Tasks that are involved in adding a node to a cluster

Task	Reference
Set up the hardware.	See "Setting up the hardware" on page 454.
Install the software manually.	See "Installing VCS filesets for a manual installation" on page 243.
Add a license key.	See "Adding a license key for a manual installation" on page 244.
Configure LLT and GAB.	See "Configuring LLT and GAB when adding a node to the cluster" on page 458.
Copy the UUID.	See "Reconfiguring VCS on the existing node" on page 446.
Add the node to the existing cluster.	See "Adding the node to the existing cluster" on page 464.
Start VCS and verify the cluster.	See "Starting VCS and verifying the cluster" on page 465.

Setting up the hardware

Figure 32-1 shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

Figure 32-1 Adding a node to a two-node cluster using two switches

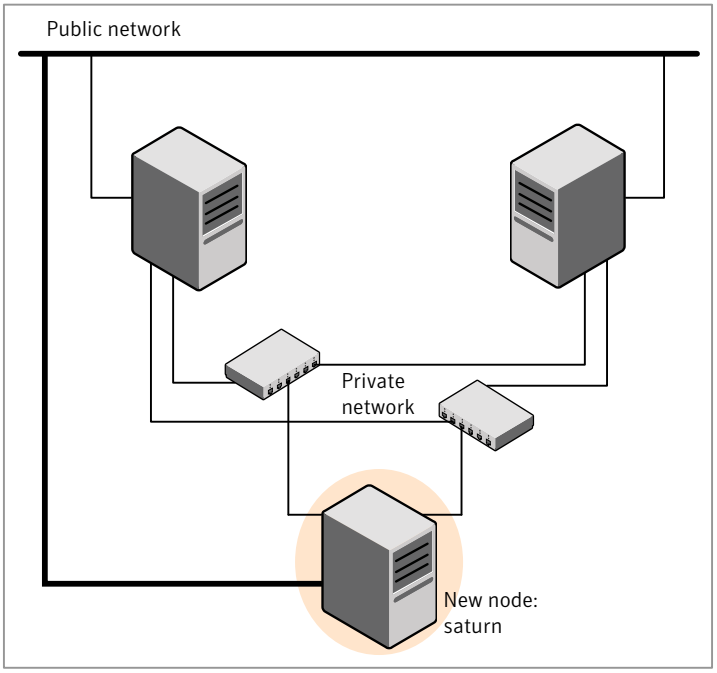

To set up the hardware

1 Connect the VCS private Ethernet controllers.

Perform the following tasks as necessary:

- When you add nodes to a two-node cluster, use independent switches or hubs for the private network connections. You can only use crossover cables for a two-node cluster, so you might have to swap out the cable for a switch or hub.
- If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

Figure 32-1 illustrates a new node being added to an existing two-node cluster using two independent hubs.

2 Connect the system to the shared storage, if required.

Installing the VCS software manually when adding a node

Install the VCS 6.2 filesets manually and add a license key.

For more information, see the following:

- See “Installing VCS software manually” on page 241.
- See “Adding a license key for a manual installation” on page 244.

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

See “Configuring LLT and GAB when adding a node to the cluster” on page 458.

Table 32-2 uses the following information for the following command examples.

Table 32-2 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  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, 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
```

Configuring LLT and GAB when adding a node to the cluster

Create the LLT and GAB configuration files on the new node and update the files on the existing nodes.

To configure LLT when adding a node to the cluster

1 Create the file /etc/llthosts on the new node. You must also update it on each of the current nodes in the cluster.

For example, suppose you add sys5 to a cluster consisting of sys1 and sys2:

- If the file on one of the existing nodes resembles:

```
0 sys1
1 sys2
```

- Update the file for all nodes, including the new one, resembling:

```
0 sys1
1 sys2
2 sys5
```

2 Create the file /etc/llttab on the new node, making sure that line beginning "`set-node`" specifies the new node.

The file /etc/llttab on an existing node can serve as a guide.

The following example describes a system where node sys2 is the new node on cluster ID number 2:

```
set-node  sys2
set-cluster 2
link en1 /dev/dlpi/en:1 - ether - -
link en2 /dev/dlpi/en:2 - ether - -
```

3 Copy the following file from one of the nodes in the existing cluster to the new node:

/etc/default/llt

4 On the new system, run the command:

```
# /sbin/lltconfig -c
```

In a setup that uses LLT over UDP, new nodes automatically join the existing cluster if the new nodes and all the existing nodes in the cluster are not separated by a router. However, if you use LLT over UDP6 link with IPv6 address and if the new node and the existing nodes are separated by a router, then do the following:

- Edit the `/etc/llttab` file on each node to reflect the link information about the new node.
- Specify the IPv6 address for UDP link of the new node to all existing nodes. Run the following command on each existing node for each UDP link:

```
# /sbin/lltconfig -a set systemid device_tag address
```

To configure GAB when adding a node to the cluster

1 Create the file /etc/gabtab on the new system.

- If the /etc/gabtab file on the existing nodes resembles:

```
/sbin/gabconfig -c
```

The file on the new node should be the same. Symantec recommends that you use the `-c -nN` option, where *N* is the total number of cluster nodes.

- If the /etc/gabtab file on the existing nodes resembles:

  ```
  /sbin/gabconfig -c -n2
  ```

  The file on all nodes, including the new node, should change to reflect the change in the number of cluster nodes. For example, the new file on each node should resemble:

  ```
  /sbin/gabconfig -c -n3
  ```

  The `-n` flag indicates to VCS the number of nodes that must be ready to form a cluster before VCS starts.

2 Copy the following file from one of the nodes in the existing cluster to the new node:

/etc/default/gab

3 On the new node, to configure GAB run the command:

```
# /sbin/gabconfig -c
```

To verify GAB

1 On the new node, run the command:

```
# /sbin/gabconfig -a
```

The output should indicate that port a membership shows all nodes including the new node. The output should resemble:

```
GAB Port Memberships
===================================
Port a gen a3640003 membership 012
```

See "Verifying GAB" on page 432.

2 Run the same command on the other nodes (sys1 and sys2) to verify that the port a membership includes the new node:

```
# /sbin/gabconfig -a
GAB Port Memberships
===================================
Port a gen a3640003 membership 012
Port h gen fd570002 membership 01
Port h gen fd570002    visible ; 2
```

Configuring I/O fencing on the new node

If the existing cluster is configured for I/O fencing, perform the following tasks on the new node:

- Prepare to configure I/O fencing on the new node.
  See "Preparing to configure I/O fencing on the new node" on page 461.
- If the existing cluster runs server-based fencing, configure server-based fencing on the new node.
  See "Configuring server-based fencing on the new node" on page 462.
  If the existing cluster runs disk-based fencing, you need not perform any additional step. Skip to the next task. After you copy the I/O fencing files and start I/O fencing, disk-based fencing automatically comes up.
- Copy the I/O fencing files from an existing node to the new node and start I/O fencing on the new node.
  See "Starting I/O fencing on the new node" on page 463.

If the existing cluster is not configured for I/O fencing, perform the procedure to add the new node to the existing cluster.

See "Adding the node to the existing cluster" on page 464.

Preparing to configure I/O fencing on the new node

Perform the following tasks before you configure and start I/O fencing on the new node.

To prepare to configure I/O fencing on the new node

1. Determine whether the existing cluster runs disk-based or server-based fencing mechanism. On one of the nodes in the existing cluster, run the following command:

   ```
   # vxfenadm -d
   ```

   If the fencing mode in the output is SCSI3, then the cluster uses disk-based fencing.

   If the fencing mode in the output is CUSTOMIZED, then the cluster uses server-based fencing.

2. In the following cases, install and configure Veritas Volume Manager (VxVM) on the new node.
   - The existing cluster uses disk-based fencing.
   - The existing cluster uses server-based fencing with at least one coordinator disk.

You need not perform this step if the existing cluster uses server-based fencing with all coordination points as CP servers.

See the *Symantec Storage Foundation and High Availability Installation Guide* for installation instructions.

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 VCS cluster, set the cluster configuration to read-write mode:

```
# haconf -makerw
```

2 Add the node sys5 to the existing vxfen group.

```
# hagrp -modify vxfen SystemList -add sys5 2
```

3 Save the configuration by running the following command from any node in the VCS cluster:

```
# haconf -dump -makero
```

Starting I/O fencing on the new node

Copy the I/O fencing files from an existing node to the new node and start I/O fencing on the new node. This task starts I/O fencing based on the fencing mechanism that is configured in the existing cluster.

To start I/O fencing on the new node

1 Copy the following I/O fencing configuration files from one of the nodes in the existing cluster to the new node:

- /etc/vxfenmode
- /etc/vxfendg—This file is required only for disk-based fencing.
- /etc/default/vxfen

2 Start I/O fencing on the new node.

```
# /etc/init.d/vxfen.rc start
```

3 Run the GAB configuration command on the new node to verify that the port b membership is formed.

```
# gabconfig -a
```

Adding the node to the existing cluster

Perform the tasks on one of the existing nodes in the cluster.

To add the new node to the existing cluster

1 Enter the command:

```
# haconf -makerw
```

2 Add the new system to the cluster:

```
# hasys -add sys1
```

3 Copy the main.cf file from an existing node to your new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \
  sys5:/etc/VRTSvcs/conf/config/
```

4 Check the VCS configuration file. No error message and a return value of zero indicates that the syntax is legal.

```
# hacf -verify /etc/VRTSvcs/conf/config/
```

5 If necessary, modify any new system attributes.

6 Enter the command:

```
# haconf -dump -makero
```

Starting VCS and verifying the cluster

Start VCS after adding the new node to the cluster and verify the cluster.

To start VCS and verify the cluster

1 Start VCS on the newly added system:

```
# hastart
```

2 Run the GAB configuration command on each node to verify that port a and port h include the new node in the membership:

```
# /sbin/gabconfig -a
  GAB Port Memberships
  ===================================
  Port a gen a3640003 membership 012
  Port h gen fd570002 membership 012
```

Adding a node using response files

Typically, you can use the response file that the installer generates on one system to add nodes to an existing cluster.

To add nodes using response files

1 Make sure the systems where you want to add nodes meet the requirements.

2 Make sure all the tasks required for preparing to add a node to an existing VCS cluster are completed.

3 Copy the response file to one of the systems where you want to add nodes.

See "Sample response file for adding a node to a VCS cluster" on page 466.

4 Edit the values of the response file variables as necessary.

See "Response file variables to add a node to a VCS cluster" on page 466.

5 Mount the product disc and navigate to the folder that contains the installation program.

6 Start adding nodes from the system to which you copied the response file. For example:

```
# ./installer -responsefile /tmp/response_file

# ./installvcs -responsefile /tmp/response_file
```

Where /tmp/*response_file* is the response file's full path name.

Depending on the fencing configuration in the existing cluster, the installer configures fencing on the new node. The installer then starts all the required Symantec processes and joins the new node to cluster. The installer indicates the location of the log file and summary file with details of the actions performed.

Response file variables to add a node to a VCS cluster

Table 32-3 lists the response file variables that you can define to add a node to an VCS cluster.

Table 32-3 Response file variables for adding a node to an VCS cluster

Variable	Description
$CFG{opt}{addnode}	Adds a node to an existing cluster. List or scalar: scalar Optional or required: required
$CFG{newnodes}	Specifies the new nodes to be added to the cluster. List or scalar: list Optional or required: required

Sample response file for adding a node to a VCS cluster

The following example shows a response file for adding a node to a VCS cluster.

```
our %CFG;

$CFG{clustersystems}=[ qw(sys1) ];
$CFG{newnodes}=[ qw(sys5) ];
$CFG{opt}{addnode}=1;
$CFG{opt}{configure}=1;d
$CFG{opt}{vr}=1;
```

```
$CFG{prod}="VCS62";
d$CFG{systems}=[ qw(sys1 sys5) ];
$CFG{vcs_allowcomms}=1;
$CFG{vcs_clusterid}=101;
$CFG{vcs_clustername}="clus1";
$CFG{vcs_lltlink1}{sys5}="en1";
$CFG{vcs_lltlink2}{sys5}="en2";

1;
```

Chapter 33

Removing a node from a VCS cluster

This chapter includes the following topics:

Removing a node from a VCS cluster

Table 33-1 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 33-1 Tasks that are involved in removing a node

Task	Reference
■ Back up the configuration file. ■ Check the status of the nodes and the service groups.	See "Verifying the status of nodes and service groups" on page 469.
■ Switch or remove any VCS service groups on the node departing the cluster. ■ Delete the node from VCS configuration.	See "Deleting the departing node from VCS configuration" on page 470.
Modify the llthosts(4) and gabtab(4) files to reflect the change.	See "Modifying configuration files on each remaining node" on page 473.
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 474.

Table 33-1 Tasks that are involved in removing a node *(continued)*

Task	Reference
On the node departing the cluster: ■ Modify startup scripts for LLT, GAB, and VCS to allow reboot of the node without affecting the cluster. ■ Unconfigure and unload the LLT and GAB utilities. ■ Remove the VCS filesets.	See "Unloading LLT and GAB and removing VCS on the departing node" on page 475.

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 VCS configuration

Before you remove a node from the cluster you need to identify the service groups that run on the node.

You then need to perform the following actions:

- Remove the service groups that other service groups depend on, or
- Switch the service groups to another node that other service groups depend on.

To remove or switch service groups from the departing node

1 Switch failover service groups from the departing node. You can switch grp3 from node sys5 to node sys2.

```
# hagrp -switch grp3 -to sys2
```

2 Check for any dependencies involving any service groups that run on the departing node; for example, grp4 runs only on the departing node.

```
# hagrp -dep
```

3 If the service group on the departing node requires other service groups—if it is a parent to service groups on other nodes—unlink the service groups.

```
# haconf -makerw
# hagrp -unlink grp4 grp1
```

These commands enable you to edit the configuration and to remove the requirement grp4 has for grp1.

4 Stop VCS on the departing node:

```
# hastop -sys sys5
```

5 Check the status again. The state of the departing node should be EXITED. Make sure that any service group that you want to fail over is online on other nodes.

```
# hastatus -summary
```

```
-- SYSTEM STATE
-- System      State          Frozen
A  sys1    RUNNING        0
A  sys2      RUNNING        0
A  sys5      EXITED         0

-- GROUP STATE
-- Group    System      Probed   AutoDisabled   State
B  grp1     sys1      Y        N              ONLINE
B  grp1     sys2      Y        N              OFFLINE
B  grp2     sys1      Y        N              ONLINE
B  grp3     sys2      Y        N              ONLINE
B  grp3     sys5    Y        Y              OFFLINE
B  grp4     sys5    Y        N              OFFLINE
```

6 Delete the departing node from the SystemList of service groups grp3 and grp4.

```
# haconf -makerw
# hagrp -modify grp3 SystemList -delete sys5
# hagrp -modify grp4 SystemList -delete sys5
```

Note: If sys5 was in the autostart list, then you need to manually add another system in the autostart list so that after reboot, the group comes online automatically.

7 For the service groups that run only on the departing node, delete the resources from the group before you delete the group.

```
# hagrp -resources grp4
    processx_grp4
    processy_grp4
# hares -delete processx_grp4
# hares -delete processy_grp4
```

8 Delete the service group that is configured to run on the departing node.

```
# hagrp -delete grp4
```

9 Check the status.

```
# hastatus -summary
    -- SYSTEM STATE
    -- System      State          Frozen
    A  sys1      RUNNING        0
    A  sys2      RUNNING        0
    A  sys5    EXITED         0

    -- GROUP STATE
    -- Group    System      Probed   AutoDisabled   State
    B  grp1     sys1      Y        N              ONLINE
    B  grp1     sys2      Y        N              OFFLINE
    B  grp2     sys1      Y        N              ONLINE
    B  grp3     sys2      Y        N              ONLINE
```

10 Delete the node from the cluster.

```
# hasys -delete sys5
```

11 Save the configuration, making it read only.

```
# haconf -dump -makero
```

Modifying configuration files on each remaining node

Perform the following tasks on each of the remaining nodes of the cluster.

To modify the configuration files on a remaining node

1 If necessary, modify the /etc/gabtab file.

No change is required to this file if the `/sbin/gabconfig` command has only the argument `-c`. Symantec recommends using the `-n`*N* option, where *N* is the number of cluster systems.

If the command has the form `/sbin/gabconfig -c -n`*N*, where *N* is the number of cluster systems, make sure that *N* is not greater than the actual number of nodes in the cluster. When *N* is greater than the number of nodes, GAB does not automatically seed.

Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`.

2 Modify /etc/llthosts file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 sys1
1 sys2
2 sys5
```

To:

```
0 sys1
1 sys2
```

Removing the node configuration from the CP server

After removing a node from a VCS cluster, perform the steps in the following procedure to remove that node's configuration from the CP server.

Note: The `cpsadm` command is used to perform the steps in this procedure. For detailed information about the `cpsadm` command, see the *Symantec Cluster Server Administrator's Guide*.

To remove the node configuration from the CP server

1. Log into the CP server as the root user.
2. View the list of VCS users on the CP server.

   If the CP server is configured to use HTTPS-based communication, run the following command:

   ```
   # cpsadm -s cp_server -a list_users
   ```

   If the CP server is configured to use IPM-based communication, run the following command:

   ```
   # cpsadm -s cp_server -p 14250 -a list_users
   ```

   Where *cp_server* is the virtual IP/ virtual hostname of the CP server.
3. Remove the VCS user associated with the node you previously removed from the cluster.

   For CP server in non-secure mode:

   ```
   # cpsadm -s cp_server  -a rm_user \
   -e cpsclient@sys5  -f cps_operator  -g vx
   ```

4. Remove the node entry from the CP server:

   ```
   # cpsadm -s cp_server -a rm_node  -h sys5 -c clus1 -n 2
   ```

5. View the list of nodes on the CP server to ensure that the node entry was removed:

   ```
   # cpsadm -s cp_server -a list_nodes
   ```

Removing security credentials from the leaving node

If the leaving node is part of a cluster that is running in a secure mode, you must remove the security credentials from node sys5. Perform the following steps.

To remove the security credentials

1 Stop the AT process.

```
# /opt/VRTSvcs/bin/vcsauth/vcsauthserver/bin/vcsauthserver.sh \
stop
```

2 Remove the credentials.

```
# rm -rf /var/VRTSvcs/vcsauth/data/
```

Unloading LLT and GAB and removing VCS on the departing node

On the node departing the cluster, unconfigure and unload the LLT and GAB utilities, and remove the VCS filesets.

See "Removing VCS filesets manually" on page 487.

You can use script-based installer to uninstall VCS on the departing node or perform the following manual steps.

If you have configured VCS as part of the Storage Foundation and High Availability products, you may have to delete other dependent filesets before you can delete all of the following ones.

To stop LLT and GAB and remove VCS

1 If you had configured I/O fencing in enabled mode, then stop I/O fencing.

```
# /etc/init.d/vxfen.rc stop
```

2 Stop GAB and LLT:

```
# /etc/init.d/gab.rc stop
# /etc/init.d/llt.rc stop
```

3 To determine the filesets to remove, enter:

```
# lslpp -L |grep VRTS
```

4 To permanently remove the VCS filesets from the system, use the `installp -u` command. Start by removing the following filesets, which may have been optionally installed, in the order shown:

```
# installp -u VRTSfcpi62
# installp -u VRTSvcswiz
# installp -u VRTSvbs
```

```
# installp -u VRTSsfmh
# installp -u VRTSvcsea
# installp -u VRTSvcsag
# installp -u VRTScps
# installp -u VRTSvcs
# installp -u VRTSamf
# installp -u VRTSvxfen
# installp -u VRTSgab
# installp -u VRTSllt
# installp -u VRTSspt
# installp -u VRTSvlic
# installp -u VRTSperl
```

5 Remove the LLT and GAB configuration files.

```
# rm /etc/llttab
# rm /etc/gabtab
# rm /etc/llthosts
```

Section 11

Uninstallation of VCS

Chapter 34

Uninstalling VCS using the installer

This chapter includes the following topics:

Preparing to uninstall VCS

Review the following prerequisites before you uninstall VCS:

- Before you remove VCS from any node in the cluster, shut down the applications that depend on VCS. For example, applications such as Java Console or any high availability agents for VCS.
- If you have manually edited any of the VCS configuration files, you need to reformat them.
  See "Reformatting VCS configuration files on a stopped cluster" on page 76.

Uninstalling VCS using the script-based installer

You must meet the following conditions to use the uninstallvcs to uninstall VCS on all nodes in the cluster at one time:

- Make sure that the communication exists between systems. By default, the uninstaller uses ssh.

- Make sure you can execute `ssh` or `rsh` commands as superuser on all nodes in the cluster.
- Make sure that the `ssh` or `rsh` is configured to operate without requests for passwords or passphrases.

If you cannot meet the prerequisites, then you must run the uninstallvcs on each node in the cluster.

The uninstallvcs removes all VCS filesets.

The following example demonstrates how to uninstall VCS using the uninstallvcs. The uninstallvcs uninstalls VCS on two nodes: sys1 sys2. The example procedure uninstalls VCS from all nodes in the cluster.

Note: If already present on the system, the uninstallation does not remove the VRTSacclib fileset.

Removing VCS 6.2 filesets

The program stops the VCS processes that are currently running during the uninstallation process.

To uninstall VCS

1. Log in as superuser from the node where you want to uninstall VCS.
2. Start uninstallvcs.

   ```
   # cd /opt/VRTS/install
   # ./uninstallvcs<version>
   ```

   Where <version> is the specific release version.

   See "About the script-based installer" on page 50.

   The program specifies the directory where the logs are created. The program displays a copyright notice and a description of the cluster:
3. Enter the names of the systems from which you want to uninstall VCS.

   The program performs system verification checks and asks to stop all running VCS processes. The installer lists all the filesets that it will remove.
4. Enter **y** to stop all the VCS processes.

   The program stops the VCS processes and proceeds with uninstalling the software.
5. Review the output as the uninstallvcs continues to do the following:

- Verifies the communication between systems
- Checks the installations on each system to determine the filesets to be uninstalled.

6 Review the output as the uninstaller stops processes, unloads kernel modules, and removes the filesets.

7 Note the location of summary, response, and log files that the uninstaller creates after removing all the filesets.

Running uninstallvcs from the VCS 6.2 disc

You may need to use the uninstallvcs on the VCS 6.2 disc in one of the following cases:

- You need to uninstall VCS after an incomplete installation.
- The uninstallvcs is not available in /opt/VRTS/install.

If you mounted the installation media to /mnt, access the uninstallvcs by changing directory to:

```
cd /mnt/cluster_server/

./uninstallvcs
```

Uninstalling VCS with the web-based installer

This section describes how to uninstall using the web-based installer.

Note: After you uninstall the product, you cannot access any file systems you created using the default disk layout version in VCS 6.2 with a previous version of VCS.

To uninstall VCS

1 Perform the required steps to save any data that you want to preserve. For example, take backups of configuration files.

2 Disable DMP native support, if it is enabled. Run the following command to disable DMP native support

```
# vxdmpadm settune dmp_native_support=off
# reboot
```

3 Start the web-based installer.

See "Starting the web-based installer" on page 186.

4 On the Select a task and a product page, select **Uninstall a Product** from the Task drop-down list.

5 Select **Symantec Cluster Server** from the Product drop-down list, and click **Next**.

6 Indicate the systems on which to uninstall. Enter one or more system names, separated by spaces. Click **Next**.

7 After the validation completes successfully, click **Next** to uninstall VCS on the selected system.

8 If there are any processes running on the target system, the installer stops the processes. Click **Next**.

9 After the installer stops the processes, the installer removes the products from the specified system.

Click **Next**.

10 After the uninstall completes, the installer displays the location of the summary, response, and log files. If required, view the files to confirm the status of the removal.

11 Click **Finish**.

Most filesets have kernel components. To ensure their complete removal, a system restart is recommended after all the filesets have been removed.

Note: If already present on the system, the uninstallation does not remove the `VRTSacclib` fileset.

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 VCS cluster (application cluster) uses the CP server that you want to unconfigure. Run the `# cpsadm -s CPS_VIP -p CPS_Port -a list_nodes` to know if any application cluster is using the CP server.

To remove the CP server configuration

1 To run the configuration removal script, enter the following command on the node where you want to remove the CP server configuration:

```
root@cps1.symantecexample.com
# /opt/VRTS/install/installvcs<version>  -configcps
```

2 Select option 3 from the menu to unconfigure the CP server.

```
[1] Configure Coordination Point Server on single node VCS system

[2] Configure Coordination Point Server on SFHA cluster

[3] Unconfigure Coordination Point Server
```

3 Review the warning message and confirm that you want to unconfigure the CP server.

```
Unconfiguring coordination point server stops the vxcpserv process.
VCS clusters using this server for coordination purpose will have
one less coordination point.
Are you sure you want to take the CP server offline? [y,n,q] (n) y
```

4 Review the screen output as the script performs the following steps to remove the CP server configuration:

- Stops the CP server
- Removes the CP server from VCS configuration
- Removes resource dependencies
- Takes the the CP server service group (CPSSG) offline, if it is online
- Removes the CPSSG service group from the VCS configuration
- Successfully unconfigured the Veritas Coordination Point Server

```
The CP server database is not being deleted on the shared storage.
It can be re-used if CP server is reconfigured on the cluster.
The same database location can be specified during CP server
configuration.
```

5 Decide if you want to delete the CP server configuration file.

```
Do you want to delete the CP Server configuration file
(/etc/vxcps.conf) and log files
(in /var/VRTScps)? [y,n,q] (n) y

Deleting /etc/vxcps.conf and log files on sys1.... Done
Deleting /etc/vxcps.conf and log files on sys2... Done
```

6 Confirm if you want to send information about this installation to Symantec to help improve installation in the future.

```
Would you like to send the information about this installation
to Symantec to help improve installation in the future? [y,n,q,?] (y)
```

Upload completed successfully.

Chapter 35

Uninstalling VCS using response files

This chapter includes the following topics:

Uninstalling VCS using response files

Typically, you can use the response file that the installer generates after you perform VCS uninstallation on one cluster to uninstall VCS on other clusters.

To perform an automated uninstallation

1. Make sure that you meet the prerequisites to uninstall VCS.
2. Copy the response file to the system where you want to uninstall VCS.

   See "Sample response file for uninstalling VCS" on page 486.

3 Edit the values of the response file variables as necessary.

See “Response file variables to uninstall VCS” on page 485.

4 Start the uninstallation from the system to which you copied the response file. For example:

```
# /opt/VRTS/install/uninstallvcs<version>
 -responsefile /tmp/response_file
```

Where <version> is the specific release version, and /tmp/*response_file* is the response file’s full path name.

See “About the script-based installer” on page 50.

Response file variables to uninstall VCS

Table 35-1 lists the response file variables that you can define to uninstall VCS.

Table 35-1 Response file variables specific to uninstalling VCS

Variable	List or Scalar	Description
CFG{opt}{uninstall}	Scalar	Uninstalls VCS filesets. (Required)
CFG{systems}	List	List of systems on which the product is to be uninstalled. (Required)
CFG{prod}	Scalar	Defines the product to be uninstalled. The value is VCS61 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 *rsh* must be used instead of ssh as the communication method between systems. (Optional)

Table 35-1 Response file variables specific to uninstalling VCS *(continued)*

Variable	List or Scalar	Description
CFG{opt}{logpath}	Scalar	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. **Note:** The installer copies the response files and summary files also to the specified *logpath* location. (Optional)

Sample response file for uninstalling VCS

Review the response file variables and their definitions.

See "Response file variables to uninstall VCS" on page 485.

```
#
# Configuration Values:
#
our %CFG;

$CFG{opt}{uninstall}=1;
$CFG{prod}="VCS62";
$CFG{systems}=[ qw(sys1 sys2) ];
1;
```

Chapter 36

Manually uninstalling VCS

This chapter includes the following topics:

Removing VCS filesets manually

You must remove the VCS filesets from each node in the cluster to uninstall VCS.

To manually remove VCS filesets on a node

1. Shut down VCS on the local system using the `hastop` command.

```
# hastop -local
```

2. Stop the CmdServer.

```
# /opt/VRTSvcs/bin/CmdServer -stop
```

3. Unconfigure the fencing, GAB, LLT, and AMF modules.

```
# /sbin/vxfenconfig -U
# /sbin/gabconfig -U
# /sbin/lltconfig -U
# /opt/VRTSamf/bin/amfconfig -U
```

4. Unload the GAB driver:

```
# /opt/VRTSgab/gabext -u
```

5 Unload the LLT driver:

```
# strload -u -d /usr/lib/drivers/pse/llt
```

6 Before you can remove VRTSveki, you need to remove these filesets if they exist: VRTSalloc, VRTSvxfs, VRTSvxvm.

```
# installp -u VRTSalloc
# installp -u VRTSvxfs
# installp -u VRTSvxvm
```

7 Remove the VCS 6.2 filesets in the following order:

```
# installp -u VRTSvcswiz
# installp -u VRTSvbs
# installp -u VRTSsfmh
# installp -u VRTSvcsea
# installp -u VRTSat(if it exists)
# installp -u VRTSvcsag
# installp -u VRTScps
# installp -u VRTSvcs
# installp -u VRTSamf
# installp -u VRTSvxfen
# installp -u VRTSgab
# installp -u VRTSllt
# installp -u VRTSveki
# installp -u VRTSspt
# installp -u VRTSsfcpi62
# installp -u VRTSperl
# installp -u VRTSvlic
```

Manually remove the CP server fencing configuration

The following procedure describes how to manually remove the CP server fencing configuration from the CP server. This procedure is performed as part of the process to stop and remove server-based IO fencing.

Note: This procedure must be performed after the VCS cluster has been stopped, but before the VCS cluster software is uninstalled.

This procedure is required so that the CP server database can be reused in the future for configuring server-based fencing on the same VCS cluster(s).

Perform the steps in the following procedure to manually remove the CP server fencing configuration.

Note: The `cpsadm` command is used in the following procedure. For detailed information about the `cpsadm` command, see the *Symantec Cluster Server Administrator's Guide*.

To manually remove the CP server fencing configuration

1 Unregister all VCS cluster nodes from all CP servers using the following command:

```
# cpsadm -s cp_server -a unreg_node -u uuid -n nodeid
```

2 Remove the VCS cluster from all CP servers using the following command:

```
# cpsadm -s cp_server -a rm_clus -u uuid
```

3 Remove all the VCS cluster users communicating to CP servers from all the CP servers using the following command:

```
# cpsadm -s cp_server -a rm_user -e user_name -g domain_type
```

4 Proceed to uninstall the VCS cluster software.

Manually deleting cluster details from a CP server

You can manually delete the cluster details from a coordination point server (CP server) using the following procedure.

To manually delete cluster details from a CP server

1 List the nodes in the CP server cluster:

```
# cpsadm -s cps1 -a list_nodes

ClusterName      UUID                                  Hostname(Node ID) Registered
===========   ====================================   ================  ===========
cluster1      {3719a60a-1dd2-11b2-b8dc-197f8305ffc0}  node0(0)          1
```

2 List the CP server users:

```
# cpsadm -s cps1 -a list_users

Username/Domain Type Cluster Name/UUID                                  Role
=================== ==================                                  =======
cpsclient@hostname/vx  cluster1/{3719a60a-1dd2-11b2-b8dc-197f8305ffc0}  Operator
```

3 Remove the privileges for each user of the cluster that is listed in step 2 from the CP server cluster. For example:

```
# cpsadm -s cps1 -a rm_clus_from_user
-c cluster1 -e cpsclient@hostname -g vx -f cps_operator
Cluster successfully deleted from user cpsclient@hostname privileges.
```

4 Remove each user of the cluster that is listed in step 2. For example:

```
# cpsadm -s cps1 -a rm_user -e cpsclient@hostname -g vx
User cpsclient@hostname successfully deleted
```

5 Unregister each node that is registered to the CP server cluster. See the output of step 1 for registered nodes. For example:

```
# cpsadm -s cps1 -a unreg_node -c cluster1 -n 0
Node 0 (node0) successfully unregistered
```

6 Remove each node from the CP server cluster. For example:

```
# cpsadm -s cps1 -a rm_node -c cluster1 -n 0
Node 0 (node0) successfully deleted
```

7 Remove the cluster.

```
# cpsadm -s cps1 -a rm_clus -c cluster1
Cluster cluster1 deleted successfully
```

8 Verify that the cluster details are removed successfully.

```
# cpsadm -s cps1 -a list_nodes

ClusterName      UUID           Hostname(Node ID) Registered
===========   ================  ================  ===========

# cpsadm -s cps1 -a list_users

Username/Domain Type Cluster Name/UUID     Role
==================== ==================    =======
```

Section 12

Installation reference

Appendix A

Services and ports

This appendix includes the following topics:

About SFHA services and ports

If you have configured a firewall, ensure that the firewall settings allow access to the services and ports used by SFHA.

Table A-1 lists the services and ports used by SFHA .

Note: The port numbers that appear in bold are mandatory for configuring SFHA.

Table A-1 SFHA services and ports

Port Number	Protocol	Description	Process
4145	TCP/UDP	VVR Connection Server VCS Cluster Heartbeats	vxio
5634	HTTPS	Symantec Storage Foundation Messaging Service	xprtld
8199	TCP	Volume Replicator Administrative Service	vras
8989	TCP	VVR Resync Utility	vxreserver

Table A-1 SFHA services and ports *(continued)*

Port Number	Protocol	Description	Process
14141	TCP	Symantec High Availability Engine Veritas Cluster Manager (Java console) (ClusterManager.exe) VCS Agent driver (VCSAgDriver.exe)	had
14144	TCP/UDP	VCS Notification	Notifier
14149	TCP/UDP	VCS Authentication	vcsauthserver
14150	TCP	Veritas Command Server	CmdServer
14155	TCP/UDP	VCS Global Cluster Option (GCO)	wac
14156	TCP/UDP	VCS Steward for GCO	steward
443	TCP	Coordination Point Server	Vxcpserv
49152-65535	TCP/UDP	Volume Replicator Packets	User configurable ports created at kernel level by `vxio.sys` file

Appendix B

VCS installation filesets

This appendix includes the following topics:

Symantec Cluster Server installation filesets

Table B-1 shows the fileset name and contents for each Symantec Cluster Server fileset.

Table B-1 Symantec Cluster Server filesets

fileset	Contents	Required/Optional
VRTSamf	Contains the binaries for the Veritas Asynchronous Monitoring Framework kernel driver functionality for all the IMF-aware agents.	Required
VRTScps	Contains the binaries for the Veritas Coordination Point Server.	Optional. Required to Coordination Point Server (CPS).
VRTSgab	Contains the binaries for Symantec Cluster Server group membership and atomic broadcast services.	Required Depends on VRTSllt.
VRTSllt	Contains the binaries for Symantec Cluster Server low-latency transport.	Required Depends on VRTSveki.
VRTSperl	Contains Perl binaries for Veritas.	Required

Table B-1 Symantec Cluster Server filesets *(continued)*

fileset	Contents	Required/Optional
VRTSsfcpi62	Product Installer The product installer fileset contains the scripts that perform the following: ■ installation ■ configuration ■ upgrade ■ uninstallation ■ adding nodes ■ removing nodes ■ etc. You can use this script to simplify the native operating system installations, configurations, and upgrades.	Required
VRTSvcswiz	Contains the wizards for Symantec Cluster Server by Symantec.	Required
VRTSspt	Contains the binaries for Veritas Software Support Tools.	Recommended fileset, optional
VRTSvcs	VRTSvcs contains the following components: ■ Contains the binaries for Symantec Cluster Server. ■ Contains the binaries for Symantec Cluster Server manual pages. ■ Contains the binaries for Symantec Cluster Server English message catalogs. ■ Contains the binaries for Symantec Cluster Server utilities. These utilities include security services.	Required Depends on VRTSperl and VRTSvlic.
VRTSvcsag	Contains the binaries for Symantec Cluster Server bundled agents.	Required Depends on VRTSvcs.

Table B-1 Symantec Cluster Server filesets *(continued)*

fileset	Contents	Required/Optional
VRTSvcsea	VRTSvcsea contains the binaries for Veritas high availability agents for DB2, Sybase, and Oracle.	Optional for VCS. Required to use VCS with the high availability agents for DB2, Sybase, or Oracle.
VRTSveki	Contains the Veritas kernel interface, which is a common set of modules that other Veritas drivers use.	Required
VRTSvlic	Contains the binaries for Symantec License Utilities.	Required
VRTSvxfen	Contains the binaries for Veritas I/O Fencing .	Required to use fencing. Depends on VRTSgab.
VRTSsfmh	Symantec Storage Foundation Managed Host Symantec Storage Foundation Managed Host is now called Veritas Operations Manager (VOM). VOM 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 fileset on a server, and add this managed host to the Central Server. The VRTSsfmcs fileset is not part of this release. You can download it separately from: http://www.symantec.com/veritas-operations-manager	Recommended

Table B-1 Symantec Cluster Server filesets *(continued)*

fileset	Contents	Required/Optional
VRTSvbs	Enables fault management and VBS command line operations on VCS nodes managed by Veritas Operations Manager. For more information, see the *Virtual Business Service–Availability User's Guide*.	Recommended Depends on VRTSsfmh. VRTSsfmh version must be 4.1 or later for VRTSvbs to get installed.

Appendix C

Installation command options

This appendix includes the following topics:

Command options for installvcs

The `installvcs` command usage takes the following form:

```
installvcs [ system1 system2 ... ]
 [ -configure | -install | -license | -upgrade | -precheck | -requirements
              | -start | -stop | -postcheck]

       [ -require installer_hot_fix_file ]
       [ -responsefile response_file ]
       [ -logpath log_path ]
       [ -tmppath tmp_path ]
       [ -tunablesfile tunables_file ]
       [ -timeout timeout_value ]
       [ -keyfile ssh_key_file ]
       [ -hostfile hostfile_path ]
       [ -pkgpath pkg_path ]
       [ -prod product_name ]
       [ -hotfix_path hotfix_path ]
       [ -hotfix2_path hotfix2_path ]
       [ -hotfix3_path hotfix3_path ]
```

```
        [ -hotfix4_path hotfix4_path ]
        [ -hotfix5_path hotfix5_path ]
        [ -nim LPP_SOURCE ]
        [ -serial | -rsh | -redirect | -installminpkgs | -installrecpkgs |
-installallpkgs | -minpkgs | -recpkgs | -allpkgs | -pkgset | -pkgtable |
-pkginfo | -makeresponsefile | -comcleanup | -comsetup | -version | -nolic |
-settunables | -tunables | -noipc | -security | -securityonenode |
-securitytrust | -fips | -addnode | -fencing | -configcps | -rolling_upgrade
-rollingupgrade_phase1 | -rollingupgrade_phase2 ]
```

Installation script options

Table C-1 shows command line options for the installation script. For an initial install or upgrade, options are not usually required. The installation script options apply to all Symantec Storage Foundation product scripts, except where otherwise noted.

See "About the script-based installer" on page 50.

Table C-1 Available command line options

Command Line Option	Function
-addnode	Adds a node to a high availability cluster.
-allpkgs	Displays all filesets required for the specified product. The filesets are listed in correct installation order. The output can be used to create scripts for command line installs, or for installations over a network.
-comcleanup	The `-comcleanup` option removes the secure shell or remote shell configuration added by installer on the systems. The option is only required when installation routines that performed auto-configuration of the shell are abruptly terminated.
-comsetup	The `-comsetup` option is used to set up the ssh or rsh communication between systems without requests for passwords or passphrases.
-configcps	The `-configcps` 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.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
–hostfile *full_path_to_file*	Specifies the location of a file that contains a list of hostnames on which to install.
-disable_dmp_native_support	Disables Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade. Retaining Dynamic Multi-pathing support for the native LVM volume groups and ZFS pools during upgrade increases fileset upgrade time depending on the number of LUNs and native LVM volume groups and ZFS pools configured on the system.
-online_upgrade	Used to perform online upgrade. Using this option, the installer upgrades the whole cluster and also supports customer's application zero down time during the upgrade procedure. Now this option only supports VCS and ApplicationHA.
-patch_path	Defines the path of a patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed .
-patch2_path	Defines the path of a second patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch3_path	Defines the path of a third patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch4_path	Defines the path of a fourth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-patch5_path	Defines the path of a fifth patch level release to be integrated with a base or a maintenance level release in order for multiple releases to be simultaneously installed.
-installallpkgs	The `-installallpkgs` option is used to select all filesets.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
-installrecpkgs	The `-installrecpkgs`option is used to select the recommended filesets set.
–installminpkgs	The `-installminpkgs`option is used to select the minimum filesets set.
-ignorepatchreqs	The `-ignorepatchreqs` option is used to allow installation or upgrading even if the prerequisite filesets or patches are missed on the system.
–keyfile *ssh_key_file*	Specifies a key file for secure shell (SSH) installs. This option passes `-I ssh_key_file` to every SSH invocation.
-license	Registers or updates product licenses on the specified systems.
–logpath *log_path*	Specifies a directory other than `/opt/VRTS/install/logs` as the location where installer log files, summary files, and response files are saved.
-makeresponsefile	Use the `-makeresponsefile` option only to generate response files. No actual software installation occurs when you use this option.
-minpkgs	Displays the minimal filesets required for the specified product. The filesets are listed in correct installation order. Optional filesets are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-nim	Produces a NIM configuration file for installing with NIM.
-noipc	Disables the installer from making outbound networking calls to Symantec Operations Readiness Tool (SORT) in order to automatically obtain patch and release information updates.
-nolic	Allows installation of product filesets without entering a license key. Licensed features cannot be configured, started, or used when this option is specified.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
–pkginfo	Displays a list of filesets and the order of installation in a human-readable format. This option only applies to the individual product installation scripts. For example, use the -pkginfo option with the installvcs script to display VCS filesets.
–pkgset	Discovers and displays the fileset group (minimum, recommended, all) and filesets that are installed on the specified systems.
-pkgtable	Displays product's filesets in correct installation order by group.
–postcheck	Checks for different HA and file system-related processes, the availability of different ports, and the availability of cluster-related service groups.
-precheck	Performs a preinstallation check to determine if systems meet all installation requirements. Symantec recommends doing a precheck before installing a product.
-prod	Specifies the product for operations.
–recpkgs	Displays the recommended filesets required for the specified product. The filesets are listed in correct installation order. Optional filesets are not listed. The output can be used to create scripts for command line installs, or for installations over a network. See `allpkgs` option.
-redirect	Displays progress details without showing the progress bar.
-require	Specifies an installer patch file.
-requirements	The `-requirements` option displays required OS version, required filesets and patches, file system space, and other system requirements in order to install the product.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
–responsefile *response_file*	Automates installation and configuration by using system and configuration information stored in a specified file instead of prompting for information. The *response_file* must be a full path name. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file.
-rolling_upgrade	Starts a rolling upgrade. Using this option, the installer detects the rolling upgrade status on cluster systems automatically without the need to specify rolling upgrade phase 1 or phase 2 explicitly.
-rollingupgrade_phase1	The `-rollingupgrade_phase1` option is used to perform rolling upgrade Phase-I. In the phase, the product kernel filesets get upgraded to the latest version.
-rollingupgrade_phase2	The `-rollingupgrade_phase2` option is used to perform rolling upgrade Phase-II. In the phase, VCS and other agent filesets upgrade to the latest version. Product kernel drivers are rolling-upgraded to the latest protocol version.
-rsh	Specify this option when you want to use RSH and RCP for communication between systems instead of the default SSH and SCP. See "About configuring secure shell or remote shell communication modes before installing products" on page 543.
-securitytrust	The `-securitytrust` 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.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
-settunables	Specify this option when you want to set tunable parameters after you install and configure a product. You may need to restart processes of the product for the tunable parameter values to take effect. You must use this option together with the `-tunablesfile` option.
-start	Starts the daemons and processes for the specified product.
-stop	Stops the daemons and processes for the specified product.
-timeout	The `-timeout` option is used to specify the number of seconds that the script should wait for each command to complete before timing out. Setting the `-timeout` option overrides the default value of 1200 seconds. Setting the `-timeout` option to 0 prevents the script from timing out. The `-timeout` option does not work with the `-serial` option
–tmppath *tmp_path*	Specifies a directory other than `/var/tmp` as the working directory for the installation scripts. This destination is where initial logging is performed and where filesets are copied on remote systems before installation.
-tunables	Lists all supported tunables and create a tunables file template.
-tunables_file *tunables_file*	Specify this option when you specify a tunables file. The tunables file should include tunable parameters.
-upgrade	Specifies that an existing version of the product exists and you plan to upgrade it.

Table C-1 Available command line options *(continued)*

Command Line Option	Function
-version	Checks and reports the installed products and their versions. Identifies the installed and missing filesets and patches where applicable for the product. Provides a summary that includes the count of the installed and any missing filesets and patches where applicable. Lists the installed patches, patches, and available updates for the installed product if an Internet connection is available.

Command options for uninstallvcs

The `uninstallvcs` command usage takes the following form:

```
uninstallvcs [ system1 system2... ]
       [ -require installer_hot_fix_file ]
       [ -responsefile response_file ]
       [ -logpath log_path ]
       [ -tmppath tmp_path ]
       [ -timeout timeout_value ]
       [ -keyfile ssh_key_file ]
       [ -hostfile hostfile_path ]
       [ -prod product_name ]
       [ -hotfix_path hotfix_path ]
       [ -hotfix2_path hotfix2_path ]
       [ -hotfix3_path hotfix3_path ]
       [ -hotfix4_path hotfix4_path ]
       [ -hotfix5_path hotfix5_path ]
       [ -serial | -rsh | -redirect | -makeresponsefile | -comcleanup |
-comsetup | -version | -noipc ]
```

For description of the `uninstallvcs` command options:

See "Installation script options" on page 500.

Appendix D

Configuration files

This appendix includes the following topics:

About the LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires /etc/llthosts and /etc/lltab files. GAB requires /etc/gabtab file.

Table D-1 lists the LLT configuration files and the information that these files contain.

Table D-1 LLT configuration files

File	Description
/etc/default/llt	This file stores the start and stop environment variables for LLT: ■ LLT_START—Defines the startup behavior for the LLT module after a system reboot. Valid values include: 1—Indicates that LLT is enabled to start up. 0—Indicates that LLT is disabled to start up. ■ LLT_STOP—Defines the shutdown behavior for the LLT module during a system shutdown. Valid values include: 1—Indicates that LLT is enabled to shut down. 0—Indicates that LLT is disabled to shut down. The installer sets the value of these variables to 1 at the end of VCS configuration. If you manually configured VCS, make sure you set the values of these environment variables to 1.

/etc/llthosts	The file `llthosts` is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file must be identical on each node in the cluster. A mismatch of the contents of the file can cause indeterminate behavior in the cluster.
	For example, the file /etc/llthosts contains the entries that resemble:
	`0 sys1`
	`1 sys2`

Table D-1 LLT configuration files *(continued)*

File	Description

/etc/llttab	The file `llttab` contains the information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the LLT network links that correspond to the specific system.
	
	For example, the file /etc/llttab contains the entries that resemble:
	
	`set-node sys1`
	`set-cluster 2`
	`link en1 /dev/dlpi/en:1 - ether - -`
	`link en2 /dev/dlpi/en:2 - ether - -`
	
	`set-node sys1`
	`set-cluster 2`
	`link en1 /dev/en:1 - ether - -`
	`link en2 /dev/en:2 - ether - -`
	
	The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines begin with the `link` command. These lines identify the two network cards that the LLT protocol uses.
	
	If you configured a low priority link under LLT, the file also includes a "link-lowpri" line.
	
	Refer to the `lltab(4)` manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the `llttab` file.

Table D-2 lists the GAB configuration files and the information that these files contain.

Table D-2 GAB configuration files

File	Description
/etc/default/gab	This file stores the start and stop environment variables for GAB: ■ GAB_START—Defines the startup behavior for the GAB module after a system reboot. Valid values include: 1—Indicates that GAB is enabled to start up. 0—Indicates that GAB is disabled to start up. ■ GAB_STOP—Defines the shutdown behavior for the GAB module during a system shutdown. Valid values include: 1—Indicates that GAB is enabled to shut down. 0—Indicates that GAB is disabled to shut down. The installer sets the value of these variables to 1 at the end of VCS configuration. If you manually configured VCS, make sure you set the values of these environment variables to 1.
/etc/gabtab	After you install VCS, the file /etc/gabtab contains a `gabconfig(1)` command that configures the GAB driver for use. The file /etc/gabtab contains a line that resembles: `/sbin/gabconfig -c -nN` The `-c` option configures the driver for use. The `-nN` specifies that the cluster is not formed until at least `N` nodes are ready to form the cluster. Symantec recommends that you set N to be the total number of nodes in the cluster. **Note:** Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. Using `-c -x` can lead to a split-brain condition. Use the `-c` option for `/sbin/gabconfig` to avoid a split-brain condition. **Note:**

About the AMF configuration files

Asynchronous Monitoring Framework (AMF) kernel driver provides asynchronous event notifications to the VCS agents that are enabled for intelligent resource monitoring.

Table D-3 lists the AMF configuration files.

Table D-3 AMF configuration files

File	Description
/etc/default/amf	This file stores the start and stop environment variables for AMF: ■ AMF_START—Defines the startup behavior for the AMF module after a system reboot or when AMF is attempted to start using the init script. Valid values include: 1—Indicates that AMF is enabled to start up. (default) 0—Indicates that AMF is disabled to start up. ■ AMF_STOP—Defines the shutdown behavior for the AMF module during a system shutdown or when AMF is attempted to stop using the init script. Valid values include: 1—Indicates that AMF is enabled to shut down. (default) 0—Indicates that AMF is disabled to shut down.
/etc/amftab	After you install VCS, the file `/etc/amftab` contains a `amfconfig`(1) command that configures the AMF driver for use. The AMF init script uses this `/etc/amftab` file to configure the AMF driver. The `/etc/amftab` file contains the following line by default: `/opt/VRTSamf/bin/amfconfig -c`

About the VCS configuration files

VCS configuration files include the following:

- main.cf
  The installer creates the VCS configuration file in the /etc/VRTSvcs/conf/config folder by default during the VCS configuration. The main.cf file contains the minimum information that defines the cluster and its nodes.
  See "Sample main.cf file for VCS clusters" on page 512.
  See "Sample main.cf file for global clusters" on page 514.
- types.cf
  The file types.cf, which is listed in the include statement in the main.cf file, defines the VCS bundled types for VCS resources. The file types.cf is also located in the folder /etc/VRTSvcs/conf/config.
  Additional files similar to types.cf may be present if agents have been added, such as OracleTypes.cf.

Note the following information about the VCS configuration file after installing and configuring VCS:

- The cluster definition includes the cluster information that you provided during the configuration. This definition includes the cluster name, cluster address, and the names of users and administrators of the cluster.
  Notice that the cluster has an attribute UserNames. The installvcs creates a user "admin" whose password is encrypted; the word "password" is the default password.
- If you set up the optional I/O fencing feature for VCS, then the UseFence = SCSI3 attribute is present.
- If you configured the cluster in secure mode, the main.cf includes "SecureClus = 1" cluster attribute.
- The installvcs creates the ClusterService service group if you configured the virtual IP, SMTP, SNMP, or global cluster options.
  The service group also has the following characteristics:
  - The group includes the IP and NIC resources.
  - The service group also includes the notifier resource configuration, which is based on your input to installvcs prompts about notification.
  - The installvcs also creates a resource dependency tree.
  - If you set up global clusters, the ClusterService service group contains an Application resource, wac (wide-area connector). This resource's attributes contain definitions for controlling the cluster in a global cluster environment. Refer to the *Symantec Cluster Server Administrator's Guide* for information about managing VCS global clusters.

Refer to the *Symantec Cluster Server Administrator's Guide* to review the configuration concepts, and descriptions of main.cf and types.cf files for AIX systems.

Sample main.cf file for VCS clusters

The following sample main.cf file is for a three-node cluster in secure mode.

```
include "types.cf"
include "OracleTypes.cf"
include "OracleASMTypes.cf"
include "Db2udbTypes.cf"
include "SybaseTypes.cf"


cluster vcs02 (
    SecureClus = 1
    )
```

```
system sysA (
    )

system sysB (
    )

system sysC (
    )

group ClusterService (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
    )

NIC csgnic (
    Device = en0
    NetworkHosts = { "10.182.13.1" }
    )

NotifierMngr ntfr (
   SnmpConsoles = { sys4" = SevereError }
   SmtpServer = "smtp.example.com"
   SmtpRecipients =  { "ozzie@example.com" = SevereError }
   )

ntfr requires csgnic

// resource dependency tree
//
//     group ClusterService
//     {
//     NotifierMngr ntfr
//         {
//         NIC csgnic
//         }
// }
```

Sample main.cf file for global clusters

If you installed VCS with the Global Cluster option, note that the ClusterService group also contains the Application resource, wac. The wac resource is required to control the cluster in a global cluster environment.

In the following main.cf file example, bold text highlights global cluster specific entries.

```
include "types.cf"

cluster vcs03 (
    ClusterAddress = "10.182.13.50"
    SecureClus = 1
    )

system sysA (
    )

system sysB (
    )

system sysC (
    )

group ClusterService (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
    )

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart -secure"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac -secure" }
    RestartLimit = 3
    )

IP gcoip (
    Device = en0
    Address = "10.182.13.50"
    NetMask = "255.255.240.0"
    )
```

```
NIC csgnic (
    Device = en0
    NetworkHosts = { "10.182.13.1" }
    )

NotifierMngr ntfr (
   SnmpConsoles = { sys4 = SevereError }
   SmtpServer = "smtp.example.com"
   SmtpRecipients =  { "ozzie@example.com" = SevereError }
   )

gcoip requires csgnic
ntfr requires csgnic
wac requires gcoip

// resource dependency tree
//
//     group ClusterService
//     {
//     NotifierMngr ntfr
//         {
//         NIC csgnic
//         }
//     Application wac
//         {
//         IP gcoip
//             {
//             NIC csgnic
//             }
//         }
//     }
```

About I/O fencing configuration files

Table D-4 lists the I/O fencing configuration files.

Table D-4 I/O fencing configuration files

File	Description
/etc/default/vxfen	This file stores the start and stop environment variables for I/O fencing: ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. The installer sets the value of these variables to 1 at the end of VCS configuration. If you manually configured VCS, you must make sure to set the values of these environment variables to 1.
/etc/vxfendg	This file includes the coordinator disk group information. This file is not applicable for server-based fencing and majority-based fencing.

Table D-4 I/O fencing configuration files *(continued)*

File	Description
/etc/vxfenmode	This file contains the following parameters: ■ vxfen_mode ■ scsi3—For disk-based fencing. ■ customized—For server-based fencing. ■ disabled—To run the I/O fencing driver but not do any fencing operations. ■ majority— For fencing without the use of coordination points. ■ vxfen_mechanism This parameter is applicable only for server-based fencing. Set the value as cps. ■ scsi3_disk_policy ■ dmp—Configure the vxfen module to use DMP devices The disk policy is dmp by default. If you use iSCSI devices, you must set the disk policy as dmp. **Note:** You must use the same SCSI-3 disk policy on all the nodes. ■ List of coordination points This list is required only for server-based fencing configuration. Coordination points in server-based fencing can include coordinator disks, CP servers, or both. If you use coordinator disks, you must create a coordinator disk group containing the individual coordinator disks. Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points and multiple IP addresses for each CP server. ■ single_cp This parameter is applicable for server-based fencing which uses a single highly available CP server as its coordination point. Also applicable for when you use a coordinator disk group with single disk. ■ autoseed_gab_timeout This parameter enables GAB automatic seeding of the cluster even when some cluster nodes are unavailable. This feature is applicable for I/O fencing in SCSI3 and customized mode. 0—Turns the GAB auto-seed feature on. Any value greater than 0 indicates the number of seconds that GAB must delay before it automatically seeds the cluster. -1—Turns the GAB auto-seed feature off. This setting is the default.

Table D-4 I/O fencing configuration files *(continued)*

File	Description

/etc/vxfentab	When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.
	Note: The /etc/vxfentab file is a generated file; do not modify this file.
	For disk-based I/O fencing, the /etc/vxfentab file on each node contains a list of all paths to each coordinator disk along with its unique disk identifier. A space separates the path and the unique disk identifier. An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:
	■ DMP disk:
	`/dev/vx/rdmp/rhdisk75 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006804E795D075`
	`/dev/vx/rdmp/rhdisk76 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006814E795D076`
	`/dev/vx/rdmp/rhdisk77 HITACHI%5F1724-100%20%20FAStT%5FDISKS%5F6 00A0B8000215A5D000006824E795D077`
	For server-based fencing, the /etc/vxfentab file also includes the security settings information.
	For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.
	This file is not applicable for majority-based fencing.

Sample configuration files for CP server

The `/etc/vxcps.conf` file determines the configuration of the coordination point server (CP server.)

See "Sample CP server configuration (/etc/vxcps.conf) file output" on page 524.

The following are example main.cf files for a CP server that is hosted on a single node, and a CP server that is hosted on an SFHA cluster.

- The main.cf file for a CP server that is hosted on a single node:
  See "Sample main.cf file for CP server hosted on a single node that runs VCS" on page 519.
- The main.cf file for a CP server that is hosted on an SFHA cluster:

See "Sample main.cf file for CP server hosted on a two-node SFHA cluster" on page 521.

Note: If you use IPM-based protocol for communication between the CP server and VCS clusters (application clusters), the CP server supports Internet Protocol version 4 or version 6 (IPv4 or IPv6 addresses). If you use HTTPS-based protocol for communication, the CP server only supports Internet Protocol version 4 (IPv4 addresses).

The example main.cf files use IPv4 addresses.

Sample main.cf file for CP server hosted on a single node that runs VCS

The following is an example of a single CP server node main.cf.

For this CP server single node main.cf, note the following values:

- Cluster name: cps1
- Node name: cps1

```
include "types.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"

// cluster name:  cps1
// CP server: cps1

cluster cps1 (
     UserNames = { admin = bMNfMHmJNiNNlVNhMK, haris = fopKojNvpHouNn,
             "cps1.symantecexample.com@root@vx" = aj,
             "root@cps1.symantecexample.com" = hq }
     Administrators = { admin, haris,
             "cps1.symantecexample.com@root@vx",
             "root@cps1.symantecexample.com" }
     SecureClus = 1
     HacliUserLevel = COMMANDROOT
     )

system cps1 (
     )

group CPSSG (
      SystemList = { cps1 = 0 }
```

```
        AutoStartList = { cps1 }
        )

        IP cpsvip1 (
              Critical = 0
              Device @cps1 = en0
              Address = "10.209.3.1"
              NetMask = "255.255.252.0"
              )

        IP cpsvip2 (
              Critical = 0
              Device @cps1 = en1
              Address = "10.209.3.2"
              NetMask = "255.255.252.0"
              )

        NIC cpsnic1 (
             Critical = 0
             Device @cps1 = en0
             PingOptimize = 0
             NetworkHosts @cps1 = { "10.209.3.10 }
             )

        NIC cpsnic2 (
             Critical = 0
             Device @cps1 = en1
             PingOptimize = 0
             )

        Process vxcpserv (
             PathName = "/opt/VRTScps/bin/vxcpserv"
             ConfInterval = 30
             RestartLimit = 3
             )

        Quorum quorum (
              QuorumResources = { cpsvip1, cpsvip2 }
              )

cpsvip1 requires cpsnic1
cpsvip2 requires cpsnic2
vxcpserv requires quorum
```

```
// resource dependency tree
//
// group CPSSG
// {
// IP cpsvip1
//     {
//     NIC cpsnic1
//     }
// IP cpsvip2
//     {
//     NIC cpsnic2
//     }
// Process vxcpserv
//     {
//     Quorum quorum
//     }
// }
```

Sample main.cf file for CP server hosted on a two-node SFHA cluster

The following is an example of a main.cf, where the CP server is hosted on an SFHA cluster.

For this CP server hosted on an SFHA cluster main.cf, note the following values:

- Cluster name: cps1
- Nodes in the cluster: cps1, cps2

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "/opt/VRTScps/bin/Quorum/QuorumTypes.cf"


// cluster: cps1
// CP servers:
// cps1
// cps2

cluster cps1 (
     UserNames = { admin = ajkCjeJgkFkkIskEjh,
            "cps1.symantecexample.com@root@vx" = JK,
```

```
            "cps2.symantecexample.com@root@vx" = dl }
      Administrators = { admin, "cps1.symantecexample.com@root@vx",
            "cps2.symantecexample.com@root@vx" }
      SecureClus = 1
      )

system cps1 (
      )

system cps2 (
      )

group CPSSG (
       SystemList = { cps1 = 0, cps2 = 1 }
       AutoStartList = { cps1, cps2 } )

       DiskGroup cpsdg (
             DiskGroup = cps_dg
             )

       IP cpsvip1 (
             Critical = 0
             Device @cps1 = en0
             Device @cps2 = en0
             Address = "10.209.81.88"
             NetMask = "255.255.252.0"
             )

       IP cpsvip2 (
             Critical = 0
             Device @cps1 = en1
             Device @cps2 = en1
             Address = "10.209.81.89"
             NetMask = "255.255.252.0"
             )

       Mount cpsmount (
             MountPoint = "/etc/VRTScps/db"
             BlockDevice = "/dev/vx/dsk/cps_dg/cps_volume"
             FSType = vxfs
             FsckOpt = "-y"
             )
```

```
        NIC cpsnic1 (
                Critical = 0
                Device @cps1 = en0
                Device @cps2 = en0
                PingOptimize = 0
                NetworkHosts @cps1 = { "10.209.81.10 }
                )

        NIC cpsnic2 (
                Critical = 0
                Device @cps1 = en1
                Device @cps2 = en1
                PingOptimize = 0
                )

        Process vxcpserv (
                 PathName = "/opt/VRTScps/bin/vxcpserv"
                 )

        Quorum quorum (
                 QuorumResources = { cpsvip1, cpsvip2 }
                 )

        Volume cpsvol (
                 Volume = cps_volume
                 DiskGroup = cps_dg
                 )

    cpsmount requires cpsvol
    cpsvip1 requires cpsnic1
    cpsvip2 requires cpsnic2
    cpsvol requires cpsdg
    vxcpserv requires cpsmount
    vxcpserv requires quorum


    // resource dependency tree
    //
    // group CPSSG
    // {
    // IP cpsvip1
    //     {
    //     NIC cpsnic1
```

```
//     }
// IP cpsvip2
//     {
//     NIC cpsnic2
//     }
// Process vxcpserv
//     {
//     Quorum quorum
//     Mount cpsmount
//         {
//         Volume cpsvol
//             {
//             DiskGroup cpsdg
//             }
//         }
//     }
// }
```

Sample CP server configuration (/etc/vxcps.conf) file output

The following is an example of a coordination point server (CP server) configuration file `/etc/vxcps.conf` output.

```
##  The vxcps.conf file determines the
## configuration for Veritas CP Server.
cps_name=cps1
vip=[10.209.81.88]
vip=[10.209.81.89]:56789
vip_https=[10.209.81.88]:55443
vip_https=[10.209.81.89]
port=14250
port_https=443
security=1
db=/etc/VRTScps/db
ssl_conf_file=/etc/vxcps_ssl.properties
```

Appendix E

Installing VCS on a single node

This appendix includes the following topics:

About installing VCS on a single node

You can install VCS 6.2 on a single node. You can subsequently add another node to the single-node cluster to form a multinode cluster. You can also prepare a single node cluster for addition into a multi-node cluster. Single node clusters can be used for testing as well.

You can install VCS onto a single node using the installer program or you can add it manually.

See "Creating a single-node cluster using the installer program" on page 526.

See "Creating a single-node cluster manually" on page 527.

Creating a single-node cluster using the installer program

Table E-1 specifies the tasks that are involved to install VCS on a single node using the installer program.

Table E-1 Tasks to create a single-node cluster using the installer

Task	Reference
Prepare for installation.	See "Preparing for a single node installation" on page 526.
Install the VCS software on the system using the installer.	See "Starting the installer for the single node cluster" on page 526.

Preparing for a single node installation

You can use the installer program to install a cluster on a single system for either of the two following purposes:

- To prepare the single node cluster to join a larger cluster
- To prepare the single node cluster to be a stand-alone single node cluster

When you prepare it to join a larger cluster, enable it with LLT and GAB. For a stand-alone cluster, you do not need to enable LLT and GAB.

For more information about LLT and GAB:

See "About LLT and GAB" on page 25.

Starting the installer for the single node cluster

When you install VCS on a single system, follow the instructions in this guide for installing VCS using the product installer.

During the installation, you need to answer two questions specifically for single node installations. When the installer asks:

```
Enter the system names separated by spaces on which to install
VCS[q,?]
```

Enter a single system name. While you configure, the installer asks if you want to enable LLT and GAB:

```
If you plan to run VCS on a single node without any need for
adding cluster node online, you have an option to proceed
```

```
without starting GAB and LLT.
Starting GAB and LLT is recommended.
Do you want to start GAB and LLT? [y,n,q,?] (y)
```

Answer **n** if you want to use the single node cluster as a stand-alone cluster.

Selecting **n** disables LLT, GAB, and I/O fencing kernel modules of VCS. So, the kernel programs are not loaded to the node.

Answer **y** if you plan to incorporate the single node cluster into a multi-node cluster in the future.

Continue with the installation.

Creating a single-node cluster manually

Table E-2 specifies the tasks that you need to perform to install VCS on a single node.

Table E-2 Tasks to create a single-node cluster manually

Task	Reference
Set the PATH variable	See "Setting the path variable for a manual single node installation" on page 527.
Install the VCS software manually and add a license key	See "Installing VCS software manually on a single node" on page 528.
Remove any LLT or GAB configuration files and rename LLT and GAB startup files. A single-node cluster does not require the node-to-node communication service, LLT, or the membership communication service, GAB.	
Start VCS and verify single-node operation.	See "Verifying single-node operation" on page 528.

Setting the path variable for a manual single node installation

Set the path variable.

See "Setting the PATH variable" on page 72.

Installing VCS software manually on a single node

Install the VCS 6.2 filesets manually and install the license key.

Refer to the following sections:

- See “Installing VCS software manually” on page 241.
- See “Adding a license key for a manual installation” on page 244.

Configuring VCS

You now need to configure VCS.

See “Configuring VCS manually” on page 254.

Verifying single-node operation

After successfully creating a single-node cluster, start VCS and verify the cluster.

To verify single-node cluster

1 Bring up VCS manually as a single-node cluster using `hastart` with the `-onenode` option:

```
# hastart -onenode
```

2 Verify that the `had` and `hashadow` daemons are running in single-node mode:

```
# ps -ef | grep had
root  285   1  0 14:49:31 ?  0:02 /opt/VRTSvcs/bin/had -onenode
root  288   1  0 14:49:33 ?  0:00 /opt/VRTSvcs/bin/hashadow
```

Appendix F

Configuring LLT over UDP

This appendix includes the following topics:

Using the UDP layer for LLT

VCS provides the option of using LLT over the UDP (User Datagram Protocol) layer for clusters using wide-area networks and routers. UDP makes LLT packets routable and thus able to span longer distances more economically.

When to use LLT over UDP

Use LLT over UDP in the following situations:

- LLT must be used over WANs
- When hardware, such as blade servers, do not support LLT over Ethernet

LLT over UDP is slower than LLT over Ethernet. Use LLT over UDP only when the hardware configuration makes it necessary.

Manually configuring LLT over UDP using IPv4

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on separate subnets. Set the broadcast address in /etc/lltab explicitly depending on the subnet for each link.

See "Broadcast address in the /etc/llttab file" on page 530.

- 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 532.
- Set the broadcast address correctly for direct-attached (non-routed) links.
  See "Sample configuration: direct-attached links" on page 534.
- 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 535.

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 /dev/xti/udp - udp 50000  -  192.168.9.1 192.168.9.255
link link2 /dev/xti/udp - udp 50001  -  192.168.10.1 192.168.10.255
```

  Verify the subnet mask using the ifconfig command to ensure that the two links are on separate subnets.

- Display the content of the /etc/llttab file on the second node sys2:

```
sys2 # cat /etc/llttab

set-node sys2
set-cluster 1
link link1 /dev/xti/udp - udp 50000  -  192.168.9.2 192.168.9.255
link link2 /dev/xti/udp - udp 50001  -  192.168.10.2 192.168.10.255
```

  Verify the subnet mask using the ifconfig command to ensure that the two links are on separate subnets.

The link command in the /etc/lltab file

Review the link command information in this section for the /etc/lltab file. See the following information for sample configurations:

- See "Sample configuration: direct-attached links" on page 534.
- See "Sample configuration: links crossing IP routers" on page 535.

Table F-1 describes the fields of the link command that are shown in the /etc/lltab file examples. Note that some of the fields differ from the command for standard LLT links.

Table F-1 Field description for link command in /etc/lltab

Field	Description
tag-name	A unique string that is used as a tag by LLT; for example link1, link2,....
device	The device path of the UDP protocol; for example /dev/xti/udp.
node-range	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
link-type	Type of link; must be "udp" for LLT over UDP.
udp-port	Unique UDP port in the range of 49152-65535 for the link. See "Selecting UDP ports" on page 532.
MTU	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the `lltstat -l` command to display the current value.
IP address	IP address of the link on the local node.
bcast-address	■ For clusters with enabled broadcasts, specify the value of the subnet broadcast address. ■ "-" is the default for clusters spanning routers.

The set-addr command in the /etc/lltab file

The `set-addr` command in the /etc/lltab file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See "Sample configuration: links crossing IP routers" on page 535.

Table F-2 describes the fields of the set-addr command.

Table F-2 Field description for set-addr command in /etc/lltab

Field	Description
node-id	The node ID of the peer node; for example, 0.
link tag-name	The string that LLT uses to identify the link; for example link1, link2,....
address	IP address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
  - Ports from the range of well-known ports, 0 to 1023
  - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -a | more
UDP
   Local Address         Remote Address      State
--------------------  --------------------  ------
      *.*                                   Unbound
      *.32771                               Idle
      *.32776                               Idle
      *.32777                               Idle
      *.name                                Idle
      *.biff                                Idle
      *.talk                                Idle
      *.32779                               Idle
      .
      .
      .
      *.55098                               Idle
      *.syslog                              Idle
      *.58702                               Idle
      *.*                                   Unbound
```

```
# netstat -a |head -2;netstat -a | grep udp
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address Foreign Address (state)
udp4       0      0  *.daytime     *.*
udp4       0      0  *.time        *.*
udp4       0      0  *.sunrpc      *.*
udp4       0      0  *.snmp        *.*
udp4       0      0  *.syslog      *.*
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the /etc/services file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use. Run the following command and answer the prompt to set the netmask:

```
# ifconfig interface_name netmask netmask
```

For example:

- For the first network interface on the node sys1:

  ```
  IP address=192.168.9.1, Broadcast address=192.168.9.255,
  Netmask=255.255.255.0
  ```

  For the first network interface on the node sys2:

  ```
  IP address=192.168.9.2, Broadcast address=192.168.9.255,
  Netmask=255.255.255.0
  ```

- For the second network interface on the node sys1:

  ```
  IP address=192.168.10.1, Broadcast address=192.168.10.255,
  Netmask=255.255.255.0
  ```

  For the second network interface on the node sys2:

  ```
  IP address=192.168.10.2, Broadcast address=192.168.10.255,
  Netmask=255.255.255.0
  ```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in /etc/lltab depending on the subnet that the links are on.

An example of a typical /etc/lltab file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

```
# cat /etc/lltab
set-node nodexyz
set-cluster 100

link link1 /dev/xti/udp - udp 50000 - 192.168.30.1
192.168.30.255
link link2 /dev/xti/udp - udp 50001 - 192.168.31.1
192.168.31.255
```

Sample configuration: direct-attached links

Figure F-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure F-1 A typical configuration of direct-attached links that use LLT over UDP

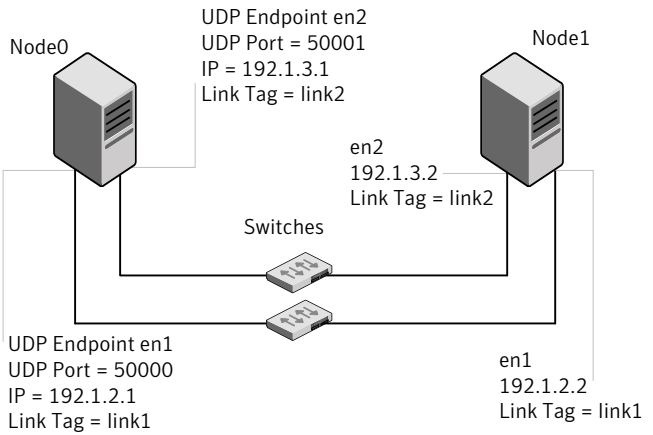

The configuration that the /etc/lltab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT sends broadcast requests to peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the /etc/lltab file using the

`set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the /etc/llttab file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /dev/xti/udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 /dev/xti/udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
# configure Links
# link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /dev/xti/udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 /dev/xti/udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

Figure F-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 F-2 A typical configuration of links crossing an IP router

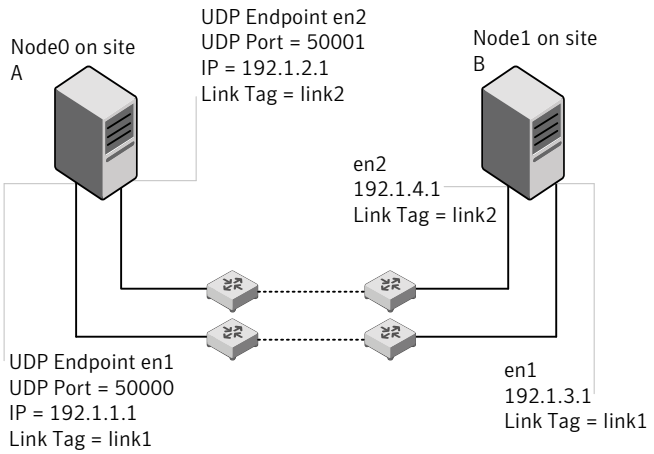

The configuration that the following `/etc/lltab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/lltab` file.

```
set-node Node1
set-cluster 1
link link1 /dev/xti/udp - udp 50000 - 192.1.3.1 -
link link2 /dev/xti/udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr        0 link1 192.1.1.1
set-addr        0 link2 192.1.2.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

The `/etc/lltab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 /dev/xti/udp - udp 50000 - 192.1.1.1 -
link link2 /dev/xti/udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr        1 link1 192.1.3.1
set-addr        1 link2 192.1.4.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

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 538.
- 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 540.

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 539.
- See "Sample configuration: links crossing IP routers" on page 540.

Note that some of the fields in Table F-3 differ from the command for standard LLT links.

Table F-3 describes the fields of the link command that are shown in the /etc/llttab file examples.

Table F-3 Field description for link command in /etc/llttab

Field	Description
tag-name	A unique string that is used as a tag by LLT; for example link1, link2,....
device	The device path of the UDP protocol; for example /dev/xti/udp6.
node-range	Nodes using the link. "-" indicates all cluster nodes are to be configured for this link.
link-type	Type of link; must be "udp6" for LLT over UDP.
udp-port	Unique UDP port in the range of 49152-65535 for the link. See "Selecting UDP ports" on page 538.

Table F-3 Field description for link command in /etc/lltab *(continued)*

Field	Description
MTU	"-" is the default, which has a value of 8192. The value may be increased or decreased depending on the configuration. Use the `lltstat -l` command to display the current value.
IPv6 address	IPv6 address of the link on the local node.
mcast-address	"-" is the default for clusters spanning routers.

The set-addr command in the /etc/lltab file

The `set-addr` command in the /etc/lltab file is required when the multicast feature of LLT is disabled, such as when LLT must cross IP routers.

See "Sample configuration: links crossing IP routers" on page 540.

Table F-4 describes the fields of the set-addr command.

Table F-4 Field description for set-addr command in /etc/lltab

Field	Description
node-id	The ID of the peer node; for example, 0.
link tag-name	The string that LLT uses to identify the link; for example link1, link2,....
address	IPv6 address assigned to the link for the peer node.

Selecting UDP ports

When you select a UDP port, select an available 16-bit integer from the range that follows:

- Use available ports in the private range 49152 to 65535
- Do not use the following ports:
  - Ports from the range of well-known ports, 0 to 1023
  - Ports from the range of registered ports, 1024 to 49151

To check which ports are defined as defaults for a node, examine the file /etc/services. You should also use the `netstat` command to list the UDP ports currently in use. For example:

```
# netstat -a | more
Active Internet connections (including servers)
Proto Recv-Q Send-Q  Local Address   Foreign Address    (state)
tcp        0      0  *.32778         *.*                LISTEN
tcp        0      0  *.32781         *.*                LISTEN
udp4       0      0  *.daytime       *.*
udp4       0      0  *.time          *.*
udp4       0      0  *.sunrpc        *.*
udp        0      0  *.snmp          *.*
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the /etc/services file, its associated name is displayed rather than the port number in the output.

Sample configuration: direct-attached links

Figure F-3 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure F-3 A typical configuration of direct-attached links that use LLT over UDP

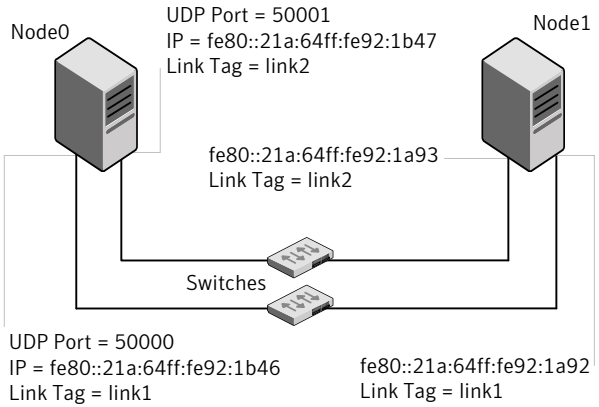

The configuration that the /etc/llttab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT uses IPv6 multicast requests for peer node address discovery. So the addresses of peer nodes do not need to be specified in the /etc/llttab file using the `set-addr` command. Use the `ifconfig -a` command to verify that the IPv6 address is set correctly.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address mcast-address
link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -
```

The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
# configure Links
# link tag-name device node-range  link-type udp port MTU \
IP-address mcast-address
link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -
```

Sample configuration: links crossing IP routers

Figure F-4 depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows two nodes of a four-node cluster.

Figure F-4 A typical configuration of links crossing an IP router

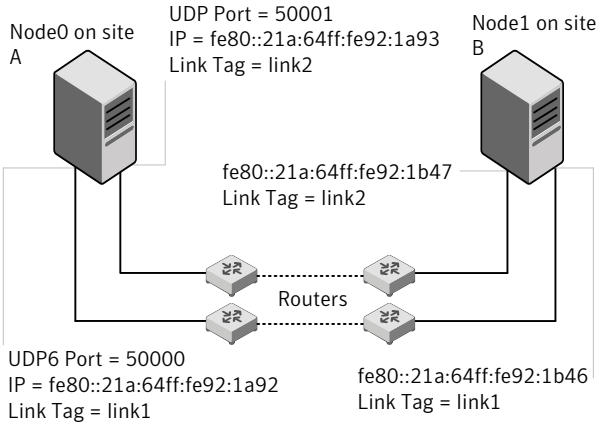

The configuration that the following `/etc/lltab` file represents for Node 1 has links crossing IP routers. Notice that IPv6 addresses are shown for each link on each peer node. In this configuration multicasts are disabled.

```
set-node Node1
set-cluster 1

link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1a92 -
link link1 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1a93 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 0 link1 fe80::21a:64ff:fe92:1b46
set-addr 0 link2 fe80::21a:64ff:fe92:1b47
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb     0
set-arp         0
```

The `/etc/llttab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 /dev/xti/udp6 - udp6 50000 - fe80::21a:64ff:fe92:1b46 -
link link2 /dev/xti/udp6 - udp6 50001 - fe80::21a:64ff:fe92:1b47 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr 1 link1 fe80::21a:64ff:fe92:1a92
set-addr 1 link2 fe80::21a:64ff:fe92:1a93
set-addr 2 link1 fe80::21a:64ff:fe92:1d70
set-addr 2 link2 fe80::21a:64ff:fe92:1d71
set-addr 3 link1 fe80::209:6bff:fe1b:1c94
set-addr 3 link2 fe80::209:6bff:fe1b:1c95

#disable LLT multicasts
set-bcasthb     0
set-arp         0
```

LLT over UDP sample /etc/lltab

The following is a sample of LLT over UDP in the etc/lltab file.

```
set-node sys1
set-cluster clus1
link en1 /dev/xti/udp - udp 50000 - 192.168.10.1 -
link en2 /dev/xti/udp - udp 50001 - 192.168.11.1 -
link-lowpri en0 /dev/xti/udp - udp 50004 - 10.200.58.205 -
set-addr 1 en1 192.168.10.2
set-addr 1 en2 192.168.11.2
set-addr 1 en0 10.200.58.206
set-bcasthb 0
set-arp 0
```

Appendix G

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

Establishing communication between nodes is required to install Symantec software from a remote system, or to install and configure a cluster. The node from which the installer is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities. You need to run the installer with superuser privileges on the systems where you plan to install Symantec software.

You can install products to remote systems using either secure shell (ssh) or remote shell (rsh). Symantec recommends that you use ssh as it is more secure than rsh.

You can set up ssh and rsh connections in many ways.

- You can manually set up the SSH and RSH connection with UNIX shell commands.
- You can run the `installer -comsetup` command to interactively set up SSH and RSH connection.
- You can run the password utility, `pwdutil.pl`.

This section contains an example of how to set up ssh password free communication. The example sets up ssh between a source system (sys1) that contains the installation directories, and a target system (sys2). This procedure also applies to multiple target systems.

Note: The script- and web-based installers support establishing passwordless communication for you.

Manually configuring passwordless ssh

The ssh program enables you to log into and execute commands on a remote system. ssh enables encrypted communications and an authentication process between two untrusted hosts over an insecure network.

In this procedure, you first create a DSA key pair. From the key pair, you append the public key from the source system to the authorized_keys file on the target systems.

Read the ssh documentation and online manual pages before enabling ssh. Contact your operating system support provider for issues regarding ssh configuration.

Visit the OpenSSH website that is located at: http://openssh.org to access online manuals and other resources.

To create the DSA key pair

1 On the source system (sys1), log in as root, and navigate to the root directory.

```
sys1 # cd /
```

2 Make sure the `/.ssh` directory is on all the target installation systems (sys2 in this example). If that directory is not present, create it on all the target systems and set the write permission to root only:

```
sys2 # mkdir /.ssh
```

Change the permissions of this directory, to secure it.

```
sys2 # chmod go-w /.ssh
```

3 To generate a DSA key pair on the source system, type the following command:

```
sys1 # ssh-keygen -t dsa
```

System output similar to the following is displayed:

```
Generating public/private dsa key pair.
Enter file in which to save the key (//.ssh/id_dsa):
```

4 Press Enter to accept the default location of `/.ssh/id_dsa`.

5 When the program asks you to enter the passphrase, press the Enter key twice.

```
Enter passphrase (empty for no passphrase):
```

Do not enter a passphrase. Press Enter.

```
Enter same passphrase again:
```

Press Enter again.

To append the public key from the source system to the authorized_keys file on the target system, using secure file transfer

1. From the source system (sys1), move the public key to a temporary file on the target system (sys2).

   Use the secure file transfer program.

   In this example, the file name `id_dsa.pub` in the root directory is the name for the temporary file for the public key.

   Use the following command for secure file transfer:

   ```
   sys1 # sftp sys2
   ```

   If the secure file transfer is set up for the first time on this system, output similar to the following lines is displayed:

   ```
   Connecting to sys2 ...
   The authenticity of host 'sys2 (10.182.00.00)'
   can't be established. DSA key fingerprint is
   fb:6f:9f:61:91:9d:44:6b:87:86:ef:68:a6:fd:88:7d.
   Are you sure you want to continue connecting (yes/no)?
   ```

2. Enter `yes`.

   Output similar to the following is displayed:

   ```
   Warning: Permanently added 'sys2,10.182.00.00'
   (DSA) to the list of known hosts.
   root@sys2 password:
   ```

3. Enter the root password of sys2.
4. At the `sftp` prompt, type the following command:

   ```
   sftp> put /.ssh/id_dsa.pub
   ```

   The following output is displayed:

   ```
   Uploading /.ssh/id_dsa.pub to /id_dsa.pub
   ```

5. To quit the SFTP session, type the following command:

   ```
   sftp> quit
   ```

6 To begin the ssh session on the target system (sys2 in this example), type the following command on sys1:

```
sys1 # ssh sys2
```

Enter the root password of sys2 at the prompt:

```
password:
```

7 After you log in to sys2, enter the following command to append the id_dsa.pub file to the authorized_keys file:

```
sys2 # cat /id_dsa.pub >> /.ssh/authorized_keys
```

8 After the id_dsa.pub public key file is copied to the target system (sys2), and added to the authorized keys file, delete it. To delete the id_dsa.pub public key file, enter the following command on sys2:

```
sys2 # rm /id_dsa.pub
```

9 To log out of the ssh session, enter the following command:

```
sys2 # exit
```

10 Run the following commands on the source installation system. If your ssh session has expired or terminated, you can also run these commands to renew the session. These commands bring the private key into the shell environment and make the key globally available to the user root:

```
sys1 # exec /usr/bin/ssh-agent $SHELL
sys1 # ssh-add

  Identity added: //.ssh/id_dsa
```

This shell-specific step is valid only while the shell is active. You must execute the procedure again if you close the shell during the session.

To verify that you can connect to a target system

1 On the source system (sys1), enter the following command:

```
sys1 # ssh -l root sys2 uname -a
```

where sys2 is the name of the target system.

2 The command should execute from the source system (sys1) to the target system (sys2) without the system requesting a passphrase or password.

3 Repeat this procedure for each target system.

Setting up ssh and rsh connection using the installer -comsetup command

You can interactively set up the ssh and rsh connections using the `installer -comsetup` command.

Enter the following:

```
# ./installer -comsetup

Input the name of the systems to set up communication:
Enter the Solaris 10 Sparc system names separated by spaces:
[q,?] sys2
Set up communication for the system sys2:

  Checking communication on sys2 ..................... Failed

CPI ERROR V-9-20-1303 ssh permission was denied on sys2. rsh
permission was denied on sys2. Either ssh or rsh is required
to be set up and ensure that it is working properly between the local
node and sys2 for communication

Either ssh or rsh needs to be set up between the local system and
sys2 for communication

Would you like the installer to setup ssh or rsh communication
automatically between the systems?
Superuser passwords for the systems will be asked. [y,n,q,?] (y) y

Enter the superuser password for system sys2:

     1)  Setup ssh between the systems
```

```
     2)  Setup rsh between the systems
     b)  Back to previous menu

Select the communication method [1-2,b,q,?] (1) 1

Setting up communication between systems. Please wait.
Re-verifying systems.

 Checking communication on sys2 ..................... Done

Successfully set up communication for the system sys2
```

Setting up ssh and rsh connection using the pwdutil.pl utility

The password utility, `pwdutil.pl`, is bundled in the 6.2 release under the `scripts` directory. The users can run the utility in their script to set up the ssh and rsh connection automatically.

```
# ./pwdutil.pl -h
Usage:

Command syntax with simple format:

    pwdutil.pl check|configure|unconfigure ssh|rsh <hostname|IP addr>
    [<user>] [<password>] [<port>]

Command syntax with advanced format:

    pwdutil.pl [--action|-a 'check|configure|unconfigure']
               [--type|-t 'ssh|rsh']
               [--user|-u  '<user>']
               [--password|-p '<password>']
               [--port|-P '<port>']
               [--hostfile|-f '<hostfile>']
               [--keyfile|-k '<keyfile>']
               [-debug|-d]
               <host_URI>

    pwdutil.pl -h | -?
```

Table G-1 Options with pwdutil.pl utility

Option	Usage
--action|-a 'check|configure|unconfigure'	Specifies action type, default is 'check'.
--type|-t 'ssh|rsh'	Specifies connection type, default is 'ssh'.
--user|-u '<user>'	Specifies user id, default is the local user id.
--password|-p '<password>'	Specifies user password, default is the user id.
--port|-P '<port>'	Specifies port number for ssh connection, default is 22
--keyfile|-k '<keyfile>'	Specifies the private key file.
--hostfile|-f '<hostfile>'	Specifies the file which list the hosts.
-debug	Prints debug information.
-h|-?	Prints help messages.
<host_URI>	Can be in the following formats: <hostname> <user>:<password>@<hostname> <user>:<password>@<hostname>: <port>

You can check, configure, and unconfigure ssh or rsh using the `pwdutil.pl`utility. For example:

- To check ssh connection for only one host:

  ```
  pwdutil.pl check ssh hostname
  ```

- To configure ssh for only one host:

  ```
  pwdutil.pl configure ssh hostname user password
  ```

- To unconfigure rsh for only one host:

  ```
  pwdutil.pl unconfigure rsh hostname
  ```

- To configure ssh for multiple hosts with same user ID and password:

```
pwdutil.pl -a configure -t ssh -u user -p password hostname1
hostname2 hostname3
```

- To configure ssh or rsh for different hosts with different user ID and password:

```
pwdutil.pl -a configure -t ssh user1:password1@hostname1
user2:password2@hostname2
```

- To check or configure ssh or rsh for multiple hosts with one configuration file:

```
pwdutil.pl -a configure -t ssh --hostfile /tmp/sshrsh_hostfile
```

- To keep the host configuration file secret, you can use the 3rd party utility to encrypt and decrypt the host file with password.
  For example:

```
### run openssl to encrypt the host file in base64 format
# openssl aes-256-cbc -a -salt -in /hostfile -out /hostfile.enc
enter aes-256-cbc encryption password: <password>
Verifying - enter aes-256-cbc encryption password: <password>

### remove the original plain text file
# rm /hostfile

### run openssl to decrypt the encrypted host file
# pwdutil.pl -a configure -t ssh `openssl aes-256-cbc -d -a
-in /hostfile.enc`
enter aes-256-cbc decryption password: <password>
```

- To use the ssh authentication keys which are not under the default $HOME/.ssh directory, you can use --keyfile option to specify the ssh keys. For example:

```
### create a directory to host the key pairs:
# mkdir /keystore

### generate private and public key pair under the directory:
# ssh-keygen -t rsa -f /keystore/id_rsa

### setup ssh connection with the new generated key pair under
the directory:
# pwdutil.pl -a configure -t ssh --keyfile /keystore/id_rsa
user:password@hostname
```

You can see the contents of the configuration file by using the following command:

```
# cat /tmp/sshrsh_hostfile
user1:password1@hostname1
user2:password2@hostname2
user3:password3@hostname3
user4:password4@hostname4

# all default: check ssh connection with local user
hostname5
The following exit values are returned:

0     Successful completion.
1     Command syntax error.
2     Ssh or rsh binaries do not exist.
3     Ssh or rsh service is down on the remote machine.
4     Ssh or rsh command execution is denied due to password is required.
5     Invalid password is provided.
255   Other unknown error.
```

Restarting the ssh session

After you complete this procedure, ssh can be restarted in any of the following scenarios:

- After a terminal session is closed
- After a new terminal session is opened
- After a system is restarted
- After too much time has elapsed, to refresh ssh

To restart ssh

1. On the source installation system (sys1), bring the private key into the shell environment.

   ```
   sys1 # exec /usr/bin/ssh-agent $SHELL
   ```

2. Make the key globally available for the user `root`

   ```
   sys1 # ssh-add
   ```

Enabling rsh for AIX

To enable `rsh`, create a `/.rhosts` file on each target system. Then add a line to the file specifying the full domain name of the source system. For example, add the line:

sysname.*domainname*.com root

Change permissions on the `/.rhosts` file to 600 by typing the following command:

```
# chmod 600 /.rhosts
```

After you complete an installation procedure, delete the `.rhosts` file from each target system to ensure security:

```
# rm -f /.rhosts
```

Appendix H

Troubleshooting VCS installation

This appendix includes the following topics:

What to do if you see a licensing reminder

In this release, you can install without a license key. In order to comply with the End User License Agreement, you must either install a license key or make the host managed by a Management Server. If you do not comply with these terms within 60 days, the following warning messages result:

```
WARNING V-365-1-1 This host is not entitled to run Symantec Storage
Foundation/Symantec Cluster Server.As set forth in the End User
License Agreement (EULA) you must complete one of the two options
set forth below. To comply with this condition of the EULA and
```

```
stop logging of this message, you have <nn> days to either:
- make this host managed by a Management Server (see
  http://go.symantec.com/sfhakeyless for details and free download),
  or
- add a valid license key matching the functionality in use on this host
  using the command 'vxlicinst' and validate using the command
'vxkeyless set NONE'.
```

To comply with the terms of the EULA, and remove these messages, you must do one of the following within 60 days:

- Install a valid license key corresponding to the functionality in use on the host. After you install the license key, you must validate the license key using the following command:

  ```
  # /opt/VRTS/bin/vxlicrep
  ```

- Continue with keyless licensing by managing the server or cluster with a management server.
  For more information about keyless licensing, see the following URL:
  http://go.symantec.com/sfhakeyless

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 Symantec products

After the installation and configuration is complete, the Symantec product installer starts the processes that the installed products use. You can use the product installer to stop or start the processes, if required.

To stop the processes

- Use the `-stop` option to stop the product installation script.

  For example, to stop the product's processes, enter the following command:

  ```
  # ./installer -stop
  ```

  or

  ```
  # /opt/VRTS/install/installvcs<version> -stop
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 50.

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/installvcs<version> -start
  ```

  Where `<version>` is the specific release version.

  See "About the script-based installer" on page 50.

Installer cannot create UUID for the cluster

The installer displays the following error message if the installer cannot find the uuidconfig.pl script before it configures the UUID for the cluster:

```
Couldn't find uuidconfig.pl for uuid configuration,
please create uuid manually before start vcs
```

You may see the error message during VCS configuration, upgrade, or when you add a node to the cluster using the installer.

Workaround: To start VCS, you must run the uuidconfig.pl script manually to configure the UUID on each cluster node.

To configure the cluster UUID when you create a cluster manually

- On one node in the cluster, perform the following command to populate the cluster UUID on each node in the cluster.

  ```
  # /opt/VRTSvcs/bin/uuidconfig.pl -clus -configure nodeA
  nodeB ... nodeN
  ```

  Where nodeA, nodeB, through nodeN are the names of the cluster nodes.

LLT startup script displays errors

If more than one system on the network has the same clusterid-nodeid pair and the same Ethernet sap/UDP port, then the LLT startup script displays error messages similar to the following:

```
LLT lltconfig ERROR V-14-2-15238 node 1 already exists
in cluster 8383 and has the address - 00:18:8B:E4:DE:27
LLT lltconfig ERROR V-14-2-15241 LLT not configured,
use -o to override this warning
LLT lltconfig ERROR V-14-2-15664 LLT could not
configure any link
LLT lltconfig ERROR V-14-2-15245 cluster id 1 is
already being used by nid 0 and has the
address - 00:04:23:AC:24:2D
LLT lltconfig ERROR V-14-2-15664 LLT could not
configure any link
```

Recommended action: Ensure that all systems on the network have unique clusterid-nodeid pair. You can use the `lltdump -f device -D` command to get the list of unique clusterid-nodeid pairs connected to the network. This utility is available only for LLT-over-ethernet.

The vxfentsthdw utility fails for Active/Passive arrays when you test disks in raw format

DMP opens the secondary (passive) paths with an exclusive flag in Active/Passive arrays. So, if you test the secondary (passive) raw paths of the disk, the vxfentsthdw command fails due to DMP's exclusive flag.

Recommended action: For Active/Passive arrays when you want to test the disks in raw format, you must use an active enabled path with the vxfentsthdw command. Run the `vxdmpadm getsubpaths dmpnodename=enclosure-based_name` command to list the active enabled paths.

The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails

While running the vxfentsthdw utility, you may see a message that resembles as follows:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node
FAILED.
Contact the storage provider to have the hardware configuration
fixed.
```

The disk array does not support returning success for a `SCSI TEST UNIT READY` command when another host has the disk reserved using SCSI-3 persistent reservations. This happens with the Hitachi Data Systems 99XX arrays if bit 186 of the system mode option is not enabled.

Issues during fencing startup on VCS cluster nodes set up for server-based fencing

Table H-1 Fencing startup issues on VCS cluster (client cluster) nodes

Issue	Description and resolution
`cpsadm` command on the VCS cluster gives connection error	If you receive a connection error message after issuing the `cpsadm` command on the VCS cluster, perform the following actions: ■ Ensure that the CP server is reachable from all the VCS cluster nodes. ■ Check the /etc/vxfenmode file and ensure that the VCS cluster nodes use the correct CP server virtual IP or virtual hostname and the correct port number. ■ For HTTPS communication, ensure that the virtual IP and ports listed for the server can listen to HTTPS requests.
Authorization failure	Authorization failure occurs when the nodes on the client clusters and or users are not added in the CP server configuration. Therefore, fencing on the VCS cluster (client cluster) node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points. To resolve this issue, add the client cluster node and user in the CP server configuration and restart fencing. See "Preparing the CP servers manually for use by the VCS cluster" on page 277.

Table H-1 Fencing startup issues on VCS cluster (client cluster) nodes *(continued)*

Issue	Description and resolution
Authentication failure	If you had configured secure communication between the CP server and the VCS cluster (client cluster) nodes, authentication failure can occur due to the following causes: ■ The client cluster requires its own private key, a signed certificate, and a Certification Authority's (CA) certificate to establish secure communication with the CP server. If any of the files are missing or corrupt, communication fails. ■ If the client cluster certificate does not correspond to the client's private key, communication fails. ■ If the CP server and client cluster do not have a common CA in their certificate chain of trust, then communication fails.

Appendix I

Sample VCS 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

The following CP server configuration diagrams can be used as guides when setting up CP server within your configuration:

- Two unique client clusters that are served by 3 CP servers:
  See Figure I-1 on page 561.
- Client cluster that is served by highly available CP server and 2 SCSI-3 disks:
- Two node campus cluster that is served be remote CP server and 2 SCSI-3 disks:
- Multiple client clusters that are served by highly available CP server and 2 SCSI-3 disks:

Two unique client clusters served by 3 CP servers

Figure I-1 displays a configuration where two unique client clusters are being served by 3 CP servers (coordination points). Each client cluster has its own unique user ID (UUID1 and UUID2).

In the `vxfenmode` file on the client nodes, vxfenmode is set to `customized` with vxfen mechanism set to `cps`.

Figure I-1 Two unique client clusters served by 3 CP servers

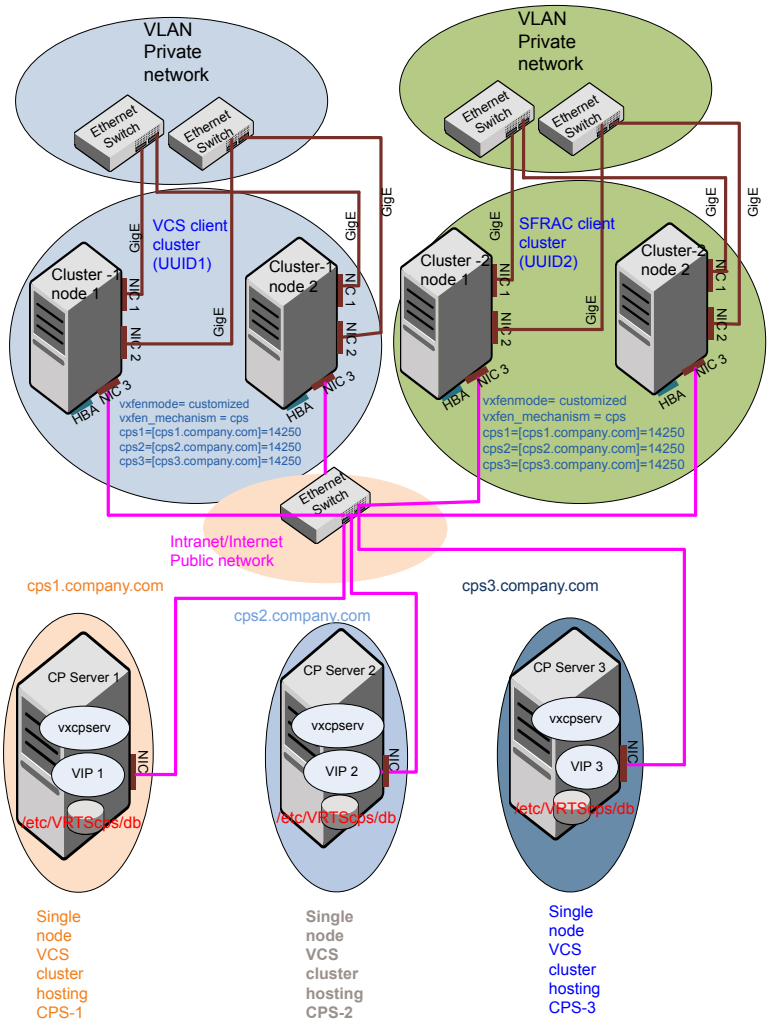

Client cluster served by highly available CPS and 2 SCSI-3 disks

Figure I-2 displays a configuration where a client cluster is served by one highly available CP server and 2 local SCSI-3 LUNs (disks).

In the `vxfenmode` file on the client nodes, `vxfenmode` is set to customized with vxfen mechanism set to `cps`.

The two SCSI-3 disks are part of the disk group vxfencoorddg. The third coordination point is a CP server hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure I-2 Client cluster served by highly available CP server and 2 SCSI-3 disks

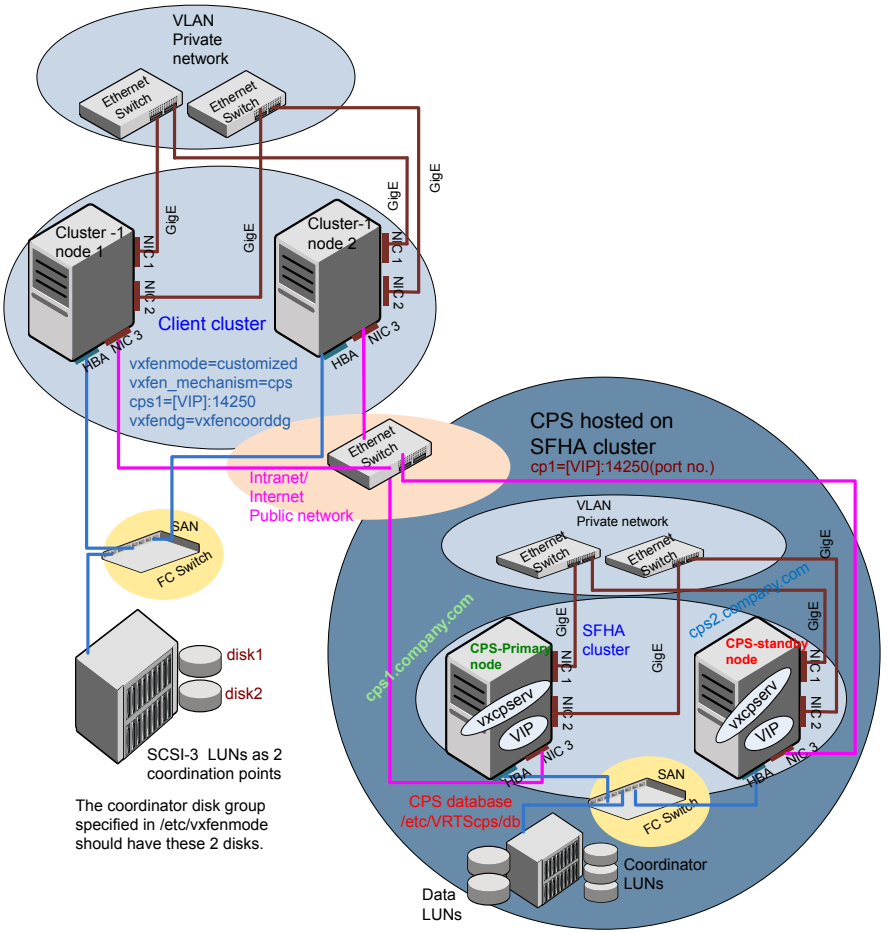

Two node campus cluster served by remote CP server and 2 SCSI-3 disks

Figure I-3 displays a configuration where a two node campus cluster is being served by one remote CP server and 2 local SCSI-3 LUN (disks).

In the `vxfenmode` file on the client nodes, vxfenmode is set to `customized` with vxfen mechanism set to `cps`.

The two SCSI-3 disks (one from each site) are part of disk group vxfencoorddg. The third coordination point is a CP server on a single node VCS cluster.

Figure I-3 Two node campus cluster served by remote CP server and 2 SCSI-3

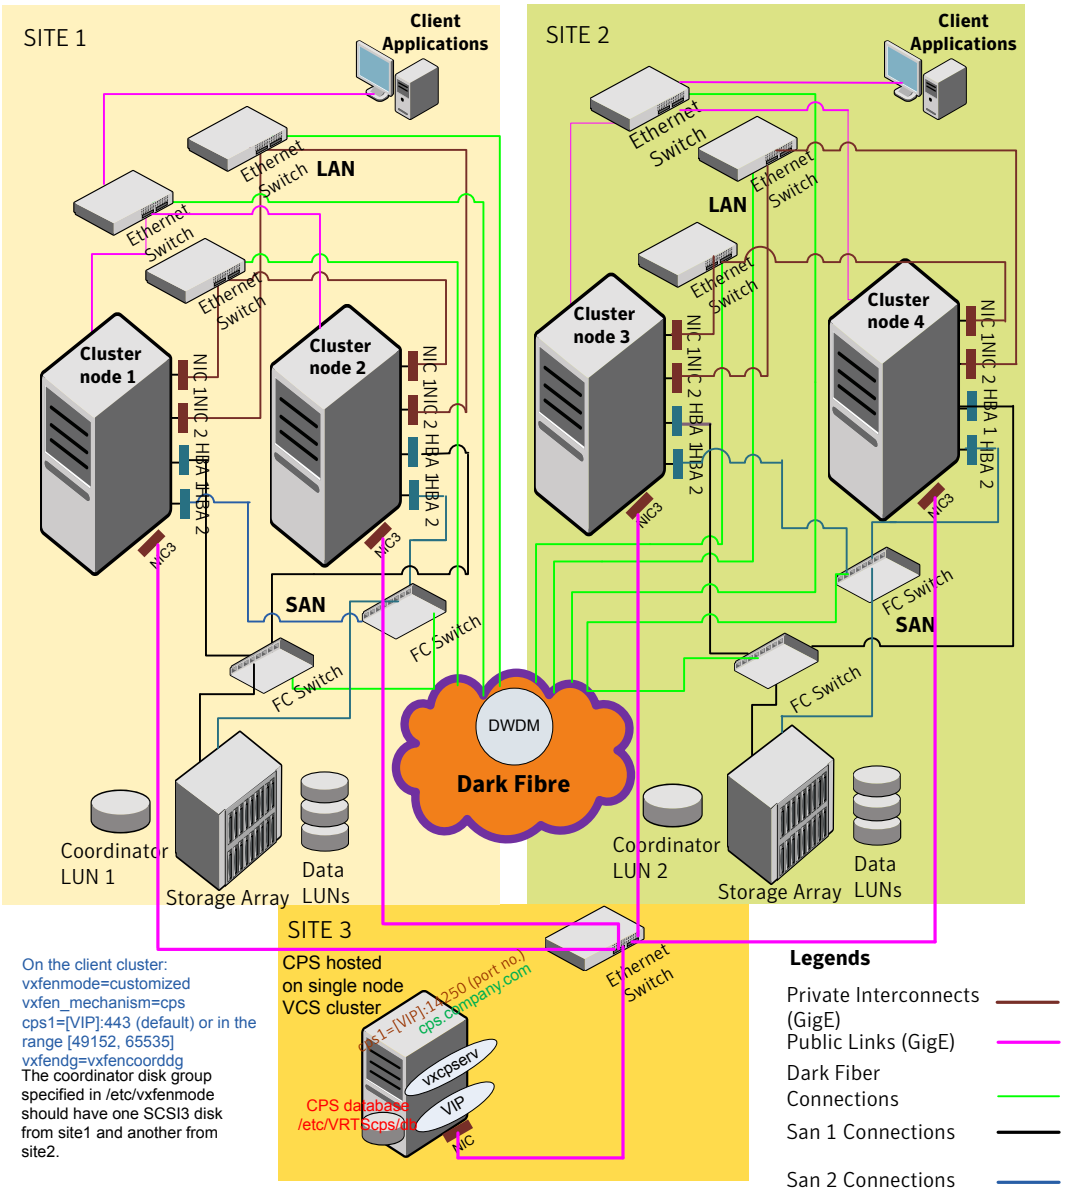

Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

Figure I-4 displays a configuration where multiple client clusters are being served by one highly available CP server and 2 local SCSI-3 LUNS (disks).

In the `vxfenmode` file on the client nodes, vxfenmode is set to `customized` with vxfen mechanism set to `cps`.

The two SCSI-3 disks are are part of the disk group vxfencoorddg. The third coordination point is a CP server, hosted on an SFHA cluster, with its own shared database and coordinator disks.

Figure I-4 Multiple client clusters served by highly available CP server and 2 SCSI-3 disks

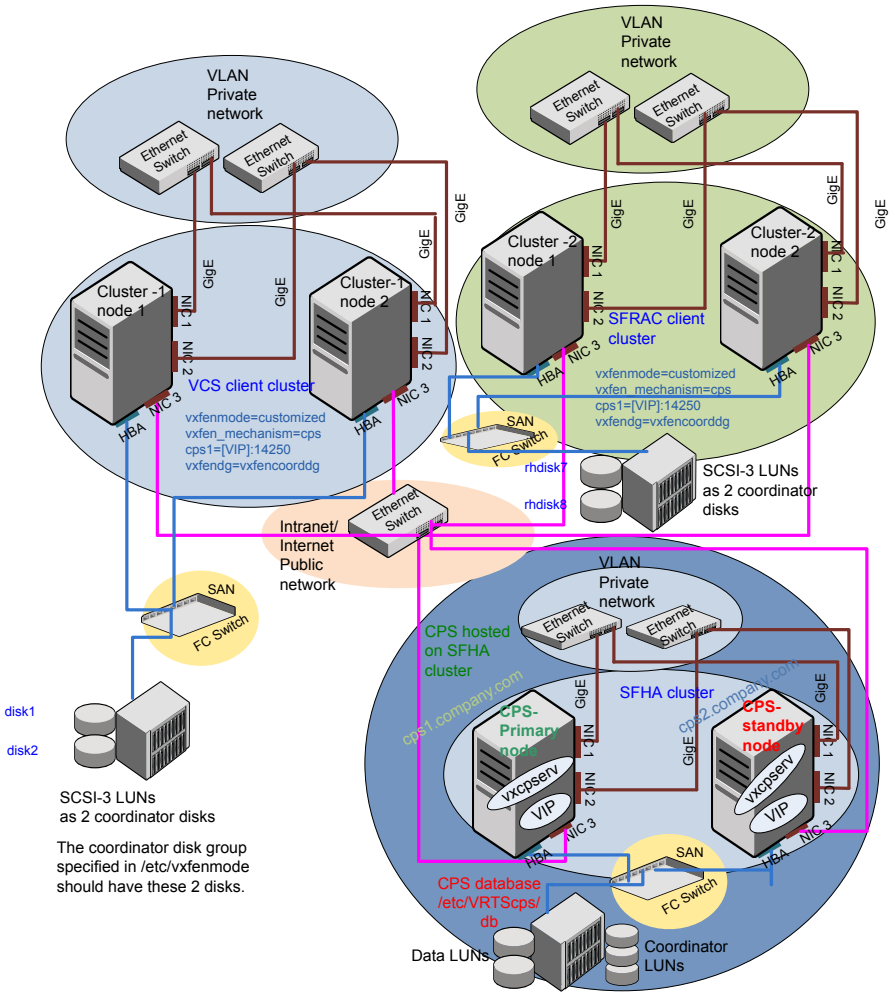

Appendix J

Changing NFS server major numbers for VxVM volumes

This appendix includes the following topics:

Changing NFS server major numbers for VxVM volumes

In a VCS cluster, block devices providing NFS service must have the same major and minor numbers on each cluster node. Major numbers identify required device drivers (such as AIX partition or VxVM volume). Minor numbers identify the specific devices themselves. NFS also uses major and minor numbers to identify the exported file system. Major and minor numbers must be verified to ensure that the NFS identity for the file system is the same when exported from each node.

Use the `haremajor` command to determine and reassign the major number that a system uses for shared VxVM volume block devices. For Veritas Volume Manager, the major number is set to the `vxio` driver number. To be highly available, each NFS server in a VCS cluster must have the same `vxio` driver number, or major number.

To list the major number currently in use on a system

- Use the command:

```
# haremajor -v
55
```

Run this command on each cluster node. If major numbers are not the same on each node, you must change them on the nodes so that they are identical.

To list the available major numbers for a system

- Use the command:

```
# haremajor -a
54,56..58,60,62..
```

The output shows the numbers that are not in use on the system where the command is issued.

To reset the major number on a system

- You can reset the major number to an available number on a system. For example, to set the major number to 75 type:

```
# haremajor -s 75
```

Appendix K

Compatibility issues when installing Symantec Cluster Server with other products

This appendix includes the following topics:

Installing, uninstalling, or upgrading Storage Foundation products when other Symantec products are present

Installing Storage Foundation when other Symantec products are installed can create compatibility issues. For example, installing Storage Foundation products when VOM, ApplicationHA, and NetBackup are present on the systems.

Installing, uninstalling, or upgrading Storage Foundation products when VOM is already present

If you plan to install or upgrade Storage Foundation products on systems where VOM has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where VOM Central Server is present, the installer skips the VRTSsfmh upgrade and leaves the VOM Central Server and Managed Host filesets as is.
- When uninstalling Storage Foundation products where VOM Central Server is present, the installer does not uninstall VRTSsfmh.
- When you install or upgrade Storage Foundation products where VOM Managed Host is present, the installer gives warning messages that it will upgrade VRTSsfmh.

Installing, uninstalling, or upgrading Storage Foundation products when NetBackup is already present

If you plan to install or upgrade Storage Foundation on systems where NetBackup has already been installed, be aware of the following compatibility issues:

- When you install or upgrade Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpbx and VRTSicsco. It does not upgrade VRTSat.
- When you uninstall Storage Foundation products where NetBackup is present, the installer does not uninstall VRTSpbx, VRTSicsco, and VRTSat.

Appendix L

Upgrading the Steward process

This appendix includes the following topics:

Upgrading the Steward process

The Steward process can be configured in both secure and non-secure mode. The following procedures provide the steps to upgrade the Steward process.

Upgrading Steward configured in secure mode from 6.1 to 6.2

To upgrade Steward configured in secure mode:

1. Log on to the Steward system as a root user.
2. Stop the Steward process.

   ```
   # steward -stop -secure
   ```

3. Uninstall the VRTSvcs and VRTSperl filesets.
4. Install the VRTSvcs and VRTSperl filesets.
5. Start the Steward process.

   ```
   # steward -start -secure
   ```

Upgrading Steward configured in non-secure mode from 6.1 to 6.2

To upgrade Steward configured in non-secure mode:

1. Log on to the Steward system as a root user.
2. Stop the Steward process.

   ```
   # steward -stop
   ```

3. Copy and replace the Steward binary from a node in the VCS 6.2 cluster to the Steward system. The file resides in the `/opt/VRTSvcs/bin/` directory. Make sure that the source from where you copy the binary is also running the same version of AIX as the destination system.
4. Start the Steward process.

   ```
   # steward -start
   ```

Refer to *About the Steward process: Split-brain in two-cluster global clusters* in the *Symantec Cluster Server Administrator's Guide* for more information.

Index

D

E

F

G

H

S

T

U

V

W