Veritas™ Cluster Server Installation Guide

Solaris

5.0 Maintenance Pack 3

Veritas 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: 5.0 MP3

Document version: 5.0MP3.0

Legal Notice

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014

http://www.symantec.com

Technical Support

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

Symantec's maintenance offerings include the following:

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

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

www.symantec.com/techsupp/

Contacting Technical Support

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

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

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

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

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

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

Licensing and registration

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

www.symantec.com/techsupp/

Customer service

Customer service information is available at the following URL:

www.symantec.com/techsupp/

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

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

Documentation feedback

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

Maintenance agreement resources

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

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

Additional enterprise services

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

Enterprise services that are available include the following:

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

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

www.symantec.com

Select your country or language from the site index.

Contents

Chapter 1

Introducing Veritas Cluster Server

This chapter includes the following topics:

About Veritas Cluster Server

Veritas™ Cluster Server by Symantec is a high-availability solution for cluster configurations. Veritas 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 shared storage devices. When a system is part of a VCS cluster, it is 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

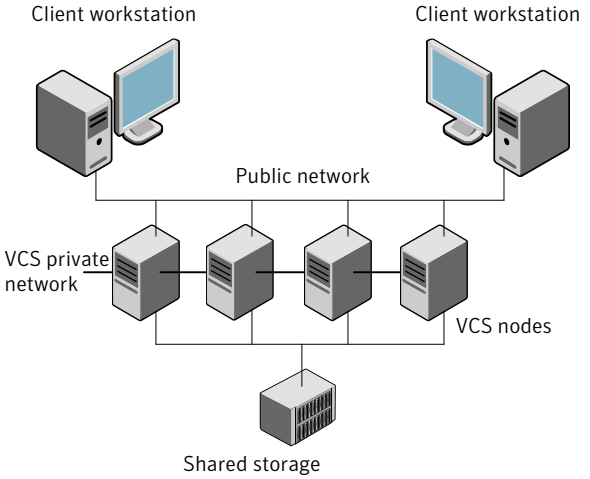

Client workstations receive service over the public network from applications running on VCS nodes. VCS monitors the nodes and their services. VCS nodes in the cluster communicate over a private network.

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, the nodes that join or leaving 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

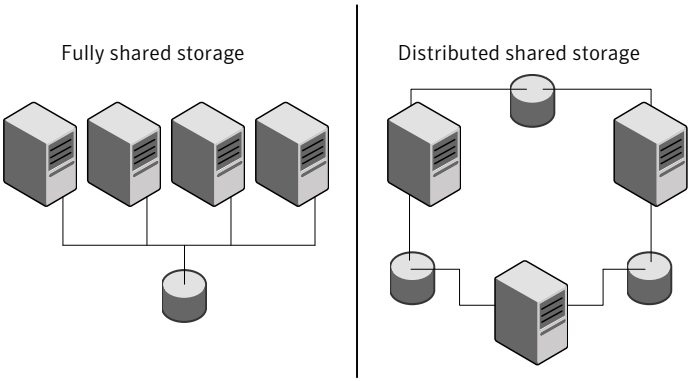

About LLT and GAB

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.

LLT configuration files are as follows:

- /etc/llthosts—lists all the nodes in the cluster
- /etc/llttab file—describes the local system's private network links to the other nodes in the cluster

GAB (Group Membership and Atomic Broadcast) provides the global message order that is required to maintain a synchronized state among the nodes. It monitors disk communications such as the VCS heartbeat utility. The /etc/gabtab file is the GAB configuration file.

See "About the LLT and GAB configuration files" on page 127.

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 Solaris 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 *Veritas Cluster Server User's Guide*.

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

Figure 1-3 Two Ethernet connections connecting two nodes

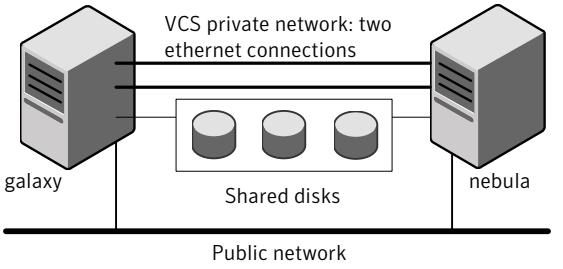

About preexisting network partitions

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

About VCS seeding

To protect your cluster from a preexisting network partition, VCS uses a seed. A seed 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.

About VCS optional features

You can use the Veritas Installation Assessment Service to assess your setup for VCS installation.

See "Veritas Installation Assessment Service" on page 17.

To configure the optional features of the VCS components, make sure to install all packages when the installation program prompts you. Review the description of the optional features and decide the features that you want to configure with VCS:

VCS notifications	See "About VCS notifications" on page 17.
VCS global clusters	See "About global clusters" on page 17.
I/O fencing	See "About I/O fencing" on page 18.

Veritas Installation Assessment Service

The Veritas Installation Assessment Service (VIAS) utility assists you in getting ready for a Veritas Storage Foundation and High Availability Solutions installation or upgrade. The VIAS utility allows the preinstallation evaluation of a configuration, to validate it prior to starting an installation or upgrade.

https://vias.symantec.com/

About VCS notifications

You can configure both SNMP and SMTP notifications for VCS. Symantec recommends you to configure 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 *Veritas Cluster Server User's Guide*.

About global clusters

Global clusters provide the ability to fail over applications between geographically distributed clusters when disaster occurs. You require a separate license to configure global clusters. You must add this license during the installation. The

installer only asks about configuring global clusters if you have used the global cluster license.

See *Veritas Cluster Server User'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.

See *Veritas Cluster Server User's Guide*.

The fencing operation determines the following:

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

This decision prevents possible data corruption. The installvcs program installs the VCS I/O fencing driver, VRTSvxfen. If you want to protect data on shared disks, you must configure I/O fencing after you install and configure VCS.

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

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

About VCS optional components

You can add the following optional components to VCS:

Symantec Product Authentication Service	See "About Symantec Product Authentication Service (AT)" on page 19.
Veritas Cluster Server Management Console	See "About Veritas Cluster Server Management Console" on page 20.
Cluster Manager (Java console)	See "About Cluster Manager (Java Console)" on page 20.

Figure 1-4 illustrates a sample VCS deployment with the optional components configured.

Figure 1-4 Typical VCS setup with optional components

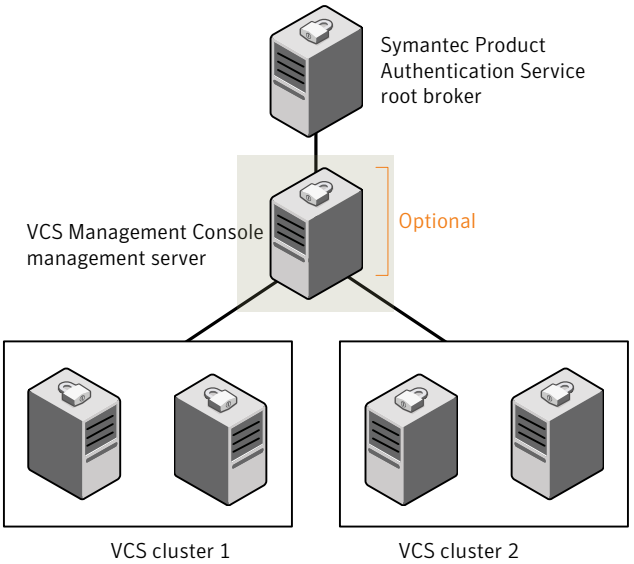

About Symantec Product Authentication Service (AT)

VCS uses Symantec Product Authentication Service (AT) to provide secure communication between cluster nodes and clients. It uses digital certificates for authentication and SSL to encrypt communication over the public network to secure communications.

AT uses the following brokers to establish trust relationship between the cluster components:

- Root broker
  A root broker serves as the main registration and certification authority; it has a self-signed certificate and can authenticate other brokers. The root broker is only used during initial creation of an authentication broker.
  A root broker can serve multiple clusters. Symantec recommends that you install a single root broker on a utility system. The utility system, such as an email server or domain controller, can be highly available.
- Authentication brokers
  Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have root-signed certificates. Each node in VCS serves as an authentication broker.

See Symantec Product Authentication Service documentation for more information.

See "Preparing to configure the clusters in secure mode" on page 25.

About Veritas Cluster Server Management Console

Veritas Cluster Server Management Console is a management interface that enables you to monitor and administer clusters from a Web console.

Veritas Cluster Server Management Console is a high availability management solution that enables monitoring and administering clusters from a single Web console.

You can configure Veritas Cluster Server Management Console to manage a single cluster, multiple clusters, or both.

See *Veritas Cluster Server Management Console Implementation Guide*.

About Cluster Manager (Java Console)

Cluster Manager (Java Console) offers complete administration capabilities for your cluster. Use the different views in the Java Console to monitor clusters and VCS objects, including service groups, systems, resources, and resource types. You can perform many administrative operations using the Java Console. You can also perform these operations using the command line interface or using the Veritas Cluster Server Management Console.

See "Installing the Java Console" on page 82.

See *Veritas Cluster Server User's Guide*.

Chapter 2

Planning to install VCS

This chapter includes the following topics:

About planning to install VCS

Every node where you want to install VCS must meet the hardware and software requirements.

For the latest information on updates, patches, and software issues, read the following Veritas Technical Support TechNote:

http://entsupport.symantec.com/docs/281987

http://entsupport.symantec.com/docs/286955

To find information on supported hardware, see the hardware compatibility list (HCL) in the following TechNote:

For Solaris SPARC:

http://entsupport.symantec.com/docs/283282

For Solaris x64:

http://entsupport.symantec.com/docs/283161

Hardware requirements

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 SPARC or x64 systems running Solaris 8 or later as appropriate.
DVD drive	One drive in a system that can communicate to all the nodes in the cluster.
Disks	Typical VCS configurations require that shared disks support the applications that migrate between systems in the cluster. The VCS I/O fencing feature requires that all data and coordinator disks support SCSI-3 Persistent Reservations (PR). See "About setting up I/O fencing" on page 113.
Disk space	To run VCS, LLT, GAB, the Web Console, and the Java Console, each VCS node 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 If you do not have enough free space in /var, then use the `installvcs` command with `tmppath` option. Make sure that the specified tmppath file system has the required free space. ■ 10 MB in / **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.
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 256 megabytes.

Supported operating systems

- On SPARC: Solaris 8, 9, and 10 (32-bit and 64-bit)
- On x64: Solaris 10 (64-bit)

Required patches

For each Solaris platform and/or architecture, Symantec recommends that you apply the latest Solaris operating system patches available from Sun. See the following site:

http://sunsolve.sun.com

Make sure you install the following patches for VCS:

- On Solaris x64: To use I/O Fencing, you must install the Solaris 10 patches 119716-10 and 118844-27.
- On systems running Solaris 10, VCS requires NFS Patch 118833-24 to enable NFS agents to function properly.
- If you are running the MultiNICB agent on Solaris 9, you must have the following patch from Sun: 116670-04.
- If you are running the MultiNICB agent on Solaris 8, you must be at a minimum level of Solaris 8 update 2.
- If you are using VCS with non-global zones and want the zone root on shared storage, use Solaris 10 Update 3 or later.

Supported software

VCS supports the following volume managers and files systems:

- Veritas Volume Manager (VxVM) with Veritas File System (VxFS)
  - VxVM 4.0 with VxFS 4.0
    (Solaris SPARC 8 and 9 only)
  - VxVM 4.1 with VxFS 4.1
  - VxVM 5.0 with VxFS 5.0
  - VxVM 5.0 MP1 with VxFS 5.0 MP1
    (Solaris SPARC only)
  - VxVM 5.0 MP3 with VxFS 5.0 MP3

Chapter 3

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.

See "About planning to install VCS" on page 21.

Preparing to configure the clusters in secure mode

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

Refer to the *Veritas Cluster Server User's Guide* for instructions to configure AT in a cluster that does not run in secure mode.

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

- A system in your enterprise is configured as root broker (RB).
  If a root broker system does not exist, install and configure root broker on a system.
  See "Installing the root broker for the security infrastructure" on page 29.
- An authentication broker (AB) account for each node in the cluster is set up on the root broker system.
  See "Creating authentication broker accounts on root broker system" on page 30.

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

The installvcs program provides the following configuration modes:

Automatic mode	The root broker system must allow rsh or ssh passwordless login to use this mode.
Semi-automatic mode	This mode requires encrypted files (BLOB files) from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See "Setting up inter-system communication" on page 38.
Manual mode	This mode requires root_hash file and the root broker information from the AT administrator to configure a cluster in secure mode. The nodes in the cluster must allow rsh or ssh passwordless login. See "Setting up inter-system communication" on page 38.

Figure 3-1 depicts the flow of configuring VCS cluster in secure mode.

Figure 3-1 Workflow to configure VCS cluster in secure mode

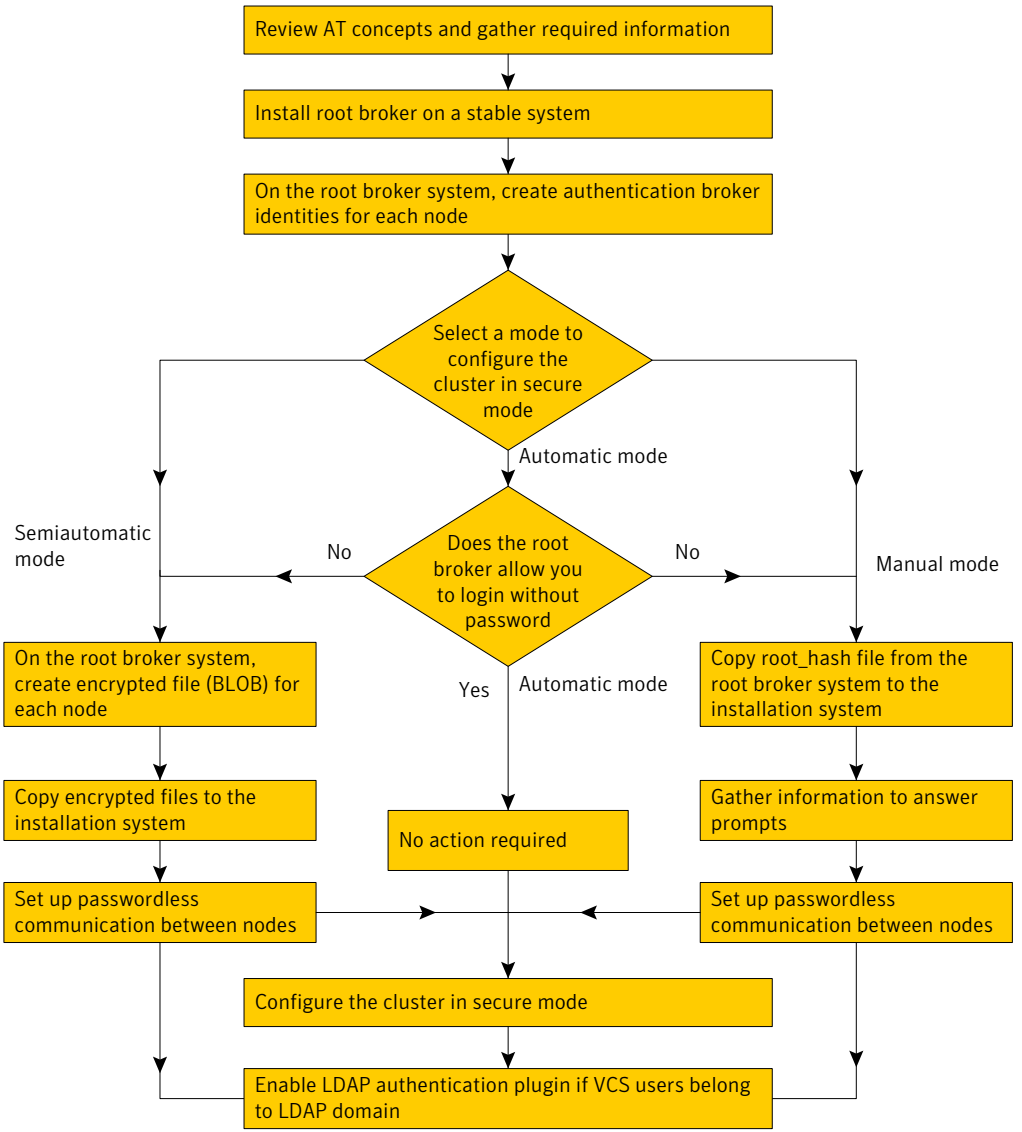

Table 3-1 lists the preparatory tasks in the order which the AT and VCS administrators must perform.

Table 3-1 Preparatory tasks to configure a cluster in secure mode

Tasks	Who performs this task
Decide one of the following configuration modes to set up a cluster in secure mode: ■ Automatic mode ■ Semi-automatic mode ■ Manual mode	VCS administrator
Install the root broker on a stable system in the enterprise. See "Installing the root broker for the security infrastructure" on page 29.	AT administrator
On the root broker system, create authentication broker accounts for each node in the cluster. See "Creating authentication broker accounts on root broker system" on page 30. AT administrator requires the following information from the VCS administrator: ■ Node names that are designated to serve as authentication brokers ■ Password for each authentication broker	AT administrator
To use the semi-automatic mode, create the encrypted files (BLOB files) for each node and provide the files to the VCS administrator. See "Creating encrypted files for the security infrastructure" on page 31. AT administrator requires the following additional information from the VCS administrator: ■ Administrator password for each authentication broker Typically, the password is the same for all nodes.	AT administrator
To use the manual mode, provide the root_hash file (/opt/VRTSat/bin/root_hash) from the root broker system to the VCS administrator.	AT administrator
Copy the files that are required to configure a cluster in secure mode to the system from where you plan to install and configure VCS. See "Preparing the installation system for the security infrastructure" on page 33.	VCS administrator

Installing the root broker for the security infrastructure

Install the root broker only if you plan to use AT to configure the cluster in secure mode. The root broker administrator must install and configure the root broker before you configure the Authentication Service for VCS. Symantec recommends that you install the root broker on a stable system that is outside the cluster. You can install the root broker on an AIX, HP-UX, Linux, or Solaris system.

See Symantec Product Authentication Service documentation for more information.

See "About Symantec Product Authentication Service (AT)" on page 19.

To install the root broker

1. Change to the directory where you can start the installvcs program:

   ```
   # cd cluster_server
   ```

2. Start the Root Broker installation program:

   ```
   # ./installvcs -security
   ```

3. Select to install the Root Broker from the three choices that the installer presents:

   ```
   3 Install Symantec Security Services Root Broker
   ```

4. Enter the name of the system where you want to install the Root Broker.

   ```
   Enter the system name on which to install VxSS: venus
   ```

5. Review the output as the installer does the following:
   - Checks to make sure that VCS supports the operating system
   - Verifies that you install from the global zone (only on Solaris)
   - Checks if the system is already configured for security
6. Review the output as the installvcs program checks for the installed packages on the system.

   The installvcs program lists the packages that the program is about to install on the system. Press Enter to continue.
7. Review the output as the installer installs the root broker on the system.
8. Enter **y** when the installer prompts you to configure the Symantec Product Authentication Service.

9 Enter a password for the root broker. Make sure the password contains a minimum of five characters.

10 Enter a password for the authentication broker. Make sure the password contains a minimum of five characters.

11 Press the Enter key to start the Authentication Server processes.

```
Do you want to start Symantec Product Authentication Service
processes now? [y,n,q] y
```

12 Review the output as the installer starts the Authentication Service.

Creating authentication broker accounts on root broker system

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

To create authentication broker accounts on root broker system

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

```
venus> # vssat showalltrustedcreds
```

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

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

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

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

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

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

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

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

Proceed to step 3.

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

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

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

Creating encrypted files for the security infrastructure

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

To create encrypted files

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

hash	The value of the root hash string, which consists of 40 characters. Execute the following command to find this value: `venus> # vssat showbrokerhash`
root_domain	The value for the domain name of the root broker system. Execute the following command to find this value: `venus> # vssat showalltrustedcreds`

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

identity	The value for the authentication broker identity, which you provided to create authentication broker principal on the root broker system. This is the value for the **--prplname** option of the `addprpl` command. See "Creating authentication broker accounts on root broker system" on page 30.
password	The value for the authentication broker password, which you provided to create authentication broker principal on the root broker system. This is the value for the **--password** option of the `addprpl` command. See "Creating authentication broker accounts on root broker system" on page 30.
broker_admin_password	The value for the authentication broker password for Administrator account on the node. This password must be at least five characters.

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

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

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

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

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

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

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

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

For example:

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

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

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

Preparing the installation system for the security infrastructure

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

To prepare the installation system for the security infrastructure

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

Automatic mode	Do the following: ■ Gather the root broker system name from the AT administrator. ■ During VCS configuration, choose the configuration option 1 when the installvcs program prompts.
Semi-automatic mode	Do the following: ■ Copy the encrypted files (BLOB files) to the system from where you plan to install VCS. Note the path of these files that you copied to the installation system. ■ During VCS configuration, choose the configuration option 2 when the installvcs program prompts.

Manual mode

Do the following:

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

Performing preinstallation tasks

Table 3-2 lists the tasks you must perform before proceeding to install VCS.

Table 3-2 Preinstallation tasks

Task	Reference
Obtain license keys.	See "Obtaining VCS license keys" on page 35.
Set up the private network.	See "Setting up the private network" on page 36.
Enable communication between systems.	See "Setting up inter-system communication" on page 38.
Set up ssh on cluster systems.	See "Setting up ssh on cluster systems" on page 39.
Set up shared storage for I/O fencing (optional)	See "Setting up shared storage" on page 42.
Set the PATH and the MANPATH variables.	See "Setting the PATH variable" on page 46. See "Setting the MANPATH variable" on page 46.
Disable the abort sequence on SPARC systems.	See "Disabling the abort sequence on SPARC systems" on page 46.

Table 3-2 Preinstallation tasks *(continued)*

Task	Reference
Review basic instructions to optimize LLT media speeds.	See "Optimizing LLT media speed settings on private NICs" on page 48.
Review guidelines to help you set the LLT interconnects.	See "Guidelines for setting the media speed of the LLT interconnects" on page 48.
Install the patches that are required for Java Run Time environment from Sun.	
Prepare zone environments	See "Preparing zone environments" on page 48.
Mount the product disc	See "Mounting the product disc" on page 49.
Verify the systems before installation	See "Performing automated pre-installation check" on page 49.

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:

https://licensing.symantec.com

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 Help link at this site to access the *License Portal User Guide* and FAQ.

The VRTSvlic package 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 currently installed licenses
vxlictest	Retrieves the features and their descriptions that are encoded in a license key

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.

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. However, Sun systems assign the same MAC address to all interfaces by default. Thus, connecting two or more interfaces to a network switch can cause problems.

For example, consider the following case where:

- The IP address is configured on one interface and LLT on another
- Both interfaces are connected to a switch (assume separate VLANs)

The duplicate MAC address on the two switch ports can cause the switch to incorrectly redirect IP traffic to the LLT interface and vice-versa. To avoid this issue, configure the system to assign unique MAC addresses by setting the `eeprom(1M)` parameter local-mac-address to true.

The following products make extensive use of the private cluster interconnects for distributed locking:

- Veritas Storage Foundation Cluster File System (CFS)
- Veritas Storage Foundation for Oracle RAC (SF Oracle RAC)

Symantec recommends network switches for the CFS and the SF Oracle RAC clusters due to their performance characteristics.

Refer to the *Veritas Cluster Server User's Guide* to review VCS performance considerations.

Figure 3-2 shows two private networks for use with VCS.

Figure 3-2 Private network setups: two-node and four-node clusters

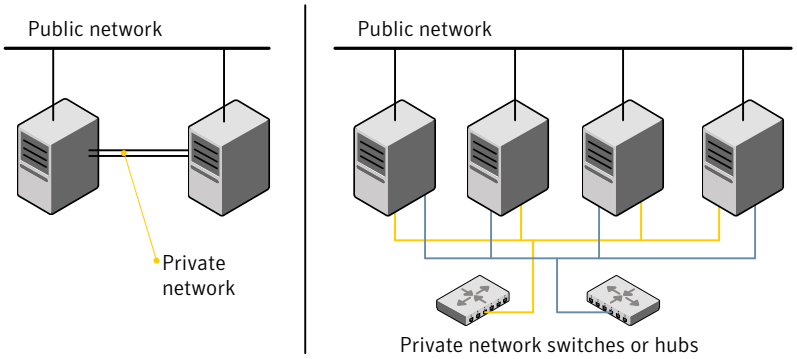

To set up the private network

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.
- The network interface card to set up private interface is not part of any aggregated interface.

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 Configure the Ethernet devices that are used for the private network such that the autonegotiation protocol is not used. You can achieve a more stable configuration with crossover cables if the auto-negotiation protocol is not used.

To achieve this stable configuration, do one of the following:

- Edit the /etc/system file to disable autonegotiation on all Ethernet devices system-wide.
- Create a qfe.conf or bge.conf file in the /kernel/drv directory to disable autonegotiation for the individual devices that are used for private network.

Refer to the Sun Ethernet driver product documentation for information on these methods.

5 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 installvcs program configures the private network in the cluster during installation.

See "About installing and configuring VCS" on page 51.

More information about configuring LLT for the private network links is in the manual installation chapter.

See "About VCS manual installation" on page 87.

Setting up inter-system communication

When you install VCS using the installvcs program, to install and configure the entire cluster at one time, make sure that communication between systems exists. By default the installer uses ssh. You must grant root privileges for the system where you run installvcs program. This privilege facilitates to issue ssh or rsh commands on all systems in the cluster. If ssh is used to communicate between systems, it must be configured in a way such that it operates without requests for passwords or passphrases. Similarly, rsh must be configured in such a way to not prompt for passwords.

If system communication is not possible between systems using ssh or rsh, you have recourse.

See "Installing VCS with a response file where ssh or rsh are disabled" on page 241.

See "About VCS manual installation" on page 87.

Setting up ssh on cluster systems

Use the Secure Shell (ssh) to install VCS on all systems in a cluster from a system outside of the cluster. Before you start the installation process, verify that ssh is configured correctly.

Use Secure Shell (ssh) to do the following:

- Log on to another system over a network
- Execute commands on a remote system
- Copy files from one system to another

The ssh shell provides strong authentication and secure communications over channels. It is intended to replace rlogin, rsh, and rcp.

Configuring ssh

The procedure to configure ssh uses OpenSSH example file names and commands.

Note: You can configure ssh in other ways. Regardless of how ssh is configured, complete the last step in the example to verify the configuration.

To configure ssh

1. Log in as root on the source system from which you want to install the Veritas product.
2. To generate a DSA key pair on the source system, type the following:

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

   System output similar to the following is displayed:

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

3. Press **Enter** to accept the default location of /.ssh/id_dsa. System output similar to the following is displayed:

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

4. Do not enter a passphrase. Press **Enter**. Enter same passphrase again:

   Press **Enter** again.

5 Make sure the /.ssh directory is on all the target installation systems. If that directory is absent, create it on the target system and set the write permission to root only:

```
# mkdir /.ssh
# chmod go-w /
# chmod 700 /.ssh
# chmod go-rwx /.ssh
```

6 Make sure the secure file transfer program (SFTP) is enabled on all the target installation systems. To enable SFTP, the /etc/ssh/sshd_config file must contain the following two lines:

```
PermitRootLogin yes
Subsystem sftp /usr/lib/ssh/sftp-server
```

7 If the lines are not there, add them and restart SSH. To restart SSH on Solaris 10, type the following command:

```
# svcadm restart ssh
```

To restart on Solaris 9, type the following commands:

```
# /etc/init.d/sshd stop
# /etc/init.d/sshd start
```

8 To copy the public DSA key, /.ssh/id_dsa.pub to each target system, type the following commands:

```
# sftp target_sys
```

If you run this step for the first time on a system, output similar to the following appears:

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

9 Enter **yes**. Output similar to the following is displayed:

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

10 Enter the root password.

11 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
```

12 To quit the SFTP session, type the following command:

```
sftp> quit
```

13 To begin the ssh session on the target system, type the following command:

```
# ssh target_sys
```

14 Enter the root password at the prompt:

```
password:
```

15 After you log in, enter the following command to append the authorization key to the id_dsa.pub file:

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

16 Delete the id_dsa.pub public key file. Before you delete this public key file, make sure to complete the following tasks:

- The file is copied to the target (host) system
- The file is added to the authorized keys file

To delete the id_dsa.pub public key file, type the following command:

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

17 To log out of the ssh session, type the following command:

```
# exit
```

18 When you install from a source system that is also an installation target, add the local system id_dsa.pub key to the local /.ssh/authorized_key file. The installation can fail if the installation source system is not authenticated.

19 Run the following commands on the source installation system. These commands bring the private key into the shell environment and makes the key globally available for the user root:

```
# exec /usr/bin/ssh-agent $SHELL
# ssh-add
Identity added: /.ssh/identity
```

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

20 To verify that you can connect to the target system, type the following command:

```
# ssh -l root target_sys uname -a
```

The commands should execute on the remote system without any requests for a passphrase or password from the system.

Setting up shared storage

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

See "About setting up I/O fencing" on page 113.

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

Setting up shared storage: SCSI disks

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

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

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

To set up shared storage

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

   Refer to the documentation that is shipped with the host adapters, the storage, and the systems.

2. With both nodes powered off, power on the storage devices.

3. Power on one system, but do not allow it to boot. If necessary, halt the system so that you can use the ok prompt.

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

4. Find the paths to the host adapters:

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

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

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

   ```
   {0} ok nvedit
   ```

   As you edit the script, note the following points:

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

   In this example, edit the nvramrc script as follows:

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

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

```
{0} ok nvstore
```

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

```
{0} ok printenv nvramrc
```

You can re-edit the file to make corrections:

```
{0} ok nvedit
```

Or, discard the changes if necessary by entering:

```
{0} ok nvquit
```

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

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

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

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

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

Permit the system to continue booting.

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

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

Permit the system to continue booting.

Setting up shared storage: Fiber channel

Perform the following steps to set up fiber channel.

To set up shared storage

1 Install the required FC-AL controllers.

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

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

3 Boot each system with the reconfigure devices option:

```
ok boot -r
```

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

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

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

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

Setting the PATH variable

Installation commands as well as other commands reside in the /sbin, /usr/sbin, /opt/VRTS/bin, and /opt/VRTSvcs/bin directories. Add these directories to your PATH environment variable.

To set the PATH variable

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

```
$ PATH=/usr/sbin:/sbin:/opt/VRTS/bin:/opt/VRTSvcs/bin: \
$PATH; export PATH
```

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

```
% setenv PATH /usr/sbin:/sbin:/opt/VRTS/bin: \
/opt/VRTSvcs/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 or ksh), type:

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

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

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

Disabling the abort sequence on SPARC systems

Most UNIX operating systems provide a method to perform a "break" or "console abort." The inherent problem when you abort a hung system is that it ceases to

heartbeat in the cluster. When other cluster members believe that the aborted node is a failed node, these cluster members may begin corrective action.

Keep the following points in mind:

- The only action that you must perform following a system abort is to reset the system to achieve the following:
  - Preserve data integrity
  - Prevent the cluster from taking additional corrective actions
- Do not resume the processor as cluster membership may have changed and failover actions may already be in progress.
- To remove this potential problem on Sun SPARC systems, you should alias the `go` function in the OpenBoot eeprom to display a message.

To alias the go function to display a message

1. At the ok prompt, enter:

   ```
   nvedit
   ```

2. Press Ctrl+L to display the current contents of the nvramrc buffer.
3. Press Ctrl+N until the editor displays the last line of the buffer.
4. Add the following lines exactly as shown. Press Return after adding each line.

   ```
   ." Aliasing the OpenBoot 'go' command! "
   : go ." It is inadvisable to use the 'go' command in a clustered
   environment. " cr
   ." Please use the 'power-off' or 'reset-all' commands instead. "
   cr
   ." Thank you, from your friendly neighborhood sysadmin. " ;
   ```

5. Press Ctrl+C to exit the nvramrc editor.
6. To verify that no errors exist, type the `nvrun` command. You should see only the following text:

   ```
   Aliasing the OpenBoot 'go' command!
   ```

7. Type the `nvstore` command to commit your changes to the non-volatile RAM (NVRAM) for use in subsequent reboots.
8. After you perform these commands, at reboot you see this output:

   ```
   Aliasing the OpenBoot 'go' command! go isn't unique.
   ```

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.

Guidelines for setting the media speed of the LLT interconnects

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

- If you have hubs or switches for LLT interconnects, Symantec recommends using the Auto_Negotiation media speed setting on each Ethernet card on each node.
- If you have hubs or switches for LLT interconnects and you do not use the Auto_Negotiation media speed setting, then do the following:
  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), set the media speed to the highest value common to both cards, typically 100_Full_Duplex.
- Symantec does not recommend using dissimilar network cards for private links.

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

Preparing zone environments

You need to keep the following items in mind when you install or upgrade VCS in a zone environment.

- When you install or upgrade VCS using the installer program, all zones are upgraded (both global and non-global) unless they are detached and unmounted.
- If you install VCS on Solaris 10 systems that run non-global zones, you need to make sure that non-global zones do not inherit the /opt directory. Run the following command to make sure that the /opt directory is not in the inherit-pkg-dir clause:

```
# zonecfg -z zone_name info
zonepath: /export/home/zone1
autoboot: false
```

```
pool: yourpool
inherit-pkg-dir:
dir: /lib
inherit-pkg-dir:
dir: /platform
inherit-pkg-dir:
dir: /sbin
inherit-pkg-dir:
dir: /usr
```

If the /opt directory appears in the output, remove the /opt directory from the zone's configuration and reinstall the zone.

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 need not be part of the cluster. The systems must be in the same subnet.

2. Insert the product disc into a DVD drive that is connected to your system.

3. If Solaris volume management software is running on your system, the software disc automatically mounts as /cdrom/cdrom0.

4. If Solaris volume management software is not available to mount the DVD, you must mount it manually. After you insert the software disc, enter:

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

   Where c0t6d0s2 is the default address for the disc drive.

Performing automated pre-installation 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 pre-installation check is:

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

You can also use the Veritas Installation Assessment Service utility for a detailed assessment of your setup.

See "Veritas Installation Assessment Service" on page 17.

To check the systems

1 Navigate to the folder that contains the installvcs program.

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

2 Start the pre-installation check:

```
# ./installvcs -precheck galaxy nebula
```

The program proceeds in a noninteractive mode to examine the systems for licenses, packages, 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.

See "About installvcs program command options" on page 57.

Chapter 4

Installing and configuring VCS

This chapter includes the following topics:

About installing and configuring VCS

You can install Veritas Cluster Server on clusters of up to 32 systems. You can install VCS using one of the following:

Veritas product installer	Use the product installer to install multiple Veritas products.
installvcs program	Use this to install just VCS.

The Veritas product installer and the installvcs program use ssh to install by default. Refer to the *Getting Started Guide* for more information.

Getting your VCS installation and configuration information ready

The VCS installation and configuration program prompts you for information about certain VCS components.

When you perform the installation, prepare the following information:

- To install VCS packages you need:

The system names where you plan to install VCS	Example: **galaxy**, **nebula**
The required license keys	Depending on the type of installation, keys 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 35.
To decide whether to install: ■ the required VCS packages ■ all the VCS packages	Install only the required packages if you do not want to configure any optional components or features. The default option is to install all packages. See "Optional VCS packages" on page 55.

- To configure Veritas Cluster Server you need:

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: **vcs_cluster27**
A unique ID number for the cluster	A number in the range of 0-65535. Within the site that contains the cluster, each cluster must have a unique ID. Example: **7**

The device names of the NICs that the private networks use among systems	A network interface card that is not part of any aggregated interface, or an aggregated interface. Do not use the network interface card that is used for the public network, which is typically hme0 for SPARC and bge0 for x64. For example on a SPARC system: **qfe0, qfe1** For example on an x64 system: **e1000g0, e1000g1**

- To configure VCS clusters in secure mode (optional), you need:

For automatic mode (default)	■ The name of the Root Broker system Example: `east` See "About Symantec Product Authentication Service (AT)" on page 19. ■ Access to the Root Broker system without use of a password.
For semiautomatic mode using encrypted files	The path for the encrypted files that you get from the Root Broker administrator. See "Creating encrypted files for the security infrastructure" on page 31.
For semiautomatic mode without using encrypted files	■ The fully-qualified hostname (FQDN) of the Root Broker . (e.g. east.symantecexample.com) The given example puts a system in the (DNS) domain symantecexample.com with the unqualified hostname east, which is designated as the Root Broker. ■ The root broker's security domain (e.g. **root@east.symantecexample.com**) ■ The root broker's port (e.g. **2821**) ■ The path to the local root hash (e.g. **/var/tmp/privatedir/root_hash**) ■ The authentication broker's principal name on each cluster node (e.g. **galaxy.symantecexample.com** and **nebula.symantecexample.com**)

- To add VCS users, which is not required if you configure your cluster in secure mode, you need:

User names	Example: **smith**

User passwords	Enter the password at the prompt.
To decide user privileges	Users have three levels of privileges: A=Administrator, O=Operator, or G=Guest. Example: **A**

- To configure SMTP email notification (optional), you need:

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: I=Information, W=Warning, E=Error, and S=SevereError. Example: **E**

- To configure SNMP trap notification (optional), you need:

The port number for the SNMP trap daemon	The default port number is 162.
The system name for each SNMP console	Example: **saturn**
To decide the minimum severity of events for SNMP trap notification	Events have four levels of severity: I=Information, W=Warning, E=Error, and S=SevereError. Example: **E**

- To configure global clusters (optional), you need:

The name of the public NIC	You must specify appropriate values for the NIC. For example for SPARC systems: **hme0** For example for x64 systems: **bge0**
The virtual IP address of the NIC	You must specify appropriate values for the virtual IP address. Example: **10.10.12.1**

The netmask for the virtual IP address	You must specify appropriate values for the netmask. Example: `255.255.240.0`

Optional VCS packages

The optional VCS packages include the following packages:

- VRTScmccc – Veritas Cluster Management Console Cluster Connector
- VRTScmcs – Veritas Cluster Management Console for Single Cluster Mode
- VRTScssim – VCS Simulator
- VRTScscm – Veritas Cluster Server Cluster Manager
- VRTSvcsmn – Manual pages for VCS commands

About the VCS installation program

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

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

- Licensing VCS
- Installing VCS packages 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
- The Symantec Product Authentication Services feature
- The wide area Global Cluster feature

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

See "About preparing to install VCS" on page 25.

The uninstallvcs program, a companion to installvcs program, uninstalls VCS packages.

See "About the uninstallvcs program" on page 199.

Optional features of the installvcs program

Table 4-1 specifies the optional actions that the installvcs program can perform.

Table 4-1 installvcs optional features

Optional action	Reference
Check the systems to verify that they meet the requirements to install VCS.	See "Performing automated pre-installation check" on page 49.
Upgrade VCS to version 5.0 MP3 if VCS currently runs on a cluster.	See "Upgrading to VCS 5.0 MP3" on page 156.
Install VCS packages without configuring VCS.	See "Installing VCS using installonly option" on page 60.
Configure or reconfigure VCS when VCS packages are already installed.	See "Configuring VCS using configure option" on page 60.
Perform secure installations using the values that are stored in a configuration file.	See "Installing VCS with a response file where ssh or rsh are disabled" on page 241.
Perform automated installations using the values that are stored in a configuration file.	See "Performing automated VCS installations" on page 233.

Interacting with the installvcs program

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

See "About the uninstallvcs program" on page 199.

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 program does not install the VCS Java Console.

See "Installing the Java Console" on page 82.

About installvcs program command options

In addition to the `-precheck`, `-responsefile`, `-installonly`, and `-configure` options, the installvcs program has other useful options.

The `installvcs` command usage takes the following form:

```
installvcs [ system1 system2... ] [ options ]
```

Table 4-2 lists the installvcs command options.

Table 4-2 installvcs options

Option and Syntax	Description
`-configure`	Configure VCS after using `-installonly` option to install VCS. See "Configuring VCS using configure option" on page 60.
`-enckeyfile` *encryption_key_file*	See the `-responsefile` and the `-encrypt` options.
`-encrypt` *password*	Encrypt password using the encryption key that is provided with the `-enckeyfile` option so that the encrypted password can be stored in response files.
`-hostfile`	Specifies the location of a file that contains the system names for the installer.
`-installonly`	Install product packages on systems without configuring VCS. See "Installing VCS using installonly option" on page 60.
`-installpkgs`	Display VCS packages in correct installation order. Output can be used to create scripts for command line installs, or for installations over a network. See the `requiredpkgs` option.

Table 4-2 installvcs options *(continued)*

Option and Syntax	Description
`-jumpstart`	Use this option to generate finish scripts that the Solaris JumpStart Server can use for Veritas products. Use complete paths when you specify the available locations to store the finish scripts.
`-keyfile` *`ssh_key_file`*	Specifies a key file for SSH. The option passes `-i` *`ssh_key_file`* with each SSH invocation.
`-license`	Register or update product licenses on the specified systems. Useful for replacing demo license.
`-logpath` *`log_path`*	Specifies that log_path, not /opt/VRTS/install/logs, is the location where installvcs log files, summary file, and response file are saved.
`-noextrapkgs`	Specifies that additional product packages such as VxVM and VxFS need not be installed. **Note:** VCS product upgrades in the future can be simplified if you do not install additional product packages.
`-nolic`	Install product packages on systems without licensing or configuration. License-based features or variants are not installed when using this option.
`-nooptionalpkgs`	Specifies that the optional product packages such as man pages and documentation need not be installed.
`-nostart`	Bypass starting VCS after completing installation and configuration.
`-patchpath` *`patch_path`*	Specifies that patch_path contains all patches that the installvcs program is about to install on all systems. The patch_path is the complete path of a directory. **Note:** You can use this option when you download recent versions of patches.
`-pkgpath` *`pkg_path`*	Specifies that pkg_path contains all packages that the installvcs program is about to install on all systems. The pkg_path is the complete path of a directory, usually NFS mounted.

Table 4-2 installvcs options *(continued)*

Option and Syntax	Description
`-precheck`	Verify that systems meet the installation requirements before proceeding with VCS installation. Symantec recommends doing a precheck before installing VCS. See "Performing automated pre-installation check" on page 49.
`-requiredpkgs`	Displays all required VCS packages in correct installation order. Optional packages are not listed. Output can be used to create scripts for command line installs, or for installations over a network. See `installpkgs` option.
`-responsefile` *`response_file`* [`-enckeyfile` *`encryption_key_file`*]	Perform automated VCS installation using the system and the configuration information that is stored in a specified file instead of prompting for information. The response_file must be a full path name. If not specified, the response file is automatically generated as installerernumber.response where number is random. You must edit the response file to use it for subsequent installations. Variable field definitions are defined within the file. The `-enckeyfile` option and encryption_key_file name are required with the `-responsefile` option when the response file contains encrypted passwords. See "Installing VCS with a response file where ssh or rsh are disabled" on page 241. See "Performing automated VCS installations" on page 233.
`-rootpath` *`root_path`*	Specifies that root_path is the root location for the installation of all packages. On Solaris, `-rootpath` passes `-I` root_path to pkgadd command.
`-rsh`	Specifies that *rsh* and `rcp` are to be used for communication between systems instead of ssh and `scp`. This option requires that systems be preconfigured such that *rsh* commands between systems execute without prompting for passwords or confirmations

Table 4-2 installvcs options *(continued)*

Option and Syntax	Description
`-security`	Enable or disable Symantec Product Authentication Service in a VCS cluster that is running. Install and configure Root Broker for Symantec Product Authentication Service. See "About Symantec Product Authentication Service (AT)" on page 19.
`-serial`	Performs the installation, uninstallation, start, and stop operations on the systems in a serial fashion. By default, the installer performs these operations simultaneously on all the systems.
`-timeout`	Specifies the timeout value (in seconds) for each command that the installer issues during the installation. The default timeout value is set to 600 seconds.
`-tmppath` *tmp_path*	Specifies that tmp_path is the working directory for installvcs program. This path is different from the /var/tmp path. This destination is where initial logging is performed and where packages are copied on remote systems before installation.
`-verbose`	Displays the details when the installer installs the packages. By default, the installer displays only a progress bar during the packages installation.

Installing VCS using installonly option

In certain situations, users may choose to install the VCS packages on a system before they are ready for cluster configuration. During such situations, the `installvcs -installonly` option can be used. The installation program licenses and installs VCS packages on the systems that you enter without creating any VCS configuration files.

Configuring VCS using configure option

If you installed VCS and did not choose to configure VCS immediately, use the `installvcs -configure` option. You can configure VCS when you are ready for cluster configuration. The installvcs program prompts for cluster information, and creates VCS configuration files without performing installation.

See "Configuring the basic cluster" on page 69.

The `-configure` option can be used to reconfigure a VCS cluster. VCS must not be running on systems when this reconfiguration is performed.

If you manually edited the main.cf file, you need to reformat the main.cf file.

See "Reformatting VCS configuration files on a stopped cluster" on page 61.

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 *Veritas Cluster Server User'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
```

Installing and configuring VCS 5.0 MP3

The example installation demonstrates how to install VCS on two systems: galaxy and nebula. The example installation chooses to install all VCS packages and configures all optional features. For this example, the cluster's name is vcs_cluster2 and the cluster's ID is 7.

Figure 4-1 illustrates the Solaris SPARC systems on which you would install and run VCS.

Figure 4-1 An example of a VCS installation on a two-node cluster

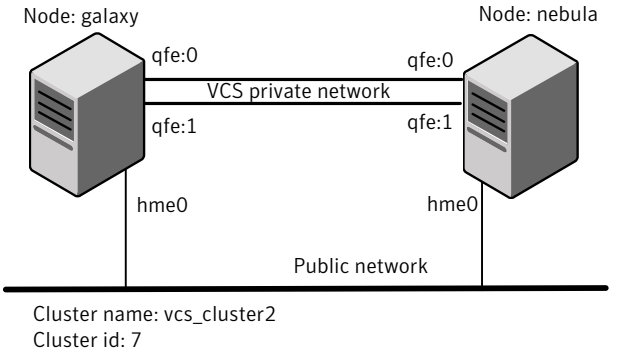

Figure 4-2 illustrates the Solaris x64 systems on which you would install and run VCS.

Figure 4-2 An example of a VCS installation on a two-node cluster

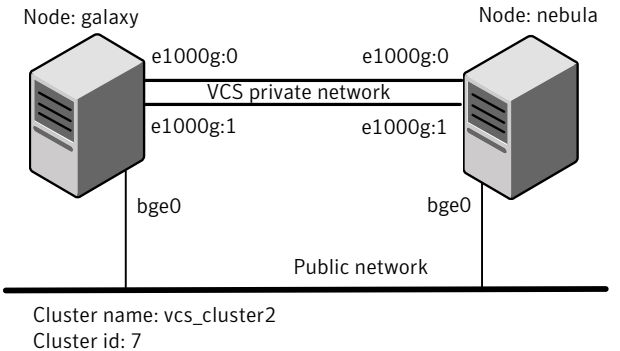

Overview of tasks

Table 4-3 lists the installation and the configuration tasks.

Table 4-3 Installation and configuration tasks

Task	Reference
License and install VCS	■ See "Starting the software installation" on page 63. ■ See "Specifying systems for installation" on page 64. ■ See "Licensing VCS" on page 65. ■ See "Choosing VCS packages for installation" on page 66. ■ See "Choosing to install VCS packages or configure VCS" on page 67. ■ See "Installing VCS packages" on page 78.
Configure the cluster and its features	■ See "Starting the software configuration" on page 68. ■ See "Specifying systems for configuration" on page 69. ■ See "Configuring the basic cluster" on page 69. ■ See "Adding VCS users" on page 73. (optional) ■ See "Configuring SMTP email notification" on page 74. (optional) ■ See "Configuring SNMP trap notification" on page 76. (optional) ■ See "Configuring global clusters" on page 77. (optional)
Create configuration files	See "Creating VCS configuration files" on page 79.
Start VCS and its components	■ See "Starting VCS" on page 79. ■ See "Completing the installation" on page 80.
For clusters that run in secure mode, enable LDAP authentication plug-in if VCS users belong to LDAP domain.	■ See "Enabling LDAP authentication for clusters that run in secure mode" on page 80.
Install language packages	■ See "Installing language packages" on page 81.
Perform the post-installation tasks	■ See "About configuring VCS clusters for data integrity" on page 111. ■ See "Installing the Java Console" on page 82.
Verify the cluster	See "Verifying the cluster after installation" on page 84.

Starting the software installation

You can install VCS using the Veritas product installer or the installvcs program.

To install VCS using the product installer

1. Confirm that you are logged in as the superuser and 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: I for "Install/Upgrade a Product."
4. From the displayed list of products to install, choose: Veritas Cluster Server.

To install VCS using the installvcs program

1. Confirm that you are logged in as the superuser and mounted the product disc.
2. Navigate to the folder that contains the installvcs program.

   ```
   # cd cluster_server
   ```

3. Start the installvcs program.

   ```
   # ./installvcs
   ```

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

Specifying systems for installation

The installer prompts for the system names on which you want to install and then performs an initial system check.

To specify system names for installation

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

   ```
   Enter the system names separated by spaces on which to install
   VCS: galaxy nebula
   ```

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

   See "Creating a single-node cluster using the installer program" on page 189.
2. Review the output as the installer verifies the systems you specify.

   The installer does the following:

- Checks that the local node running the installer can communicate with remote nodes
  If the installer finds *ssh* binaries, it confirms that *ssh* can operate without requests for passwords or passphrases.
- Makes sure the systems use the proper operating system
- Makes sure the systems install from the global zone
- 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 5.0 MP3.

Licensing VCS

The installer checks whether VCS license keys are currently in place on each system. If license keys are not installed, the installer prompts you for the license keys.

See "Checking licensing information on the system" on page 84.

To license VCS

1. Review the output as the utility checks system licensing and installs the licensing package.
2. Enter the license key for Veritas Cluster Server as the installer prompts for each node.

```
Enter a VCS license key for galaxy: [?] XXXX-XXXX-XXXX-XXXX-XXX
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on galaxy
VCS license registered on galaxy
```

3 Enter keys for additional product features.

```
Do you want to enter another license key for galaxy? [y,n,q,?]
(n) y

Enter a VCS license key for galaxy: [?] XXXX-XXXX-XXXX-XXXX-XXX
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on galaxy

Do you want to enter another license key for galaxy? [y,n,q,?]
(n)
```

4 Review the output as the installer registers the license key on the other nodes. Enter keys for additional product features on the other nodes when the installer prompts you.

```
XXXX-XXXX-XXXX-XXXX-XXX successfully registered on nebula
VCS license registered on nebula

Do you want to enter another license key for nebula? [y,n,q,?]
(n)
```

Choosing VCS packages for installation

The installer verifies for any previously installed packages and then based on your choice installs all the VCS packages or only the required packages.

To install VCS packages

1 Review the output as the installer checks the packages that are already installed.

2 Review the output as the installer makes sure that the required OS patches are available on all nodes.

If the installer reports that any of the patches are not available, install the patches on the node before proceeding with the VCS installation.

3 Choose the VCS packages that you want to install.

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

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

1 Installs only the required VCS packages.

2 Installs all the VCS packages.

You must choose this option to configure any optional VCS feature. Note that this option is the default if you already installed the SF HA packages.

3 Installs all the VCS and the SF HA packages. (default option)

If you already installed the SF HA packages, the installer does not list this option.

4 View the list of packages that the installer would install on each node.

If the current version of a package is on a system, the installer removes it from the package installation list for the system.

Choosing to install VCS packages or configure VCS

While you must configure VCS before you can use VCS, you can do one of the following:

- Choose to install and configure VCS now.
  See "Configuring the basic cluster" on page 69.
- Install packages on the systems and leave the cluster configuration steps for later.

To install VCS packages now and configure VCS later

1 If you do not want to configure VCS now, enter n at the prompt.

```
Are you ready to configure VCS? [y,n,q] (y) n
```

The utility 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 utility stops and indicates the actions required to proceed with the process.

2 Review the output as the installer uninstalls any previous versions and installs the VCS 5.0 MP3 packages.

3 Configure the cluster later.

See "Configuring VCS using configure option" on page 60.

Starting the software configuration

You can configure VCS using the Veritas product installer or the installvcs program.

To configure VCS using the product installer

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

2 Start the installer.

```
# ./installer
```

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

3 From the opening Selection Menu, choose: C for "Configure an Installed Product."

4 From the displayed list of products to configure, choose: Veritas Cluster Server.

To configure VCS using the installvcs program

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

2 Navigate to the folder that contains the installvcs program.

```
# cd /cluster_server
```

3 Start the installvcs program.

```
# ./installvcs -configure
```

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

Specifying systems for configuration

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

To specify system names for installation

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

```
Enter the system names separated by spaces on which to configure
VCS: galaxy nebula
```

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

The installer does the following tasks:

- Checks that the local node running the installer can communicate with remote nodes
  If the installer finds ssh binaries, it confirms that ssh can operate without requests for passwords or passphrases.
- Makes sure the systems use the proper operating system
- Makes sure the systems install from the global zone
- Checks whether VCS is installed
- Exits if VCS 5.0 MP3 is not installed

Configuring the basic cluster

Enter the cluster information when the installer prompts you.

To configure the cluster

1 Review the configuration instructions that the installer presents.

2 Enter the unique cluster name and cluster ID.

```
Enter the unique cluster name: [?] vcs_cluster2
Enter the unique Cluster ID number between 0-65535: [b,?] 7
```

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

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

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

You must choose the network interface cards or the aggregated interfaces that the installer discovers and reports. If you want to use aggregated interfaces that the installer has not discovered, then you must manually edit the /etc/lltab file before you start VCS when the installer prompts after product configuration.

See "Starting VCS" on page 79.

You must not enter the network interface card that is used for the public network (typically hme0.)

Answer the following prompts based on architecture:

- For Solaris SPARC:

```
Enter the NIC for the first private heartbeat NIC on galaxy:
[b,?] qfe0
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat NIC on galaxy:
[b,?] qfe1
Would you like to configure a third private heartbeat link?
[y,n,q,b,?](n)
Do you want to configure an additional low priority heartbeat
link? [y,n,q,b,?] (n)
```

- For Solaris x64:

```
Enter the NIC for the first private heartbeat NIC on galaxy:
[b,?] e1000g0
```

```
Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)
Enter the NIC for the second private heartbeat NIC on galaxy:
[b,?] e1000g1
Would you like to configure a third private heartbeat link?
[y,n,q,b,?](n)
Do you want to configure an additional low priority heartbeat
link? [y,n,q,b,?] (n)
```

5 Choose whether to use the same NIC details to configure private heartbeat links on other systems.

```
Are you using the same NICs for private heartbeat links on all
systems? [y,n,q,b,?] (y)
```

If you want to use the NIC details that you entered for galaxy, make sure the same NICs are available on each system. Then, enter **y** at the prompt.

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

6 Verify and confirm the information that the installer summarizes.

Configuring the cluster in secure mode

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

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

See "Preparing to configure the clusters in secure mode" on page 25.

To configure the cluster in secure mode

1 Choose whether to configure VCS to use Symantec Product Authentication Service.

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

- If you want to configure the cluster in secure mode, make sure you meet the prerequisites and enter **y**.
- If you do not want to configure the cluster in secure mode, enter **n**.
  You must add VCS users when the configuration program prompts.

See "Adding VCS users" on page 73.

2 Select one of the options to enable security.

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

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

`Option 1`. Automatic configuration	Enter the name of the Root Broker system when prompted. Requires a remote access to the Root Broker. Review the output as the installer verifies communication with the Root Broker system, checks vxatd process and version, and checks security domain.
`Option 2`. Semiautomatic configuration	Enter the path of the encrypted file (BLOB file) for each node when prompted.

`Option 3`. Manual configuration

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

```
Enter root Broker name:
east.symantecexample.com
Enter root broker FQDN: [b]
(symantecexample.com)
symantecexample.com
Enter root broker domain: [b]
(root@east.symantecexample.com)
root@east.symantecexample.com
Enter root broker port: [b] (2821) 2821
Enter path to the locally accessible
root hash [b] (/var/tmp/
installvcs-1Lcljr/root_hash)
     /root/root_hash
```

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

```
Enter authentication broker principal name on
galaxy [b]
(galaxy.symantecexample.com)
galaxy.symantecexample.com
Enter authentication broker password on galaxy:
Enter authentication broker principal name on
nebula [b]
(nebula.symantecexample.com)
nebula.symantecexample.com
Enter authentication broker password on nebula:
```

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

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

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

Adding VCS users

If you have enabled Symantec Product Authentication Service, you do not need to add VCS users now. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

To add VCS users

1 Review the required information to add VCS users.

2 Reset the password for the Admin user, if necessary.

```
Do you want to set the password for the Admin user
(default password='password')? [y,n,q] (n) y

Enter New 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: [?] smith
Enter New Password:*******

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

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

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

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

Configuring SMTP email notification

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

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

To configure SMTP email notification

1 Review the required information to configure the SMTP email notification.

2 Specify whether you want to configure the SMTP notification.

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

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

See "Configuring SNMP trap notification" on page 76.

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,?] smtp.example.com
```

- Enter the email address of each recipient.

```
Enter the full email address of the SMTP recipient
(example: user@yourcompany.com): [b,?] 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,?] 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,?] 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,?] 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 *Veritas Cluster Server User's Guide* for more information.

To configure the SNMP trap notification

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

2 Specify whether you want to configure the SNMP notification.

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

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

See "Configuring global clusters" on page 77.

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,?] (162)
```

- Enter the SNMP console system name.

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

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

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

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

Configuring global clusters

If you had installed a valid HA/DR license, the installer provides you an option to configure this cluster as global cluster. If not, the installer proceeds to configure VCS based on the configuration details you provided.

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. Note that you can also run the gcoconfig utility in each cluster later to update the VCS configuration file for global cluster.

See *Veritas Cluster Server User's Guide* for instructions to set up VCS global clusters.

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

To configure the global cluster option

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

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

   If you skip this option, the installer proceeds to configure VCS based on the configuration details you provided.
3. Provide information to configure this cluster as global cluster.

   The installer prompts you for a NIC, a virtual IP address, and value for the netmask.
4. Verify and confirm the configuration of the global cluster.

   ```
   Global Cluster Option configuration verification:

   NIC: hme0
   IP: 10.10.12.1
   Netmask: 255.255.240.0

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

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

Installing VCS packages

After the installer gathers all the configuration information, the installer installs the packages on the cluster systems. If you already installed the packages and chose to configure or reconfigure the cluster, the installer proceeds to create the configuration files.

See "Creating VCS configuration files" on page 79.

The utility 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 utility stops and indicates the actions that are required to proceed with the process. Review the output as the installer uninstalls any previous versions and installs the VCS 5.0 MP3 packages.

Creating VCS configuration files

After you install the packages and provide the configuration information, the installer continues to create configuration files and copies them to each system:

```
Creating Cluster Server configuration files ..... Done
Copying configuration files to galaxy..... Done
Copying configuration files to nebula..... Done
Cluster Server configured successfully.
```

If you chose to configure the cluster in secure mode, the installer also configures the Symantec Product Authentication Service.

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

- Creates the security principal
- Executes the encrypted file to create security principal on each node in the cluster

The installer then does the following before the installer starts VCS in secure mode:

- Creates the VxSS service group
- Creates the Authentication Server credentials on each node in the cluster
- Creates the Web credentials for VCS users
- Sets up trust with the root broker

Starting VCS

You can now start VCS and its components on each system. If you chose to configure the cluster in secure mode, the installer also starts the Authentication Service processes on each node in the cluster.

Note: To use aggregated interfaces that the installer has not discovered for private heartbeats, do not opt to start VCS.

See "Configuring the basic cluster" on page 69.

To start VCS

1 Confirm to start VCS and its components on each node.

Enter **y** if you want to start VCS.

If you want to use aggregated interfaces that the installer has not discovered for private heartbeats, enter **n**. Skip to step 2

```
Do you want to start Veritas Cluster Server processes now?
[y,n,q] (y) n
```

2 Do the following to use aggregated interfaces for private heartbeats:

- Edit the /etc/llttab file to replace the names of NICs with the names of the aggregated interfaces.
- Reboot the system for the configuration changes to take effect.

Completing the installation

After VCS 5.0 MP3 installation completes successfully, the installer creates summary, log, and response files. The files provide the useful information that can assist you with the installation and can also assist future installations.

Review the location of the installation log files, summary file, and response file that the installer displays.

Table 4-4 specifies the files that are created at the end of the installation.

Table 4-4 File description

File	Description
summary file	■ Lists the packages that are installed on each system. ■ Describes the cluster and its configured resources. ■ Provides the information for managing the cluster.
log file	Details the entire installation.
response file	Contains the configuration information that can be used to perform secure or unattended installations on other systems. See "Example response file" on page 234.

Enabling LDAP authentication for clusters that run in secure mode

Symantec Product Authentication Service (AT) supports LDAP (Lightweight Directory Access Protocol) user authentication through a plug-in for the

authentication broker. AT supports all common LDAP distributions such as Sun Directory Server, Netscape, OpenLDAP, and Windows Active Directory.

See the *Symantec Product Authentication Service Administrator's Guide*.

The LDAP schema and syntax for LDAP commands (such as, ldapadd, ldapmodify, and ldapsearch) vary based on your LDAP implementation.

Before adding the LDAP domain in Symantec Product Authentication Service, note the following information about your LDAP environment:

- The type of LDAP schema used (the default is RFC 2307)
  - UserObjectClass (the default is posixAccount)
  - UserObject Attribute (the default is uid)
  - User Group Attribute (the default is gidNumber)
  - Group Object Class (the default is posixGroup)
  - GroupObject Attribute (the default is cn)
  - Group GID Attribute (the default is gidNumber)
  - Group Membership Attribute (the default is memberUid)
- URL to the LDAP Directory
- Distinguished name for the user container (for example, UserBaseDN=ou=people,dc=comp,dc=com)
- Distinguished name for the group container (for example, GroupBaseDN=ou=group,dc=comp,dc=com)

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. To enable LDAP authentication plug-in, you must verify the LDAP environment, add the LDAP domain in AT, and then verify LDAP authentication. The AT component packaged with VCS requires you to manually edit the VRTSatlocal.conf file to enable LDAP authentication.

Refer to the *Symantec Product Authentication Service Administrator's Guide* for instructions.

If you have not already added VCS users during installation, you can add the users later.

See *Veritas Cluster Server User's Guide* for instructions to add VCS users.

Installing language packages

Before you install the language packages, do the following:

- Make sure `install_lp` command uses the *ssh* or *rsh* commands as root on all systems in the cluster.
- Make sure that permissions are granted for the system on which `install_lp` is run.

To install the language packages

1. Insert the language disc into the drive.

   The Solaris volume-management software automatically mounts the disc as /cdrom/cdrom0.

2. Change to the /cdrom/cdrom0 directory.

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

3. Install the language packages:

   ```
   # ./install_lp
   ```

To install the language patches

1. Insert the language disc into the drive.

   The Solaris volume-management software automatically mounts the disc as /cdrom/cdrom0.

2. Change to the /cdrom/cdrom0 directory.

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

3. Install the language patches:

   ```
   # ./installmlp
   ```

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 NT/2000 Professional/XP/2003 system, or Solaris system with X-Windows. 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.

When you install the Java Console on the Solaris system, make sure a printer is configured to that system. If you print the online JavaHelp on a system that does not have a printer that is configured, the Java Console might hang.

Review the information about using the Cluster Manager and the Configuration Editor components of the Java Console. For more information, refer to the *Veritas Cluster Server User's Guide*.

Installing the Java Console on Solaris

Review the procedure to install the Java console.

To install Java console on Solaris

1. Create a directory for installation of the Java Console:

   ```
   # mkdir /tmp/install
   ```

2. Insert the software disc with the VCS software into a drive that is connected to the system. The Solaris volume-management software automatically mounts the disc as /cdrom/cdrom0. Type the command:

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

3. Copy the compressed package files from the software disc to the temporary directory:

   ```
   # cp -r cluster_server/pkgs/VRTScscm* /tmp/install
   ```

4. If your system does not have the gunzip utility, copy it from the disc:

   ```
   # cp /cdrom_path/gnu/gunzip /tmp/install
   ```

5. Go to the temporary directory and unzip the compressed package file:

   ```
   # cd /tmp/install
   # gunzip VRTScscm.tar.gz
   ```

   The file VRTScscm.tar is now present in the temporary directory.

6. Extract the compressed file from the tar file:

   ```
   # tar -xvf VRTScscm.tar
   ```

7. Install the software:

   ```
   # pkgadd -d . VRTScscm
   ```

8. Answer `Yes` if prompted.

Installing the Java Console on a Windows system

You can install the VCS Java Console (Cluster Manager) on a Windows NT/2000 Professional/XP/2003 system to administer the cluster.

To install the Java Console on a Windows system

1. Insert the software disc with the VCS software into a drive on your Windows system. For supported languages other than English, insert the language pack media disc into the drive.
2. Using Windows Explorer, select the disc drive.
3. Go to \windows\VCSWindowsInstallers\ClusterManager.
4. Open the language folder of your choice, for example EN or JA.
5. Double-click setup.exe.
6. The Veritas Cluster Manager Install Wizard guides you through the installation process.

Verifying the cluster after installation

When you have used installvcs program and chosen to configure and start VCS, VCS and all components are properly configured and can start correctly. You must verify that your cluster operates properly after the installation.

See "About verifying the VCS installation" on page 127.

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.

Checking licensing information on the system

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

To check licensing information

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

```
# cd /opt/VRTS/bin
# ./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                  = 1249
License Type                   = PERMANENT
OEM ID                         = 478

Features :=
Platform                       = Solaris
Version                        = 5.0
Tier                           = 0
Reserved                       = 0
Mode                           = VCS
```

Updating product licenses using vxlicinst

You can use the `vxlicinst` command 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 85.

To update product licenses

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

```
# cd /opt/VRTS/bin
# ./vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```

Replacing a VCS demo license with a permanent license

When a VCS demonstration 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:

   ```
   # cd /opt/VRTS/bin
   # ./vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
   ```

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

   ```
   # cd /opt/VRTS/bin
   # ./vxlicrep
   ```

5. Start VCS on each node:

   ```
   # hastart
   ```

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.

Chapter 5

Manually installing and configuring VCS

This chapter includes the following topics:

About VCS manual installation

You can manually install and configure VCS instead of using the installvcs program.

Review the following criteria for a manual installation:

- You want to install a single VCS package.
- You want to install VCS to one system in a cluster that runs VCS 5.0 MP3.
- You cannot install on a system over the network, which can occur when you do not have remote root user access.

A manual installation takes a lot of time, patience, and care. Symantec recommends that you use the installvcs program instead of the manual installation when possible.

Requirements for installing VCS

Review the following requirements and verify that you are ready to install the VCS software:

- See "Hardware requirements" on page 21.
- See "Supported operating systems" on page 22.
- See "Supported software" on page 23.

Installing VCS software manually

Table 5-1 lists the tasks that you must perform when you manually install and configure VCS 5.0 MP3.

Table 5-1 Manual installation tasks for VCS 5.0 MP3

Task	Reference
Install VCS software manually on each node in the cluster.	See "Installing VCS packages for a manual installation" on page 90.
Install VCS language pack software manually on each node in the cluster.	See "Installing Japanese language packages in a manual installation" on page 93.
Add a license key.	See "Adding a license key for a manual installation" on page 97.
Restore the configuration files from your previous VCS installation.	See "Upgrading the configuration files" on page 97.
Copy the installation guide to each node.	
Configure LLT and GAB.	■ See "Configuring LLT for a manual installation" on page 97. ■ See "Configuring GAB for a manual installation" on page 100.
Configure VCS.	See "Configuring VCS" on page 100.
Start LLT, GAB, and VCS services.	See "Starting LLT, GAB, and VCS for a manual installation" on page 102.
Modify the VCS configuration.	See "Modifying the VCS configuration" on page 103.

Table 5-1 Manual installation tasks for VCS 5.0 MP3 *(continued)*

Task	Reference
Replace demo license with a permanent license.	See "Replacing a VCS demo license with a permanent license for manual installations" on page 103.

Preparing for a manual installation

Before you install, log in as the superuser. Mount the disc, and copy the files in a temporary folder for installation.

See "Mounting the product disc" on page 49.

To prepare for installation

1 Copy the compressed package and patch files from the software disc to the temporary directory:

```
# cp -r cluster_server/pkgs/* /tmp/install
# cp -r cluster_server/patches/* /tmp/install
```

2 Go to the temporary directory and unzip the compressed package files:

```
# cd /tmp/install
# gunzip *.gz
```

3 If your system does not have the `gunzip` utility, copy it from the disc:

```
# cp /cdrom/cdrom0/gnu/gunzip /tmp/install
```

4 List the files for the following architectures:

- For SPARC:

```
# ls /tmp/install
.               137338-01.tar     VRTSacclib.tar  VRTSperl.tar
..              README.123207-03  VRTSat.tar      VRTSspt.tar
123207-03.tar   README.123208-03  VRTScmccc.tar   VRTSvcs.tar
123208-03.tar   README.123209-03  VRTScmcs.tar    VRTSvcsag.tar
123209-03.tar   README.123210-03  VRTScscm.tar    VRTSvcsmg.tar
123210-03.tar   README.123211-03  VRTScscw.tar    VRTSvcsmn.tar
123211-03.tar   README.123722-01  VRTScssim.tar   VRTSvlic.tar
123722-01.tar   README.123983-01  VRTScutil.tar   VRTSvxfen.tar
123983-01.tar   README.123984-01  VRTSgab.tar     VRTSweb.tar
123984-01.tar   README.123985-01  VRTSicsco.tar   info
123985-01.tar   README.125150-07  VRTSjre15.tar
```

```
125150-07.tar  README.127333-01  VRTSllt.tar
127333-01.tar  README.137338-01  VRTSpbx.tar
```

- For x64:

```
# ls /tmp/install
128048-03.tar     README.128050-03  VRTScscw.tar      VRTSspt.tar
128049-03.tar     README.137339-01  VRTScssim.tar     VRTSvcs.tar
128050-03.tar     README.137384-03  VRTScutil.tar     VRTSvcsag.tar
137339-01.tar     README.137388-01  VRTSgab.tar       VRTSvcsmg.tar
137384-03.tar     VRTSacclib.tar    VRTSicsco.tar     VRTSvcsmn.tar
137388-01.tar     VRTSat.tar        VRTSjre15.tar     VRTSvlic.tar
info              VRTScmccc.tar     VRTSllt.tar       VRTSvxfen.tar
README.128048-03  VRTScmcs.tar      VRTSpbx.tar       VRTSweb.tar
README.128049-03  VRTScscm.tar      VRTSperl.tar
```

Installing VCS packages for a manual installation

VCS has both required and optional packages. Install the required packages first. All packages are installed in the /opt directory.

When you select the optional packages, note the following information:

- Symantec recommends that you install the packages for VCS manual pages (VRTSvcsmn).
- The I/O fencing package can be used only with the shared disks that support SCSI-3 Persistent Reservations (PR). See the *Veritas Cluster Server User's Guide* for a conceptual description of I/O fencing. You need to test shared storage for SCSI-3 PR and to implement I/O fencing.
  See "About setting up I/O fencing" on page 113.
- The VCS configuration wizard (VRTScscw) package includes wizards for the installation and configuration of Veritas products that require VCS configuration.
- To use the Java Console with VCS Simulator, you must install the VRTScssim and VRTScscm packages.

Perform the steps to install VCS packages on each node in the cluster.

To install VCS packages on a node

1. Extract the following required and optional VCS packages and patches from the compressed files:
   - Extract the required packages.

```
# tar -xvf VRTSperl.tar
# tar -xvf VRTSvlic.tar
# tar -xvf VRTSicsco.tar
# tar -xvf VRTSpbx.tar
# tar -xvf VRTSat.tar
# tar -xvf VRTSspt.tar
# tar -xvf VRTSllt.tar
# tar -xvf VRTSgab.tar
# tar -xvf VRTSvxfen.tar
# tar -xvf VRTSvcs.tar
# tar -xvf VRTSvcsmg.tar
# tar -xvf VRTSvcsag.tar
# tar -xvf VRTSjre15.tar
# tar -xvf VRTScutil.tar
# tar -xvf VRTScscw.tar
# tar -xvf VRTSweb.tar
# tar -xvf VRTSacclib.tar
```

- Extract the optional packages. Omit the packages that you do not want.

```
# tar -xvf VRTSvcsmn.tar
# tar -xvf VRTScscm.tar
# tar -xvf VRTScssim.tar
# tar -xvf VRTScmcs.tar
# tar -xvf VRTScmccc.tar
```

- For SPARC 8, extract the patches:

```
# tar -xvf 127333-01.tar
# tar -xvf 137338-01.tar
# tar -xvf 123722-01.tar
# tar -xvf 123207-03.tar
# tar -xvf 125150-07.tar
# tar -xvf 123984-01.tar
# tar -xvf 123983-01.tar
```

- For SPARC 9, extract the patches:

```
# tar -xvf 127333-01.tar
# tar -xvf 137338-01.tar
# tar -xvf 123722-01.tar
# tar -xvf 123208-03.tar
```

```
# tar -xvf 125150-07.tar
# tar -xvf 123984-01.tar
# tar -xvf 123983-01.tar
```

- For SPARC 10, extract the patches:

```
# tar -xvf 127333-01.tar
# tar -xvf 137338-01.tar
# tar -xvf 123722-01.tar
# tar -xvf 123209-03.tar
# tar -xvf 123210-03.tar
# tar -xvf 125150-07.tar
# tar -xvf 123211-03.tar
# tar -xvf 123983-01.tar
```

- For x64, extract the patches:

```
# tar -xvf 137388-01.tar
# tar -xvf 137339-01.tar
# tar -xvf 128049-03.tar
# tar -xvf 128048-03.tar
# tar -xvf 137384-03.tar
# tar -xvf 128050-03.tar
```

2 Install the following required and optional VCS packages from the compressed files:

- Install the following required packages in the order shown:

```
# pkgadd -d . VRTSperl VRTSvlic VRTSicsco VRTSpbx VRTSat
VRTSspt VRTSllt VRTSgab VRTSvxfen VRTSvcs VRTSvcsmg VRTSvcsag
VRTSjre15 VRTScutil VRTScscw VRTSweb VRTSacclib
```

- Install the optional packages, in the order shown. Omit the packages that you do not want.

```
# pkgadd -d . VRTSvcsmn VRTScscm VRTScssim VRTScmcs
VRTScmcc
```

3 Install the following patches from the compressed files:

- For SPARC 8, install the patches:

```
# patchadd 127333-01
# patchadd 137338-01
```

```
# patchadd 123722-01
# patchadd 123207-03
# patchadd 125150-07
# patchadd 123984-01
# patchadd 123983-01
```

- For SPARC 9, install the patches:

```
# patchadd 127333-01
# patchadd 137338-01
# patchadd 123722-01
# patchadd 123208-03
# patchadd 125150-07
# patchadd 123984-01
# patchadd 123983-01
```

- For SPARC 10, install the patches:

```
# patchadd 127333-01
# patchadd 137338-01
# patchadd 123722-01
# patchadd 123209-03
# patchadd 123210-03
# patchadd 125150-07
# patchadd 123211-03
# patchadd 123984-01
# patchadd 123983-01
```

- For x64, install the patches:

```
# patchadd 137388-01
# patchadd 137339-01
# patchadd 128049-03
# patchadd 128048-03
# patchadd 137384-03
# patchadd 128050-03
```

Installing Japanese language packages in a manual installation

Install the language packages that VCS requires after you install the base VCS packages. The Japanese language packages are as follows:

Required packages	
VRTSmulic	Multi Language Symantec License Utilities
VRTSjaico	Symantec Infrastructure Core Services Common Japanese Language
VRTSjapbx	Symantec Private Branch Exchange Japanese Language
VRTSatJA	Symantec Product Authentication Service Software Japanese Language Kit
VRTSjacs	Japanese Veritas Cluster Server Message Catalogs by Symantec
VRTSjacsj	Japanese VERITAS Cluster Server Cluster Manager
VRTSjacsu	Japanese Symantec Veritas Cluster Utility
VRTSjaweb	Japanese Symantec Web Server Language Pack
Optional packages	
VRTSjacmc	Veritas Cluster Management Console Japanese Localization
VRTSjacsm	Japanese VERITAS Cluster Server Simulator

Before you install, make sure that you are logged on as superuser and that you have mounted the language disc.

See "Mounting the product disc" on page 49.

Perform the steps on each node in the cluster to install Japanese language packages.

To install the language packages and patches on a node

1 Copy the compressed package and patch files from the software disc to the temporary directory.

```
# cp -r ja/cluster_server/pkgs/* /tmp
# cp -r ja/cluster_server/patches/* /tmp
```

2 If your system does not have the gunzip utility, you can copy it from the base product's disc.

```
# cp /cdrom_path/gnu/gunzip /tmp/install
```

3 Go to the temporary directory and unzip the compressed package files.

```
# cd /tmp
# gunzip *.gz
```

4 List the files in the temporary directory.

```
# ls /tmp/install
123680-05.tar  info              README.123982-01  VRTSjacsu.tar
123975-03.tar  README.123680-05  VRTSatJA.tar      VRTSjaico.tar
123976-03.tar  README.123975-03  VRTSjacmc.tar     VRTSjapbx.tar
123977-03.tar  README.123976-03  VRTSjacsj.tar     VRTSjaweb.tar
123978-03.tar  README.123977-03  VRTSjacsm.tar     VRTSmulic.tar
123982-01.tar  README.123978-03  VRTSjacs.tar      VRTSperl.tar
```

5 Extract the compressed packages:

- Extract the required files:

```
# tar -xvf VRTSmulic.tar
# tar -xvf VRTSjaico.tar
# tar -xvf VRTSjapbx.tar
# tar -xvf VRTSatJA.tar
# tar -xvf VRTSjacs.tar
# tar -xvf VRTSjacsj.tar
# tar -xvf VRTSjacsu.tar
# tar -xvf VRTSjaweb.tar
```

- Extract the optional packages:

```
# tar -xvf VRTSjacmc.tar
# tar -xvf VRTSjacsm.tar
```

6 Extract the compressed patches:

- Extract the required patches for Solaris 10:

```
# tar -xvf 123680-05.tar
# tar -xvf 123977-03.tar
# tar -xvf 123978-03.tar
```

- Extract the required patches for Solaris 9:

```
# tar -xvf 123680-05.tar
# tar -xvf 123976-03.tar
```

- Extract the required patches for Solaris 8:

```
# tar -xvf 123680-05.tar
# tar -xvf 123975-03.tar
```

- Extract the patch for the VRTSjacmc package.

```
# tar -xvf 123982-01.tar
```

7 Install the following required and optional VCS packages from the compressed files:

- Install the following required packages in the order shown:

```
# pkgadd -d . VRTSmulic VRTSjaico VRTSjapbx VRTSatJA
VRTSjacs VRTSjacsj VRTSjacsu VRTSjaweb
```

- Install the optional packages, in the order shown. Omit the packages that you do not want.

```
# pkgadd -d . VRTSjacmc VRTSjacsm
```

8 Install the following required and optional VCS patches from the compressed files:

- Install the following required patches for Solaris 10:

```
# patchadd 123680-05
# patchadd 123977-03
# patchadd 123978-03
```

- Install the following required patches for Solaris 9:

```
# patchadd 123680-05
# patchadd 123976-03
```

- Install the following required patches for Solaris 8:

```
# patchadd 123680-05
# patchadd 123975-03
```

- Install the following patch for the VRTSjacmc package.

```
# patchadd 123982-01
```

Adding a license key for a manual installation

After you have installed all packages on each cluster node, use the `vxlicinst` command to add the VCS license key on each system:

```
# cd /opt/VRTS/bin
# ./vxlicinst -k XXXX-XXXX-XXXX-XXXX-XXXX-XXX
```

Checking licensing information on the system for a manual installation

Use the `vxlicrep` utility to display information about all Veritas licenses on a system. For example, enter:

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

Upgrading the configuration files

You need to restore the configuration files from your previous VCS installation if you manually added 5.0 MP3 packages to upgrade your cluster to VCS.

Configuring LLT for a manual installation

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, set up two files: /etc/llthosts and /etc/lltab on each node in the cluster.

Setting up /etc/llthosts for a manual installation

The file llthosts(4) is a database. It contains one entry for the system that links the LLT system ID (in the first column) with the LLT host name. You must create an identical file on each node in the cluster.

Use vi or another editor, to create the file /etc/llthosts that contains the entries that resemble:

```
0 north
1 south
```

Setting up /etc/lltab for a manual installation

The /etc/lltab file must specify the system's ID number (or its node name), 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 "LLT directives for a manual installation" on page 99.

Use vi or another editor to create the file /etc/lltab that contains the entries that resemble the following:

- For SPARC:

```
set-node north
set-cluster 2
link qfe0 qfe:0 - ether - -
link qfe1 qfe:1 - ether - -
```

- For x64:

```
set-node north
set-cluster 2
link e1000g0 e1000g:0 - ether - -
link e1000g1 e1000g:1 - ether - -
```

The first line must identify the system where the file exists. In the example, the value for `set-node` can be: north, 0, or the file name /etc/nodename. The file needs to contain the name of the system (north in this example) to use these choices. 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 file/opt/VRTSllt/sample-lltab.

LLT directives for a manual installation

For more information about LLT directives, refer to the `lltab`(4) manual page.

Table 5-2 contains the LLT directives for a manual installation.

Table 5-2 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 0-31. 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.
`link`	Attaches LLT to a network interface. At least one link is required, and up to eight are supported. 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 `lltab` 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 `lltab(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.
`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-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. Note that LLT distributes network traffic evenly across all available network connections. In addition to enabling VCS communication, broadcasts heartbeats to monitor each network connection.

For more information about LLT directives, refer to the `lltab`(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.

By default, Sun systems assign the same MAC address to all interfaces. Thus, connecting two or more interfaces to a network switch can cause problems. Consider the following example. You configure an IP on one public interface and LLT on another. Both interfaces are connected to a switch. The duplicate MAC address on the two switch ports can cause the switch to incorrectly redirect IP traffic to the LLT interface and vice-versa. To avoid this issue, configure the system to assign unique MAC addresses by setting the eeprom(1M) parameter local-mac-address? to true.

Configuring GAB for a manual installation

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, use vi or another editor to set up an /etc/gabtab configuration file on each node in the cluster. 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. By default, N is the number of systems in the cluster.

Note: Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. Using `-c -x` dramatically increases configuration time for the Gigabit Ethernet controller and can lead to a split-brain condition.

Configuring VCS

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 resources.
- The name of the cluster.
- The name of the systems that make up the cluster.

Editing the main.cf file

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 *Veritas Cluster Server User's Guide*.

To edit the main.cf file

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.
3. Save and close the file.

Example, main.cf

An example main.cf for a two-node cluster:

```
include "types.cf"
cluster VCSCluster2 ( )
system north
system south
```

An example main.cf for a single-node cluster:

```
include "types.cf"
cluster VCSCluster1 ( )
system sn1
```

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 resources. During new installations, the types.cf file is automatically copied in to the /etc/VRTSvcs/conf/config directory.

Starting LLT, GAB, and VCS for a manual installation

Start LLT, GAB, and VCS.

To start LLT

- On each node, type:

```
# /etc/rc2.d/S70llt start
```

If LLT is configured correctly on each node, the console output resembles:

```
Apr  5 14:46:18 north llt: LLT:10009: LLT Protocol available
```

See "Verifying LLT" on page 135.

To start GAB

- On each node, type:

```
# /etc/rc2.d/S92gab start
```

If GAB is configured correctly on each node, the console output resembles:

```
Apr  5 14:46:29 north gab: GAB:20021: GAB available
Apr  5 14:51:50 north gab: GAB:20026: Port a registration
waiting for seed port membership
```

See "Verifying GAB" on page 138.

To start VCS

- On each node, type:

```
# /etc/rc3.d/S99vcs start
```

If VCS is configured correctly on each node, the console output resembles:

```
Apr  5 14:52:02 north gab: GAB:20036: Port h gen 3972a201
      membership 01
```

See "Verifying the cluster" on page 139.

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 Cluster Server Management Console, or the Cluster Manager (Java Console). For information on management tools, refer to the *Veritas Cluster Server User's Guide*.

You can also edit the main.cf file directly. For information on the structure of the main.cf file, refer to the *Veritas Cluster Server User'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 Cluster Management Console, 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. Presented in this guide is a reference configuration example of a system with a ClusterService group.

See "Sample main.cf file for VCS clusters" on page 130.

Replacing a VCS demo license with a permanent license for manual installations

When a VCS demonstration key license expires, you can replace it with a permanent license using the `vxlicinst` program.

See "Checking licensing information on the system" on page 84.

Installing VCS using JumpStart

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

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

Tasks for a JumpStart installation of VCS

For detailed instructions, follow the JumpStart documentation that came with your operating system. These steps are provided as a summary only.

To install the Veritas packages on a JumpStart server

1. Add a client (register to the JumpStart server).

   See the JumpStart documentation that came with your operating system for details.

2. Copy the compressed Veritas Cluster Server packages and patch files to a temporary directory, uncompress the files, and extract the packages from the tar files.

   See "Copying and unzipping the VCS packages and patches" on page 105.

3. Copy the packages and patches to the JumpStart server under a shared directory on the network. Note the subdirectory with the relevant packages for your installation.

4. Determine the installation order and modify the rules file.

   See "Establishing the order of installation" on page 106.

5. Write the JumpStart start and finish scripts.

   The pkgadd operation and the patchadd operation to install the packages and patches must be coded in a script that the JumpStart server can use.

   Run the installvcs command with the jumpstart option to create a sample finish file. To create the sample finish file, perform the following task:

   ```
   # cd cluster_server
   # ./installvcs -jumpstart dir_path
   ```

   Where *dir_path* indicates the path to the directory in which to create the finish file.

   For the language pack, add lines for the language packages and patches in the finish script.

   See "Adding language pack information to the finish file" on page 107.

6. Add packages and patches to the appropriate location and set up the JumpStart environment.

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

7 Create the response files for the VRTS packages.

See "Creating the JumpStart response files" on page 109.

8 Run JumpStart to install the packages.

JumpStart may restart the system after the packages have been installed.

9 Run the installvcs command from the disc directory to configure the Veritas software.

```
# /cdrom/cdrom0/cluster_server/installvcs -configure
```

Copying and unzipping the VCS packages and patches

Before you can install the packages, you must unzip them, and extract them from the tar file.

To unzip the packages

1 Log on as superuser (root).

2 Create a directory for installation.

```
# mkdir /parent_directory/install
```

3 Insert the product disc into a drive that is connected to the system.

The Solaris volume management software automatically mounts the disc as /cdrom/cdrom0. Type the command:

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

4 Copy the compressed package files and patches from the software disc to the temporary directory.

If you download the software, you need to unzip and untar the downloaded file to the temporary directory.

```
# cp -r cluster_server/pkgs/* /parent_directory/install/pkgs
# cp -r cluster_Server/patches/* /parent_directory/install/patches
```

For the language pack, copy the compressed package and patch files from the language pack disc.

```
# cp -r ja/cluster_server/pkgs/* /parent_directory/install/pkgs
# cp -r ja/cluster_server/patches/* /parent_directory/install/patches
```

5 If your system does not have the gunzip utility, copy it from the disc:

```
# cp /cdrom/cdrom0/gnu/gunzip /parent_directory/install
```

6 Go to the temporary directory and unzip the compressed package files and patches.

```
# cd /parent_directory/install/pkgs
# gunzip VRTS*.gz
# cd /parent_directory/install/patches
# gunzip *.gz
```

7 Decompress and extract each package.

```
# cd /parent_directory/install/pkgs
# tar -xvf package_name.tar
# tar -xvf package_name.tar
# tar -xvf package_name.tar
.
.
```

8 Decompress and extract each patch.

```
# cd /parent_directory/install/patches
# tar -xvf patch_name.tar
# tar -xvf patch_name.tar
# tar -xvf patch_name.tar
.
.
```

9 List the files in the temporary directory.

```
# ls /parent_directory/install/pkgs
# ls /parent_directory/install/patches
```

10 Use these directories to provide the packages and patches for the manual installation procedure.

Establishing the order of installation

You must install the packages in the correct order. For example, some packages must be installed before other packages because of various product dependencies.

The list of the available packages has descriptions of each package. To get the installation order for Veritas Cluster Server patches and packages, use the option -requiredpkgs or -installpkgs with the scripts from the disc.

The requiredpkgs option displays only the required packages and the installpkgs option displays all packages.

To get package installation order for Veritas Cluster Server

1 Change to the Cluster Server installation directory. Use the following command:

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

2 Run the script with the -requiredpkgs option or the -installpkgs option.

```
# ./installvcs -requiredpkgs

or

# ./installvcs -installpkgs
```

3 For the language pack, the package order is:

```
VRTSmulic VRTSjaico VRTSjapbx VRTSjaweb
VRTSjacs VRTSjacsu VRTSjacsj VRTSjacsm VRTSatJA VRTSjacmc
```

4 For the language pack, the patch order follows:

- For Solaris 10:

```
123680-05 123977-03 123978-03 123982-01
```

- For Solaris 9:

```
123680-05 123976-03 123982-01
```

- For Solaris 8:

```
123680-05 123975-03 123982-01
```

Adding language pack information to the finish file

For the language pack, add lines for the language packages and patches in the finish script. If the finish file resembles:

```
. . .
for PKG in VRTSperl VRTSvlic VRTSicsco . . .

do
.
.
.
done

for PATCH in 127333-01 137338-01 . . .
do
.
done
. . .
```

Add the following lines for the language pack after the patch information for VCS. Copy the command syntax between the "do" and "done" lines and add that for the language pack lines as well. Note that the line that starts "for PKG" is on two lines in this guide, but should be on a single line in the file.

```
. . .
for PKG in VRTSmulic VRTSjaico VRTSjapbx VRTSjaweb VRTSjacs
VRTSjacsu VRTSjacsj VRTSjacsm VRTSatJA VRTSjacmc

do
.
.
.
done

for PATCH in xxxxxx-xx yyyyyy-yy. . .

do
.
done
. . .
```

Where *xxxxxx-xx* and *yyyyyy-yy* are place holders for patch IDs.

In the "for PATCH" line, you must replace the patch IDs in operating system-specific order as follows:

- For Solaris 10:

  ```
  for PATCH in 123680-05 123977-03 123978-03 123982-01
  ```

- For Solaris 9:

```
for PATCH 123680-05 123976-03 123982-01
```

- For Solaris 8:

```
for PATCH 123680-05 123975-03 123982-01
```

Creating the JumpStart response files

Use the following instructions to create the response files for JumpStart and add the relevant parameters to the finish file.

To create the response files for JumpStart

1 The VRTSjre15, VRTSjre, and VRTScutil packages require an empty response file.

Add the following lines to your scripts to create the empty response files. You can then perform a `pkgadd` command for the VRTSjre15, VRTSjre, and VRTScutil packages.

```
touch responsefile
# pkgadd -r responsefile package_name
```

2 The VRTScssim response files must contain the following settings:

```
BASEDIR=/opt
```

3 Copy the supplied VRTS admin files, and modify them if needed.

```
# cp storage_foundation/scripts/VRTS* \
    /parent_directory/install/pkgs
```

4 Specify the -a adminfile option to the pkgadd command. This adminfile must be created in the current directory, and contain the following entries:

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

Chapter 6

Configuring VCS clusters for data integrity

This chapter includes the following topics:

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.

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
Coordination points	Act as a global lock during membership changes

I/O fencing in VCS involves coordination points and data disks. Each component has a unique purpose and uses different physical disk devices. The fencing driver, known as vxfen, directs VxVM as necessary to carry out actual fencing operations at the disk group level.

About data disks

Data disks are standard disk devices for data storage and are either physical disks or RAID Logical Units (LUNs). These disks must support SCSI-3 PR and are part of standard VxVM or CVM disk groups.

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

About coordination points

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

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

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. Users cannot store data on these disks or include the disks in a disk group for user data. The coordinator disks can be any three disks that support SCSI-3 PR. Coordinator disks cannot be the special devices that array vendors use. For example, you cannot use EMC gatekeeper devices as coordinator disks.

Symantec recommends using the smallest possible LUNs for coordinator disks. Because coordinator disks do not store any data, cluster nodes need to only register with them and do not need to reserve them.

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

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

You can use iSCSI devices as coordinator disks for I/O fencing. However, I/O fencing supports iSCSI devices only when you use DMP disk policy. If you use iSCSI devices as coordinator disks, make sure that the /etc/vxfenmode file has the disk policy set to DMP.

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

http://entsupport.symantec.com/docs/283161

About setting up I/O fencing

Figure 6-1 illustrates the tasks involved to configure I/O fencing.

Figure 6-1 Workflow to configure I/O fencing

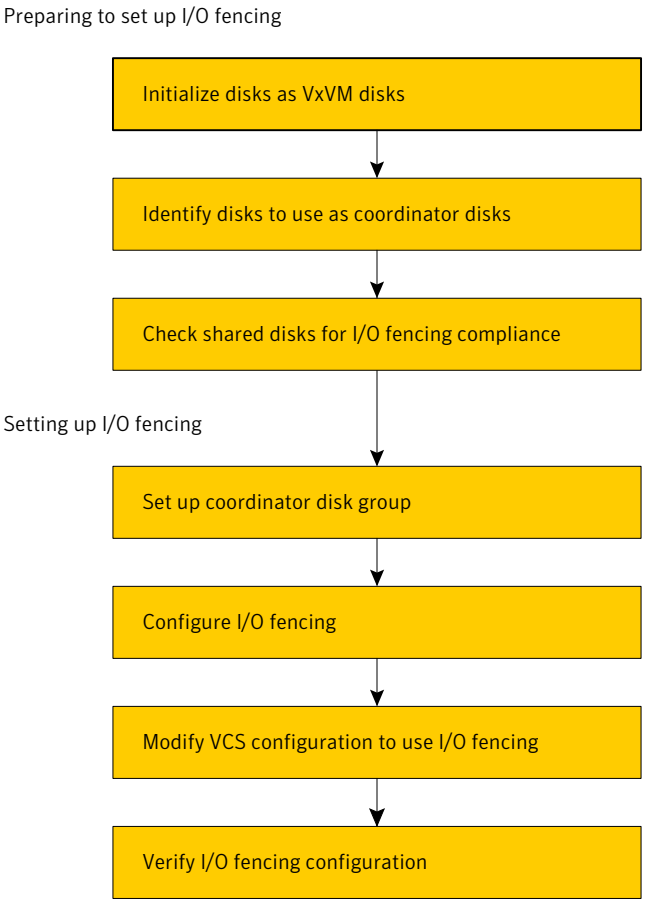

See "Preparing to configure I/O fencing" on page 116.

See "Setting up I/O fencing" on page 121.

I/O fencing requires the coordinator disks be configured in a disk group. The coordinator disks must be accessible to each node in the cluster. These disks enable the vxfen driver to resolve potential split-brain conditions and prevent data corruption.

Review the following requirements for coordinator disks:

- You must have three coordinator disks.
- Each of the coordinator disks must use a physically separate disk or LUN.

- Each of the coordinator disks should exist on a different disk array, if possible.
- You must initialize each disk as a VxVM disk.
- The coordinator disks must support SCSI-3 persistent reservations.
- The coordinator disks must exist in a disk group (for example, vxfencoorddg).
- Symantec recommends using hardware-based mirroring for coordinator disks.

The I/O fencing configuration files include:

/etc/vxfendg	You must create this file to include the coordinator disk group information.
/etc/vxfenmode	You must set the I/O fencing mode to SCSI-3. You can configure the vxfen module to use either DMP devices or the underlying raw character devices. Note that you must use the same SCSI-3 disk policy on all the nodes. The SCSI-3 disk policy can either be raw or dmp. The policy is raw by default.

/etc/vxfentab	When you run the vxfen startup file to start I/O fencing, the script creates this /etc/vxfentab file on each node with a list of all paths to each coordinator disk. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files.
	Thus any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all paths to the coordinator disks.
	Note: The /etc/vxfentab file is a generated file; do not modify this file.
	An example of the /etc/vxfentab file on one node resembles as follows:
	■ Raw disk:
	`/dev/rdsk/c1t1d0s2`
	`/dev/rdsk/c2t1d0s2`
	`/dev/rdsk/c3t1d2s2`
	■ DMP disk:
	`/dev/vx/rdmp/c1t1d0s2`
	`/dev/vx/rdmp/c2t1d0s2`
	`/dev/vx/rdmp/c3t1d0s2`

In some cases you must remove disks from or add disks to an existing coordinator disk group.

Warning: If you remove disks from an existing coordinator disk group, then be sure to remove the registration and reservation keys from these disks before you add the disks to another disk group.

Preparing to configure I/O fencing

Make sure you performed the following tasks before configuring I/O fencing for VCS:

- Install the correct operating system.
- Install the VRTSvxfen package when you installed VCS.
- Install a version of Veritas Volume Manager (VxVM) that supports SCSI-3 persistent reservations (SCSI-3 PR).
  Refer to the installation guide that comes with the Storage Foundation product that you use.

The shared storage that you add for use with VCS software must support SCSI-3 persistent reservations, a functionality that enables the use of I/O fencing.

Perform the following preparatory tasks to configure I/O fencing:

Initialize disks as VxVM disks	See "Initializing disks as VxVM disks" on page 116.
Identify disks to use as coordinator disks	See "Identifying disks to use as coordinator disks" on page 118.
Check shared disks for I/O fencing The tasks involved in checking the shared disks for I/O fencing are as follows: ■ Verify that the nodes have access to the same disk ■ Test the disks using the vxfentsthdw utility	See "Checking shared disks for I/O fencing" on page 118.

Initializing disks as VxVM disks

Install the driver and HBA card. Refer to the documentation from the vendor for instructions.

After you physically add shared disks to the nodes, you must do the following:

- Initialize them as VxVM disks
- Verify that all the nodes see the same disk

See the *Veritas Volume Manager Administrator's Guide* for more information on how to add and configure disks.

To initialize disks as VxVM disks

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

   ```
   # devfsadm
   ```

2. 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.

3. 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
   ============================================================

   libvxCLARiiON.so            DGC
   libvxcscovrts.so            CSCOVRTS
   libvxemc.so                 EMC
   ```

4. 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.

5. To initialize the disks as VxVM disks, use one of the following methods:

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

```
vxdisksetup -i device_name
```

The example specifies the CDS format:

```
# vxdisksetup -i c2t13d0
```

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

Identifying disks to use as coordinator disks

After you add and initialize disks, 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 list
```

2 Pick three SCSI-3 PR compliant shared disks as coordinator disks.

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.

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 *Veritas Cluster Server User's Guide*.

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

- Verifying that nodes have access to the same disk
  See "Verifying that the nodes have access to the same disk" on page 119.
- Testing the shared disks for SCSI-3

See "Testing the disks using vxfentsthdw utility" on page 119.

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.

   ```
   /sbin/vxfenadm -i diskpath
   ```

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

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

   From node A, enter:

   ```
   # /sbin/vxfenadm -i /dev/rdsk/c1t1d0s2

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

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

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

   ```
   # /sbin/vxfenadm -i /dev/rdsk/c3t1d2s2

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

Testing the disks using vxfentsthdw utility

This procedure uses the /dev/rdsk/c1t1d0s2 disk in the steps.

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

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

```
The disk /dev/rdsk/c1t1d0s2 is ready to be configured for I/O Fencing on
node galaxy
```

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

To test the disks using vxfentsthdw utility

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

   See "Setting up inter-system communication" on page 38.

   After you complete the testing process, remove permissions for communication and restore public network connections.

   See "Removing permissions for communication" on page 126.

2. From one node, start the utility.

   Do one of the following:

   - If you use `ssh` for communication:

     ```
     # /opt/VRTSvcs/vxfen/bin/vxfentsthdw
     ```

   - If you use `rsh` for communication:

     ```
     # /opt/VRTSvcs/vxfen/bin/vxfentsthdw -n
     ```

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

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

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

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

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

```
Enter the disk name to be checked for SCSI-3 PGR on node
galaxy in the format: /dev/rdsk/cxtxdxsx
/dev/rdsk/c2t13d0s2
Enter the disk name to be checked for SCSI-3 PGR on node
nebula in the format: /dev/rdsk/cxtxdxsx
Make sure it's the same disk as seen by nodes galaxy and nebula
/dev/rdsk/c2t13d0s2
```

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

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

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

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

ALL tests on the disk /dev/rdsk/c1t1d0s2 have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
```

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

Setting up I/O fencing

Make sure you completed the preparatory tasks before you set up I/O fencing.

Tasks that are involved in setting up I/O fencing include:

Table 6-1 Tasks to set up I/O fencing

Action	Description
Setting up coordinator disk groups	See "Setting up coordinator disk groups" on page 122.
Configuring I/O fencing	See "Configuring I/O fencing" on page 122.
Modifying VCS configuration to use I/O fencing	See "Modifying VCS configuration to use I/O fencing" on page 123.
Verifying I/O fencing configuration	See "Verifying I/O fencing configuration" on page 125.

Setting up coordinator disk groups

From one node, create a disk group named vxfencoorddg. This group must contain three disks or LUNs. If you use VxVM 5.0 or later, you must also set the coordinator attribute for the coordinator disk group. VxVM uses this attribute to prevent the reassignment of coordinator disks to other disk groups.

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

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

The following example procedure assumes that the disks have the device names c1t1d0s2, c2t1d0s2, and c3t1d0s2.

To create the vxfencoorddg disk group

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

```
# vxdg init vxfencoorddg c1t1d0s2 c2t1d0s2 c3t1d0s2
```

2 If you use VxVM 5.0 or later, set the coordinator attribute value as "on" for the coordinator disk group.

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

Configuring I/O fencing

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. Deport the coordinator disk group:

   ```
   # vxdg deport vxfencoorddg
   ```

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

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

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

   ```
   # vxdg deport vxfencoorddg
   ```

4. 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.

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

   - For DMP configuration:

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

   - For raw device configuration:

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

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

   ```
   # more /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

/etc/VRTSvcs/conf/config/main.cf. If you reset this attribute to UseFence = None, VCS does not make use of I/O fencing abilities while failing over service groups. However, I/O fencing needs to be disabled separately.

To modify VCS configuration to enable I/O fencing

1 Save the existing configuration:

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

2 Stop VCS on all nodes:

```
# hastop -all
```

3 If the I/O fencing driver vxfen is already running, stop the I/O fencing driver.

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

4 Make a backup copy of the main.cf file:

```
# cd /etc/VRTSvcs/conf/config
# cp main.cf main.orig
```

5 On one node, use vi or another text editor to edit the main.cf file. To modify the list of cluster attributes, add the UseFence attribute and assign its value as SCSI3.

```
cluster clus1
UserNames = { admin = "cDRpdxPmHpzS." }
Administrators = { admin }
HacliUserLevel = COMMANDROOT
CounterInterval = 5
UseFence = SCSI3
)
```

6 Save and close the file.

7 Verify the syntax of the file /etc/VRTSvcs/conf/config/main.cf:

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

8 Using rcp or another utility, copy the VCS configuration file from a node (for example, galaxy) to the remaining cluster nodes.

For example, on each remaining node, enter:

```
# rcp galaxy:/etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config
```

9 Start the I/O fencing driver and VCS. Perform the following steps on each node:

- Start the I/O fencing driver.
  The vxfen startup script also invokes the `vxfenconfig` command, which configures the vxfen driver to start and use the coordinator disks that are listed in /etc/vxfentab.

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

- Start VCS.

  ```
  # /opt/VRTS/bin/hastart
  ```

Verifying I/O fencing configuration

Verify from the vxfenadm output that the SCSI-3 disk policy reflects the configuration in the /etc/vxfenmode file.

To verify I/O fencing configuration

- On one of the nodes, type:

```
# vxfenadm -d

I/O Fencing Cluster Information:
================================

Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:

   * 0 (galaxy)
   1 (nebula)

RFSM State Information:
   node 0 in state 8 (running)
   node 1 in state 8 (running)
```

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 7

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 LLT and GAB configuration files

Low Latency Transport (LLT) and Group Membership and Atomic Broadcast (GAB) are VCS communication services. LLT requires /etc/llthosts and /etc/llttab files. GAB requires /etc/gabtab file.

The information that these LLT and GAB configuration files contain is as follows:

- The /etc/llthosts file

The file `llthosts` is a database that contains one entry per system. This file links the LLT system ID (in the first column) with the LLT host name. This file is identical on each node in the cluster.
For example, the file /etc/llthosts contains the entries that resemble:

```
0       galaxy
1       nebula
```

- The /etc/llttab file

  The file `llttab` contains the information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the private network links that correspond to the specific system. For example, the file /etc/llttab contains the entries that resemble the following:

  - For Solaris SPARC:

    ```
    set-node galaxy
    set-cluster 2
    link qfe0 qfe:0 - ether - -
    link qfe1 qfe:1 - ether - -
    ```

  - For Solaris x64:

    ```
    set-node galaxy
    set-cluster 2
    link e1000g0 e1000g:0 - ether - -
    link e1000g1 e1000g:1 - 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.
  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.

- The /etc/gabtab file
  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. By default, `N` is the number of nodes in the cluster.

Note: The use of the `-c -x` option for `/sbin/gabconfig` is not recommended. The Gigabit Ethernet controller does not support the use of `-c -x`.

About the VCS configuration file main.cf

The VCS configuration file /etc/VRTSvcs/conf/config/main.cf is created during the installation process.

See "Sample main.cf file for VCS clusters" on page 130.

See "Sample main.cf file for global clusters" on page 132.

The main.cf file contains the minimum information that defines the cluster and its nodes. In addition, the file types.cf, which is listed in the include statement, defines the VCS bundled types for VCS resources. The file types.cf is also located in the directory /etc/VRTSvcs/conf/config after installation.

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 program 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 that you added is present.
- If you configured the cluster in secure mode, the main.cf includes the VxSS service group and "SecureClus = 1" cluster attribute.
- The installvcs program creates the ClusterService service group. The group includes the IP, NIC, and VRTSWebApp resources.
  The service group also has the following characteristics:
  - The service group also includes the notifier resource configuration, which is based on your input to installvcs program prompts about notification.
  - The installvcs program also creates a resource dependency tree.

- If you set up global clusters, the ClusterService service group contains an Application resource, wac (wide-area connector). This resource's attributes contain definitions for controlling the cluster in a global cluster environment.
  Refer to the *Veritas Cluster Server User's Guide* for information about managing VCS global clusters.

Refer to the *Veritas Cluster Server User's Guide* to review the configuration concepts, and descriptions of main.cf and types.cf files for Solaris systems.

Sample main.cf file for VCS clusters

The following sample main.cf file is for a three-node secure cluster that the Cluster Management Console manages locally.

```
include "types.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 = hme0
    )

NotifierMngr ntfr (
    SnmpConsoles = { vcslab4079 = SevereError }
    SmtpServer = "smtp.veritas.com"
```

```
    SmtpRecipients = { "johndoe@veritas.com" = SevereError }
    )

ntfr requires csgnic

// resource dependency tree
//
//     group ClusterService
//     {
//     NotifierMngr ntfr
//         {
//         NIC csgnic
//         }
// }

group VxSS (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    Parallel = 1
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
    )

Phantom phantom_vxss (
    )

ProcessOnOnly vxatd (
    IgnoreArgs = 1
    PathName = "/opt/VRTSat/bin/vxatd"
    )

// resource dependency tree
//
//     group VxSS
//     {
//     Phantom phantom_vxss
//     ProcessOnOnly vxatd
//     }
```

Sample main.cf file for global clusters

If you installed 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"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
    )

IP gcoip (
    Device = hme0
    Address = "10.182.13.50"
    NetMask = "255.255.240.0"
    )
```

```
NIC csgnic (
    Device = hme0
    )

NotifierMngr ntfr (
    SnmpConsoles = { vcslab4079 = SevereError }
    SmtpServer = "smtp.veritas.com"
    SmtpRecipients = { "johndoe@veritas.com" = SevereError }
    )

gcoip requires csgnic
ntfr requires csgnic
wac requires gcoip

// resource dependency tree
//
//     group ClusterService
//     {
//     NotifierMngr ntfr
//         {
//         NIC csgnic
//         }
//     Application wac
//         {
//         IP gcoip
//             {
//             NIC csgnic
//             }
//         }
//     }

group VxSS (
    SystemList = { sysA = 0, sysB = 1, sysC = 2 }
    Parallel = 1
    AutoStartList = { sysA, sysB, sysC }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
    )

Phantom phantom_vxss (
    )
```

```
ProcessOnOnly vxatd (
    IgnoreArgs = 1
    PathName = "/opt/VRTSat/bin/vxatd"
    )

// resource dependency tree
//
//     group VxSS
//     {
//     Phantom phantom_vxss
//     ProcessOnOnly vxatd
//     }
```

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 127.

   See "About the VCS configuration file main.cf" on page 129.

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 46.

3. If you use Sun SCI adapters for your private network, move the scripts `S70llt` and `S92gab` from the directory /etc/rc2.d to directory /etc/rc3.d. so that they are run after the `S19sci` and `S23scid` scripts.
4. Verify LLT operation.

   See "Verifying LLT" on page 135.

5. Verify GAB operation.

   See "Verifying GAB" on page 138.

6. Verify the cluster operation.

   See "Verifying the cluster" on page 139.

Verifying LLT

Use the `lltstat` command to verify that links are active for LLT. If LLT is configured correctly, this command shows all the nodes in the cluster. The command also returns information about the links for LLT for the node on which you typed the command.

Refer to the `lltstat(1M)` manual page for more information.

To verify LLT

1. Log in as superuser on the node galaxy.
2. Run the `lltstat` command on the node galaxy to view the status of LLT.

   ```
   lltstat -n
   ```

   The output on galaxy resembles:

   ```
   LLT node information:
       Node             State            Links
       *0 galaxy          OPEN             2
        1 nebula          OPEN             2
   ```

   Each node has two links and each node is in the OPEN state. The asterisk (*) denotes the node on which you typed the command.

3 Log in as superuser on the node nebula.

4 Run the `lltstat` command on the node nebula to view the status of LLT.

```
lltstat -n
```

The output on nebula resembles:

```
LLT node information:
    Node             State           Links
     0 galaxy          OPEN           2
    *1 nebula          OPEN           2
```

5 To view additional information about LLT, run the `lltstat -nvv` command on each node.

For example, run the following command on the node galaxy in a two-node cluster:

```
lltstat -nvv | more
```

The output on galaxy resembles the following:

- For Solaris SPARC:

```
Node          State      Link    Status      Address
*0 galaxy      OPEN
                         qfe:0 UP     08:00:20:93:0E:34
                         qfe:1 UP     08:00:20:93:0E:34
1 nebula       OPEN
                         qfe:0 UP     08:00:20:8F:D1:F2
                         qfe:1 DOWN
2              CONNWAIT
                         qfe:0 DOWN
                         qfe:1 DOWN
3              CONNWAIT
                         qfe:0 DOWN
                         qfe:1 DOWN
.
.
.

31             CONNWAIT
                         qfe:0 DOWN
                         /dev/qfe:1 DOWN
```

- For Solaris x64:

```
Node            State       Link    Status      Address
*0 galaxy       OPEN
                            e1000g:0 UP     08:00:20:93:0E:34
                            e1000g:1 UP     08:00:20:93:0E:34
1 nebula        OPEN
                            e1000g:0 UP     08:00:20:8F:D1:F2
                            e1000g:1 DOWN
2               CONNWAIT
                            e1000g:0 DOWN
                            e1000g:1 DOWN
3               CONNWAIT
                            e1000g:0 DOWN
                            e1000g:1 DOWN
.
.
.

31              CONNWAIT
                            e1000g:0 DOWN
                            e1000g:1 DOWN
```

Note that the output lists 32 nodes. The command reports the status on the two nodes in the cluster, galaxy and nebula, along with the details for the non-existent nodes.

For each correctly configured node, the information must show the following:

- A state of OPEN
- A status for each link of UP
- An address for each link

However, the output in the example shows different details for the node nebula. The private network connection is possibly broken or the information in the /etc/llttab file may be incorrect.

6 To obtain information about the ports open for LLT, type `lltstat -p` on any node.

For example, type `lltstat -p` on the node galaxy in a two-node cluster:

```
lltstat -p
```

The output resembles:

```
LLT port information:
  Port  Usage       Cookie
  0     gab          0x0
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects:    0 1
  7     gab          0x7
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects:    0 1
  31    gab          0x1F
        opens:       0 2 3 4 5 6 7 8 9 10 11 ... 28 29 30 31
        connects:    0 1
```

Verifying 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 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 *Veritas Cluster Server User'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:

```
GAB Port Memberships
===================================
Port a gen a36e0003 membership 01
Port h gen fd570002 membership 01
```

- 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 *Veritas Cluster Server User's Guide* for a description of system states and the transitions between them.

To verify the cluster

1 To verify the status of the cluster, type the following command:

```
hastatus -summary
```

The output resembles:

```
-- SYSTEM STATE
-- System               State                Frozen

A  galaxy                RUNNING               0
A  nebula                RUNNING               0

-- GROUP STATE
-- Group           System        Probed  AutoDisabled   State

B  ClusterService  galaxy         Y       N              ONLINE
B  ClusterService  nebula         Y       N              OFFLINE
```

2 Review the command output for the following information:

- The system state
  If the value of the system state is RUNNING, VCS is successfully installed and started.
- The ClusterService group state
  In the sample output, the group state lists the ClusterService group, which is ONLINE on galaxy and OFFLINE on nebula.

Verifying the cluster nodes

Verify the information of the cluster systems using the `hasys -display` command. The information for each node in the output should be similar.

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

Refer to the *Veritas Cluster Server User's Guide* for information about the system attributes for VCS.

To verify the cluster nodes

- On one of the nodes, type the `hasys -display` command:

```
hasys -display
```

The example shows the output when the command is run on the node galaxy. The list continues with similar information for nebula (not shown) and any other nodes in the cluster.

Note: The following example is for SPARC. x64 clusters have different command output.

```
#System    Attribute             Value

galaxy     AgentsStopped         0

galaxy     AvailableCapacity     100

galaxy     CPUBinding            BindTo None CPUNumber 0

galaxy     CPUUsage              0

galaxy     CPUUsageMonitoring    Enabled 0 ActionThreshold 0
                                 ActionTimeLimit 0 Action NONE
                                 NotifyThreshold 0 NotifyTimeLimit 0

galaxy     Capacity              100

galaxy     ConfigBlockCount      130

galaxy     ConfigCheckSum        46688

galaxy     ConfigDiskState       CURRENT

galaxy     ConfigFile            /etc/VRTSvcs/conf/config

galaxy     ConfigInfoCnt         0

galaxy     ConfigModDate         Fri May 26 17:22:48 2006

galaxy     ConnectorState        Down

galaxy     CurrentLimits

galaxy     DiskHbStatus

galaxy     DynamicLoad           0

galaxy     EngineRestarted       0

galaxy     EngineVersion         5.0.30.0

galaxy     Frozen                0
```

```
#System    Attribute             Value

galaxy     GUIIPAddr

galaxy     LLTNodeId             0

galaxy     LicenseType           DEMO

galaxy     Limits

galaxy     LinkHbStatus          qfe:0 UP qfe:1 UP

galaxy     LoadTimeCounter       0

galaxy     LoadTimeThreshold     600

galaxy     LoadWarningLevel      80

galaxy     NoAutoDisable         0

galaxy     NodeId                0

galaxy     OnGrpCnt              1

galaxy     ShutdownTimeout       120

galaxy     SourceFile            ./main.cf

galaxy     SysInfo               Solaris:galaxy,Generic_
                                 118558-11,5.9,sun4u

galaxy     SysName               galaxy

galaxy     SysState              RUNNING

galaxy     SystemLocation

galaxy     SystemOwner

galaxy     TFrozen               0

galaxy     TRSE                  0

galaxy     UpDownState           Up

galaxy     UserInt               0

galaxy     UserStr

galaxy     VCSFeatures           DR

galaxy     VCSMode               VCS
```

Chapter 8

Upgrading VCS

This chapter includes the following topics:

About VCS 5.0 MP3 upgrade

Upgrade to VCS 5.0 MP3 with the installvcs program or the installmp programs depending on the version of VCS that you use. You also have the option to do a typical upgrade or a minimal downtime upgrade. This chapter presents both types of upgrades.

See "VCS supported upgrade paths" on page 145.

You can also find information on a special upgrade scenario when the zone root is on shared storage.

See "Special upgrading scenario" on page 167.

VCS supported upgrade paths

Review the following information to help you decide the programs that you need to use to upgrade VCS.

Table 8-1 lists the supported upgrade paths for Solaris SPARC.

Table 8-1 Supported upgrade paths for Solaris SPARC

From	To	Upgrade program to use
VCS 4.0 MP2 VCS 4.1 MP2	VCS 5.0 MP3	installvcs program
Language support for the following: ■ VCS 4.0 MP2 ■ VCS 4.1 MP2	VCS 5.0 MP3 with language support	installvcs program, install_lp, and then installmlp
VCS 5.0 VCS 5.0 MP1	VCS 5.0 MP3	installmp
Language support for the following: ■ VCS 5.0 ■ VCS 5.0 MP1	VCS 5.0 MP3 with language support	installmp and then installmlp

Table 8-2 lists the supported upgrade paths for the Solaris x64 Platform Edition.

Table 8-2 Supported upgrade paths for Solaris x64 Platform Edition

From	To	Upgrade program to use
VCS 4.1 Phase 2	VCS 5.0 MP3	installvcs program
VCS 5.0	VCS 5.0 MP3	installmp

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). Run the installvcs program on each node to upgrade the cluster to VCS 5.0 MP3. On the first node, the program updates the configuration and stops the cluster before you upgrade the system. On the other nodes, it uninstalls the previous version and installs VCS 5.0 MP3. After the last node is upgraded and started, the upgrade is complete.

About minimal downtime upgrade

Use a minimal downtime upgrade to upgrade VCS. This procedure minimizes downtime for the cluster that you want to upgrade. Depending on the situation, you can calculate the approximate downtime as follows:

You can fail over all your service groups to the nodes that are up.	Downtime equals the time that is taken to offline and online the service groups.
You have a service group that you cannot fail over to a node that runs during upgrade.	Downtime for that service group equals the time that is taken to perform an upgrade and restart the node.

Prerequisites for a minimal downtime upgrade

Before you start the upgrade, confirm that you have licenses for all the nodes that you plan to upgrade.

Planning for the minimal downtime upgrade

Plan out the movement of the service groups from node-to-node to minimize the downtime for any particular service group.

Some rough guidelines follow:

- Split the cluster in half. If the cluster has an odd number of nodes, calculate (n+1)/2, and start the upgrade with the even number of nodes.
- Split the cluster so that your high priority service groups remain online during the upgrade of the first subcluster.

Minimal downtime upgrade limitations

The following limitations primarily describe not to tamper with configurations or service groups during the minimum downtime upgrade:

- While you perform the upgrades, do not choose any configuration options.
- 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.

- 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.

Minimal downtime upgrade example

In this example, you have four nodes: node01, node02, node03, and node04. You also have four service groups: sg1, sg2, sg3, and sg4. Each service group is running on one node as follows:

- node01 runs sg2.
- node02 runs sg1.
- node03 runs sg4.
- node04 runs sg3.

In your system list, you have each service group that fails over to one other node as follows:

- sg1 can fail over between node01 and node02.
- sg2 can fail over between node01 and node03.
- sg3 can fail over between node01 and node04.
- sg4 can fail over between node03 and node04.

Figure 8-1 shows four nodes, four service groups, and their failover paths.

Figure 8-1 Four nodes, four service groups, and their failover paths

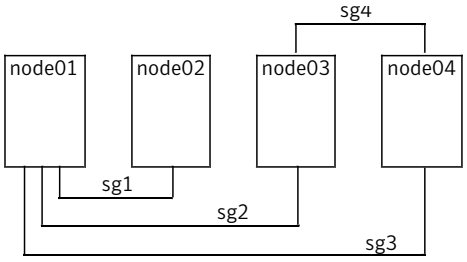

Minimal downtime example overview

This example presumes that you have at least one service group (in this case sg3), that cannot stay online on both nodes during the upgrade. In this situation, sg3 must be a low-priority service group. The cluster is split with node02 and node03 together for the first upgrade, and node01 and node04 together for the next upgrade.

You switch sg1 to run on node01. Switch sg4 to run on node04. You then perform the upgrade on node02 and node03. When you finish the upgrade on node02 and node03, you need to upgrade node01 and node04.

Your cluster is down when you stop HAD on node01 and node04, but have not yet started node02 and node03.

You have to take your service groups offline manually on node01 and node04. When you start node02 and node03, the service groups come online. Restart node01 and node04 when the upgrade completes. They then rejoin the cluster and you can balance the load on systems by switching service groups.

About changes to VCS bundled agents

Review the changes to VCS bundled agents if you upgrade to VCS 5.0 MP3.

Deprecated agents

The following agents are no longer supported:

- CampusCluster
- CFSQlogckd
- ClusterMonitorConfig
- Disk
- DiskReservation
- NFSLock—Use the NFSRestart agent to provide high availability to NFS record locks.
- Service group heartbeat (ServiceGroupHB)—VCS does not support service group heartbeats in this release. Symantec recommends using I/O fencing.

Removing deprecated resource types

With VCS 5.0 MP3, certain resource type definitions are no longer used. Before you start the upgrade process, you must remove the resources of the deprecated resource types from your cluster configuration.

The list of resource types that are not used in VCS 5.0 MP3 are as follows:

- CampusCluster
- ClusterMonitorConfig
- Disk
- DiskReservation
- NFSLock
- ServiceGroupHB

Note: The ClusterConnectorConfig resource type has replaced the ClusterMonitorConfig resource type.

Review the changes to VCS agents in version 5.0 MP3.

See "About changes to VCS bundled agents" on page 149.

If you use the resource type ServiceGroupHB, Symantec recommends the use of I/O fencing.

VCS 5.0 does not support gabdiskhb. So, the installvcs program removes the gabdiskhb entry from the /etc/gabtab file.

Note: Make sure you start VCS on the local node before starting on the other nodes. This standard ensures that HAD reads the configuration from the local node and updates it on the remaining nodes.

Perform the following steps to remove the deprecated resource types.

To remove the deprecated resource types

1 Save the VCS configuration and stop the VCS engine.

```
# haconf -dump -makero
# hastop -all -force
```

2 Back up the configuration file, main.cf to a location on the cluster node.

3 Edit the main.cf located under /etc/VRTSvcs/conf/config.

Perform the following instructions:

- Remove the resource of the deprecated resource types.
  You must modify the resource dependencies to ensure that the configuration works properly.

- Save the main.cf.
- Reformat the main.cf file.

```
# hacf -cftocmd config
# hacf -cmdtocf config
```

4 Verify the configuration.

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

5 Start VCS on the local node.

6 Start VCS on other nodes.

New agents

The following new agents are in the 5.0 MP3 release:

- DiskGroupSnap—Verifies the configuration and the data integrity in a campus cluster environment.
- LDom—Monitors and manages logical domains on Solaris SPARC.
- Zpool—Monitors ZFS storage pools.
- SambaServer—Monitors the smbd process.
- SambaShare—Use to make a Samba Share highly available or to monitor it.
- NetBios—Use to make the nmbd process highly available or to monitor it.

The following new agents were added in the 5.0 release:

- Apache (now bundled on all platforms)—Provides high availability to an Apache Web server.
- NFSRestart—Provides high availability for NFS record locks.
- ProcessOnOnly—Starts and monitors a user-specified process.
- RemoteGroup—Monitors and manages a service group on another system.
- SANVolume—Monitors volumes in a SAN environment managed using Storage Foundation Volume Server. This agent is not supported.

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for more information on these new agents.

New and modified attributes for VCS 5.0 MP3 agents

Table 8-3 lists the attributes that VCS adds or modifies when you upgrade to VCS 5.0 MP3 from VCS 5.0.

Table 8-3 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 5.0

Agent	New and modified attributes	Default value
Apache		
New attribute		
	SupportedActions	"checkconffile.vfd"
	ContainerType	Zone
	PidFile	
	ContainerName	
	IntentionalOffline	0
DNS		
New attributes		
	SupportedActions	"dig.vfd", "keyfile.vfd", "master.vfd"
	ResRecord	
	CreatePTR	0
	OffDelRR	0
DiskGroup		
New attributes		
	SupportedActions	"license.vfd", "disk.vfd", "udid.vfd", "verifyplex.vfd", "checkudid", "campusplex", "numdisks", "joindg", "splitdg", "getvxvminfo", "volinuse"
	UmountVolumes	0
Mount		
New attribute		

Table 8-3 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 5.0 *(continued)*

Agent	New and modified attributes	Default value
	VxFSMountLock	1
Modified attribute		
	SupportedActions	"mountpoint.vfd", "mounted.vfd", "vxfslic.vfd", "chgmntlock", "mountentry.vfd"
NFSRestart		
New attributes		
	SupportedActions	"lockdir.vfd", "nfsconf.vfd"
Share		
New attributes		
	SupportedActions	"direxists.vfd"

Table 8-4 lists the attributes that VCS adds or modifies when you upgrade to VCS 5.0 MP3 from VCS 4.1.

Table 8-4 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 4.1

Agent	New and modified attributes	Default value
Application		
New attributes		
	SupportedActions	program.vfd, user.vfd, cksum.vfd, getcksum
DiskGroup		
New attributes		
	SupportedActions	"license.vfd", "disk.vfd", "udid.vfd", "verifyplex.vfd", "checkudid", "campusplex", "numdisks", "joindg", "splitdg", "getvxvminfo", "volinuse"
	PanicSystemOnDGLoss	1

Table 8-4 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 4.1 *(continued)*

Agent	New and modified attributes	Default value
	DiskGroupType	Private
	UmountVolumes	0
Modified attributes		
	tempUseFence	Invalid
DNS		
New attributes		
	SupportedActions	"dig.vfd", "keyfile.vfd", "master.vfd"
	ResRecord	
	CreatePTR	0
	OffDelRR	0
IP		
New attributes		
	SupportedActions	"device.vfd", "route.vfd"
	ContainerName	
Modified attribute		
	IfconfigTwice	
IPMultiNIC		
New attributes		
	ContainerName	
Modified attribute		
	IfconfigTwice	
IPMultiNICB		
New attributes		

Table 8-4 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 4.1 *(continued)*

Agent	New and modified attributes	Default value
	ToleranceLimit	1
	MonitorInterval	30
	ContainerName	
Modified attribute		
	DeviceChoice	0
Mount		
New attributes		
	SupportedActions	"mountpoint.vfd", "mounted.vfd", "vxfslic.vfd", "chgmntlock", "mountentry.vfd"
	VxFSMountLock	1
	ContainerName	
Modified attribute		
	SnapUmount	
MultiNICA		
Modified attribute		
	IfconfigTwice	
MultiNICB		
New attributes		
	GroupName	
Modified attributes		
	NoBroadcast	
	Failback	
NFS		

Table 8-4 New and modified attributes for VCS 5.0 MP3 agents for upgrades from VCS 4.1 *(continued)*

Agent	New and modified attributes	Default value
New attributes		
	LockFileTimeout	180
NIC		
New attributes		
	SupportedActions	"device.vfd"
Process		
New attribute		
	SupportedActions	"program.vfd", getcksum
Share		
New attribute		
	SupportedActions	"direxists.vfd"

Upgrading to VCS 5.0 MP3

You must upgrade the following to version 5.0 MP3:

VCS

Depending on your upgrade path, use installvcs or installmp program to upgrade VCS.

- See "Upgrading from VCS 4.x" on page 157.
- See "Upgrading from VCS 5.x or later" on page 159.

See "Performing a minimal downtime upgrade to VCS 5.0 MP3" on page 161.

VCS Cluster Manager

See "Upgrading the Cluster Manager (Java Console)" on page 166.

VCS Simulator

See "Upgrading the VCS Simulator" on page 167.

Upgrading from VCS 4.x

If you have a VCS cluster with a version earlier than 5.0, run the installvcs program to upgrade to VCS 5.0 MP3.

See "VCS supported upgrade paths" on page 145.

To perform pre-upgrade tasks

1. Review required patches.

   See "Required patches" on page 23.

2. Log on as superuser on one of the systems for installation.
3. Before you upgrade VCS, Symantec recommends that you back up the types.cf and main.cf configuration files.
4. Stop the application agents that are installed on the VxVM disk (for example the NBU agent). Perform the following steps to stop the application agents.
   - Take the resources offline on all systems that you want to upgrade.

     ```
     # hares -offline resname -sys sysname
     ```

   - Stop the application agents that are installed on VxVM disk on all the systems.

     ```
     # haagent -stop agentname -sys sysname
     ```

   - Ensure that the agent processes are not running.

     ```
     # ps -ef | grep Agent
     ```

   This command does not list any processes in the VxVM installation directory.
5. If you have a zone environment, prepare it for upgrade.

   See "Preparing zone environments" on page 48.

6 Remove the deprecated resource types.

See "Deprecated agents" on page 149.

7 Make sure that LLT, GAB, and VCS are running on all of the nodes in the cluster. The installvcs program cannot proceed unless these processes are running.

```
# lltconfig
LLT is running
# gabconfig -a
=======================================================
Port a gen    cc701 membership 01
Port h gen    cc704 membership 01
```

If any of these are not running, refer to the *Veritas Cluster Server User's Guide* for instructions on how to start them.

To upgrade to VCS 5.0 MP3 using installvcs

1 Insert the disc that contains the 5.0 MP3 software into the disc drive of one of the cluster nodes.

2 Mount the disc on a suitable mount point.

3 Navigate to the directory that contains the installvcs program.

4 Make sure that you have saved any changes to your configuration.

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

5 Enter the following command to start the VCS upgrade:

```
# ./installvcs [-rsh]
```

6 After the initial system checks and the requirements checks are complete, press Return to start upgrading the packages.

7 When the installation is complete, note the locations of the summary, log, and response files.

8 If you see a message that states, "some processes failed to start," or that "GAB or LLT cannot be stopped or unloaded successfully," disregard them.

Execute the following command to restart your systems:

```
# /usr/sbin/shutdown -y -i6 -g0
```

To perform post-upgrade tasks

- If you used the AllowNativeCliUsers attribute before you upgraded VCS, you must use the halogin utility now.

Upgrading the language pack from 4.x

You can now run the install_lp and installmlp programs to upgrade the language pack.

To upgrade language support packages and patches

1. Insert the language disc into the disc drive.
2. Change to the /cdrom/cdrom0 directory:

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

3. Install the language packages:

```
# ./install_lp
```

4. Install the language patches:

```
# ./installmlp
```

Upgrading from VCS 5.x or later

If you are currently running a VCS cluster with VCS 5.0 or later, you must run the installmp program to upgrade to VCS 5.0 MP3.

See "VCS supported upgrade paths" on page 145.

To perform pre-upgrade tasks

1. Review required patches.

   See "Required patches" on page 23.
2. Log on as superuser on one of the systems for installation.
3. Before you upgrade VCS, Symantec recommends that you back up the types.cf and main.cf configuration files.
4. Stop the application agents that are installed on the VxVM disk (for example, NBU agent). Perform the following steps to stop the application agents:
   - Take the resources offline on all systems that you want to upgrade.

     ```
     # hares -offline resname -sys sysname
     ```

- Stop the application agents that are installed on VxVM disk on all the systems.

  ```
  # haagent -stop agentname -sys sysname
  ```

- Ensure that the agent processes are not running.

  ```
  # ps -ef | grep Agent
  ```

  This command does not list any processes in the VxVM installation directory.

5 On each node that you want to upgrade, run the `had -v` command. Run the command to confirm that each node that you plan to upgrade is version 5.0 or later.

```
# had -v
```

Review the output, and confirm that it is 5.0 or later.

6 If you have a zone environment, prepare it for upgrade.

See "Preparing zone environments" on page 48.

To upgrade to VCS 5.0 MP3 using installmp

1 Insert the disc that contains the 5.0 MP3 software into the disc drive of one of the cluster nodes.

2 Mount the disc on a suitable mount point.

3 Navigate to the directory that contains the installmp program.

4 Make sure that you have saved any changes to your configuration.

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

5 Enter the following command to start the VCS upgrade:

```
# ./installmp [-rsh]
```

6 When you are prompted, enter the names of the nodes that you want to upgrade.

7 After the initial system checks and the requirements checks are complete, press Return to start upgrading the packages.

8 When the installation is complete, note the locations of the summary, log, and the response files.

9 Execute the following command to restart your nodes:

```
# /usr/sbin/shutdown -y -i6 -g0
```

Upgrading the language pack from 5.x

You can now run the installmlp program to upgrade the language pack.

To upgrade language support packages and patches

1 Insert the language disc into the disc drive.

2 Change to the /cdrom/cdrom0 directory:

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

3 Install the language patches:

```
# ./installmlp
```

Performing a minimal downtime upgrade to VCS 5.0 MP3

Perform a minimal downtime upgrade in the following phases:

- Select a first group of one or more cluster nodes as target nodes to upgrade now. Leave a group of one or more nodes online to upgrade later.

  Upgrade the target nodes as follows:

  - Switch the service groups from the nodes that you plan to upgrade now to the nodes that you plan to upgrade later.
  - Install the maintenance patches.
  - Restart the target nodes.
  - Upgrade the remaining nodes in the second group.
  - Bring the service groups online on the nodes that you upgraded last.

Performing the pre-upgrade tasks

Perform the following procedure to prepare for the upgrade.

To perform pre-upgrade tasks

1. Select a node or a group of nodes in the cluster as the nodes that you want to upgrade first.
2. Log on as superuser on one of the target nodes for the upgrade.
3. Verify that /opt/VRTS/bin is set in your PATH environment variable to execute all product commands.
4. Back up the llttab, llthosts, gabtab, types.cf, and main.cf files.

```
# cp /etc/llttab /etc/llttab.bkp
# cp /etc/llthosts /etc/llthosts.bkp
# cp /etc/gabtab /etc/gabtab.bkp
# cp /etc/VRTSvcs/conf/config/main.cf \
/etc/VRTSvcs/conf/config/main.cf.bkp
# cp /etc/VRTSvcs/conf/config/types.cf \
/etc/VRTSvcs/conf/config/types.cf.bkp
```

5. Establish where the service groups are online. On one of the nodes in the cluster at the prompt, enter:

```
# hagrp -state
#Group Attribute System Value
sg1 State node01 |ONLINE|
sg1 State node02 |OFFLINE|
sg2 State node01 |ONLINE|
sg2 State node02 |OFFLINE|
```

   Where you plan to upgrade node01.

6 Switch the service groups to the remaining nodes where you plan to upgrade VCS later.

```
# hagrp -switch service_group -to nodename
```

For example:

```
# hagrp -switch sg1 -to node02
```

7 Verify that the service groups are offline on the target nodes for upgrade.

```
# hagrp -state
#Group Attribute System Value
sg1 State node01 |OFFLINE|
sg1 State node02 |ONLINE|
sg2 State node01 |OFFLINE|
sg2 State node02 |ONLINE|
```

Performing the minimal downtime upgrade

You now perform the upgrade on the selected nodes.

To perform the minimal downtime upgrade

1 Insert the VCS 5.0 MP3 software disc into the disc drive of one of the nodes.

2 Mount the disc on a suitable mount point.

3 Navigate to the directory that contains the installvcs program or the installmp program.

4 Upgrade to VCS 5.0 MP3.

Depending on what version you are upgrade from, use the installvcs program or the installmp program. You must specify the target nodes for the upgrade or the program upgrades all the nodes in the cluster:

- If you want to perform an upgrade from VCS 4.x to VCS 5.0 MP3, use the installvcs program.

```
# ./installvcs [-rsh] node01 node02 ...
```

Where node01 and node02 are the names of the nodes that you want to upgrade.

- If you want to perform an upgrade from VCS 5.x to VCS 5.0 MP3, use the installmp program.

```
# ./installmp [-rsh] node01 node02 ...
```

Where node01 and node02 are the names of the nodes that you want to upgrade.

5 After the initial system checks and the requirements checks are complete, press Return to start the upgrade.

6 When the installation is complete, note the locations of the summary, log, and the response files.

If you see a message that states, "some processes failed to start," or that "GAB or LLT cannot be stopped or unloaded successfully," disregard them.

7 Change the cluster ID in the /etc/llttab file on the nodes that were upgraded. Find the line that contains “set-cluster” and change the cluster ID that follows this keyword. Make sure that the new cluster ID is unique within the LAN.

8 On one of the upgraded nodes, edit the main.cf file to freeze all the service groups. Add the “Frozen=1” line to all the service group definitions. Copy the updated main.cf to all the upgraded nodes.

For example, if the original group's definition is:

```
Group oracle_sg (
SystemList = { node01 = 0, node02 = 1 }
AutoStartList = { node01, node02 }
```

The new group definition, after you add “Frozen = 1" is:

```
Group oracle_sg (
SystemList = { node01 = 0, node02 = 1 }
AutoStartList = { node01, node02 }
Frozen = 1
```

9 Restart the target nodes.

10 Manually seed the cluster with the `gabconfig -cx` command. From one of the nodes that you have upgraded, run the following command:

```
# gabconfig -cx
```

Run the `gabconfig -a` command to see if port a and port h are seeded. Output resembles the following:

```
# gabconfig -a
GAB Port Memberships

===============================

Port a gen   1ebf01 membership 0

Port h gen   1ebf03 membership 0
```

11 While the target nodes come up, upgrade the remaining node or set of nodes in the cluster.

Repeat the step 1 to step 7 on the remaining nodes.

Unfreezing and bringing service groups online

Perform the following on the nodes that you have previously upgraded while you upgrade the final node or set of nodes.

To unfreeze the service groups and bring the online

1 On one of the nodes that you have previously upgraded, make the configuration writable.

```
# haconf -makerw
```

2 Unfreeze the service groups. On an upgraded node, run the following command.

```
# hagrp -unfreeze service_group -persistent
```

3 Bring all the service groups online on the upgraded nodes. On an upgraded node, for each service group run the following command.

```
# hagrp -online service_group -sys nodename
```

4 Save the configuration.

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

5 When the upgrade is complete on the final nodes, restart them.

Execute the following command to restart the nodes:

```
# /usr/sbin/shutdown -y -i6 -g0
```

Upgrading the language pack after a minimal downtime upgrade

You can now run the install_lp and installmlp programs to upgrade the language pack.

To upgrade language support packages and patches

1 Insert the language disc into the disc drive.

2 Change to the /cdrom/cdrom0 directory:

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

3 Install the language packages:

```
# ./install_lp
```

4 Install the language patches:

```
# ./installmlp
```

Upgrading the Cluster Manager (Java Console)

This release includes updates for Cluster Manager (Java Console).

To upgrade the Java Console on a Windows client

1 Stop Cluster Manager if it is running.

2 Remove Cluster Manager from the system using Add or Remove Programs.

3 Insert the software disc into a drive on your Windows system.

4 Start the installer:

- For English, use the following path:
  \windows\VCSWindowsInstallers\ClusterManager\EN\setup.exe

- For supported languages other than English, start the installer from the following path on the language disc:
  \ja\wndows\VCSWindowsInstallers\ClusterManager\JA\setup.exe

5 Follow the wizard instructions to complete the installation.

Upgrading the VCS Simulator

This release includes updates for 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 using Add or Remove Programs.

4 Insert the software disc into a drive on your Windows system.

5 Start the installer:

- For English, use the following path:
  \windows\VCSWindowsInstallers\Simulator \EN\vrtsvcssim.msi
- For supported languages other than English, start the installer from the following path on the language disc:
  \ja\wndows\VCSWindowsInstallersSimulator\JA\setup.exe

6 Follow the wizard instructions to complete the installation.

Special upgrading scenario

When you use installvcs or installmp to upgrade, the installer program handles zone states and ensures a smooth upgrade. Symantec recommends that you use these programs to upgrade.

For a manual upgrade of nodes that have their zone root on Veritas File System (VxFS) shared storage, the proper zone state is mandatory. Failure to have the proper zone state for the nodes causes patchadd to fail. This failure can result in different versions of VCS running in the cluster, which results in an unsupported configuration.

Limitations, prerequisites, and definitions

Review the following information.

Limitation

This upgrade is an unsupported upgrade from VCS 5.x that can result in an unsupported configuration.

Prerequisites

All nodes must run Solaris 10 Update 3 or later.

You should have good to excellent knowledge of VCS and zones to perform these tasks.

Definitions

In the following procedures:

- *inactive_local_zoneroot_mountpoint*
  The name of a mount point in a non-global zone on a file system that you can use for the duration of the upgrade. It is in the installed state on a node and does not have access to zone root on shared storage.
- *inactive_local_zonename*
  The name of the non-global zone. This non-global zone is in the installed state on a node and does not have access to zone root on shared storage.

Upgrading VCS when the zone root is on Veritas File System shared storage

The following procedures are to make one active non-global zone upgradeable with the zone root on shared storage. The corresponding non-global zones on the other nodes in the cluster are then detached from shared storage. They are detached to prevent them from being upgraded one at a time.

Stopping the cluster and upgrading nodeA

Stop the cluster and upgrade nodeA.

To stop the cluster and upgrade nodeA

1 Stop the cluster. On nodeA in the cluster, run the following command:

```
# hastop -all
```

2 On nodeA, bring up the volumes and the file systems that are related to the zone root.

Note: For a faster upgrade, you can boot the zones to bring them into the running state.

---

3 Use the patchadd command to upgrade nodeA.

```
# patchadd nnnnnn-nn
# patchadd xxxxxx-xx
.
.
```

Where *nnnnnn-nn* and *xxxxxx-xx* are the IDs of the patches that you add.

Detaching the zones on nodeB - nodeN

Use a mount point as a temporary zone root directory. You then detach the non-global zones in the cluster that are in the installed state. Detach them to prevent the operating system from trying to upgrade these zones and failing.

To detach an inactive non-global zone

1 Change a temporary zone root file system directory's permission to 700 (if it is not already 700).

```
# chmod 700 inactive_local_zoneroot_mountpoint
```

2 Detach the inactive non-global zone.

```
# zoneadm -z inactive_local_zonename detach
```

3 Change the temporary zone root file system directory's permission to 700 and detach all other non-global zones that are in the installed state.

Upgrading nodeB - nodeN

Use the `patchadd` command to upgrade nodeB - nodeN (the global zones).

To upgrade nodeB - nodeN

- Use the `patchadd` command to upgrade nodeB - nodeN (the global zones).

```
# patchadd nnnnnn-nn
# patchadd xxxxxx-xx
.
.
```

Where *nnnnnn-nn* and *xxxxxx-xx* are the IDs of the patches that you add.

Removing files from the temporary file system that was used for the zone root

After you restart the nodes, mount the temporary file system for the zone root.

To manually mount the Veritas File System after restart

1. Reattach the inactive non-global zone.

```
# zoneadm -z inactive_local_zonename attach -F
```

2. Remove the files that are in the *inactive_local_zoneroot_mountpoint* directory.
3. Verify the VCS configuration and its status is the same as before upgrade.

```
# hastatus -sum
```

Chapter 9

Adding and removing cluster nodes

This chapter includes the following topics:

About adding and removing nodes

After you install VCS and create a cluster, you can add and remove nodes from the cluster. You can create a cluster of up to 32 nodes.

Adding a node to a cluster

The system you add to the cluster must meet the hardware and software requirements.

See "Hardware requirements" on page 21.

Table 9-1 specifies the tasks that are involved in adding a cluster. The example demonstrates how to add a node saturn to already existing nodes, galaxy and nebula.

Table 9-1 Tasks that are involved in adding a node to a cluster

Task	Reference
Set up the hardware.	See "Setting up the hardware" on page 172.

Table 9-1 Tasks that are involved in adding a node to a cluster *(continued)*

Task	Reference
Install the software manually and add a license key.	See "Installing the VCS software manually when adding a node" on page 173.
For a cluster that is running in secure mode, verify the existing security setup on the node.	See "Setting up the node to run in secure mode" on page 174.
Configure LLT and GAB.	See "Configuring LLT and GAB" on page 176.
Add the node to the existing cluster.	See "Adding the node to the existing cluster" on page 178.
Start VCS and verify the cluster.	See "Starting VCS and verifying the cluster" on page 179.

Setting up the hardware

Figure 9-1 shows that before you configure a new system on an existing cluster, you must physically add the system to the cluster.

Figure 9-1 Adding a node to a three-node cluster using two independent hubs

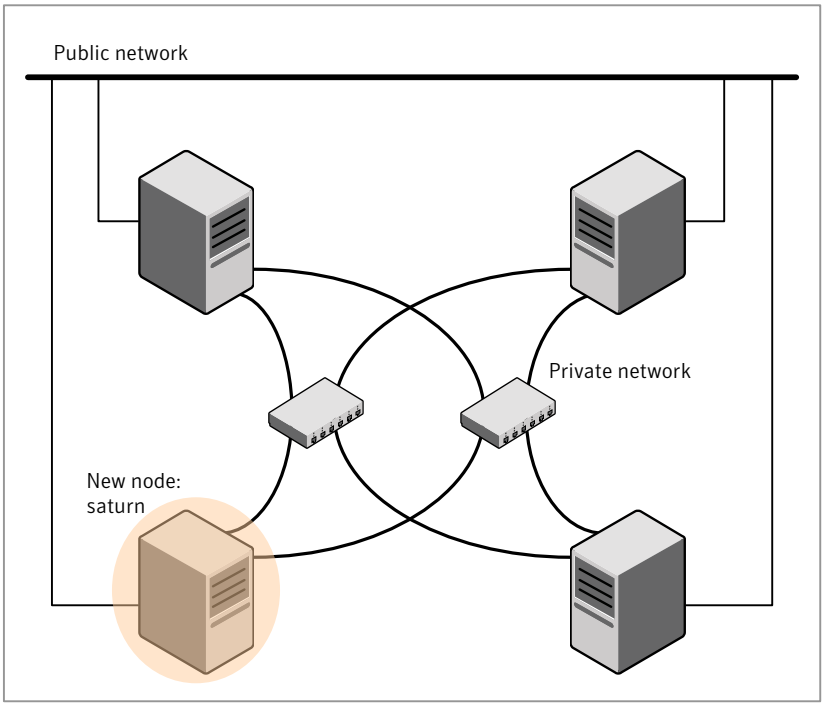

To set up the hardware

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 9-1 illustrates a new node being added to an existing three-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 5.0 MP3 packages manually and add a license key.

For more information, see the following:

- See "Installing VCS software manually" on page 88.
- See "Adding a license key for a manual installation" on page 97.

Setting up the node to run in secure mode

You must follow this procedure only if you are adding a node to a cluster that is running in secure mode. If you are adding a node to a cluster that is not running in a secure mode, proceed with configuring LLT and GAB.

See "Configuring LLT and GAB" on page 176.

Table 9-2 uses the following information for the following command examples.

Table 9-2 The command examples definitions

Name	Fully-qualified host name (FQHN)	Function
saturn	saturn.nodes.example.com	The new node that you are adding to the cluster.
RB1	RB1.brokers.example.com	The root broker for the cluster
RB2	RB2.brokers.example.com	Another root broker, not the cluster's RB

To verify the existing security setup on the node

1. If node saturn is configured as an authentication broker (AB) belonging to a root broker, perform the following steps. Else, proceed to configuring the authentication broker on node saturn.

   See "Configuring the authentication broker on node saturn" on page 175.

2. Find out the root broker to which the node saturn belongs using the following command.

```
# vssregctl -l -q -b \
"Security\Authentication\Authentication Broker" \
 -k "BrokerName"
```

3 If the node saturn already belongs to root broker RB1, it is configured as part of the cluster. Proceed to setting up VCS related security configuration.

See "Setting up VCS related security configuration" on page 176.

4 If the node saturn belongs to a different root broker (for example RB2), perform the following steps to remove the security credentials from node saturn.

- Kill /opt/VRTSat/bin/vxatd process.
- Remove the credential that RB2 has given to AB on node saturn.

```
# vssat deletecred --domain type:domainname \
--prplname prplname
```

For example:

```
# vssat deletecred --domain vx:root@RB2.brokers.example.com \
--prplname saturn.nodes.example.com
```

Configuring the authentication broker on node saturn

Configure a new authentication broker (AB) on node saturn. This AB belongs to root broker RB1.

To configure the authentication broker on node saturn

1 Create a principal for node saturn on root broker RB1. Execute the following command on root broker RB1.

```
# vssat addprpl --pdrtype root --domain domainname \
 --prplname prplname --password password \
 --prpltype service
```

For example:

```
# vssat addprpl --pdrtype root \
--domain root@RB1.brokers.example.com \
--prplname saturn.nodes.example.com \
--password flurbdicate --prpltype service
```

2 Ensure that there is no clock skew between the times on node saturn and RB1.

3 Copy the /opt/VRTSat/bin/root_hash file from RB1 to node saturn.

4 Configure AB on node saturn to talk to RB1.

```
# vxatd -o -a -n prplname -p password -x vx -y domainname -q \
rootbroker -z 2821 -h roothash_file_path
```

For example:

```
# vxatd -o -a -n saturn.nodes.example.com -p flurbdicate \
-x vx -y root@RB1.brokers.example.com -q RB1 \
-z 2821 -h roothash_file_path
```

5 Verify that AB is configured properly.

```
# vssat showbrokermode
```

The command should return 1, indicating the mode to be AB.

Setting up VCS related security configuration

Perform the following steps to configure VCS related security settings.

Setting up VCS related security configuration

1 Start /opt/VRTSat/bin/vxatd process.

2 Create HA_SERVICES domain for VCS.

```
# vssat createpd --pdrtype ab --domain HA_SERVICES
```

3 Add VCS and webserver principal to AB on node saturn.

```
# vssat addprpl --pdrtype ab --domain HA_SERVICES --prplname
webserver_VCS_prplname --password new_password --prpltype
service --can_proxy
```

4 Create /etc/VRTSvcs/conf/config/.secure file.

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

Configuring LLT and GAB

Create the LLT and GAB configuration files on the new node and update the files on the existing nodes.

To configure LLT

1 Create the file /etc/llthosts on the new node. You must also update it on each of the current nodes in the cluster.

For example, suppose you add saturn to a cluster consisting of galaxy and nebula:

- If the file on one of the existing nodes resembles:

```
0 galaxy
1 nebula
```

- Update the file for all nodes, including the new one, resembling:

```
0 galaxy
1 nebula
2 saturn
```

2 Create the file /etc/llttab on the new node, making sure that line beginning "`set-node`" specifies the new node.

The file /etc/llttab on an existing node can serve as a guide.

The following example describes a system where node saturn is the new node on cluster number 2:

- For Solaris SPARC:

```
set-node saturn
set-cluster 2
link qfe0 qfe:0 - ether - -
link qfe1 qfe:1 - ether - -
```

- For Solaris x64:

```
set-node saturn
set-cluster 2
link e1000g0 e1000g:0 - ether - -
link e1000g1 e1000g:1 - ether - -
```

3 On the new system, run the command:

```
# /sbin/lltconfig -c
```

To configure GAB

1 Create the file /etc/gabtab on the new system.

- If the /etc/gabtab file on the existing nodes resembles:

```
/sbin/gabconfig -c
```

The file on the new node should be the same. Symantec recommends that you use the `-c -n`*N* option, where *N* is the 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 On the new node, run the command, to configure GAB:

```
# /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 138.

2 Run the same command on the other nodes (galaxy and nebula) to verify that the port a membership includes the new node:

```
# /sbin/gabconfig -a
GAB Port Memberships
===================================
Port a gen a3640003 membership 012
Port h gen fd570002 membership 01
Port h gen fd570002    visible ; 2
```

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 saturn
```

3 Stop VCS on the new node:

```
# hastop -sys saturn
```

4 Copy the main.cf file from an existing node to your new node:

```
# rcp /etc/VRTSvcs/conf/config/main.cf \
  saturn:/etc/VRTSvcs/conf/config/
```

5 Start VCS on the new node:

```
# hastart
```

6 If necessary, modify any new system attributes.

7 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 From the new system, start VCS with the new system added to the cluster:

```
# 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
```

Removing a node from a cluster

Table 9-3 specifies the tasks that are involved in removing a node from a cluster. In the example procedure, the cluster consists of nodes A, B, and C; node C is to leave the cluster.

Table 9-3 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 181.
■ 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 182.
Modify the llthosts and gabtab files to reflect the change.	See "Modifying configuration files on each remaining node" on page 185.
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 185.

Table 9-3 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 packages.	See "Unloading LLT and GAB and removing VCS on the departing node" on page 186.

Verifying the status of nodes and service groups

Start by issuing the following commands from one of the nodes to remain, node A or node B.

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  A           RUNNING        0
A  B           RUNNING        0
A  C           RUNNING        0

-- GROUP STATE
-- Group    System      Probed   AutoDisabled   State
B  grp1     A           Y        N              ONLINE
B  grp1     B           Y        N              OFFLINE
B  grp2     A           Y        N              ONLINE
B  grp3     B           Y        N              OFFLINE
B  grp3     C           Y        N              ONLINE
B  grp4     C           Y        N              ONLINE
```

The example output from the `hastatus` command shows that nodes A, B, and C are the nodes in the cluster. Also, service group grp3 is configured to run on node B and node C, the departing node. Service group grp4 runs only on node C. Service groups grp1 and grp2 do not run on node C.

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 C to node B.

```
# hagrp -switch grp3 -to B
```

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 C
```

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  A           RUNNING        0
    A  B           RUNNING        0
    A  C           EXITED         0

    -- GROUP STATE
    -- Group    System      Probed   AutoDisabled   State
    B  grp1     A           Y        N              ONLINE
    B  grp1     B           Y        N              OFFLINE
    B  grp2     A           Y        N              ONLINE
    B  grp3     B           Y        N              ONLINE
    B  grp3     C           Y        Y              OFFLINE
    B  grp4     C           Y        N              OFFLINE
```

6 Delete the departing node from the SystemList of service groups grp3 and grp4.

```
# hagrp -modify grp3 SystemList -delete C
# hagrp -modify grp4 SystemList -delete C
```

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  A           RUNNING        0
    A  B           RUNNING        0
    A  C           EXITED         0

    -- GROUP STATE
    -- Group    System      Probed   AutoDisabled   State
    B  grp1     A           Y        N              ONLINE
    B  grp1     B           Y        N              OFFLINE
    B  grp2     A           Y        N              ONLINE
    B  grp3     B           Y        N              ONLINE
```

10 Delete the node from the cluster.

```
# hasys -delete C
```

11 Save the configuration, making it read only.

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

Modifying configuration files on each remaining node

Perform the following tasks on each of the remaining nodes of the cluster.

To modify the configuration files on a remaining node

1 If necessary, modify the /etc/gabtab file.

No change is required to this file if the `/sbin/gabconfig` command has only the argument `-c`. Symantec recommends using the `-nN` option, where *N* is the number of cluster systems.

If the command has the form `/sbin/gabconfig -c -nN`, where *N* is the number of cluster systems, make sure that *N* is not greater than the actual number of nodes in the cluster. When *N* is greater than the number of nodes, GAB does not automatically seed.

Note: Symantec does not recommend the use of the `-c -x` option for `/sbin/gabconfig`. The Gigabit Ethernet controller does not support the use of `-c -x`.

---

2 Modify /etc/llthosts file on each remaining nodes to remove the entry of the departing node.

For example, change:

```
0 A
1 B
2 C
```

To:

```
0 A
1 B
```

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 C. Perform the following steps.

To remove the security credentials

1 Kill /opt/VRTSat/bin/vxatd process.

2 Remove the root credentials on node C.

```
# vssat deletecred --domain type:domainname --prplname prplname
```

Unloading LLT and GAB and removing VCS on the departing node

Perform the tasks on the node that is departing the cluster.

To unconfigure and unload LLT and GAB and remove VCS

1 Unconfigure GAB and LLT:

```
# /sbin/gabconfig -U
# /sbin/lltconfig -U
```

2 Unload the GAB and LLT modules from the kernel.

- Determine the kernel module IDs:

```
# modinfo | grep gab
# modinfo | grep llt
```

  The module IDs are in the left-hand column of the output.

- Unload the module from the kernel:

```
# modunload -i gab_id
# modunload -i llt_id
```

3 Rename the startup files to prevent LLT, GAB, or VCS from starting up in the future:

```
# mv /etc/rc2.d/S70llt /etc/rc2.d/s70llt
# mv /etc/rc2.d/S92gab /etc/rc2.d/s92gab
# mv /etc/rc3.d/S99vcs /etc/rc3.d/s99vcs
```

4 To determine the packages to remove, enter:

```
# pkginfo | grep VRTS
```

5 To permanently remove the VCS patches from the system, use the `patchrm` command for the following versions of Solaris:

- For Solaris 8, perform the following commands:

```
# patchrm 123983-01
# patchrm 123984-01
# patchrm 125150-07
# patchrm 123207-03
# patchrm 123722-01
# patchrm 137338-01
```

```
# patchrm 127333-01
```

- For Solaris 9, perform the following commands:

```
# patchrm 123983-01
# patchrm 123984-01
# patchrm 125150-07
# patchrm 123208-03
# patchrm 123722-01
# patchrm 137338-01
# patchrm 127333-01
```

- For Solaris 10, perform the following commands:

```
# patchrm 123983-01
# patchrm 123984-01
# patchrm 123211-03
# patchrm 123210-03
# patchrm 123209-03
# patchrm 123722-01
# patchrm 137338-01
# patchrm 127333-01
```

- For x64, perform the following commands:

```
# patchrm 128050-03
# patchrm 137384-03
# patchrm 128048-03
# patchrm 128049-03
# patchrm 137339-01
# patchrm 137388-01
```

6 To permanently remove the VCS packages from the system, use the `pkgrm` command. Start by removing the following packages, which may have been optionally installed, in the order shown:

```
# pkgrm VRTScmccc VRTScmcs VRTSacclib VRTScssim
VRTScscm VRTSweb VRTScscw VRTScutil VRTSjre15
VRTSvcsmn VRTSvcsag VRTSvcsmg VRTSvcs VRTSvxfen
VRTSgab VRTSllt VRTSspt VRTSat VRTSpbx
VRTSicsco VRTSvlic VRTSperl
```

7 Remove the LLT and GAB configuration files.

```
# rm /etc/llttab
# rm /etc/gabtab
# rm /etc/llthosts
```

8 Remove the language packages and patches.

See "Removing VCS packages manually" on page 202.

Chapter 10

Installing VCS on a single node

This chapter includes the following topics:

About installing VCS on a single node

You can install VCS 5.0 MP3 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 189.

See "Creating a single-node cluster manually" on page 191.

Creating a single-node cluster using the installer program

Table 10-1 specifies the tasks that are involved to install VCS on a single node using the installer program.

Table 10-1 Tasks to create a single-node cluster using the installer

Task	Reference
Prepare for installation.	See "Preparing for a single node installation" on page 190.
Install the VCS software on the system using the installer.	See "Starting the installer for the single node cluster" on page 190.

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, install 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 15.

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.

See "Starting the software installation" on page 63.

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:
```

Enter a single system name. The installer now 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,?] (n)
```

Answer `n` if you want to use the single node cluster as a stand-alone cluster.

Answer y if you plan to incorporate the single node cluster into a multi-node cluster in the future.

Continue with the installation.

See "Licensing VCS" on page 65.

Creating a single-node cluster manually

Table 10-2 specifies the tasks that you need to perform to install VCS on a single node.

Table 10-2 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 191.
Install the VCS software manually and add a license key	See "Installing the VCS software manually on a single node" on page 191.
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.	See "Renaming the LLT and GAB startup files" on page 192.
Create and modify the VCS configuration files.	See "Configuring VCS" on page 192.
Start VCS and verify single-node operation.	See "Verifying single-node operation" on page 192.

Setting the path variable for a manual single node installation

Set the path variable.

See "Setting the PATH variable" on page 46.

Installing the VCS software manually on a single node

Install the VCS 5.0 MP3 packages and patches manually and install the license key.

Refer to the following sections:

- See "Installing VCS software manually" on page 88.

- See "Adding a license key for a manual installation" on page 97.

Renaming the LLT and GAB startup files

You may need the LLT and GAB startup files to upgrade the single-node cluster to a multiple-node cluster at a later time.

To rename the LLT and GAB startup files

- Rename the LLT and GAB startup files.

```
# mv /etc/rc2.d/S70llt /etc/rc2.d/X70llt
# mv /etc/rc2.d/S92gab /etc/rc2.d/X92gab
```

Configuring VCS

You now need to configure VCS.

See "Configuring VCS" on page 100.

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 ha
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
```

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 10-3 specifies the activities that you need to perform to add nodes to a single-node cluster.

Table 10-3 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 193.
■ 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 194.
Connect both nodes to shared storage.	See "Configuring the shared storage" on page 195.
■ Bring up VCS on Node A. ■ Edit the configuration file.	See "Bringing up the existing node" on page 195.
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 196.
Edit the configuration files on Node B.	See "Creating configuration files" on page 196.
Start LLT and GAB on Node B.	See "Starting LLT and GAB" on page 196.
■ Start LLT and GAB on Node A. ■ Restart VCS on Node A. ■ Modify service groups for two nodes.	See "Reconfiguring VCS on the existing node" on page 197.
■ Start VCS on Node B. ■ Verify the two-node cluster.	See "Verifying configuration on both nodes" on page 198.

Setting up a node to join the single-node cluster

The new node to join the existing single node running VCS must run the same version of operating system and patch level.

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 packages and configuration files.
  See "Removing a node from a cluster" on page 180.
- If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS.
  See "Removing VCS packages manually" on page 202.
- If you renamed the LLT and GAB startup files, remove them.
  See "Renaming the LLT and GAB startup files" on page 192.

2 If necessary, install VxVM and VxFS.

See "Installing VxVM or VxFS if necessary" on page 194.

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 36.

To install and configure Ethernet cards for private network

1 Shut down VCS on Node A.

```
# hastop -local
```

2 Shut down the node to get to the OK prompt:

```
# sync;sync;init 0
```

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 42.

Bringing up the existing node

Bring up the node.

To bring up the node

1 On Node A, enter the command:

```
ok boot -r
```

2 Log in as superuser.

3 Make the VCS configuration writable.

```
# haconf -makerw
```

4 Display the service groups currently configured.

```
# hagrp -list
```

5 Freeze the service groups.

```
# hagrp -freeze group -persistent
```

Repeat this command for each service group in step 4.

6 Make the configuration read-only.

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

7 Stop VCS on Node A.

```
# hastop -local -force
```

8 Rename the GAB and LLT startup files so they can be used.

```
# mv /etc/rc2.d/X92gab /etc/rc2.d/S92gab
# mv /etc/rc2.d/X70llt /etc/rc2.d/S70llt
```

Installing the VCS software manually when adding a node to a single node cluster

Install the VCS 5.0 MP3 packages manually and install the license key.

Refer to the following sections:

- See "Installing VCS software manually" on page 88.
- See "Adding a license key for a manual installation" on page 97.

Creating configuration files

Create the configuration files for your cluster.

To create the configuration files

1 Create the file /etc/llttab that lists both the nodes.

See "Setting up /etc/llttab for a manual installation" on page 98.

2 Create the file /etc/llthosts. Set up /etc/llthosts for a two-node cluster.

See "Setting up /etc/llthosts for a manual installation" on page 98.

3 Create the file /etc/gabtab.

See "Configuring GAB for a manual installation" on page 100.

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 start
```

2 Start GAB on Node B.

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

3 Start GAB on Node A.

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

4 Check the membership of the cluster.

```
# gabconfig -a
```

5 Start VCS on Node A.

```
# hastart
```

6 Make the VCS configuration writable.

```
# haconf -makerw
```

7 Add Node B to the cluster.

```
# hasys -add sysB
```

8 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. Implement the new two-node configuration.

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

Chapter 11

Uninstalling VCS

This chapter includes the following topics:

About the uninstallvcs program

You can uninstall VCS from all nodes in the cluster or from specific nodes in the cluster using the uninstallvcs program. The uninstallvcs program does not automatically uninstall VCS enterprise agents, but offers uninstallation if proper packages dependencies on VRTSvcs are found.

If uninstallvcs program does not remove an enterprise agent, see the documentation for the specific enterprise agent for instructions on how to remove it.

Prerequisites for using the uninstallvcs program

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.
- Before you remove VCS from fewer than all nodes in a cluster, stop the service groups on the nodes from which you uninstall VCS. You must also reconfigure VCS on the remaining nodes.
  See "About adding and removing nodes" on page 171.

- 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 61.

Uninstalling VCS 5.0 MP3

You must meet the following conditions to use the uninstallvcs program 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 program on each node in the cluster.

The example demonstrates how to uninstall VCS using the uninstallvcs program. The uninstallvcs program uninstalls VCS on two nodes: galaxy nebula. The example procedure uninstalls VCS from all nodes in the cluster.

Removing VCS 5.0 MP3 packages

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 program.

```
# cd /opt/VRTS/install
# ./uninstallvcs
```

The program specifies the directory where the logs are created. The program displays a copyright notice and a description of the cluster:

```
VCS configuration files exist on this system with the following
information:

      Cluster Name: VCS_cluster2
      Cluster ID Number: 7
      Systems: galaxy nebula
      Service Groups: ClusterService groupA groupB
```

3 Answer the prompt to proceed with uninstalling the software.

Select one of the following:

- To uninstall VCS on all nodes, press `Enter`.
- To uninstall VCS only on specific nodes, enter `n`.

```
Do you want to uninstall VCS from these systems? [y,n,q] (y)
```

4 If the uninstallvcs program prompts, enter a list of nodes from which you want to uninstall VCS.

The uninstallvcs program prompts this information in one of the following conditions:

- You enter `n` .
- The program finds no VCS configuration files on the local node.

5 Review the output as the uninstallvcs program continues to do the following:

- Verifies the communication between systems
- Checks the installations on each system to determine the packages to be uninstalled

6 If packages, such as enterprise agents, are found to be dependent on a VCS package, the uninstaller prompt you on whether you want them removed. Enter `y` to remove the designated packages.

7 Review the uninstaller report after the verification.

8 Press Enter to uninstall the VCS packages.

```
Are you sure you want to uninstall VCS packages? [y,n,q] (y)
```

9 Review the output as the uninstaller stops processes, unloads kernel modules, and removes the packages.

10 Note the location of summary and log files that the uninstaller creates after removing all the packages.

Running uninstallvcs from the VCS 5.0 MP3 disc

You may need to use the uninstallvcs program on the VCS 5.0 MP3 disc in one of the following cases:

- You need to uninstall VCS after an incomplete installation.
- The uninstallvcs program is not available in /opt/VRTS/install.

Removing VCS packages manually

You must remove the VCS packages from each node in the cluster to uninstall VCS.

To manually remove VCS packages on a node

1 Shut down VCS on the local system using the `hastop` command.

```
# hastop -local
```

2 Unconfigure the GAB and the LLT utilities.

```
# /sbin/gabconfig -U
# /sbin/lltconfig -U
```

3 Determine the GAB kernel module ID:

```
# modinfo | grep gab
```

The module ID is in the left-hand column of the output.

4 Unload the GAB module from the kernel:

```
# modunload -i gab_id
```

5 Determine the LLT kernel module ID:

```
# modinfo | grep llt
```

The module ID is in the left-hand column of the output.

6 Unload the LLT module from the kernel:

```
# modunload -i llt_id
```

7 For the language pack, remove the VCS patches:

- Remove the following patch for the VRTSjacmc package.

```
# patchrm 123982-01
```

- Remove the following required patches for Solaris 10:

```
# patchrm 123978-03
# patchrm 123977-03
# patchrm 123680-05
```

- Remove the following required patches for Solaris 9:

```
# patchrm 123976-03
# patchrm 123680-05
```

- Remove the following required patches for Solaris 8:

```
# patchrm 123975-03
# patchrm 123680-05
```

8 For the language pack, remove the following required and optional VCS packages:

- Remove the optional packages, in the order shown.

```
# pkgrm VRTSjacsm VRTSjacmc
```

- Remove the following packages in the order shown:

```
# pkgrm VRTSjaweb VRTSjacsu VRTSjacsj VRTSjacs
VRTSatJA VRTSjapbx VRTSjaico VRTSmulic
```

9 Remove the VCS patches. Depending on the version and architecture, use the following:

- For SPARC 8, remove the patches:

```
# patchrm 123983-01
# patchrm 123984-01
# patchrm 125150-07
# patchrm 123207-03
# patchrm 123722-01
# patchrm 137338-01
# patchrm 127333-01
```

- For SPARC 9, remove the patches:

```
# patchrm 123983-01
# patchrm 123984-01
# patchrm 125150-07
# patchrm 123208-03
# patchrm 123722-01
# patchrm 137338-01
# patchrm 127333-01
```

- For SPARC 10, remove the patches:

```
# patchrm 123983-01
# patchrm 123984-01
# patchrm 123211-03
# patchrm 123210-03
# patchrm 123209-03
# patchrm 123722-01
# patchrm 137338-01
# patchrm 127333-01
```

- For x64, remove the patches:

```
# patchrm 128050-03
# patchrm 137384-03
# patchrm 128048-03
# patchrm 128049-03
```

```
# patchrm 137339-01
# patchrm 137388-01
```

10 Remove the VCS 5.0 MP3 packages in the following order:

```
# pkgrm VRTScmcc VRTScmcs VRTScssim VRTScscm
VRTSvcsmn VRTSacclib VRTSweb VRTScscw VRTScutil
VRTSjre15 VRTSvcsag VRTSvcsmg VRTSvcs VRTSvxfen
VRTSgab VRTSllt VRTSspt VRTSat VRTSpbx
VRTSicsco VRTSvlic VRTSperl
```

Appendix A

Upgrading the operating system

This appendix includes the following topics:

Upgrading Solaris versions

An operating system upgrade can take hours to finish. When you upgrade, you typically upgrade one node at a time. Coordinate with your system administrator to plan for the down time of each system. Plan ahead to move service groups to running nodes, while the nodes that you upgrade are down. Planning ahead reduces downtime and ensures availability of services for your customers.

When you upgrade the operating system, you must remove the GAB, LLT, and fencing packages and patches before you upgrade the operating system. Reinstall fencing, GAB, and LLT after upgrading the operating system.

Note: Be sure that you have the Symantec software disc with the VCS software on hand before you begin.

You must upgrade the operating system on each node in the cluster to completely upgrade the Solaris versions in the cluster.

Upgrading Solaris on a node

The tasks that you need to perform when upgrading the Solaris operating system include the following:

- Stopping VCS
- Stopping GAB and LLT and unloading the kernel modules
- Removing packages and patches
- Upgrading Solaris operating system
- Reinstalling fencing, GAB, and LLT from the software disc
- Restarting VCS

To stop VCS

1. Make the VCS configuration writable. On the first system, type:

   ```
   # haconf -makerw
   ```

2. Move all service groups from the node you are plan to upgrade to another system. Keep services from failing over to this server. On the system where you plan to upgrade, type:

   ```
   # hasys -freeze -persistent -evacuate upgrade_server
   ```

3. Check if all service groups and resources are offline on the system and online on the other system. Type:

   ```
   # hastatus -summary
   ```

4 Close the configuration and unload the VCS services on the system that you plan to upgrade. On the system that you plan to upgrade, type:

```
# haconf -dump -makero
# hastop -local
```

5 Confirm that VCS has stopped. On the upgrading system, type:

```
# gabconfig -a
```

Output resembles:

```
GAB Port Memberships
===================================
Port a gen 23dc0001 membership 01
```

Note that the output shows no membership for port h.

To stop GAB and LLT and unload the kernel modules

1 Unconfigure GAB. Type:

```
# gabconfig -U
```

2 Unload the GAB module from the kernel:

- Determine the kernel module ID:

```
# modinfo | grep gab
```

- Unload the module from the kernel:

```
# modunload -i gab_id
```

3 Unconfigure LLT. On each system, type:

```
# lltconfig -U
```

The following message is displayed on the console:

```
lltconfig: this will attempt to stop and reset LLT.
Confirm (y/n)?
```

4 Type `Y` on each system in response to the message.

5 Unload the LLT module from the kernel:

- Determine the kernel module ID:

```
# modinfo | grep llt
```

The module IDs are in the left-hand column of the output.

- Unload the module from the kernel:

```
# modunload -i llt_id
```

To remove the fencing, GAB, and LLT packages and patches

1. On each node, use the `pkgrm` command to remove the fencing, GAB, and LLT packages.

```
# pkgrm VRTSvxfen VRTSgab VRTSllt
```

2. On each node, use the `patchrm` command to remove the patches:

- For Solaris 9:

```
# patchrm 123208-03
```

- For Solaris 8:

```
# patchrm 123207-03
```

To upgrade the operating system

1. Follow the Sun installation guide to upgrade the operating system kernel to the new version of Solaris.
2. As the system comes up, enter single-user mode.

To reinstall fencing, GAB, and LLT from the software disc and restart

1. In single-user mode, log on as superuser on the system that you have upgraded.
2. Check whether the /tmp directory is mounted.

```
# mount
```

3. If the /tmp directory is not mounted, then enter:

```
# mount /tmp
```

4. Create a directory for installation:

```
# mkdir /tmp/install
```

5 Insert the software disc with the VCS software into a system drive where you have upgraded. The Solaris volume-management software automatically mounts the disc as /cdrom/cdrom0. Type the command:

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

6 Copy the compressed package files from the software disc to the temporary directory:

```
# cp -r cluster_server/pkgs/VRTSllt.tar.gz /tmp/install
# cp -r cluster_server/pkgs/VRTSgab.tar.gz /tmp/install
# cp -r cluster_server/pkgs/VRTSvxfen.tar.gz /tmp/install
# cp -r cluster_server/patches/1232* /tmp/install
```

7 If your system does not have the gunzip utility, copy it from the disc:

```
# cp /cdrom_path/gnu/gunzip /tmp/install
```

8 Go to the temporary directory and unzip the compressed package files:

```
# cd /tmp/install
# gunzip VRTSllt.tar.gz
# gunzip VRTSgab.tar.gz
# gunzip VRTSvxfen.tar.gz
# gunzip 1232*.tar.gz
```

The following files are now present in the temporary directory:

```
VRTSgab.tar
VRTSllt.tar
VRTSvxfen.tar
123207-03.tar
123208-03.tar
123209-03.tar
123210-03.tar
123211-03.tar
```

9 Extract the required VCS files from the compressed files:

```
# tar -xvf VRTSllt.tar
# tar -xvf VRTSgab.tar
# tar -xvf VRTSvxfen.tar
# tar -xvf 123207-03.tar
# tar -xvf 123208-03.tar
# tar -xvf 123209-03.tar
# tar -xvf 123210-03.tar
# tar -xvf 123211-03.tar
```

10 Install the LLT, GAB, and fencing packages and patches. As you enter the command, be sure to install the packages and patches in the order shown:

- Install the packages.

  ```
  # pkgadd -d . VRTSllt VRTSgab VRTSvxfen
  ```

- For Solaris 10:

  ```
  # patchadd 123209-03
  # patchadd 123210-03
  # patchadd 123211-03
  ```

- For Solaris 9:

  ```
  # patchadd 123208-03
  ```

- For Solaris 8:

  ```
  # patchadd 123207-03
  ```

11 Bring up the system in multi-user mode:

```
# cd /
# init 3
```

To restart VCS

1 Verify that VCS services are running on the upgraded server. On the upgraded server, type:

```
# ps -ef | grep ha
root   576   1  0 16:54:12 ?    0:02 /opt/VRTSvcs/bin/had
root   578   1  0 16:54:13 ?    0:00 /opt/VRTSvcs/bin/hashadow
```

2 If the VCS services are not running, reload the VCS services. Type:

```
# hastart
```

3 Unfreeze the upgraded server and save the configuration. On the upgraded server, type:

```
# hasys -unfreeze -persistent upgraded_server
# haconf -dump -makero
```

Live Upgrade for VCS

Use Solaris Live Upgrade to perform an operating system upgrade from one disk to another disk on a single node. When you use Solaris Live Upgrade, you can keep a node operational while you upgrade its operating system. You move VCS from one disk to another after you perform the Live Upgrade. When you use Solaris Live Upgrade to upgrade the OS, downtime for your node, and VCS, is the time it takes for a reboot.

When you use VCS with Solaris Live Upgrade, you must pay attention to VCS kernel components such as VRTSllt, VRTSgab, and VRTSvxfen. These components are unique for each operating system version. Before you reboot the target disk, you must remove these packages and patches and re-install them for the upgraded version of the operating system.

Requirements

Before you perform the operating system upgrade, you must have:

- VCS installed and running on Solaris 8, 9, or 10.
- An alternate target boot disk of equal or greater size than your current source boot disk.

Performing Live Upgrade for VCS

The general procedure is to install the Solaris Live Upgrade packages and patches on the running disk. Clone the current operating system onto the new disk. Upgrade the operating system for the clone. Mount the new disk. Remove and re-install the kernel-level packages and patches for the version of the operating system. Migrate VCS to the new disk.

To install the Solaris Live Upgrade packages and patches

1. Format the target disk to have the same size and the same partition as the source disk.
2. Install the Solaris Live Upgrade packages on the current source operating system disk.

   The version of the Solaris Live Upgrade packages must match the version of the operating system that you are upgrading to. For example, if your current source disk has Solaris 9 and you want to upgrade the target boot disk to Solaris 10, install the Live Upgrade packages from the Solaris 10 disk onto your Solaris 9 operating system.

To create and populate the new boot environment by cloning the current operating environment

1. Make sure that your second alternate book disk is the same size and has the same partitions as your current boot environment.
2. Execute the `lucreate` command with the following options to create a new boot environment for the alternate boot disk.

   The `-c` option assigns the specified name to the current boot environment.

   The `-m` option specifies the root slice (/) that you plan to copy to /dev/dsk/c0t1d0s0.

   The `-n` option specifies the name of boot environment.

   For example:

   ```
       # lucreate -c sol_9 -m /:/dev/dsk/c0t1d0s0:ufs -m \
   -:/dev/dsk/c0t1d0s1:swap -n sol_10
   ```

   Output from this command includes naming, analyzing, creating, checking, and populating the new environment.
3. After you set up the disk and create the boot environment, upgrade the operating system in the new boot environment.

To upgrade the new boot environment

1 Execute the `luupgrade` command with following options to upgrade the new boot environment.

The `-u` option specifies upgrading the operating system for the boot environment.

The `-n` option specifies the boot environment to upgrade.

The `-s` option specifies the source for the operating system's image or flash image.

For example:

```
# luupgrade -u -n sol_10 -s /net/vcsinstall/netinstall/2.10
```

Output from this command includes validation and upgrade messages.

2 You now need to remove operating system specific kernel-level packages and patches. You must then re-install the packages and patches that match the operating system version.

To mount the new environment and remove and re-install the kernel-level packages and patches

1 Create a new directory and mount the new disk.

```
# mkdir /tmp/a
# mount /dev/dsk/c0t1d0s0 /tmp/a
```

2 Remove the old packages and patches.

- Remove the old packages.

```
# pkgrm -R /tmp/a VRTSvxfen VRTSgab VRTSllt
```

- Remove the old patches for Solaris 10.

```
# patchrm -R /tmp/a 123211-03
# patchrm -R /tmp/a 123210-03
# patchrm -R /tmp/a 123209-03
```

- Remove the old patches for Solaris 9.

```
# patchrm -R /tmp/a 123208-03
```

- Remove the old patches for Solaris 8.

```
# patchrm -R /tmp/a 123207-03
```

3 Copy the compressed package files from the software disc to the temporary directory:

```
# cp -r cluster_server/pkgs/VRTSllt.tar.gz /tmp/install
# cp -r cluster_server/pkgs/VRTSgab.tar.gz /tmp/install
# cp -r cluster_server/pkgs/VRTSvxfen.tar.gz /tmp/install
# cp -r cluster_server/patches/1232* /tmp/install
```

4 If your system does not have the gunzip utility, copy it from the disc.

```
# cp /cdrom_path/gnu/gunzip /tmp/install
```

5 Go to the temporary directory and unzip the compressed package files.

```
# cd /tmp/install
# gunzip VRTSllt.tar.gz
# gunzip VRTSgab.tar.gz
# gunzip VRTSvxfen.tar.gz
# gunzip 1232*.tar.gz
```

The following files are now present in the temporary directory.

```
VRTSgab.tar
VRTSllt.tar
VRTSvxfen.tar
123207-03.tar
123208-03.tar
123209-03.tar
123210-03.tar
123211-03.tar
```

6 Extract the required VCS files from the compressed files.

```
# tar -xvf VRTSllt.tar
# tar -xvf VRTSgab.tar
# tar -xvf VRTSvxfen.tar
# tar -xvf 123207-03.tar
# tar -xvf 123208-03.tar
# tar -xvf 123209-03.tar
# tar -xvf 123210-03.tar
# tar -xvf 123211-03.tar
```

7 Add the packages and patches that match the version on the operating system on the new boot environment.

- Add the packages.

  ```
  # pkgadd -d . -R /tmp/a VRTSllt VRTSgab VRTSvxfen
  ```

- Add the patches for Solaris 10.

  ```
  # patchadd -R /tmp/a 123209-03
  # patchadd -R /tmp/a 123210-03
  # patchadd -R /tmp/a 123211-03
  ```

- Add the patches for Solaris 9.

  ```
  # patchadd -R /tmp/a 123208-03
  ```

- Add the patches for Solaris 8.

  ```
  # patchadd -R /tmp/a 123207-03
  ```

To stop VCS, boot the new environment for VCS, and verify the new environment

1 Stop VCS.

```
# hastop -all
```

2 Enter the `init` command to select the new disk and start the node, for example:

```
# init 0
OK boot disk1
```

3 Enter the `luactivate` command to verify that your current boot environment is the one that you want.

```
# luactivate
```

In the example, the output is:

```
Sol_10
```

Appendix B

Advanced VCS installation topics

This appendix includes the following topics:

Reconciling major/minor numbers for NFS shared disks

Your configuration may include disks on the shared bus that support NFS. You can configure the NFS file systems that you export on disk partitions or on Veritas Volume Manager volumes. An example disk partition name is `/dev/dsk/c1t1d0s3`. An example volume name is `/dev/vx/dsk/shareddg/vol3`. Each name represents the block device on which the file system is to be mounted.

In a VCS cluster, block devices providing NFS service must have the same major and minor numbers on each cluster node. Major numbers identify required device drivers (such as a Solaris partition or a VxVM volume). Minor numbers identify the specific devices themselves. NFS also uses major and minor numbers to identify the exported file system.

Major and minor numbers must be verified to ensure that the NFS identity for the file system is the same when exported from each node.

Checking major and minor numbers for disk partitions

The following sections describe checking and changing, if necessary, the major and minor numbers for disk partitions used by cluster nodes.

To check major and minor numbers on disk partitions

- Use the following command on all nodes exporting an NFS file system. This command displays the major and minor numbers for the block device.

  ```
  # ls -lL block_device
  ```

  The variable *block_device* refers to a partition where a file system is mounted for export by NFS. Use this command on each NFS file system. For example, type:

  ```
  # ls -lL /dev/dsk/c1t1d0s3
  ```

  Output on Node A resembles:

  ```
  crw-r-----  1 root  sys  32,1 Dec 3 11:50 /dev/dsk/c1t1d0s3
  ```

  Output on Node B resembles:

  ```
  crw-r-----  1 root  sys  32,1 Dec 3 11:55 /dev/dsk/c1t1d0s3
  ```

  Note that the major numbers (32) and the minor numbers (1) match, satisfactorily meeting the requirement for NFS file systems.

To reconcile the major numbers that do not match on disk partitions

1. Reconcile the major and minor numbers, if required. For example, if the output in the previous section resembles the following, perform the instructions beginning step 2:

   Output on Node A:

   ```
   crw-r-----  1 root  sys  32,1 Dec 3 11:50 /dev/dsk/c1t1d0s3
   ```

   Output on Node B:

   ```
   crw-r-----  1 root  sys  36,1 Dec 3 11:55 /dev/dsk/c1t1d0s3
   ```

2. Place the VCS command directory in your path. For example:

   ```
   # export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin
   ```

3 Attempt to change the major number on System B (now 36) to match that of System A (32). Use the command:

```
# haremajor -sd major_number
```

For example, on Node B, enter:

```
# haremajor -sd 32
```

4 If the command succeeds, go to step 8.

5 If the command fails, you may see a message resembling:

```
Error: Preexisiting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

6 Notice that the number 36 (the major number on Node A) is not available on Node B. Run the `haremajor` command on Node B and change it to 128,

```
# haremajor -sd 128
```

7 Run the same command on Node A. If the command fails on Node A, the output lists the available numbers. Rerun the command on both nodes, setting the major number to one available to both.

8 Reboot each system on which the command succeeds.

9 Proceed to reconcile the major numbers for your next partition.

To reconcile the minor numbers that do not match on disk partitions

1 In the example, the minor numbers are `1` and `3` and are reconciled by setting to 30 on each node.

2 Type the following command on both nodes using the name of the block device:

```
# ls -1 /dev/dsk/c1t1d0s3
```

Output from this command resembles the following on Node A:

```
lrwxrwxrwx  1 root  root  83 Dec 3 11:50
 /dev/dsk/c1t1d0s3          -> ../../
 devices/sbus@1f,0/QLGC,isp@0,10000/sd@1,0:d,raw
```

The `device name` (in bold) includes the slash following the word `devices`, and continues to, but does not include, the colon.

3 Type the following command on both nodes to determine the instance numbers that the SCSI driver uses:

```
# grep sd /etc/path_to_inst | sort -n -k 2,2
```

Output from this command resembles the following on Node A:

```
"/sbus@1f,0/QLGC,isp@0,10000/sd@0,0" 0 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@1,0" 1 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@2,0" 2 "sd"
"/sbus@1f,0/QLGC,isp@0,10000/sd@3,0" 3 "sd"
.
.
"/sbus@1f,0/SUNW,fas@e,8800000/sd@d,0" 27 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@e,0" 28 "sd"
"/sbus@1f,0/SUNW,fas@e,8800000/sd@f,0" 29 "sd"
```

In the output, the instance numbers are in the second field.

The instance number that is associated with the device name that matches the name for Node A displayed in step 2, is "1."

4 Compare instance numbers for the device in the output on each node.

After you review the instance numbers, perform one of the following tasks:

- If the instance number from one node is unused on the other— it does not appear in the output of step 3—edit `/etc/path_to_inst`.
  You edit this file to make the second node's instance number similar to the number of the first node.
- If the instance numbers in use on both nodes, edit `/etc/path_to_inst` on both nodes. Change the instance number that is associated with the device name to an unused number. The number needs to be greater than the highest number that other devices use. For example, the output of step 3 shows the instance numbers that all devices use (from 0 to 29). You edit the file `/etc/path_to_inst` on each node and reset the instance numbers to 30.

5 Type the following command to reboot each node on which `/etc/path_to_inst` was modified:

```
# reboot -- -rv
```

Checking the major and minor number for VxVM volumes

The following sections describe checking and changing, if necessary, the major and minor numbers for the VxVM volumes that cluster systems use.

To check major and minor numbers on VxVM volumes

1. Place the VCS command directory in your path. For example:

   ```
   # export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin
   ```

2. To list the devices, use the `ls -lL` *block_device* command on each node:

   ```
   # ls -lL /dev/vx/dsk/shareddg/vol3
   ```

   On Node A, the output may resemble:

   ```
   brw-------   1 root  root  32,43000 Mar 22 16:4 1
   /dev/vx/dsk/shareddg/vol3
   ```

   On Node B, the output may resemble:

   ```
   brw-------   1 root  root  36,43000 Mar 22 16:4 1
   /dev/vx/dsk/shareddg/vol3
   ```

3. Import the associated shared disk group on each node.

4 Use the following command on each node exporting an NFS file system. The command displays the major numbers for `vxio` and `vxspec` that Veritas Volume Manager uses . Note that other major numbers are also displayed, but only `vxio` and `vxspec` are of concern for reconciliation:

```
# grep vx /etc/name_to_major
```

Output on Node A:

```
vxdmp 30
vxio 32
vxspec 33
vxfen 87
vxglm 91
```

Output on Node B:

```
vxdmp 30
vxio 36
vxspec 37
vxfen 87
vxglm 91
```

5 To change Node B's major numbers for `vxio` and `vxspec` to match those of Node A, use the command:

```
haremajor -vx major_number_vxio major_number_vxspec
```

For example, enter:

```
# haremajor -vx 32 33
```

If the command succeeds, proceed to step 8. If this command fails, you receive a report similar to the following:

```
Error: Preexisiting major number 32
These are available numbers on this system: 128...
Check /etc/name_to_major on all systems for
available numbers.
```

6 If you receive this report, use the `haremajor` command on Node A to change the major number (`32/33`) to match that of Node B (`36/37`). For example, enter:

```
# haremajor -vx 36 37
```

If the command fails again, you receive a report similar to the following:

```
Error: Preexisiting major number 36
These are available numbers on this node: 126...
Check /etc/name_to_major on all systems for
available numbers.
```

7 If you receive the second report, choose the larger of the two available numbers (in this example, `128`). Use this number in the `haremajor` command to reconcile the major numbers. Type the following command on both nodes:

```
# haremajor -vx 128 129
```

8 Reboot each node on which `haremajor` was successful.

9 If the minor numbers match, proceed to reconcile the major and minor numbers of your next NFS block device.

10 If the block device on which the minor number does not match is a volume, consult the `vxdg`(1M) manual page. The manual page provides instructions on reconciling the Veritas Volume Manager minor numbers, and gives specific reference to the `reminor` option.

Node where the vxio driver number have been changed require rebooting.

Using the UDP layer for LLT

VCS 5.0 MP3 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.

Note: LLT over UDP is not supported on IPv6.

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 that LLT over Ethernet. Only use LLT over UDP when the hardware configuration makes it necessary.

Configuring LLT over UDP

The following checklist is to configure LLT over UDP:

- Make sure that the LLT private links are on different physical networks.
  If the LLT private links are not on different physical networks, then make sure that the links are on separate subnets. Set the broadcast address in /etc/llttab explicitly depending on the subnet for each link.
  See "Broadcast address in the /etc/llttab file" on page 226.
- 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 228.
- Set the broadcast address correctly for direct-attached (non-routed) links.
- 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 231.

Broadcast address in the /etc/llttab file

The broadcast address is set explicitly for each link in the following example.

```
# cat /etc/llttab
set-node Node0
set-cluster 1

link link1 /dev/udp - udp  50000  -  10.20.30.1  10.20.30.255
link link2 /dev/udp - udp  50001  -  10.20.31.1  10.20.31.255
```

Verify the subnet mask using the ifconfig command to ensure that the two links are on separate subnets.

The link command in the /etc/llttab file

Review the link command information in this section for the /etc/llttab file. See the following information for sample configurations:

- See "Sample configuration: direct-attached links" on page 229.

- See "Sample configuration: links crossing IP routers" on page 231.

Note that some of the fields in Table B-1 on page 227. differ from the command for standard LLT links.

Table B-1 describes the fields of the link command that are shown in the /etc/llttab file examples.

Table B-1 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/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 228.
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 displays 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/llttab file

The `set-addr` command in the /etc/llttab file is required when the broadcast feature of LLT is disabled, such as when LLT must cross IP routers.

See "Sample configuration: links crossing IP routers" on page 231.

Table B-2 describes the fields of the set-addr command.

Table B-2 Field description for set-addr command in /etc/llttab

Field	Description
node-id	The ID of the cluster node; for example, 0.

Table B-2 Field description for set-addr command in /etc/llttab *(continued)*

Field	Description
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
-------------------- -------------------- -------
      *.sunrpc                            Idle
      *.*                                 Unbound
      *.32771                             Idle
      *.32776                             Idle
      *.32777                             Idle
      *.name                              Idle
      *.biff                              Idle
      *.talk                              Idle
      *.32779                             Idle
.
.
.
      *.55098                             Idle
      *.syslog                            Idle
      *.58702                             Idle
      *.*                                 Unbound
```

Look in the UDP section of the output; the UDP ports that are listed under Local Address are already in use. If a port is listed in the /etc/services file, its associated name is displayed rather than the port number in the output.

Configuring the netmask for LLT

For nodes on different subnets, set the netmask so that the nodes can access the subnets in use.

For example, with the following interfaces:

- For first network interface

```
IP address=192.168.30.1, Broadcast address=192.168.30.255,
Netmask=255.255.255.0
```

- For second network interface

```
IP address=192.168.31.1, Broadcast address=192.168.31.255,
 Netmask=Mask:255.255.255.0
```

Configuring the broadcast address for LLT

For nodes on different subnets, set the broadcast address in /etc/llttab depending on the subnet that the links are on.

An example of a typical /etc/llttab file when nodes are on different subnets. Note the explicitly set broadcast address for each link.

```
# cat /etc/llttab
set-node nodexyz
set-cluster 100

link link1 /dev/udp - udp 50000 - 192.168.30.1 192.168.30.255
link link2 /dev/udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: direct-attached links

Figure B-1 depicts a typical configuration of direct-attached links employing LLT over UDP.

Figure B-1 A typical configuration of direct-attached links that use LLT over UDP

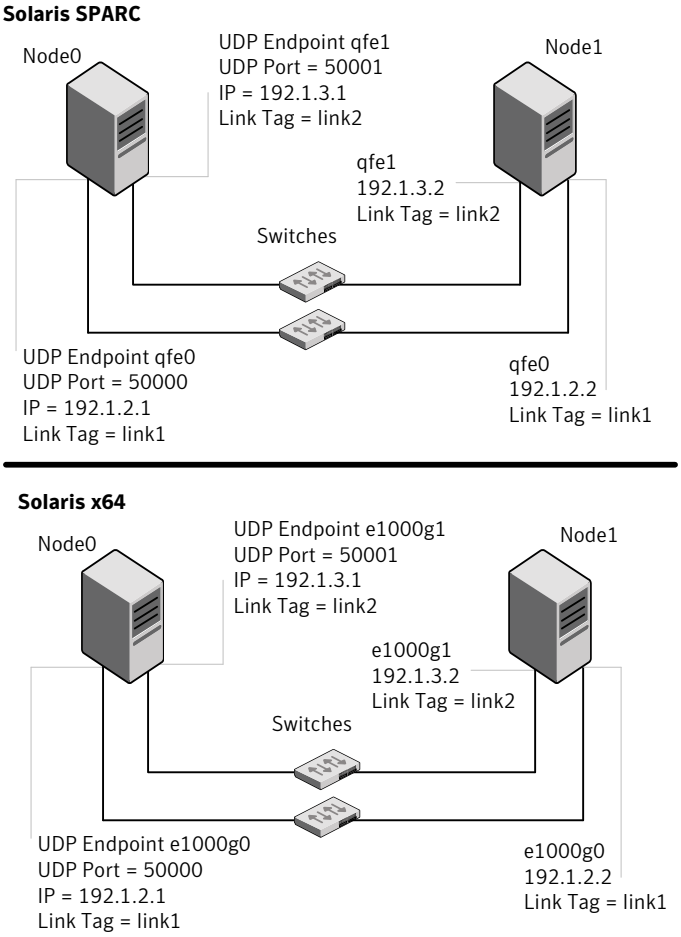

The configuration that the /etc/llttab file for Node 0 represents has directly attached crossover links. It might also have the links that are connected through a hub or switch. These links do not cross routers.

LLT broadcasts requests peer nodes to discover their addresses. So the addresses of peer nodes do not need to be specified in the /etc/llttab file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of

the links in the /etc/llttab file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /dev/udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 /dev/udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 resembles:

```
set-node Node1
set-cluster 1
#configure Links
#link tag-name device node-range link-type udp port MTU \
IP-address bcast-address
link link1 /dev/udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 /dev/udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: links crossing IP routers

Figure B-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 B-2 A typical configuration of links crossing an IP router

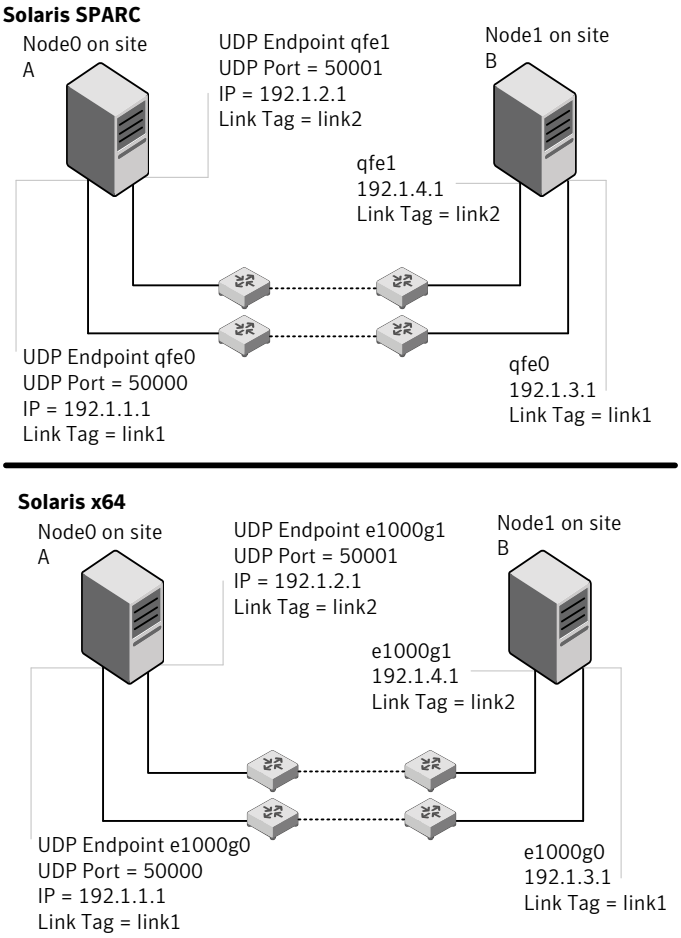

The configuration that the following `/etc/lltab` file represents for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the `link` command of the `/etc/lltab` file.

```
set-node Node1
set-cluster 1
```

```
link link1 /dev/udp - udp 50000 - 192.1.3.1 -
link link2  /dev/udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr        0 link1 192.1.1.1
set-addr        0 link2 192.1.2.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

The `/etc/llttab` file on Node 0 resembles:

```
set-node Node0
set-cluster 1

link link1 /dev/udp - udp 50000 - 192.1.1.1 -
link link2 /dev/udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr node-id link tag-name address
set-addr        1 link1 192.1.3.1
set-addr        1 link2 192.1.4.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

Performing automated VCS installations

Using installvcs program with the `-responsefile` option is useful not only for installing and configuring VCS within a secure environment. This option is also useful for conducting unattended installations to other clusters as well. Typically, you can use the response file generated during the installation of VCS on one

cluster to install VCS on other clusters. You can copy the file to a system in another cluster and manually edit the file to contain appropriate values.

When the systems are set up and meet the requirements for installation, you can perform an unattended installation. You perform the installation from one of the cluster systems where you have copied the response file.

To perform unattended installation

1. Navigate to the folder containing the installvcs program.

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

2. Start the installation from one of the cluster systems where you have copied the response file.

   ```
   # ./installvcs -responsefile /tmp/response_file
   ```

   Where /tmp/response_file is the response file's full path name.

Syntax in the response file

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

For example,

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

or, in the case of an integer value:

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

or, in the case of a list:

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

Example response file

The example response file resembles the file that installvcs creates after the example VCS installation. The file is a modified version of the response file generated on vcs_cluster2 that you can use to install VCS on vcs_cluster3. Review the variables that are required for installation.

Note: For Solaris x64 Platform Edition, replace hme0, qfe0, qfe1 with e1000g0, e1000g2, and e1000g3 in the following response file.

See "Response file variable definitions" on page 235.

```
#
# installvcs configuration values:
#
$CPI::CFG{AT_ROOTDOMAIN}="root\@east.symantecexample.com";
$CPI::CFG{CMC_CC_CONFIGURED}=1;
$CPI::CFG{CMC_CLUSTERID}{east}=1146235600;
$CPI::CFG{CMC_MSADDR}{east}="mgmtserver1";
$CPI::CFG{CMC_MSADDR}{west}="mgmtserver1";
$CPI::CFG{CMC_MS_ROOT_HASH}="758a33dbd6fae716...3deb54e562fe98";
$CPI::CFG{CMC_SERVICE_PASSWORD}="U2FsdVkX18v...n0hTSWwodThc+rX";
$CPI::CFG{ENCRYPTED}="U2FsdGVkX1+k2DHcnW7b6...ghdh+zW4G0WFIJA=";
$CPI::CFG{KEYS}{east}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{KEYS}{west}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{OBC_IGNOREWARNINGS}=0;
$CPI::CFG{OBC_MODE}="STANDALONE";
$CPI::CFG{OPT}{INSTALL}=1;
$CPI::CFG{OPT}{NOEXTRAPKGS}=1;
$CPI::CFG{OPT}{RSH}=1;
$CPI::CFG{SYSTEMS}=[ qw(east west) ];
$CPI::CFG{UPI}="VCS";
$CPI::CFG{VCS_ALLOWCOMMS}="Y";
$CPI::CFG{VCS_CLUSTERID}=13221;
$CPI::CFG{VCS_CLUSTERNAME}="vcs_cluster3";
$CPI::CFG{VCS_CSGNETMASK}="255.255.240.0";
$CPI::CFG{VCS_CSGNIC}{ALL}="hme0";
$CPI::CFG{VCS_CSGVIP}="10.10.12.1";
$CPI::CFG{VCS_LLTLINK1}{east}="qfe0";
$CPI::CFG{VCS_LLTLINK1}{west}="qfe0";
$CPI::CFG{VCS_LLTLINK2}{east}="qfe1";
$CPI::CFG{VCS_LLTLINK2}{west}="qfe1";

$CPI::CFG{VCS_SMTPRECP}=[ qw(earnie@symantecexample.com) ];
$CPI::CFG{VCS_SMTPRSEV}=[ qw(SevereError) ];
$CPI::CFG{VCS_SMTPSERVER}="smtp.symantecexample.com";
$CPI::CFG{VCS_SNMPCONS}=[ qw(neptune) ];
$CPI::CFG{VCS_SNMPCSEV}=[ qw(SevereError) ];
$CPI::CFG{VCS_SNMPPORT}=162;
```

Response file variable definitions

Note that some optional variables make it necessary to define other optional variables. For example, all the variables that are related to the cluster service

group (CSGNIC, CSGVIP, and CSGNETMASK) must be defined if any are defined. The same is true for the SMTP notification (SMTPSERVER, SMTPRECP, and SMTPRSEV), the SNMP trap notification (SNMPPORT, SNMPCONS, and SNMPCSEV), and the Global Cluster Option (CGONIC, GCOVIP, and GCONETMASK).

Table B-3 lists the variables that the response file uses and the variable definitions.

Table B-3 Response file variables

Variable	Description
$CPI::CFG{OPT}{INSTALL}	Installs and configures VCS. List or scalar: scalar Optional or required: required
$CPI::CFG{OPT}{INSTALLONLY}	Installs VCS packages. Configuration can be performed at a later time using the `-configure` option. List or scalar: scalar Optional or required: optional
$CPI::CFG{SYSTEMS}	List of systems on which the product is to be installed, uninstalled, or configured. List or scalar: list Optional or required: required
$CPI::CFG{SYSTEMSCFG}	List of systems to be recognized in configuration if secure environment prevents all systems from being installed at once. List or scalar: list Optional or required: optional
$CPI::CFG{UPI}	Defines the product to be installed, uninstalled, or configured. List or scalar: scalar Optional or required: required
$CPI::CFG{OPT}{KEYFILE}	Defines the location of an ssh keyfile that is used to communicate with all remote systems. List or scalar: scalar Optional or required: optional

Table B-3 Response file variables *(continued)*

Variable	Description
$CPI::CFG{OPT}{LICENSE}	Licenses VCS only. List or scalar: scalar Optional or required: optional
$CPI::CFG{OPT}{NOLIC}	Installs the product without any license. List or scalar: scalar Optional or required: optional
$CPI::CFG{AT_ROOTDOMAIN}	Defines the name of the system where the root broker is installed. List or scalar: list Optional or required: optional
$CPI::CFG{OPT}{PATCHPATH}	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. List or scalar: scalar Optional or required: optional
$CPI::CFG{OPT}{PKGPATH}	Defines a location, typically an NFS mount, from which all remote systems can install product depots. The location must be accessible from all target systems. List or scalar: scalar Optional or required: optional
$CPI::CFG{OPT}{TMPPATH}	Defines the location where a working directory is created to store temporary files and the depots that are needed during the install. The default location is /var/tmp. List or scalar: scalar Optional or required: optional
$CPI::CFG{OPT}{RSH}	Defines that *rsh* must be used instead of ssh as the communication method between systems. List or scalar: scalar Optional or required: optional

Table B-3 Response file variables *(continued)*

Variable	Description
$CPI::CFG{DONOTINSTALL} {PACKAGE}	Instructs the installation to not install the optional packages in the list. List or scalar: list Optional or required: optional
$CPI::CFG{DONOTREMOVE} {PACKAGE}	Instructs the uninstallation to not remove the optional packages in the list. List or scalar: list Optional or required: optional
$CPI::CFG{VCS_CLUSTERNAME}	Defines the name of the cluster. List or scalar: scalar Optional or required: required
$CPI::CFG{VCS_CLUSTERID}	An integer between 0 and 65535 that uniquely identifies the cluster. List or scalar: scalar Optional or required: required
$CPI::CFG{KEYS} {SYSTEM}	List of keys to be registered on the system. List or scalar: scalar Optional or required: optional
$CPI::CFG{OPT_LOGPATH}	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs. List or scalar: scalar Optional or required: optional
$CPI::CFG{CONFIGURE}	Performs the configuration if the packages are already installed using the `-installonly` option. List or scalar: scalar Optional or required: optional

Table B-3 Response file variables *(continued)*

Variable	Description
$CPI::CFG{VCS_LLTLINK#} {SYSTEM}	Defines the NIC to be used for a private heartbeat link on each system. Two LLT links are required per system (LLTLINK1 and LLTLINK2). Up to four LLT links can be configured. List or scalar: scalar Optional or required: required
$CPI::CFG{VCS_LLTLINKLOWPRI} {SYSTEM}	Defines a low priority heartbeat link. Typically, LLTLINKLOWPRI is used on a public network link to provide an additional layer of communication. List or scalar: scalar Optional or required: optional
$CPI::CFG{VCS_CSGNIC}	Defines the NIC for Cluster Management Console to use on a system. 'ALL' can be entered as a system value if the same NIC is used on all systems. List or scalar: scalar Optional or required: optional
$CPI::CFG{CSGVIP}	Defines the virtual IP address that the Cluster Management Console uses. List or scalar: scalar Optional or required: optional
$CPI::CFG{VCS_CSGNETMASK}	Defines the Netmask of the virtual IP address that the Cluster Management Console uses. List or scalar: scalar Optional or required: optional
$CPI::CFG{VCS_SMTPSERVER}	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for Web notification. List or scalar: scalar Optional or required: optional

Table B-3 Response file variables *(continued)*

Variable	Description
$CPI::CFG{VCS_SMTPRECP}	List of full email addresses (example: user@symantecexample.com) of SMTP recipients. List or scalar: list Optional or required: optional
$CPI::CFG{VCS_SMTPRSEV}	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. List or scalar: list Optional or required: optional
$CPI::CFG{VCS_SNMPPORT}	Defines the SNMP trap daemon port (default=162). List or scalar: scalar Optional or required: optional
$CPI::CFG{VCS_SNMPCONS}	List of SNMP console system names List or scalar: list Optional or required: optional
$CPI::CFG{VCS_SNMPCSEV}	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. List or scalar: list Optional or required: optional
$CPI::CFG{VCS_GCONIC} {SYSTEM}	Defines the NIC for the Virtual IP that the Global Cluster Option uses. 'ALL' can be entered as a system value if the same NIC is used on all systems. List or scalar: scalar Optional or required: optional

Table B-3 Response file variables *(continued)*

Variable	Description
$CPI::CFG{VCS_GCOVIP}	Defines the virtual IP address to that the Global Cluster Option uses. List or scalar: scalar Optional or required: optional
$CPI::CFG{VCS_GCONETMASK}	Defines the Netmask of the virtual IP address that the Global Cluster Option uses. List or scalar: scalar Optional or required: optional
$CPI::CFG{VCS_USERENPW}	List of encoded passwords for users List or scalar: list Optional or required: optional
$CPI::CFG{VCS_USERNAME}	List of names of users List or scalar: list Optional or required: optional
$CPI::CFG{VCS_USERPRIV}	List of privileges for users List or scalar: list Optional or required: optional
$CPI::CFG{OPT}{UNINSTALL}	List of systems where VCS must be uninstalled. List or scalar: scalar Optional or required: optional

Installing VCS with a response file where ssh or rsh are disabled

In secure enterprise environments, ssh or rsh communication is not allowed between systems. In such cases, the installvcs program can install and configure VCS only on systems with which it can communicate—most often the local system only. When installation is complete, VCS creates a response file.

See "Example response file" on page 234.

The response file that the installvcs program generates contains descriptions and explanations of the variables and their values. You copy this file to the other

systems in the cluster, and edit it to reflect the current local system. You can use the installation program with the `-responsefile` option to install and configure VCS identically on each system without being prompted.

To use installvcs in a secure environment

1. On one node in the cluster, start VCS installation using the installvcs program.

   See "Starting the software installation" on page 63.

2. Review the output as the installer performs the initial system checks.

   The installer detects the inability to communicate between systems.

3. Press the Enter key to install VCS on one system and create a response file with which you can install on other systems.

   ```
   Would you like to install Cluster Server on systems galaxy only
   and create a responsefile for systems nebula? [y,n,q] (y)
   ```

4. Enter all cluster information. Proceed with the installation and configuration tasks.

   See "Installing and configuring VCS 5.0 MP3" on page 62.

   The installvcs program installs and configures VCS on systems where communication is possible.

5. After the installation is complete, review the installer report.

   The installer stores the installvcs-*universaluniqueidentifier* response file in the /opt/VRTS/install/logs/installvcs-*universaluniqueidentifier*/.response directory where *universaluniqueidentifier* is a variable to uniquely identify the file.

6. If you start VCS before VCS is installed and started on all nodes in the cluster, you see the output similar to:

   ```
   VCS:11306:Did not receive cluster membership, manual
   intervention may be needed for seeding
   ```

7. Use a method of your choice (for example, by using NFS, ftp, or a floppy disk). Place a copy of the response file in a directory such as /tmp on the next system to install VCS.

8 On the next system, edit the response file.

For the variables in the example, change the name of the system to reflect the current local system:

```
.
$CFG{SYSTEMS} = ["east"];
.
.
$CFG{KEYS}{east} = ["XXXX-XXXX-XXXX-XXXX-XXXX-XXX"];
.
```

For demo or site licenses, the license key need not be changed. When license keys are "node-locked" to specific cluster nodes, you must edit the license key.

9 On the next system, perform the following:

- Mount the product disc.
  See "Mounting the product disc" on page 49.
- Start the software installation using the `installvcs -responsefile` option.

  ```
  # ./installvcs -responsefile /tmp/installvcs-uui.response
  ```

  Where uui is the Universal Unique Identifier that the installer automatically assigned to the response file.
  See "Starting the software installation" on page 63.

10 Repeat step 7 through step 9 until VCS has been installed on all nodes in the cluster.

Index

L

M

N

O

P