# Veritas™ Cluster Server Installation Guide

Solaris x64 Platform Edition

5.0

symantec.

Veritas Cluster Server
Installation Guide

Third-party legal notices

Third-party software may be recommended, distributed, embedded, or bundled with this Symantec product. Such third-party software is licensed separately by its copyright holder. All third-party copyrights associated with this product are listed in the accompanying release notes.

Opteron is a registered trademark of Advanced Micro Devices.
Solaris is a trademark of Sun Microsystems, Inc.

Licensing and registration

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.

Technical support

Visit http://www.symantec.com/enterprise/support/assistance_care.jsp for technical assistance. Use the Knowledge Base search feature to access resources such as TechNotes, product alerts, software downloads, hardware compatibility lists, and our customer email notification service. If you encounter an error when using a product, include the error number preceding the message when contacting Technical Services.You can also use the error number to search for information in TechNotes or documents on the website.

Contents

Chapter 3 Installing and configuring VCS

Chapter 4   Manually installing and configuring VCS

Chapter 6   Verifying the VCS installation

Chapter 7   Upgrading to VCS 5.0

Introducing VCS

This chapter contains the following topics:

- About VCS
- VCS basics

About VCS

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.

VCS basics

A single VCS cluster consists of multiple systems 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.

Client application continue operation 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 connected to shared storage. 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.

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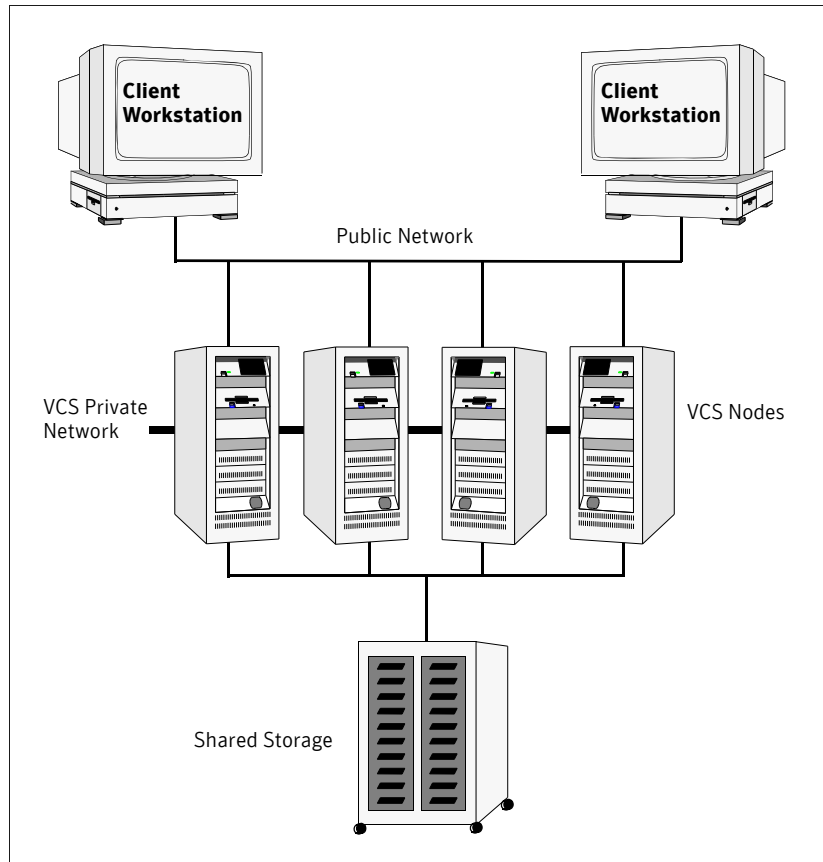
Multiple nodes

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 and to recognize active nodes, nodes that are joining or leaving the cluster, and failed nodes. The private network requires two communication channels to guard against network partitions.

Shared storage

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

Figure 1-2 illustrates the flexibility of VCS shared storage configurations. VCS nodes can only access physically-attached storage.

Figure 1-2   Two examples of shared storage configurations
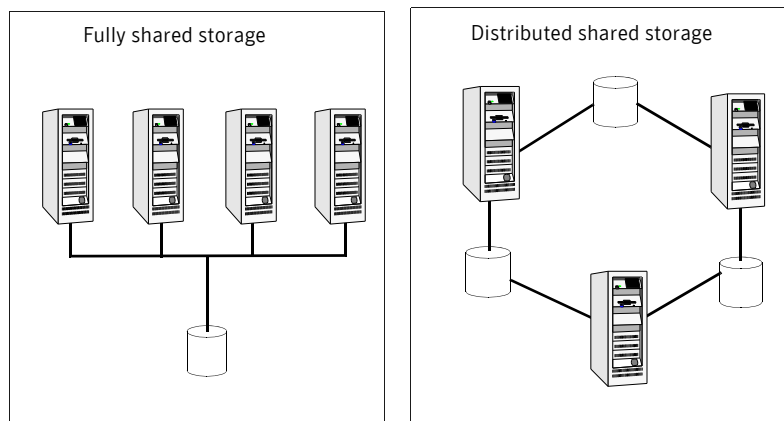
LLT and GAB

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. The system administrator configures LLT by creating the configuration files /etc/llthosts, which lists all the nodes in the cluster, and /etc/llttab, which 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 required to maintain a synchronized state among the nodes. It monitors disk communications such as the VCS heartbeat utility. The system administrator configures the GAB driver by creating a configuration file (/etc/gabtab).

See "Verifying LLT and GAB configuration files" on page 129.

Network channels for heartbeating

For the VCS private network, two network channels must be available for heartbeating. These network connections are also used for transmitting information.

Each Solaris x64 Platform Edition cluster configuration requires at least two network channels between the systems. The requirement for two channels protects your cluster against network partitioning. Refer to the *Veritas Cluster Server User's Guide* for more information on network partitioning.

Figure 1-3 illustrates a two-node VCS cluster where sysA and sysB have two private network connections.

Figure 1-3   Two nodes connected by two ethernet connections
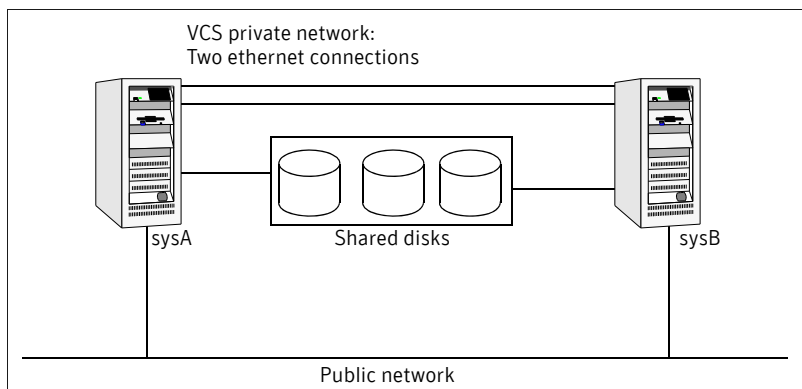
Preexisting network partitions

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

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 when:

■   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.

You need to perform a manual seed to run VCS from a cold start (all systems down) when one or more systems of the cluster are unavailable. VCS does not start service groups on a system until it has a seed.

Preparing to install and configure VCS

This chapter contains the following topics:

- About preparing to install VCS 5.0
- About VCS 5.0 features
- Preparing to install VCS 5.0
- Performing pre-installation tasks
- Getting your VCS installation and configuration information ready

About preparing to install VCS 5.0

Before you install any product, read the following Veritas Technical Support TechNote for the latest information on updates, patches, and software issues regarding this release:

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

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

http://entsupport.symantec.com/docs/287089

About VCS 5.0 features

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:

- Symantec Product Authentication Service
- SMTP email notification for VCS events
- SNMP trap notification for VCS events
- Global clusters using VCS
- I/O fencing

Symantec Product Authentication Service

Symantec Product Authentication Service secures communication between cluster nodes and clients, including the Java console, by using digital certificates for authentication and SSL to encrypt communication over the public network. For more information about the Authentication Service, see the *Veritas Cluster Server User's Guide.*

To configure the cluster in secure mode, VCS requires you to configure a system in your enterprise as root broker and all nodes in the cluster as authentication brokers.

- 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.

- Authentication brokers
  Authentication brokers serve as intermediate registration and certification authorities. Authentication brokers have certificates that are signed by the root. Each node in VCS serves as an authentication broker.

You can set up Authentication Service for the cluster during the installation or after installation. Refer to the *Veritas Cluster Server User's Guide* to configure the cluster in secure mode after the installation and configuration process.

See "Configuring the cluster in secure mode" on page 54.

Figure 2-4 depicts the flow of configuring VCS in secure mode.

Figure 2-4   Secure VCS cluster configuration flowchart
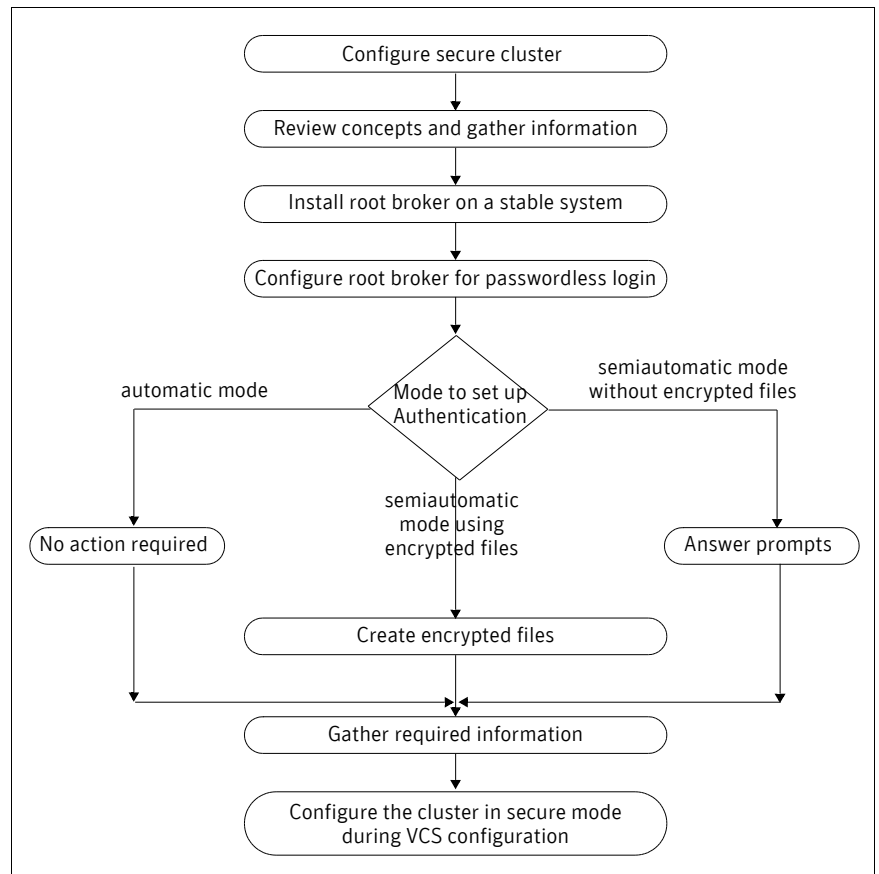
If you decide to enable Authentication Service, the root broker administrator must perform the following preparatory tasks:

■   Install the root broker on another stable system.

    The root broker is the main registration and certification authority and can serve multiple clusters. Symantec recommends that you install a single root broker on a utility computer such as an email server or domain controller, which can be highly available.

    See "Installing root broker for Symantec Product Authentication Service" on page 27.

■   Configure the root broker system for a passwordless login when you want to use the automatic mode.

The installvcs program provides the following modes to enable Symantec Product Authentication Service:

■ In the automatic mode, the installer configures Authentication Service automatically without any user intervention.
You must provide the name of the root broker system.

■ In the semiautomatic modes, the installer provides you an option to use encrypted files or answer the installer prompts to enable security. The semiautomatic mode requires the root broker administrator to set up the basic authentication environment and create principals for authentication brokers. You must complete the following preparatory tasks to configure security in the semiautomatic mode:

With encrypted file	■	The root broker administrator must create an encrypted file for each node in the cluster. See "Creating encrypted files for Symantec Product Authentication Service" on page 28.
	■	You must fetch the encrypted files from the root broker administrator and copy the encrypted files to the installation node. Make a note of the path of these encrypted files.
Without encrypted file	■	You must gather the following information from the root broker administrator: - Root broker name - Root broker domain name - Root broker port (Default is 2821) - Authentication broker principal name for each node - Authentication broker password for each Authentication broker
	■	You must fetch the root_hash file from the root broker system and copy the root_hash file to a directory in the installation node. Make a note of the path of this root_hash file.

Note: Make sure that the system clocks of the root broker and authentication brokers are in sync.

SMTP email notification for VCS events

You have the option to configure SMTP email notification of VCS events by the VCS Notifier component. If you choose SMTP notification, be ready to answer prompts for the following information:

■ The domain-based address of the SMTP server that is to send notification email about the events within the cluster, for example: smtp.symantecexample.com.

■ The email address of each SMTP recipient to be notified, for example: john@symantecexample.com.

■ The minimum severity of events for SMTP email notification. Events have four levels of severity: Information, Warning, Error, and SevereError.

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

SNMP trap notification for VCS events

You have the option to configure SNMP trap notification of VCS events by the VCS Notifier component. If you choose SNMP notification, be ready to answer prompts for the following information:

■ The port number, 162 by default, for the SNMP trap daemon.

■ The system name for each SNMP console.

■ The minimum severity of events for SNMP trap notification. Events have four levels of severity: Information, Warning, Error, and SevereError.

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

Global clusters using VCS

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.

If you choose to configure global clusters, the installer enables you to choose whether or not to use the same NIC, virtual IP address, and netmask as are configured for the ClusterService group, which are the defaults. If you choose not to use the same networking information, you must specify appropriate values for the NIC, virtual IP address, and netmask when you are prompted.

I/O fencing

I/O fencing protects the data on shared disks. When nodes in a cluster detect a change in cluster membership that could indicate a split brain condition, the fencing operation proceeds to determine which nodes are to retain access to the shared storage and which nodes are to be ejected from the cluster, thus preventing possible data corruption. The *Veritas Cluster Server User's Guide* describes I/O fencing concepts in detail. The installvcs program installs the VCS I/O fencing driver, VRTSvxfen.

Note: Symantec strongly recommends that you use VCS I/O fencing to deter potential split brain scenarios in your cluster.

See "Setting up I/O fencing" on page 97.

Preparing to install VCS 5.0

Each node on which you want to install VCS must meet the hardware and software requirements.

- "Hardware requirements" on page 26
- "Supported operating systems" on page 26
- "Supported software" on page 27

After planning the VCS features that you want to configure, you must prepare to configure these features. Figure 2-5 represents the major tasks and decisions required to install and configure VCS.

Figure 2-5   Workflow for fresh install of VCS 5.0
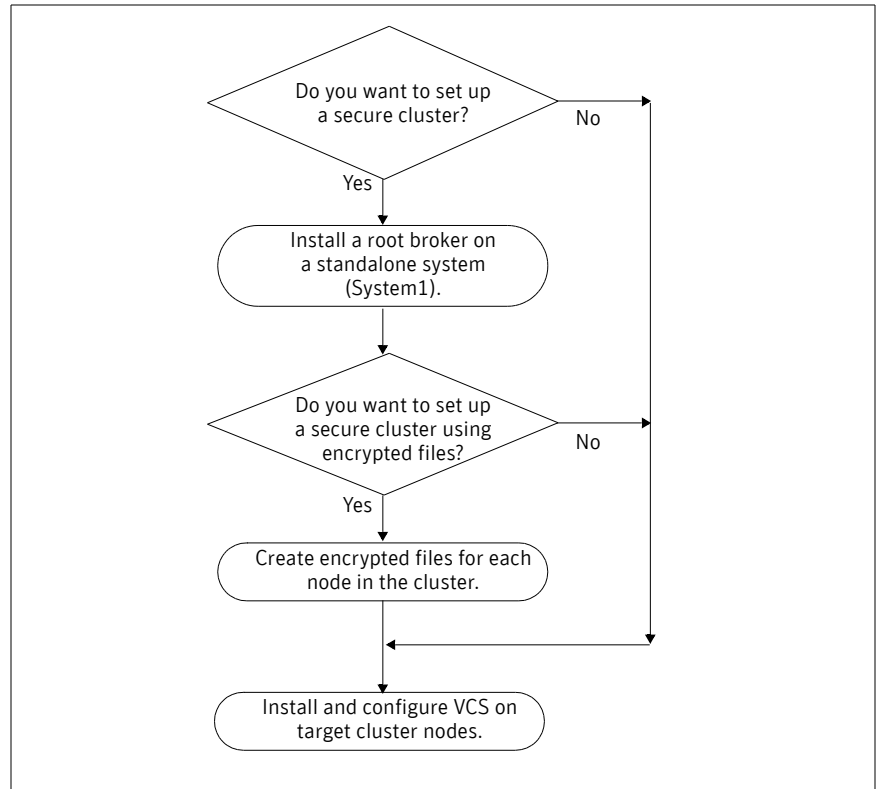
Complete the following preparatory tasks based on the VCS features you want to configure:

■   "Installing root broker for Symantec Product Authentication Service" on page 27

■   "Creating encrypted files for Symantec Product Authentication Service" on page 28

Hardware requirements

Make sure that you meet the following requirements.

Table 2-1   Hardware requirements for a cluster

Item	Description
VCS systems	Sun Fire Opteron servers running Solaris 10 x64 Platform Edition.
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 applications that migrate between systems in the cluster. The VCS I/O fencing feature requires that all disks used as data disks or as coordinator disks must support SCSI-3 Persistent Reservations (PR).
	The coordinator disk does not store data, so configure the disk as the smallest possible LUN on a disk array to avoid wasting space.
	See "Setting up I/O fencing" on page 97.

Disk space	To run VCS, LLT, GAB, the Web Console, and the Java Console, each VCS system 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 /

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.
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 system requires at least 256 megabytes.

Supported operating systems

Run VCS 5.0 on these operating systems at the suggested patch levels.

■ Solaris 10 (64-bit) operating systems

For each platform, Symantec recommends applying the latest cumulative operating system patches available from Sun. Visit the Sun website for more information.

Supported software

- Veritas Volume Manager (VxVM) (Solaris 10 only), 4.1, and 5.0
- Veritas File System (VxFS) (Solaris 10 only), 4.1, and 5.0

Note: If you plan to set up VCS I/O fencing in a shared storage environment, Symantec recommends using VxVM versions 4.1 or 5.0.

Installing root broker for Symantec Product Authentication Service

Install the root broker only if you plan on using Symantec Product Authentication Service. 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 *Veritas Cluster Server User's Guide* for more information. You can configure the Authentication Service during or after VCS installation.

See "Symantec Product Authentication Service" on page 20.

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: **east**

5   Review the output as the installer:
    - checks to make sure that the VCS supports the operating system
    - verifies that you are installing 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 will be installed 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 Enter 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.

13 If you plan to configure the Authentication Service during VCS installation, choose to configure the cluster in secure mode when the installer prompts you.
    See "Installing and configuring VCS 5.0" on page 48.

Creating encrypted files for
Symantec Product Authentication Service

Create encrypted files only if you plan on choosing the semiautomatic mode that uses an encrypted file to configure the Authentication Service. The encrypted files must be created by the administrator on the root broker node. The administrator must create encrypted files for each node that would be a part of the cluster before you configure the Authentication Service for VCS. See *Veritas Cluster Server User's Guide* for more information. You can configure the Authentication Service during or after VCS installation.

See "Symantec Product Authentication Service" on page 20.

To create encrypted files

1   Determine the root broker domain name. Enter the following command on the root broker system:
    ```
    east> # vssat showalltrustedcreds
    ```
    For example, the domain name would resemble
    "Domain Name: root@east.symantecexample.com" in the output.

2 For each node in the cluster, make sure that you have created an account on root broker system.

For example, to verify on node north:

```
east> # vssat showprpl --pdrtype root \
--domain root@east.symantecexample.com --prplname north
```

■ 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:

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

■ If the output displays an error similar to "Failed To Get Attributes For Principal," then the account for given authentication broker is not created on this root broker. Proceed to step 3.

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

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

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

4 Make a note of the following information that is required for the input file for the encrypted file.

■ hash - The root hash string that consists of 40 characters, as shown by the command:

```
east> # vssat showbrokerhash
```

■ identity - Authentication broker identity
The value that you provide for **--prplname** in step 3 (for example, north).

■ password - Authentication broker password
The value that you provide for **--password** in step 3.

■ root_domain - Domain name of the root broker system
The value that you determined in step 1.

■ broker_admin_password - Authentication broker password for Administrator account on the node
Provide a password of at least five characters long.

5 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 north would resemble:

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

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

6 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.

7 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:
   ```
   east> # vssat createpkg --in /tmp/north.blob.in \
   --out /tmp/north.blob.out --host_ctx north
   ```
   Note that this command creates a encrypted file even if you provide wrong password for "password=" entry, but the encrypted file will fail to install on authentication broker node.

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

9 If you plan to configure the Authentication Service during VCS installation, choose to configure the cluster in secure mode when the installer prompts you.

   See "Installing and configuring VCS 5.0" on page 48.

Performing pre-installation tasks

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

Table 2-2   Pre-installation tasks

Task	Reference
Set the PATH and MANPATH variables.	"Setting the PATH variable" on page 31
	"Setting the MANPATH variable" on page 32
Set up the private network.	"Setting up the private network" on page 32
Set up shared storage for I/O fencing (optional)	"Setting up shared storage" on page 34
Enable communication between systems.	"Enabling communication between systems" on page 35
Review basic instructions to optimize LLT media speeds.	"Optimizing LLT media speed settings on private NICs" on page 35
Review guidelines to help you set the LLT interconnects.	"Guidelines for setting the media speed of the LLT interconnects" on page 35
Set up ssh on cluster systems.	"Setting up ssh on cluster systems" on page 36
Obtain license keys.	"Obtaining VCS license keys" on page 38
Install the patches required for Java Run Time Environment from Sun.	"Patches required for Java Run Time Environment from Sun" on page 39
Mount the product disc	"Mounting the product disc" on page 39

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

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
```

Setting up the private network

VCS requires you to setup a private network between the systems that will form part of a cluster.

To set up the private network

1 Install the required network interface cards.

2 Connect the VCS private Ethernet controllers on each system.

3 Use cross-over Ethernet cables (supported only on two systems), or independent hubs, for each VCS communication network. Ensure that power to the hubs comes from separate sources. On each system, use two independent network cards to provide redundancy.

   During the process of setting up heartbeat connections, note that a chance for data corruption exists if a failure removes all communications between the systems and still leaves the systems running and capable of accessing shared storage.

Figure 2-6   Private network setups: two-node and four-node clusters
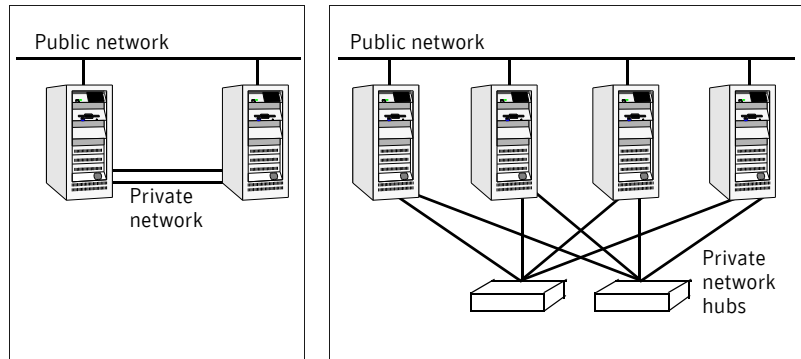
4   Configure the Ethernet devices used for the private network such that the
     auto-negotiation protocol is not used. This ensures a more stable
     configuration with cross-over cables.

     You can achieve this either by editing the /etc/system file to disable auto-
     negotiation on all Ethernet devices system-wide or by creating a
     e1000g.conf file in the /kernel/drv directory to disable auto-negotiation for
     the individual devices used for private network. Refer to the Sun Ethernet
     driver product documentation for information on these methods.

5   Test network connections by temporarily assigning network addresses and
     use `telnet` or `ping` to verify communications.

     LLT uses its own protocol, and does not use TCP/IP. Therefore, to ensure the
     private network connections are used only for LLT communication and not
     for TCP/IP traffic, unplumb and unconfigure the temporary addresses after
     testing.

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

     See "Installing and configuring VCS" on page 45.

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

     See "Manually installing and configuring VCS" on page 83.

Using network switches

You can use network switches instead of hubs. By default, Sun systems assign
the same MAC address to all interfaces. Therefore, connecting two or more
interfaces to a network switch can cause problems. For example, if IP is
configured on one interface and LLT on another, and both interfaces are
connected to a switch (assuming 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, configure the system to assign unique MAC addresses by setting the `eeprom(1M)` parameter local-mac-address to true.

Note: Because of their performance characteristics, network switches are recommended for clusters supporting the Veritas Storage Foundation Cluster File System and Veritas Storage Foundation for Oracle RAC, which make extensive use of the private cluster interconnects for distributed locking. See the *Veritas Cluster Server User's Guide* and review the chapter on VCS performance considerations.

Setting up shared storage

The following sections describe setting up SCSI and Fibre 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 "Setting up I/O fencing" on page 97.

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 storage shared between nodes, the SCSI address, or SCSI initiator ID, of each node must be unique. For x64 Opteron systems, the SCSI initiator or host ID is set using the system's SCSI BIOS configuration utility. Access this utility at boot time by pressing the Ctrl+C keys. For more information on using the SCSI BIOS configuration utility, see *Veritas Cluster Server x64 User's Guide.*

Setting up shared storage: Fibre channel

1   Install the required FC-AL controllers.

2   Connect the FC-AL controllers and the shared storage devices to the same hub or switch. If a fibre switch is being used, be sure that no zoning is implemented, which would prevent all systems from seeing all shared devices required to run the critical application.

3   Create the /reconfigure file that the OS reads when you reboot the system.
    ```
    # touch /reconfigure
    ```

4   Reboot each system.

> **Note:** If the controller is the supported default, the system finds and registers the device automatically. To update the system, instead of rebooting, you can run the `devfsadm` command.

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

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

- If Volume Manger is not being used, the same number of external disk devices must appear and device nodes must be identical for all devices on all systems.

Enabling communication between systems

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 permissions for the system where you run installvcs program to issue `ssh` or `rsh` commands as root 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 "Performing VCS installation in a secure environment" on page 68.

See "Manually installing and configuring VCS" on page 83.

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 switches or hubs used for the interconnects 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

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, set the hub or switch port to the same setting as that 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.

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. Verify that ssh is configured correctly before starting the installation process.

Secure Shell (ssh) is a program to log on to another computer over a network, to execute commands on a remote system, and to copy files from one system to another. The ssh 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.

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 missing, 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
```

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

```
# sftp target_sys
```

If this is the first time this step is run on a system, output similar to the following displays:

```
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 logging in, enter the following command to append the authorization key to the id_dsa.pub file:

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

16 To delete the id_dsa.pub public key file after it is copied to the target (host) system and added to the authorized keys file, type the following command:

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

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

```
# exit
```

18 When installing 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 is a shell-specific step 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 the system requesting a passphrase or password.

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

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 features and their descriptions encoded in a license key

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

Patches required for Java Run Time Environment from Sun

The GUI modules for VCS use the Java Run Time Environment from Sun Microsystems. You need to obtain and install the latest Solaris x64 Platform Edition specific patches to enable the modules to function properly. Visit the Sun website to download the packages.

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 connected by the network to the systems where you are installing VCS. The system that you are using to install VCS need not be part of the cluster.

2   Insert the product disc into a DVD drive 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 inserting the software disc, enter:
     # **mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom**
     Where c0t6d0s2 is the default address for the disc drive.

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: **north**, **south**
The required license keys	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 38.
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 44.

■ To configure the Veritas Cluster Server you need:

The name of 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 used by the private networks among systems	Do not use the network interface card that is used for the public network, which is typically bge0. Example: **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 "Symantec Product Authentication Service" on page 20.
	■ 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 Symantec Product Authentication Service" on page 28.
For semiautomatic mode without using encrypted files	■ The fully-qualified hostname (FQDN) of the Root Broker. (e.g. east.symantecexample.com) The above example given posits 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 the each cluster node (e.g. north.symantecexample.com and south.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 the Cluster Management Console to locally manage this cluster (optional), you need:

The name of the public NIC for each node in the cluster	The device name for the NIC that provides public network access. Example: **bge0**
A virtual IP address of the NIC for the Cluster Management Console	This virtual IP address becomes a resource for use by the ClusterService group that includes the Cluster Management Console. The "Cluster Virtual IP address" can fail over to another cluster system, making the Web Console highly available. Example: **10.10.12.1**
The netmask for the virtual IP address	The subnet used with the virtual address. Example: **255.255.240.0**

■ To configure the Cluster Management Console cluster connector (optional), you need:

The management server network address for Cluster Management Console	The Cluster Management Console cluster connector requires the management server network address. Example: **mgmtserver1.symantecexample.com**
A Cluster Management Console service account password	You must have set this account password while installing the management server.
The root hash of the management server	You can use `vssat showbrokerhash` command and copy the root hash of the management server.

■ 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 can use the same NIC that you configured for the ClusterService group. Otherwise, specify appropriate values for the NIC.
	Example: **bge0**
The virtual IP address of the NIC	You can use the same virtual IP address that you configured for the ClusterService group. Otherwise, specify appropriate values for the virtual IP address.
	Example: **10.10.12.1**
The netmask for the virtual IP address	You can use the same netmask as configured for the ClusterService group. Otherwise, specify appropriate values for the netmask.
	Example: **255.255.240.0**

■ To configure I/O fencing:

| The names of the three disks that form the coordinator disk group | Example: **c1t1d0s0, c2t1d0s0, c3t1d0s0** |
| The DMP nodes names for each disk in the coordinator disk group (if using DMP) | Example: **/dev/vx/dmp** |

Optional VCS packages

The optional VCS packages include:

- VRTScmccc — Veritas Cluster Management Console Cluster Connector

- VRTScmcs — Veritas Cluster Management Console

- VRTScssim — VCS Simulator

- VRTScscm — Veritas Cluster Server Cluster Manager

- VRTSvcsdc - VCS documentation

- VRTSvcsmn - Manual pages for VCS commands

Installing and configuring VCS

This chapter contains the following topics:

- About installing and configuring VCS
- About the VCS installation program
- Installing and configuring VCS 5.0
- Installing VCS using installonly option
- Configuring VCS using configure option
- Performing VCS installation in a secure environment
- Performing automated installations
- Checking licensing information on the system
- Updating product licenses using vxlicinst
- About installvcs command options
- About the uninstallvcs program
- Uninstalling VCS 5.0

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	Offers a high-level approach to installing multiple Veritas products.

installvcs program   Offers a direct approach to installing VCS.

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

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, creating several detailed configuration files on each system
- Starting VCS processes

You can choose to configure the optional: Web-based Cluster Management Console, SNMP and SMTP notification features in the cluster, Symantec Product Authentication Services feature, or the wide area Global Cluster feature. Review the highlights of the information for which installvcs program prompts you as you proceed to configure.

See "Preparing to install and configure VCS" on page 19.

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

See "About the uninstallvcs program" on page 80.

Optional features of the installvcs program

Table 3-3 specifies the optional actions that the installvcs program can perform.

Table 3-3   installvcs optional features

Optional action	Reference
Check the systems to verify that they meet the requirements to install VCS.	See "Checking the systems for installation" on page 49.
Upgrade VCS to version 5.0 if VCS currently runs on a cluster.	"Upgrading VCS using installvcs program" on page 143
Install VCS packages without configuring VCS.	See "Installing VCS using installonly option" on page 67.

Table 3-3   installvcs optional features

Optional action	Reference
Configure or reconfigure VCS when VCS packages are already installed.	See "Configuring VCS using configure option" on page 67.
Perform secure installations using values stored in a configuration file.	See "Performing VCS installation in a secure environment" on page 68.
Perform automated installations using values stored in a configuration file.	See "Performing automated installations" on page 69.

Interacting with the installvcs program

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

Note: Installation of VCS packages takes place only after you have confirmed the information. However, you must remove the partially installed VCS files before running the installvcs program again.
See "Uninstalling VCS 5.0" on page 81.

At some points during the installation, the installer prompts you to type information and 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.

When the installer prompts you to answer a series of questions related to a configuration activity, 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 re-enter 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.
See "Installing the VCS Java Console" on page 62.

Installing and configuring VCS 5.0

Figure 3-7 illustrates the systems on which you would install and run VCS. The example installation demonstrates how to install VCS on two systems: north and south. 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 3-7   An example of a VCS installation on a two-node cluster
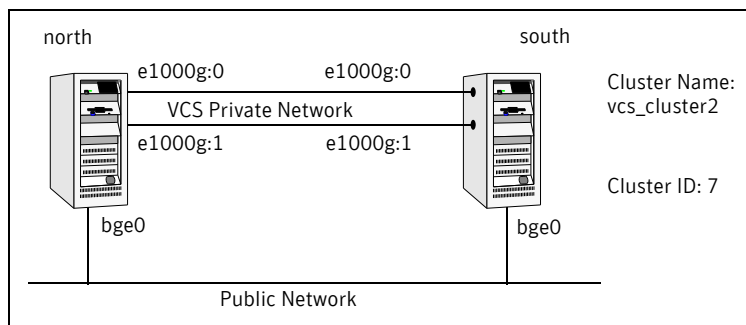
Overview of tasks

Table 3-4 lists the installation and configuration tasks.

Table 3-4   Installation and configuration tasks

Task	Reference
Start the installation process and choose the installation	■ "Checking the systems for installation" on page 49 (optional) ■ "Starting the software installation" on page 50 ■ "Specifying systems for installation" on page 50 ■ "Licensing VCS" on page 51 ■ "Choosing VCS packages" on page 52 ■ "Choosing to install VCS packages or configure VCS" on page 53

Table 3-4   Installation and configuration tasks

Task	Reference
Configure the cluster and optional features	■ "Configuring the cluster" on page 53 ■ "Configuring the cluster in secure mode" on page 54 (optional) ■ "Adding VCS users" on page 56 (optional) ■ "Configuring cluster connector" on page 56 (optional) ■ "Configuring SMTP email notification" on page 57 (optional) ■ "Configuring SNMP trap notification" on page 58 (optional) ■ "Configuring global clusters" on page 60 (optional)
Install the packages and create configuration files	■ "Installing the VCS packages" on page 60 ■ "Creating VCS configuration files" on page 61
Start VCS and its components	■ "Starting VCS" on page 61 ■ "Completing the installation" on page 61
Install language packages	■ "I" on page 62
Perform the post-installation tasks	■ "Copying the installation guide to each node" on page 62 ■ "Setting up I/O fencing" on page 97 ■ "Installing the VCS Java Console" on page 62 ■ "Establishing cluster communication with the management server" on page 64 ■ "Installing cluster connector" on page 65
Verify the cluster	■ "Verifying the cluster after installation" on page 67

Checking the systems for installation

Before beginning 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 ...
```

To check the systems

1   Navigate to the folder containing the installvcs program.

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

2   Start the pre-installation check:

     # **./installvcs -precheck north south**
     The program proceeds in a non-interactive mode, examining 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 command options" on page 77.

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 containing 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: north south
     ```
     For a single node installation, enter one name for the system.
     See "Starting the installer for the single node cluster" on page 257.

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 are installing 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.
          See "Upgrading VCS using installvcs program" on page 143.

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 76.

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 north: [?] XXXX-XXXX-XXXX-XXXX-XXX
     XXXX-XXXX-XXXX-XXXX-XXX successfully registered on north
     VCS license registered on north
     ```

3   Enter keys for additional product features.

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

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

```
Do you want to enter another license key for north? [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 south
VCS license registered on south

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

Choosing VCS packages

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 to be installed.

```
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 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 cluster" on page 53.

- 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 packages.

3   Configure the cluster later.
    See "Configuring VCS using configure option" on page 67.

Configuring the cluster

The installer provides you an option to configure VCS and its optional features.

Note: You can use `installvcs -configure` command to configure the cluster later and enter the system names where you want to configure VCS when the installer prompts you.

To configure the cluster

1   Enter **y** or press **Enter** at the prompt to configure VCS.
    ```
    It is optional to configure VCS now. If you choose to configure
    VCS later, you can either do so manually or run the
    installvcs -configure command.
    Are you ready to configure VCS?
    [y,n,q] (y) y
    ```

2   Review the configuration requirements that the installer lists.

3   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
    ```

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

5   Enter the details for the private heartbeat links.

     You must not enter the network interface card that is used for the public
     network (typically bge0.)

```
Enter the NIC for the first private heartbeat NIC on north:
[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 north:
[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)
```

6   Choose whether to use the same NICs on all nodes.

     ■   If you want to use the same NICs for private heartbeat links on all
         nodes, make sure the same NICs are available on each system and enter
         y.

     ■   Enter **n** to use NICs with different device names on some of the nodes.

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

7   Verify and confirm the information that the installer summarizes.

Configuring the cluster in secure mode

Before you configure a cluster in a secure mode, make sure to meet the
requirements for automatic or semiautomatic mode of configuration. You can
also enable Symantec Product Authentication Service later.

See "Symantec Product Authentication Service" on page 20.

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 have
         installed the root broker and enter **y**.

     ■   If you do not want to configure the cluster in secure mode and want to
         proceed to adding VCS users, enter **n**.
         See "Adding VCS users" on page 56.

2   Select one of the options to enable security.

```
Select the Security option you would like to perform [1-3,q,?]
```
Based on the mode of configuration you want to use, enter one of the following:

Option	Tasks
1. Automatic configuration	Enter the name of the Root Broker system when prompted.
	Requires 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.
2 . Semi-automatic using encrypted files	Enter the path of the file for each node when prompted.
3. Semi-automatic entering authentication information at installer prompts	Enter the following Root Broker information as the installer prompts you:

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

3 After configuring the cluster in secure mode, proceed to configure the Cluster Management Console cluster connector.
See "Configuring cluster connector" on page 56.

Adding VCS users

If you have enabled Symantec Product Authentication Service, you do not need to add VCS users now.

Proceed to configure the Cluster Management Console cluster connector. Otherwise, on systems operating under an English locale, you can add VCS users at this time.

See "Configuring the cluster in secure mode" on page 54.

See "Configuring cluster connector" on page 56.

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 cluster connector

If you configured the Cluster Management Console management server to centrally manage this cluster, you can now configure cluster connector for the buffering feature. If a firewall exists between the management server and this cluster, then you must configure cluster connector to enable centralized management. Make sure you meet the prerequisites to configure cluster connector.

To configure cluster connector

1   Review the information to configure Cluster Management Console.

2   Choose whether to configure cluster connector or not. Do one of the following:

■   To configure cluster connector on the systems, press Enter.

```
Do you want this cluster to be managed by a management
server? Enter 'y' if you have set up a management server.
[y,n,q] (y) y
```

■   To skip configuring cluster connector and advance to configuring Cluster Management Console for local cluster management, enter **n**.

3   Review the required information to configure cluster connector.

4   Enter the management server network address for the Cluster Management Console.

```
Enter the network address used by the management server [?]
(north) mgmtserver1.symantecexample.com
```

5   Verify and confirm the management server information.

6   Enter the following information that is required to securely communicate with the management server.

■   Password for the service account that is created during the management server installation

■   Hash of the Cluster Management Console management server's root broker

7   Verify and confirm the information.

Configuring SMTP email notification

You can choose to configure VCS to send event notifications to SMTP e-mail services. You need to provide the SMTP server name and e-mail addresses of people to be notified. Note that it is also possible to configure 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 one of the following:

■   To configure SMTP notification, press **Enter**.

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

- ■ To skip configuring SMTP notification and advance to configuring SNMP notification, enter **n**.

  See "Configuring SNMP trap notification" on page 58.

3 Provide information to configure SMTP notification.

- ■ 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 it is also possible to configure 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 one of the
    following:

    ■   To configure SNMP notification, press **Enter**.

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

    ■   To skip configuring SNMP notification and advance to configuring
        global clustering option, enter **n**.
        See "Configuring global clusters" on page 60.

3   Provide information to configure SNMP trap notification.

    ■   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

You can configure global clusters to link clusters at separate locations and enable wide-area failover and disaster recovery. Note that you must have entered a valid license key for VCS global clusters.

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 one of the following:

   ■   To configure global cluster option, press **Enter**.

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

   ■   To skip configuring global cluster option and advance to installing VCS packages, enter **n**.
      See "Installing the VCS packages" on page 60.

3   Provide information to configure the Global Cluster option.
   If you configured Cluster Management Console to manage this cluster locally, the installer discovers and displays the virtual IP address and netmask used by the Cluster Management Console. You can use the same virtual IP address and netmask.
   Do one of the following:

   ■   If you want to use the default values, press **Enter**.

   ■   If you do not want to use the default value, enter another IP address. The installer prompts you for a NIC and value for the netmask.

      ```
      Enter the Virtual IP address for Global Cluster Option:
      [b,?] (10.10.12.1)
      ```

4   Verify and confirm the configuration of the global cluster.

   ```
   Global Cluster Option configuration verification:

   NIC: bge0
   IP: 10.10.12.1
   Netmask: 255.255.240.0

   Matching Cluster Management Console Virtual IP configuration

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

Installing the 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 61.

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

Creating VCS configuration files

After installing the packages, the installer continues to create configuration files and copies them to each system:

```
Creating Cluster Server configuration files ............ Done
Copying configuration files to north.................... Done
Copying configuration files to south.................... 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 creates security principal or executes the encrypted file to create security principal on each node in the cluster. The installer creates the VxSS service group, creates Authentication Server credentials on each node in the cluster, and Web credentials for VCS users, and sets up trust with the root broker. Then, the installer proceeds to start VCS in secure mode.

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.

To start VCS

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

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

Completing the installation

After VCS 5.0 installation completes successfully, the installer creates summary, log, and response files. The files provide useful information that can assist you with the installation and can also assist future installations. Table 3-5 specifies the files created at the end of the installation.

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

Table 3-5   File description

File	Description
summary file	■ Lists packages installed on each system. ■ Describes the cluster and its configured resources. ■ Provides information for managing the cluster.
log file	Details the entire installation.
response file	Contains configuration information that can be used to perform secure or unattended installations on other systems. See "Example response file" on page 70.

Copying the installation guide to each node

After you install VCS, Symantec recommends that you copy the PDF version of this guide from the installation disc (cluster_server/docs/vcs_install.pdf) to the directory /opt/VRTS/docs on each node to make it available for reference.

Setting up I/O fencing

Symantec recommends you to set up the I/O fencing feature to prevent data corruption in the event of a communication breakdown in the cluster. Make sure that you do the following before you set up I/O fencing:

■ Install a version of Veritas Volume Manager (VxVM) that supports SCSI-3 persistent reservations.

■ Verify that the disks you intend to use for shared data storage and for coordinator disks support SCSI-3 PR (Persistent Reservations).

See "Setting up I/O fencing" on page 97.

Installing the VCS 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, Windows 2000 Professional system, Windows XP, or Windows 2003 system, or Solaris x64 Platform Edition system with X-Windows. The system from which you run the Java Console can be a system in the cluster or a remote workstation; the latter enables each system in the cluster to be administered remotely.

When installing the Java Console on a Solaris x64 Platform Edition system, make sure a printer is configured to that system. On a system without a configured printer, printing from the online JavaHelp could cause the Java Console to hang.

For information about using the Cluster Manager and the Configuration Editor components of the Java Console, see the applicable chapter in the *Veritas Cluster Server User's Guide*.

Installing the Java Console on Solaris x64 Platform Edition

To install Java console on Solaris x64 Platform Edition

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 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, Windows 2000 Professional, Windows XP, or Windows 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.

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.

5   Double-click setup.exe.

6   The Veritas Cluster Manager Install Wizard guides you through the installation process.

Establishing cluster communication with the management server

You can also set up multiple-cluster management feature after VCS installation and configuration.

Use the following list to prepare clusters for administration and management through the Cluster Management Console.

■   Ensure that all clusters that you want to manage run a supported version of VCS.

■   Decide which clusters are to use cluster connector to communicate with the management server, and then install cluster connector on each cluster.
    Cluster connector is a process agent. You must use cluster connector if a firewall exists between the management server and any clusters. You can install cluster connector when you install VCS 5.0.
    See "Installing cluster connector" on page 65.

■   Decide which clusters are to use a direct connection to communicate with the management server. If these clusters run a supported version of VCS, they require no further preparation.

After you prepare the clusters for management server administration, start the Cluster Management Console and use it to configure the management server to connect to clusters using direct connection. Clusters using cluster connector connect to the management server automatically.

Refer to the Veritas Cluster Management Console documentation for more information.

Installing cluster connector

Note: You must stop all VCS Web consoles, VCS Java consoles, and agent wizards that are running on any cluster nodes before you install cluster connector.

For all cluster connector installations, cluster connector installs or upgrades Symantec Product Authentication Service (version 4.3) on target systems that do not already have it.

■ You are installing cluster connector (locally or remotely) from a system running the authentication broker.

■ You are installing cluster connector (locally or remotely) from a cluster node and that cluster node is in your install list.

For example, assume that you have nodes A, B, and C each in their own cluster; each have the authentication broker installed. You also have system X with no authentication broker. You cannot install cluster connector from X. You can install cluster connector from A to B and C to other nodes.

Perform this procedure to use cluster connector for management server communications when the cluster is a supported VCS cluster. You can also use this procedure if you want to install or configure cluster connector after installing VCS 5.0 on a cluster.

To install cluster connector on a Solaris x64 Platform Edition system

1 Insert the software disc into the drive on the local system.

2 Make sure that the disc is mounted.

3 Navigate to the folder containing the installcmc program.

4 Run the installcmc program.

   # **./installcmc -rsh**

5 Enter **2** if you are prompted to select a product component.

   ```
   Enter '1' to install the Management Server, '2' to install the
   Cluster Connector: [1-2,q] (1) 2
   ```
   The installer presents a message stating that it will install cluster connector.

6 Enter the name of one system in each cluster to be managed. Separate the system names with spaces.

   ```
   Enter the name of a system in each cluster that you want the
   management server to manage. Separate system names with spaces:
   system1 system2 system3
   ```
   The installer detects the systems that you enter, performs an initial check of those systems, and then checks for installed packages on those systems.

If these checks are satisfactory, the installer lists the packages to be installed.

7 Enter **y** to verify that the information up to this point is correct.

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

The installer performs an initial system check of the local system and checks for installed packages on the local system. If these checks are satisfactory, the installer program lists the packages to be installed.

8 Press Enter.

You may install Cluster Management Console packages without performing configuration. The installer program gives you the option to configure Cluster Management Console now, and provides instructions for configuring Cluster Management Console later.

9 Enter **y** to configure Cluster Management Console.

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

10 Enter the fully-qualified management server network address, such as:

`Enter the network address used by the management server [?]`
`mgmtserver1.symantec.com`

11 Enter **y** to verify that the information up to this point is correct.

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

12 Enter a password for the cluster connector service account.

The password is the password that was entered for the cluster connector service account during management server installation.

`Enter the password for the CMC service account:`**`xxxxx`**

13 Enter the root hash of the authentication broker installed on the management server, which you can get from the Root Broker administrator.

`Enter the hash of the Management Server's root broker [?]`

- On Windows:

  `\program files\veritas\security\authentication\bin\vssat`
  `showbrokerhash`

- On UNIX systems:

  `/opt/VRTSat/bin/vssat showbrokerhash`

  The output of this command looks similar to the following:

  `Root Hash:      9dfde3d9aaebee084f8e35819c1fed7e6b01d2ae`

  Enter the alphanumeric string (the string you receive is different from the one shown).

14 Enter **y** to verify that the information up to this point is correct.

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

The installer presents:

- Installation progress percentages

- Authentication status messages

■   Cluster connector configuration status messages

15   Enter **y** to start Veritas Cluster Management Console processes.

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

The installer presents startup progress percentages and, if successful, displays the following message:

```
Startup completed successfully on all systems
```

16   Enter an encryption key of at least five characters.

This key is used to encrypt passwords in the response file. It must be referred to using the `-enckeyfile` option if the generated installation response file is to be used again.

```
A string of five or more characters is required to encrypt
passwords in the responsefile
Enter five or more characters to be used an encryption key:
```
xxxxx
```
Press [Return] to continue:
```

17   Press Enter.

Record the location that the installer program provides for the installation log files, summary file, and response file.

Verifying the cluster after installation

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

See "Verifying the VCS installation" on page 129.

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 entered 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 to 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 cluster" on page 53.

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

Performing VCS installation in a secure environment

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, a "response" file is created.

See "Example response file" on page 70.

Note that a response file generated by the installvcs program contains descriptions and explanations of the variables and their values. By copying this file to the other systems in the cluster and editing 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 50.

2   Review the output as the installer performs the initial system checks.
    The installer detects the inability to communicate between systems.

3   Press Enter 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 north only
    and create a responsefile for systems south? [y,n,q] (y)
    ```

4   Enter all cluster information. Proceed with the installation and configuration tasks.
    See "Installing and configuring VCS 5.0" on page 48.
    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 response file within the file /opt/VRTS/install/logs/ installvcs-*universaluniqueidentifier/*installvcs-*universaluniqueidentifier*.response.

6   If you start VCS before VCS is installed and started on all nodes in the
    cluster, you will see the output similar to:

    ```
    VCS:11306:Did not receive cluster membership, manual
    intervention may be needed for seeding
    ```

7   Using 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 described 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:

    ■   Mount the product disc.
        See "Mounting the product disc" on page 39.

    ■   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 50.

10 Repeat step 7 through step 9 until VCS has been installed on all nodes in the
    cluster.

Performing automated installations

Using installvcs program with the -responsefile option is useful not only for
installing and configuring VCS within a secure environment, but 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.

Assuming the systems are set up and meet the requirements for installation, you can perform unattended 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 used in response file

The syntax of Perl statements included in the response file varies, depending on whether "Scalar" or "List" values are required by the variables.

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 created by installvcs after the example VCS installation. It 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 required for installation.

See Table 3-6, "Response file variables."

```
#
# installer configuration values:
#
$CPI::CFG{AT_ROOTDOMAIN}="root\@sxvcs505.cdc.veritas.com";
$CPI::CFG{CMC_CC_CONFIGURED}=1;
$CPI::CFG{CMC_CLUSTERID}{IFRT01}=1169779551;
$CPI::CFG{CMC_MSADDR}{IFRT01}="sxvcs003.cdc.veritas.com";
$CPI::CFG{CMC_MSADDR}{sxvcs003}="sxvcs003.cdc.veritas.com";
$CPI::CFG{CMC_MS_ROOT_HASH}="wgtl0912";
$CPI::CFG{CMC_SERVICE_PASSWORD}="U2FsdGVkX1/W/
hFxuqvTgtyQhBBEjDeE";
$CPI::CFG{ENCRYPTED}="U2FsdGVkX18tyezIGKwtmQNRME+a4L3K2Uyx8nyEH
ns=";
$CPI::CFG{KEYS}{IFRT01}=[ qw(PZZV-T3KB-K2SX-OFA6-38CZ-8C) ];
$CPI::CFG{KEYS}{sxvcs003}=[ qw(PZZV-T3KB-K2SX-OFA6-38CZ-8C) ];
```

```
$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(IFRT01 sxvcs003) ];
$CPI::CFG{UPI}="VCS";
$CPI::CFG{VCS_ALLOWCOMMS}="Y";
$CPI::CFG{VCS_CLUSTERID}=8;
$CPI::CFG{VCS_CLUSTERNAME}="vcs_cluster3";
$CPI::CFG{VCS_CSGNETMASK}="255.255.248.0";
$CPI::CFG{VCS_CSGNIC}{ALL}="e1000g0";
$CPI::CFG{VCS_CSGVIP}="10.198.90.223";
$CPI::CFG{VCS_GCONETMASK}="255.255.248.0";
$CPI::CFG{VCS_GCOVIP}="10.198.90.223";
$CPI::CFG{VCS_LLTLINK1}{IFRT01}="e1000g2";
$CPI::CFG{VCS_LLTLINK1}{sxvcs003}="e1000g2";
$CPI::CFG{VCS_LLTLINK2}{IFRT01}="e1000g3";
$CPI::CFG{VCS_LLTLINK2}{sxvcs003}="e1000g3";
```

Response file variable definitions

Table 3-6 lists the variables used in the response file and their definitions. Note that while some variables are labeled as required and others as optional, some of the optional variables, if used, make it necessary to define other optional variables. For example, all variables 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), SNMP trap notification (SNMPPORT, SNMPCONS, and SNMPCSEV), and the Global Cluster Option (CGONIC, GCOVIP, and GCONETMASK).

Table 3-6   Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{OPT}{INSTALL}	Scalar	Req'd	List of systems where VCS must be installed and configured.
$CPI::CFG{OPT}{INSTALLONLY}	Scalar	Opt'l	List of systems where VCS packages must be installed. Configuration can be performed at a later time using the -configure option.

Table 3-6   Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{SYSTEMS}	List	Req'd	List of systems on which the product is to be installed, uninstalled, or configured.
$CPI::CFG{SYSTEMSCFG}	List	Opt'l	List of systems to be recognized in configuration if secure environment prevents all systems from being installed at once.
$CPI::CFG{UPI}	Scalar	Req'd	Defines the product to be installed, uninstalled, or configured.
$CPI::CFG{OPT}{KEYFILE}	Scalar	Opt'l	Defines the location of an ssh keyfile that is used to communicate with all remote systems.
$CPI::CFG{OPT}{LICENSE}	Scalar	Opt'l	Licenses VCS only.
$CPI::CFG{OPT}{NOLIC}	Scalar	Opt'l	Installs the product without any license.
$CPI::CFG{AT_ROOTDOMAIN}	List	Opt'l	Defines the name of the system where the root broker is installed.
$CPI::CFG{OPT}{PATCHPATH}	Scalar	Opt'l	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.
$CPI::CFG{OPT}{PKGPATH}	Scalar	Opt'l	Defines a location, typically an NFS mount, from which all remote systems can install product packages. The location must be accessible from all target systems.

Table 3-6 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{OPT}{TMPPATH}	Scalar	Opt'l	Defines the location where a working directory is created to store temporary files and packages needed during the install. The default location is /var/tmp.
$CPI::CFG{OPT}{RSH}	Scalar	Opt'l	Defines that `rsh` must be used instead of ssh as the communication method between systems.
$CPI::CFG{DONOTINSTALL} {*PACKAGE*}	List	Opt'l	Instructs the installation to not install the optional packages designated in the list.
$CPI::CFG{DONOTREMOVE} {*PACKAGE*}	List	Opt'l	Instructs the uninstallation to not remove the optional packages designated in the list.
$CPI::CFG{VCS_CLUSTERNAME}	Scalar	Req'd	Defines the name of the cluster.
$CPI::CFG{VCS_CLUSTERID}	Scalar	Req'd	An integer between 0 and 65535 that uniquely identifies the cluster.
$CPI::CFG{KEYS}{*SYSTEM*}	Scalar	Opt'l	List of keys to be registered on the system during installation.
$CPI::CFG{OPT_LOGPATH}	Scalar	Opt'l	Mentions the location where the log files are to be copied. following an installation. The default location is /opt/VRTS/install/logs.
$CPI::CFG{CONFIGURE}	Scalar	Opt'l	Performs configuration if the packages are already installed using the `-installonly` option.

Table 3-6   Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{VCS_LLTLINK#} {*SYSTEM*}	Scalar	Req'd	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.
$CPI::CFG{VCS_LLTLINKLOWPRI} {*SYSTEM*}	Scalar	Opt'l	Defines a low priority heartbeat link. Typically, LLTLINKLOWPRI is used on a public network link to provide an additional layer of communication.
$CPI::CFG{VCS_CSGNIC}	Scalar	Opt'l	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.
$CPI::CFG{CSGVIP}	Scalar	Opt'l	Defines the virtual IP address to be used by the Cluster Manager (Web Console).
$CPI::CFG{VCS_CSGNETMASK}	Scalar	Opt'l	Defines the Netmask of the virtual IP address to be used by the Cluster Manager (Web Console).
$CPI::CFG{VCS_SMTPSERVER}	Scalar	Opt'l	Defines the domain-based hostname (example: smtp.symantecexample.com) of the SMTP server to be used for web notification.
$CPI::CFG{VCS_SMTPRECP}	List	Opt'l	List of full email addresses (example: user@symantecexample.com) of SMTP recipients.

Table 3-6 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{VCS_SMTPRSEV}	List	Opt'l	Defines 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.
$CPI::CFG{VCS_SNMPPORT}	Scalar	Opt'l	Defines the SNMP trap daemon port (default=162).
$CPI::CFG{VCS_SNMPCONS}	List	Opt'l	List of SNMP console system names
$CPI::CFG{VCS_SNMPCSEV}	List	Opt'l	Defines 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.
$CPI::CFG{VCS_GCONIC} {*SYSTEM*}	Scalar	Opt'l	Defines the NIC for the Virtual IP used for the Global Cluster Option. 'ALL' can be entered as a system value if the same NIC is used on all systems.
$CPI::CFG{VCS_GCOVIP}	Scalar	Opt'l	Defines the virtual IP address to be used by the Global Cluster Option.
$CPI::CFG{VCS_GCONETMASK}	Scalar	Opt'l	Defines the Netmask of the virtual IP address to be used by the Global Cluster Option
$CPI::CFG{VCS_USERENPW}	List	Opt'l	List of encoded passwords for users
$CPI::CFG{VCS_USERNAME}	List	Opt'l	List of names of users
$CPI::CFG{VCS_USERPRIV}	List	Opt'l	List of privileges for users

Table 3-6   Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{OPT}{UNINSTALL}	Scalar	Opt'l	Defines the list of systems, where VCS must be uninstalled.

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 output to determine:

    ■   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.

```
License Key                        = xxx-xxx-xxx-xxx-xxx
    Product Name                       = Veritas Cluster Server
    Serial Number                      = 1006
    License Type                       = PERMANENT
    OEM ID                             = 2006
    Editions Product = YES

Features :=
    Platform                           = Solaris_x86
    Version                            = 5.0
    Tier                               = Unused
    Reserved                           = 0
    Mode                               = VCS
    Global Cluster Option = Enabled
```

From the output, you can determine the license key, the type of license, the product for which it applies, and its expiration date, if any. Demo keys have expiration dates, permanent keys and sites do not.

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 are using a demo license, you can replace the demo license.

See "Replacing a VCS demo license with a permanent license" on page 77.

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 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.

5   Start VCS on each node:
```
# hastart
```

About installvcs command options

Table 3-7 lists the `installvcs` 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 3-7   installvcs options

Option and Syntax	Description
-configure	Configure VCS after using -installonly option to install VCS. See "Configuring VCS using configure option" on page 67.
-enckeyfile encryption_key_file	See the -responsefile and the -encrypt options.
-encrypt password	Encrypt password using the encryption key provided with the -enckeyfile option so that the encrypted password can be stored in response files.
-installonly	Install product packages on systems without configuring VCS. See "Installing VCS using installonly option" on page 67.
-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.
-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.

Table 3-7   installvcs options

Option and Syntax	Description
`-nostart`	Bypass starting VCS after completing installation and configuration.
`-patchpath` *patch_path*	Specifies that *patch_path* contains all patches to be installed by installvcs program on all systems; *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 to be installed by installvcs program on all systems; *pkg_path* is the complete path of a directory, usually NFS mounted.
`-precheck`	Verify that systems meet the installation requirements before proceeding with VCS installation. Symantec recommends doing a precheck before installing VCS. See "Checking the systems for installation" 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 system and configuration information 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 installerer*number*.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 "Performing VCS installation in a secure environment" on page 68. See "Performing automated installations" on page 69.

Table 3-7   installvcs options

Option and Syntax	Description
`-rootpath` *root_path*	Specifies that *root_path* is the root location for the installation of all packages. On Solaris x64 Platform Edition, `-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 pre-configured such that `rsh` commands between systems execute without prompting for passwords or confirmations
`-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 "Symantec Product Authentication Service" on page 20.
`-tmppath` *tmp_path*	Specifies that *tmp_path*, not /var/tmp, is the working directory for installvcs program. This destination is where initial logging is performed and where packages are copied on remote systems before installation.

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 package 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 removing it.

Prerequisites

■   Before removing VCS from any node in the cluster, you must shut down applications such as Java Console or any VCS enterprise agents that depend on VCS.

■   If you do not plan to remove VCS from all nodes in the cluster, make sure that no service groups are running on the nodes from which VCS is uninstalled. You must also reconfigure VCS on the remaining nodes.
    See "Adding and removing cluster nodes" on page 155.

Uninstalling VCS 5.0

The example demonstrates how to uninstall VCS on two nodes: north and south.

Removing VCS 5.0 packages

The program stops the VCS processes that are currently running during the uninstallation process.

To uninstall VCS

1   Do one of the following to begin uninstalling:

    ■   If you can execute commands as superuser on the remote nodes in the cluster using ssh or rsh without supplying a password, run uninstallvcs program on one node to uninstall VCS on all nodes in the cluster.

    ■   If you cannot execute commands as superuser on remote nodes in the cluster using ssh or rsh, you must run uninstallvcs program on each node in the cluster.

2   Start uninstallvcs program.
    ```
    # cd /opt/VRTS/install
    # ./uninstallvcs
    ```
    The program specifies the directory where the logs are created and displays a copyright notice followed by 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: north south
        Service Groups: ClusterService groupA groupB
    ```

3   Answer the prompt to proceed with uninstalling the software.

    ■   To uninstall VCS on all nodes, press **Enter**.

    ■   To uninstall VCS only on specific nodes, enter **n**.
        Note that if you enter **n** or if no VCS configuration files are found on the local node, the uninstallvcs program prompts you to enter a list of nodes from which you want to uninstall VCS.
        ```
        Do you want to uninstall VCS from these systems? [y,n,q] (y)
        ```

4   Review the output as the uninstallvcs program continues to verify communication between systems and check the installations on each system to determine the packages to be uninstalled.

5   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.

6   Review the uninstaller report after the verification.

7   Press Enter to uninstall the VCS packages.

    ```
    Are you sure you want to uninstall VCS packages? [y,n,q] (y)
    ```

8   Review the output as the uninstaller stops processes, unloads kernel modules, and removes the packages.

9   Note the location of summary and log files that the uninstaller creates after removing all the packages.

Running uninstallvcs from the VCS 5.0 disc

If you need to uninstall VCS after an incomplete installation, or if the uninstallvcs program is not available in /opt/VRTS/install, you may need to use the uninstallvcs program on the VCS 5.0 disc.

Manually installing and configuring VCS

This chapter contains the following topics:

- About VCS manual installation
- Requirements for installing VCS
- JumpStart
- Installing VCS software manually
- Removing VCS packages manually

About VCS manual installation

You can manually install and configure VCS instead of using the installvcs program. Perform a manual installation when:

- You are installing a single VCS package.
- You are installing VCS to one system in a cluster already running VCS 5.0.
- You are unable to install on a system over the network. This can occur when you do not have remote root user access.

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

Requirements for installing VCS

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

See "Preparing to install and configure VCS" on page 19.

JumpStart

VCS is JumpStart compliant. When configuring the JumpStart server, make sure to install the following sections, and the packages and patches therein, in this order:

■ "Installing VCS packages for a manual installation" on page 86

For more information on using JumpStart, refer to the appropriate Sun Solaris documentation.

Installing VCS software manually

Table 4-1 lists the tasks involved in manually installing and configuring VCS 5.0.

Table 4-1   Tasks involved in manually installing VCS 5.0

Task	Reference
Install VCS software manually on each node in the cluster.	"Installing VCS packages for a manual installation" on page 86
Add a license key.	"Adding a license key" on page 87
Restore the configuration files from your previous VCS installation.	"Upgrading the configuration files" on page 88
Install the VCS cluster manager.	"Installing the Cluster Manager" on page 88
Copy the installation guide to each node.	"Copying the installation guide to each node" on page 88
Configure LLT and GAB.	"Configuring LLT and GAB" on page 88
Configure VCS.	"Configuring VCS" on page 91
Start LLT, GAB, and VCS services.	"Starting LLT, GAB, and VCS" on page 92
Modify the VCS configuration.	"Modifying the VCS configuration" on page 93
Replace demo license with a permanent license.	"Replacing a VCS demo license with a permanent license" on page 93

Preparing for a manual installation

Before you install, log in as the superuser. You then mount the disc and put the files in a temporary folder for installation.

See "Mounting the product disc" on page 39.

To prepare for installation

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

   # `cp -r cluster_server/pkgs/* /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:

```
# ls /tmp/install
info                VRTScscw.tar        VRTSllt.tar
VRTSvcsag.tar       SYMClma.tar         VRTScssim.tar
VRTScutil.tar       VRTSvcsdc.tar       VRTSacclib.tar
VRTSat.tar          VRTSpbx.tar         VRTSvcsmg.tar
VRTSvcsmn.tar       VRTSgab.tar         VRTSperl.tar
VRTSsmf.tar         VRTScmccc.tar       VRTSicsco.tar
VRTSjre.tar         VRTSvlic.tar        VRTScmcs.tar
VRTScscm.tar        VRTSspt.tar         VRTSvxfen.tar
VRTSweb.tar         VRTSjre15.tar       VRTSvcs.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 selecting the optional packages, note:

■   Symantec recommends that you install the packages for VCS manual pages (VRTSvcsmn) and VCS documentation (VRTSvcsdc). Install the documentation package on nodes where you want access to the documentation.

■   The I/O fencing package (VCSvxfen) can be used only with 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 "Setting up I/O fencing" on page 97.

■   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 required VCS files from the compressed files:

```
# tar -xvf VRTSvlic.tar
# tar -xvf VRTSperl.tar
```

```
# tar -xvf VRTSicsco.tar
# tar -xvf VRTSpbx.tar
# tar -xvf VRTSsmf.tar
# tar -xvf VRTSat.tar
# tar -xvf VRTSspt.tar
# tar -xvf SYMClma.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 VRTSjre.tar
# tar -xvf VRTSjre15.tar
# tar -xvf VRTScutil.tar
# tar -xvf VRTScscw.tar
# tar -xvf VRTSweb.tar
# tar -xvf VRTSacclib.tar
```

2   Install the required packages in the order shown:

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

3   Extract the optional packages:

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

4   Install the optional packages, in the order shown. Omit the packages that you do not want to install.

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

Adding a license key

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

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 license key, the type of license, the product for which it applies, and its expiration date, if any. 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 packages to upgrade your cluster to VCS.

Installing the Cluster Manager

If you did not install Cluster Manager (the VCS Java-based graphical user interface package), VRTScscm, you can do it later.

Copying the installation guide to each node

After you install VCS, Symantec recommends that you copy the PDF version of this guide from the installation disc (cluster_server/docs/vcs_install.pdf) to the directory /opt/VRTS/docs on each node to make it available for reference.

Configuring LLT and GAB

VCS uses LLT and GAB to replace the functions of TCP/IP for VCS private network communications. LLT and GAB provide the performance and reliability required by VCS for these and other functions.

LLT and GAB must be configured as described in the following sections.

Configuring low latency transport (LLT)

To configure LLT, set up two files: /etc/llthosts and /etc/llttab on each node in the cluster.

Setting up /etc/llthosts

The file llthosts(4) is a database, containing one entry per 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 entries that resemble:

```
0 north
1 south
```

Setting up /etc/llttab

The /etc/llttab 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 llttab file in /opt/VRTSllt.

See "LLT directives" on page 89.

Using vi or another editor, create the file /etc/llttab that contains entries that resemble:

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

The first line must identify the system on which the file exists. In the example above, the value for set-node could be north, 0, or the file name /etc/nodename, provided the file contains the name of the system (north in this example). 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/llttab.

LLT directives

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

Table 4-2   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 listed in /etc/llthosts file.
	Note that LLT fails to operate if any systems share the same ID.

Table 4-2   LLT directives

Directive	Description
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 llttab to set optional static MAC addresses.
	The second argument to link is the device name of the network interface. Its format is *device_name:device_instance_number*. The remaining four arguments to link are defaults; these arguments should be modified only in advanced configurations. There should be one link directive for each network interface. LLT uses an unregistered Ethernet SAP of 0xCAFE. If the SAP is unacceptable, refer to the llttab(4) manual page for information on how to customize SAP. Note that IP addresses do not need to be assigned to the network device; LLT does not use IP addresses.
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 and, in addition to enabling VCS communication, broadcasts heartbeats to monitor each network connection.

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

Additional considerations for LLT

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. For example, if IP is configured on one public interface and LLT on another, and 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, configure the system to assign unique MAC addresses by setting the eeprom(1M) parameter local-mac-address? to true.

Configuring group membership and atomic broadcast (GAB)

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 a simple /etc/gabtab file:

```
/sbin/gabconfig -c -nN
```

Where the `-c` option configures the driver for use and `-n`*N* 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"
```

To edit the main.cf file

1   Log in as superuser, and move to the directory containing the configuration file:

   # **cd /etc/VRTSvcs/conf/config**

2   Using vi, or another text editor, edit the main.cf file, defining your cluster name and system names. Refer to the following example.

3   Save and close the file.

Refer to the *Veritas Cluster Server User's Guide* for a full description of the main.cf file, how to edit it and verify it.

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 a file named types.cf. 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

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 138.

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 by using the command line, the Cluster Management Console, or Cluster Manager (the VCS Java GUI). Refer to the *Veritas Cluster Server User's Guide* for information on using the Cluster Management Console and the Java Console.

You can also edit the main.cf file directly. See the *Veritas Cluster Server User's Guide* for information on the structure of the main.cf file.

Configuring the ClusterService group

When you have successfully installed VCS, and verified that LLT, GAB, and VCS work correctly, you can create a service group to include the optional features including 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 example of a system configured with a ClusterService group.

See the "Example main.cf, for clusters without the GCO option" on page 132.

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.

See "Checking licensing information on the system" on page 76.

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(1m)` command.

    # **hastop -local**

2   Unconfigure the GAB and 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** *gab_id*
    # **modunload -i** *llt_id*

7   Remove the VCS 5.0 packages in the following order:

    # **pkgrm VRTScmccc**
    # **pkgrm VRTScmcs**
    # **pkgrm VRTSacclib**
    # **pkgrm VRTScssim**
    # **pkgrm VRTScscm**
    # **pkgrm VRTSweb**
    # **pkgrm VRTScscw**
    # **pkgrm VRTScutil**
    # **pkgrm VRTSjre15**
    # **pkgrm VRTSjre**
    # **pkgrm VRTSvcsdc**
    # **pkgrm VRTSvcsmn**
    # **pkgrm VRTSvcsag**
    # **pkgrm VRTSvcsmg**
    # **pkgrm VRTSvcs**
    # **pkgrm VRTSvxfen**
    # **pkgrm VRTSgab**
    # **pkgrm VRTSllt**
    # **pkgrm SYMClma**
    # **pkgrm VRTSspt**
    # **pkgrm VRTSat**
    # **pkgrm VRTSsmf**

```
# pkgrm VRTSpbx
# pkgrm VRTSicsco
# pkgrm VRTSperl
# pkgrm VRTSvlic
```

Setting up I/O fencing

This chapter contains the following topics:

- About I/O fencing
- Preparing to configure I/O fencing
- Setting up I/O fencing for VCS
- Additional I/O fencing information
- How I/O fencing works in different event scenarios
- About the vxfenadm utility
- Troubleshooting I/O fencing

About I/O fencing

I/O Fencing is a feature that prevents data corruption in the event of a communication breakdown in a cluster.

Note: Symantec strongly recommends that you use VCS I/O fencing to deter potential split brain scenarios in your cluster.

Preventing data corruption with I/O fencing

To provide high availability, the cluster must be capable of taking corrective action when a node fails. In this situation, VCS configures its components to reflect the altered membership.

Problems arise when the mechanism that detects the failure breaks down because symptoms appear identical to those of a failed node. For example, if a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects and the remaining node takes corrective action. However,

the failure of private interconnects (instead of the actual nodes) would present identical symptoms and cause each node to determine its peer has departed. This situation typically results in data corruption because both nodes attempt to take control of data storage in an uncoordinated manner.

In addition to a broken set of private networks, other scenarios can generate this situation. If a system is so busy that it appears to stop responding or "hang," the other nodes could declare it as dead. This declaration may also occur for nodes using 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 even though the system later returns and begins write operations.

VCS uses a technology called I/O fencing to remove the risk associated with split brain. I/O fencing allows write access for members of the active cluster and blocks access to storage from non-members; even a node that is alive is unable to cause damage.

SCSI-3 persistent reservations

SCSI-3 Persistent Reservations (SCSI-3 PR) are required for I/O fencing and resolve the issues of using SCSI reservations in a clustered SAN environment. SCSI-3 PR enables access for multiple nodes to a device and simultaneously blocks access for other nodes.

SCSI-3 reservations are persistent across SCSI bus resets and support multiple paths from a host to a disk. In contrast, only one host can use SCSI-2 reservations with one path. If the need arises to block access to a device because of data integrity concerns, only one host and one path remain active. The requirements for larger clusters, with multiple nodes reading and writing to storage in a controlled manner, make SCSI-2 reservations obsolete.

SCSI-3 PR uses a concept of registration and reservation. Each system registers its own "key" with a SCSI-3 device. Multiple systems registering keys form a membership and establish a reservation, typically set to "Write Exclusive Registrants Only." The WERO setting enables only registered systems to perform write operations. For a given disk, only one reservation can exist amidst numerous registrations.

With SCSI-3 PR technology, blocking write access is as simple as removing a registration from a device. Only registered members can "eject" the registration of another member. A member wishing to eject another member issues a "preempt and abort" command. Ejecting a node is final and atomic; an ejected node cannot eject another node. In VCS, a node registers the same key for all paths to the device. A single preempt and abort command ejects a node from all paths to the storage device.

I/O fencing components

Fencing in VCS involves coordinator disks and data disks. Each component has a unique purpose and uses different physical disk devices. The fencing driver is vxfen.

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 added to a disk group are automatically fenced, as are new paths discovered to a device.

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 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 only register with them and do not need to reserve them.

These disks 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 coordinator disks before it can fence the peer from the data drives. This concept of racing for control of the coordinator disks to gain the ability to fence data disks is key to understanding prevention of split brain through fencing.

Dynamic Multipathing devices with I/O fencing

DMP allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. You can configure coordinator disks to use Veritas Volume Manager Dynamic Multipathing (DMP) feature.

For more information on using DMP, see the *Veritas Volume Manager Administrator's Guide*.

See "Updating /etc/vxfenmode file" on page 107

I/O fencing operations

I/O fencing, provided by the kernel-based fencing module (`vxfen`), performs identically on node failures and communications failures. When the fencing module on a node is informed of a change in cluster membership by the GAB module, it immediately begins the fencing operation. The node attempts to eject the key for departed nodes from the coordinator disks using the preempt and abort command. When the node successfully ejects the departed nodes from the coordinator disks, it ejects the departed nodes from the data disks. In a split brain scenario, both sides of the split would race for control of the coordinator disks. The side winning the majority of the coordinator disks wins the race and fences the loser. The loser then panics and reboots the system.

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 accompanying the Storage Foundation product that you are using.

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.

Checking shared disks for I/O fencing

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	Stores shared data
Coordinator disks	Act as a global lock during membership changes.
	Coordinator disks are small LUNs (typically three per cluster)

See "Setting up shared storage" on page 34.

Perform the following checks for I/O fencing disks:

- Identify three SCSI-3 PR compliant shared disks as coordinator disks.
  List the disks on each node and pick three disks as coordinator disks.
  For example, execute the following commands to list the disks:

```
# lsdev -Cc disk
```

■ Test the shared disks using the vxfentsthdw script.
See "Testing the shared disks for SCSI-3" on page 101.

Testing the shared disks for SCSI-3

Use the vxfentsthdw utility to test the shared storage arrays support SCSI-3 persistent reservations and I/O fencing. Review the guidelines to run vxfentsthdw program, verify that the systems see the same disk, and proceed to test the disks. Make sure to test disks serving as coordinator disks.

See "Setting up coordinator disk groups" on page 104.

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

Review these guidelines for using vxfentsthdw

■ Verify the connection of the shared storage for data to two of the nodes on which you installed VCS.

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

■ The two nodes must have ssh (default) or rsh communication. If you use rsh, launch the vxfentsthdw utility with the -n option.
See "Enabling communication between systems" on page 35.
After completing the testing process, remove permissions for communication and restore public network connections.
See "Removing permissions for communication" on page 109.

■ To ensure both nodes are connected to the same disk during the testing, use the vxfenadm -i *diskpath* command to verify the disk serial number.
See "Verifying the nodes see the same disk" on page 101.

Verifying the nodes see the same disk

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 to verify that the same serial number for the LUN is returned on all paths to the LUN.

For example, an EMC disk is accessible by the /dev/rdsk/c2t13d0s2 path on node A and the /dev/rdsk/c2t11d0s2 path on node B.

From node A, enter:

```
# vxfenadm -i /dev/rdsk/c2t13d0s2
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/c2t11d0s2 path.

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

```
# vxfenadm -i /dev/rdsk/c2t0d2s2
Vendor id        : HITACHI
Product id       : OPEN-3      -SUN
Revision         : 0117
Serial Number    : 0401EB6F0002
```

Refer to the vxfenadm(1M) manual page.

Testing the disks using vxfentsthdw script

This procedure uses the /dev/rdsk/c2t13d0s2 disk in the steps.

If the utility does not show a message stating a disk is ready, verification has failed. Failure of verification can be the result of an improperly configured disk array. It can also be caused by 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/c2t13d0s2 is ready to be configured for I/O
Fencing on node north
```

See "Adding or removing coordinator disks" on page 125.

To test disks using vxfentsthdw script

1 Make sure system-to-system communication is functioning properly. See "Enabling communication between systems" on page 35.

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 After reviewing the overview and warning that the tests overwrite data on the disks, confirm to continue the process and enter the node names.

```
******** 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: north
Enter the second node of the cluster: south
```

4 Enter the names of the disks you are checking. For each node, the same disk may be known by a different name:

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

If the disk names 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
north
ALL tests on the disk /dev/rdsk/c2t13d0s2 have PASSED
The disk is now ready to be configured for I/O Fencing on node
north
```

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

Setting up I/O fencing for VCS

Tasks involved in setting up I/O fencing include:

■ Initializing disks

■ Setting up coordinator disk groups

■ Stopping VCS on all nodes

■ Configuring /etc/vxfendg disk group for I/O fencing

■ Updating /etc/vxfenmode file

■ Starting I/O fencing

■ Modifying VCS configuration to use I/O fencing

■ Verifying I/O fencing configuration

Initializing 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 initialize them as VxVM disks and verify that all the nodes see the same disk. Use the example procedure; see the *Veritas Volume Manager Administrator's Guide* for more information on adding and configuring disks.

To initialize disks

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

    ```
    # lsdev -Cc disk
    ```

2   If the Array Support Library (ASL) for the array you are adding is not installed, obtain and install it on each node before proceeding. The ASL for the supported storage device you are adding 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 adding and configuring 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 format=cdsdisk
        ```

        The example specifies the CDS format:

        ```
        # vxdisksetup -i c2t13d0s2 format=cdsdisk
        ```

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

Setting up coordinator disk groups

I/O fencing requires coordinator disks that are configured in a disk group and accessible to each node in the cluster. These disks enables the vxfen driver to resolve potential split brain conditions and prevent data corruption. Make sure to meet the requirements for coordinator disks and then create the coordinator disk group.

Requirements for coordinator disks

After adding and initializing disks for use as coordinator disks, make sure coordinator disks meet the following requirements:

■   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.
     See "Testing the shared disks for SCSI-3" on page 101.

■   The coordinator disks must exist in a disk group (for example, vxfencoorddg).
     See "Creating the coordinator disk group and setting the coordinator attribute" on page 105.

■   Symantec recommends using hardware-based mirroring for coordinator disks.

Creating the coordinator disk group and setting the coordinator attribute

From one node, create a disk group named vxfencoorddg. This group must contain three disks or LUNs.

You must also set the coordinator attribute for the coordinator disk group. VxVM uses this attribute to prevent the reassignment of coordinator disks to other disk groups.

Note that if you create a coordinator disk group as a regular disk group, you can turn on the coordinator attribute in Volume Manager. You do this with a `vxdg set coordinator=on` command.

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

The example procedure assumes that the disks have the device names c1t1d0s0, c2t1d0s0, and c3t1d0s0.

To create the vxfencoorddg disk group

1   On any node, create the disk group by specifying the device name of the disks:

     # **vxdg -o coordinator=on init vxfencoorddg c1t1d0s0**

2   Add the other two disks to the disk group:

     # **vxdg -g vxfencoorddg adddisk c2t1d0s0**
     # **vxdg -g vxfencoorddg adddisk c3t1d0s0**

Stopping VCS on all nodes

Before configuring the coordinator disk for use, you must stop VCS on all nodes.

To stop VCS on all nodes

◆   On one node, enter:

     # **hastop -all**

Configuring /etc/vxfendg disk group for I/O fencing

After setting up the coordinator disk group, configure it for use.

To configure the disk group for fencing

1   Deport the 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 all 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.

     Based on the contents of the /etc/vxfendg and /etc/vxfenmode files, the rc script creates the /etc/vxfentab file for use by the vxfen driver when the system starts. The rc script also invokes the vxfenconfig command, which configures the vxfen driver to start and use the coordinator disks listed in /etc/vxfentab. The /etc/vxfentab file is a generated file; do not modify this file.

Example /etc/vxfentab file

The /etc/vxfentab file gets created when you start the I/O fencing driver.

See "Starting I/O fencing" on page 107.

An example of the /etc/vxfentab file on one node resembles:

■ Raw disk

```
/dev/rdsk/c1t1d0s2
/dev/rdsk/c2t1d0s2
/dev/rdsk/c3t1d0s2
```

■ 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.

See "Adding or removing coordinator disks" on page 125.

Updating /etc/vxfenmode file

You must update the /etc/vxfenmode file to operate in SCSI-3 mode. 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, either raw or dmp, on all the nodes.

To update /etc/vxfenmode file

◆ On all cluster nodes, depending on the SCSI-3 mechanism you have chosen, type:

■ 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
```

Starting I/O fencing

You now need to start I/O fencing on each node. VxFEN, the I/O fencing driver, may already be running, so you need to restart the driver for the new configuration to take effect.

To stop I/O fencing on a node

◆ Stop the I/O fencing driver.

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

To start I/O fencing on a node

◆ Start the I/O fencing driver.

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

Modifying VCS configuration to use I/O fencing

After adding coordinator disks and configuring 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   Make a backup copy of the main.cf file:

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

4   On one node, use vi or another text editor to edit the main.cf file. Modify the
list of cluster attributes by adding the UseFence attribute and assigning its
value of SCSI3.

```
cluster rac_cluster101
UserNames = { admin = "cDRpdxPmHpzS." }
Administrators = { admin }
HacliUserLevel = COMMANDROOT
CounterInterval = 5
UseFence = SCSI3
)
```

5   Save and close the file.

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

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

7   Using rcp or another utility, copy the VCS configuration file from a node (for
example, north) to the remaining cluster nodes.
For example, on each remaining node, enter:

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

8   On each node enter the following sequence of commands. These commands
brings up VCS processes:

```
# /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: raw
Cluster Members:

    * 0 (north)
    1 (south)

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

Removing permissions for communication

After completing the installation of VCS and verification of disk support for I/O fencing, if you used rsh, remove the temporary rsh access permissions 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.

Additional I/O fencing information

Review additional information about I/O fencing, including an extended description of the vxfentsthdw command, vxfenadm command, and a description of I/O fencing behavior to protect data in certain scenarios.

vxfentsthdw options

Table 5-3 describes the methods the utility provides to test storage devices.

Table 5-3   vxfentsthdw options

vxfentsthdw option	Description	When to use
-n	Utility uses `rsh` for communication.	Use when `rsh` is used for communication.
-r	Non-destructive testing. Testing of the disks for SCSI-3 persistent reservations occurs in a non-destructive way; that is, there is only testing for reads, not writes. May be used with -m, -f, or -g options.	Use during non-destructive testing.
-t	Testing of the return value of `SCSI TEST UNIT (TUR)` command under SCSI-3 reservations. A warning is printed on failure of TUR testing.	When you want to perform TUR testing.
-d	Use DMP devices. May be used with -c or -g options.	By default, the script picks up the OS paths for disks in the disk group. If you want the script to use the DMP path, use the -d option.
-c	Utility tests the coordinator disk group prompting for systems and devices, and reporting success or failure.	For testing disks in coordinator disk group.
-m	Utility runs manually, in interactive mode, prompting for systems and devices, and reporting success or failure. May be used with -r and -t options. -m is the default option.	For testing a few disks or for sampling disks in larger arrays.

Table 5-3   vxfentsthdw options

vxfentsthdw option	Description	When to use
-f *filename*	Utility tests system/device combinations listed in a text file. May be used with -r and -t options.	For testing several disks.
-g *disk_group*	Utility tests all disk devices in a specified disk group. May be used with -r and -t options.	For testing many disks and arrays of disks. Disk groups may be temporarily created for testing purposes and destroyed (ungrouped) after testing.

Testing the coordinator disk group using vxfentsthdw -c

Use the vxfentsthdw utility to verify disks are configured to support I/O fencing. In this procedure, the vxfentsthdw utility tests the three disks one disk at a time from each node.

- From the node north, the disks are /dev/rdsk/c1t1d0s2, /dev/rdsk/c2t1d0s2, and /dev/rdsk/c3t1d0s2.

- From the node south, the same disks are seen as /dev/rdsk/c4t1d0s2, /dev/rdsk/c5t1d0s2, and /dev/rdsk/c6t1d0s2.

Note: To test the coordinator disk group using the vxfentsthdw utility, the utility requires that the coordinator disk group, vxfencoorddg, be accessible from two nodes.

To test the coordinator disk group using vxfentsthdw -c

1   Use the `vxfentsthdw` command with the `-c` option. For example:

     # **/opt/VRTSvcs/vxfen/bin/vxfentsthdw -c vxfencoorddg**

2   Enter the nodes you are using to test the coordinator disks:

     ```
     Enter the first node of the cluster:
     north
     Enter the second node of the cluster:
     south
     ```

3   Review the output of the testing process for both nodes for all disks in the coordinator disk group. Each disk should display output that resembles:

     ```
     ALL tests on the disk /dev/rdsk/c1t1d0s2 have PASSED.
     The disk is now ready to be configured for I/O Fencing on node
     north as a COORDINATOR DISK.
     ```

```
ALL tests on the disk /dev/rdsk/c4t1d0s2 have PASSED.
The disk is now ready to be configured for I/O Fencing on node
south as a COORDINATOR DISK.
```

4   After you test all disks in the disk group, the vxfencoorddg disk group is
    ready for use.

Removing and replacing a failed disk

If a disk in the coordinator disk group fails verification, remove the failed disk or
LUN from the vxfencoorddg disk group, replace it with another, and retest the
disk group.

If you need to replace a disk in an active coordinator disk group, refer to the
troubleshooting procedure.

See "Adding or removing coordinator disks" on page 125.

To remove and replace a failed disk

1   Use the vxdiskadm utility to remove the failed disk from the disk group.
    Refer to the *Veritas Volume Manager Administrator's Guide*.

2   Add a new disk to the node, initialize it, and add it to the coordinator disk
    group.
    See "Initializing disks" on page 103.
    See "Setting up coordinator disk groups" on page 104.

3   Retest the disk group.

Using the -r option for non-destructive testing

To test disk devices containing data you want to preserve, you can use the -r
option with the -m, -f, or -g options, which are described in the following
sections. For example, to use the -m option and the -r option, you can run the
utility by entering:

    # **/opt/VRTSvcs/vxfen/bin/vxfentsthdw -rm**

When invoked with the -r option, the utility does not use tests that write to the
disks. Therefore, it does not test the disks for all of the usual conditions of use.

Using the -m option

Review the procedure to test the shared disks. The utility uses the -m option.

See "Testing the shared disks for SCSI-3" on page 101.

Using the -f option

Use the -f option to test disks that are listed in a text file. For example, you can create a file to test two disks shared by systems north and south that might resemble:

```
north /dev/rdsk/c2t2d1s2 south /dev/rdsk/c3t2d1s2
north /dev/rdsk/c2t2d1s2 south /dev/rdsk/c3t2d1s2
```

where the first disk is listed in the first line and is seen by north as /dev/rdsk/c2t2d1s2 and by south as /dev/rdsk/c3t2d1s2. The other disk, in the second line, is seen as /dev/rdsk/c2t2d2s2 from north and /dev/rdsk/c3t2d2s2 from south. Typically, the list of disks could be extensive.

Suppose you created the file named disks_blue. To test the disks, you would enter:

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw -f disks_blue
```

The utility reports the test results one disk at a time, just as for the -m option.

You can redirect the test results to a text file. Precede the command with "yes" to acknowledge that the testing destroys any data on the disks to be tested.

Caution: Be advised that by redirecting the command's output to a file, a warning that the testing destroys data on the disks cannot be seen until the testing is done.

For example:

```
# yes | /opt/VRTSvcs/vxfen/bin/vxfentsthdw -f disks_blue >
blue_test.txt
```

Using the -g option

Use the -g option to test all disks within a disk group. For example, you create a temporary disk group consisting of all disks in a disk array and test the group.

Note: Do not import the test disk group as shared; that is, do not use the -s option.

The utility reports the test results one disk at a time. You can redirect the test results to a text file for review.

```
# /opt/VRTSvcs/vxfen/bin/vxfentsthdw -g red_disks_dg >
redtest.txt
```

After testing, destroy the disk group and put the disks into disk groups as you need.

Testing a disk with existing keys

If the utility detects that a coordinator disk has existing keys, you see a message that resembles:

```
There are Veritas I/O Fencing keys on the disk. Please make sure
that I/O Fencing is shut down on all nodes of the cluster before
continuing.

        ******** WARNING!!!!!!!! ********

THIS SCRIPT CAN ONLY BE USED IF THERE ARE NO OTHER ACTIVE NODES
IN THE CLUSTER!  VERIFY ALL OTHER NODES ARE POWERED OFF OR
INCAPABLE OF ACCESSING SHARED STORAGE.

If this is not the case, data corruption will result.

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

The utility prompts you with a warning before proceeding. You may continue as long as I/O fencing is not yet configured.

About VXFEN tunable parameters

On each node, edit the file /kernel/drv/vxfen.conf to change the value of the vxfen driver tunable global parameter, vxfen_max_delay and vxfen_min_delay. You must restart the system to put change into effect.

Table 5-4 describes tunable parameters for the VXFEN driver.

Table 5-4   VXFEN tunable parameters

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
vxfen_debug_sz	Size of debug log in bytes ■ Values Default: 65536 Minimum: 65536 Maximum: 256K
vxfen_max_delay and vxfen_min_delay (See below.)	In the event of a network partition, the smaller cluster delays before racing for the coordinator disks. The time delayed allows a larger sub-cluster to win the race for the coordinator disks. The vxfen_max_delay and vxfen_min_delay parameters define the delay in seconds.

Table 5-4 VXFEN tunable parameters

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
vxfen_max_delay	Specifies the maximum number of seconds that the smaller sub-cluster waits before racing with larger clusters for control of the coordinator disks. This value must be greater than the vxfen_min_delay value. ■ Values Default: 60 Minimum: 0 Maximum: 600
vxfen_min_delay	Specifies the minimum number of seconds that the smaller sub-cluster waits before racing with larger clusters for control of the coordinator disks. This value must be smaller than the vxfen_max_delay value. ■ Values Default: 1 Minimum: 0 Maximum: 600

Example of implementing the vxfen_max_delay parameter

The following is an example of changing the vxfen_max_delay parameter.

The default file /kernel/drv/vxfen.conf looks like this:

```
#
# VXFEN configuration file
#
name="vxfen" parent="pseudo" instance=0 vxfen_debug_sz=65536
vxfen_max_delay=60 vxfen_min_delay=1;
```

Change the min_delay value to 30.

```
#
# VXFEN configuration file
#
name="vxfen" parent="pseudo" instance=0 vxfen_debug_sz=65536
vxfen_max_delay=60 vxfen_min_delay=30;
```

Close and save the file. For the changes to take effect, either restart the system, or reconfigure the VXFEN module.

See "Configuring the VXFEN parameters" on page 115.

Configuring the VXFEN parameters

For the parameter changes to take effect, reconfigure the VXFEN module.

To reconfigure the VXFEN module

1 Unconfigure the VXFEN module.

 # **/sbin/vxfenconfig -U**

2 Determine the VXFEN module ID:

 # **/usr/sbin/modinfo | grep -i vxfen**
 The module ID is the number in the first column of the output.

3 Unload the VXFEN module, using the module ID you determined:

 # **/usr/sbin/modunload -i *module_ID***

4 Configure the VXFEN module:

 # **/sbin/vxfenconfig -c**

5 Start VCS.

 # hastart

6 Bring the service groups online.

 # **hagrp -online oragrp -sys north**

How I/O fencing works in different event scenarios

Table 5-5 describes how I/O fencing works to prevent data corruption in
different failure event scenarios. For each event, corrective operator actions are
indicated.

Table 5-5   I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
Both private networks fail.	Node A races for majority of coordinator disks. If Node A wins race for coordinator disks, Node A ejects Node B from the shared disks and continues.	Node B races for majority of coordinator disks. If Node B loses the race for the coordinator disks, Node B removes itself from the cluster.	When Node B is ejected from cluster, repair the private networks before attempting to bring Node B back.
Both private networks function again after event above.	Node A continues to work.	Node B has crashed. It cannot start the database since it is unable to write to the data disks.	Restart Node B after private networks are restored.

Table 5-5   I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
One private network fails.	Node A prints message about an IOFENCE on the console but continues.	Node B prints message about an IOFENCE on the console but continues.	Repair private network. After network is repaired, both nodes automatically use it.
Node A hangs.	Node A is extremely busy for some reason or is in the kernel debugger. When Node A is no longer hung or in the kernel debugger, any queued writes to the data disks fail because Node A is ejected. When Node A receives message from GAB about being ejected, it removes itself from the cluster.	Node B loses heartbeats with Node A, and races for a majority of coordinator disks. Node B wins race for coordinator disks and ejects Node A from shared data disks.	Verify private networks function and restart Node A.

Table 5-5 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action

Nodes A and B and private networks lose power. Coordinator and data disks retain power.
 	
Power returns to nodes and they restart, but private networks still have no power.
	Node A restarts and I/O fencing driver (vxfen) detects Node B is registered with coordinator disks. The driver does not see Node B listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node A from joining the cluster. Node A console displays:
	
	`Potentially a`
	`preexisting split`
	`brain. Dropping`
	`out of the`
	`cluster. Refer to`
	`the user`
	`documentation for`
	`steps required to`
	`clear preexisting`
	`split brain.`
		Node B restarts and I/O fencing driver (vxfen) detects Node A is registered with coordinator disks. The driver does not see Node A listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node B from joining the cluster. Node B console displays:
		
		`Potentially a`
		`preexisting`
		`split brain.`
		`Dropping out of`
		`the cluster.`
		`Refer to the user`
		`documentation`
		`for steps`
		`required to`
		`clear`
		`preexisting`
		`split brain.`
			Resolve preexisting split brain condition.
			
			See "System panics to prevent potential data corruption" on page 123.

Table 5-5 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
Node A crashes while Node B is down. Node B comes up and Node A is still down.	Node A is crashed.	Node B restarts and detects Node A is registered with the coordinator disks. The driver does not see Node A listed as member of the cluster. The I/O fencing device driver prints message on console: `Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.`	Resolve preexisting split brain condition. See "System panics to prevent potential data corruption" on page 123.

Table 5-5   I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
The disk array containing two of the three coordinator disks is powered off.	Node A continues to operate as long as no nodes leave the cluster.	Node B continues to operate as long as no nodes leave the cluster.	
Node B leaves the cluster and the disk array is still powered off.	Node A races for a majority of coordinator disks. Node A fails because only one of three coordinator disks is available. Node A removes itself from the cluster.	Node B leaves the cluster.	Power on failed disk array and restart I/O fencing driver to enable Node A to register with all coordinator disks.

About the vxfenadm utility

Administrators can use the `vxfenadm` command to troubleshoot and test fencing configurations. The command's options for use by administrators are:

-g	read and display keys
-i	read SCSI inquiry information from device
-m	register with disks
-n	make a reservation with disks
-p	remove registrations made by other systems
-r	read reservations
-x	remove registrations

Registration key formatting

The key defined by VxVM associated with a disk group consists of seven bytes maximum. This key becomes unique among the systems when the VxVM prefixes it with the ID of the system. The key used for I/O fencing, therefore, consists of eight bytes.

0   7

Node ID	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined	VxVM Defined

The keys currently assigned to disks can be displayed by using the vxfenadm command.

For example, from the system with node ID 1, display the key for the disk /dev/rdsk/c2t1d0s2 by entering:

```
# vxfenadm -g /dev/rdsk/c2t1d0s2
Reading SCSI Registration Keys...
Device Name: /dev/rdsk/c2t1d0s2
Total Number of Keys: 1
key[0]:
        Key Value [Numeric Format]: 65,45,45,45,45,45,45,45
        Key Value [Character Format]: A-------
```

The -g option of vxfenadm displays all eight bytes of a key value in two formats. In the numeric format, the first byte, representing the Node ID, contains the system ID plus 65. The remaining bytes contain the ASCII values of the letters of the key, in this case, "-------". In the next line, the node ID 0 is expressed as "A;" node ID 1 would be "B."

Troubleshooting I/O fencing

Headings indicate likely symptoms or procedures required for a solution.

Node is unable to join cluster while another node is being ejected

A cluster that is currently fencing out (ejecting) a node from the cluster prevents a new node from joining the cluster until the fencing operation is completed. The following are example messages that appear on the console for the new node:

```
...VCS FEN ERROR V-11-1-25 ... Unable to join running cluster
...VCS FEN ERROR V-11-1-25 ... since cluster is currently
fencing
...VCS FEN ERROR V-11-1-25 ... a node out of the cluster.

...VCS GAB.. Port b closed
```

If you see these messages when the new node is booting, the vxfen startup script on the node makes up to five attempts to join the cluster. If this is not sufficient to allow the node to join the cluster, restart the new node or attempt to restart vxfen driver with the command:

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

vxfentsthdw fails when SCSI TEST UNIT READY command fails

If you see a message resembling:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node
FAILED.
Contact the storage provider to have the hardware configuration
fixed.
The disk array does not support returning success for a SCSI
TEST UNIT READY command when another host has the disk reserved
using SCSI-3 persistent reservations. This happens with Hitachi
Data Systems 99XX arrays if bit 186 of the system mode option is
not enabled.
```

Removing existing keys from disks

Review the procedure to remove the registration and reservation keys created by another node from a disk.

To remove the registration and reservation keys from disk

1   Create a file to contain the access names of the disks:

```
# vi /tmp/disklist
```
For example:
```
/dev/rdsk/c1t0d11s2
```

2   Read the existing keys:

```
# vxfenadm -g all -f /tmp/disklist
```

The output from this command displays the key:

```
Device Name: /dev/rdsk/c1t0d11s2
Total Number Of Keys: 1
key[0]:
        Key Value [Numeric Format]:   65,49,45,45,45,45,45,45
        Key Value [Character Format]: A1------
```

3   If you know on which node the key was created, log in to that node and enter
    the following command:

```
# vxfenadm -x -kA1 -f /tmp/disklist
```

The key is removed.

4   If you do not know on which node the key was created, follow step 5 through
    step 7 to remove the key.

5   Register a second key "A2" temporarily with the disk:

```
# vxfenadm -m -k A2 -f /tmp/disklist
Registration completed for disk path /dev/rdsk/c1t0d11s2
```

6   Remove the first key from the disk by preempting it with the second key:

```
# vxfenadm -p -kA2 -f /tmp/disklist -vA1
key: A2------ prempted the key: A1------ on disk
/dev/rdsk/c1t0d11s2
```

7   Remove the temporary key assigned in step 5.

```
# vxfenadm -x -kA2 -f /tmp/disklist
Deleted the key : [A2------] from device /dev/rdsk/c1t0d11s2
```

No registration keys exist for the disk.

System panics to prevent potential data corruption

When a node experiences a split brain condition and is ejected from the cluster,
it panics and displays the following console message:

```
VXFEN:vxfen_plat_panic: Local cluster node ejected from cluster
to prevent potential data corruption.
```

How vxfen driver checks for pre-existing split brain condition

The vxfen driver functions to prevent an ejected node from rejoining the cluster
after the failure of the private network links and before the private network
links are repaired.

For example, suppose the cluster of system 1 and system 2 is functioning
normally when the private network links are broken. Also suppose system 1 is
the ejected system. When system 1 restarts before the private network links are
restored, its membership configuration does not show system 2; however, when
it attempts to register with the coordinator disks, it discovers system 2 is

registered with them. Given this conflicting information about system 2, system 1 does not join the cluster and returns an error from `vxfenconfig` that resembles:

```
vxfenconfig: ERROR: There exists the potential for a preexisting
split-brain. The coordinator disks list no nodes which are in
the current membership. However, they also list nodes which are
not in the current membership.

I/O Fencing Disabled!
```

Also, the following information is displayed on the console:

```
<date> <system name> vxfen: WARNING: Potentially a preexisting
<date> <system name> split-brain.
<date> <system name> Dropping out of cluster.
<date> <system name> Refer to user documentation for steps
<date> <system name> required to clear preexisting split-brain.
<date> <system name>
<date> <system name> I/O Fencing DISABLED!
<date> <system name>
<date> <system name> gab: GAB:20032: Port b closed
```

However, the same error can occur when the private network links are working and both systems go down, system 1 restarts, and system 2 fails to come back up. From the view of the cluster from system 1, system 2 may still have the registrations on the coordinator disks.

Case 1: system 2 up, system 1 ejected (actual potential split brain)

Determine if system1 is up or not. If it is up and running, shut it down and repair the private network links to remove the split brain condition. restart system 1.

Case 2: system 2 down, system 1 ejected (apparent potential split brain)

1   Physically verify that system 2 is down.

2   Verify the systems currently registered with the coordinator disks. Use the following command:

    # **`vxfenadm -g all -f /etc/vxfentab`**

    The output of this command identifies the keys registered with the coordinator disks.

3   Clear the keys on the coordinator disks as well as the data disks using the command `/opt/VRTSvcs/rac/bin/vxfenclearpre`.

    See "Clearing keys after split brain using vxfenclearpre command" on page 125.

4   Make any necessary repairs to system 2 and restart.

Clearing keys after split brain using vxfenclearpre command

When you have encountered a split brain condition, use the `vxfenclearpre` command to remove SCSI-3 registrations and reservations on the coordinator disks as well as on the data disks in all shared disk groups.

To clear keys after split brain

1   Shut down all other nodes in the cluster that have access to the shared storage. This prevents data corruption.

2   Start the script:
```
# cd /opt/VRTSvcs/vxfen/bin
# ./vxfenclearpre
```

3   Read the script's introduction and warning. Then, you can choose to let the script run.
```
Do you still want to continue: [y/n] (default : n) y
```
Informational messages resembling the following may appear on the console of one of the nodes in the cluster when a node is ejected from a disk/LUN:
```
<date> <system name> scsi: WARNING: /sbus@3,0/lpfs@0,0/
sd@0,1(sd91):
<date> <system name> Error for Command: <undecoded cmd 0x5f>
Error Level: Informational
<date> <system name> scsi: Requested Block: 0 Error Block 0
<date> <system name> scsi: Vendor: <vendor> Serial Number:
0400759B006E
<date> <system name> scsi: Sense Key: Unit Attention
<date> <system name> scsi: ASC: 0x2a (<vendor unique code
0x2a>), ASCQ: 0x4, FRU: 0x0
```
These informational messages may be ignored.
```
Cleaning up the coordinator disks...

Cleaning up the data disks for all shared disk groups...

Successfully removed SCSI-3 persistent registration and
reservations from the coordinator disks as well as the shared
data disks.

Reboot the server to proceed with normal cluster startup...
#
```

4   Restart all nodes in the cluster.

Adding or removing coordinator disks

Review the following information to:

■   Replace coordinator disk in the coordinator disk group

- Destroy a coordinator disk group

Note: Adding or removing coordinator disks requires all services be shut down.

Note the following about the procedure:

- A coordinator disk group requires three disks/LUNs.
- When adding a disk, add the disk to the disk group `vxfencoorddg` and retest the group for support of SCSI-3 persistent reservations.
- You can destroy the coordinator disk group such that no registration keys remain on the disks. The disks can then be used elsewhere.

To remove and replace a disk in the coordinator disk group

1 Log in as superuser on one of the cluster nodes.

2 If VCS is running, shut it down:
   # **hastop -all**

3 Stop I/O fencing on all nodes:
   # **/etc/init.d/vxfen stop**
   This removes any registration keys on the disks.

4 Import the coordinator disk group. The file /etc/vxfendg includes the name of the disk group (typically, `vxfencoorddg`) that contains the coordinator disks, so use the command:
   # **vxdg -tfC import `cat /etc/vxfendg`**
   where:
   -t specifies that the disk group is imported only until the node restarts.
   -f specifies that the import is to be done forcibly, which is necessary if one or more disks is not accessible.
   -C specifies that any import blocks are removed.

5 To remove disks from the disk group, use the VxVM disk administrator utility, `vxdiskadm`.
   You may also destroy the existing coordinator disk group. For example:
   # **vxdg destroy vxfencoorddg**

6 Add the new disk to the node, initialize it as a VxVM disk, and add it to the `vxfencoorddg` disk group.
   See "Creating the coordinator disk group and setting the coordinator attribute" on page 105.

7 Test the recreated disk group for SCSI-3 persistent reservations compliance.
   See "Testing the coordinator disk group using vxfentsthdw -c" on page 111.

8 After replacing disks in a coordinator disk group, deport the disk group:

vxdg deport `cat /etc/vxfendg`

9   On each node, start the I/O fencing driver:

/etc/init.d/vxfen start

10   If necessary, restart VCS on each node:

hastart

Verifying the VCS installation

This chapter contains the following topics:

- About verifying the VCS installation
- Verifying LLT and GAB configuration files
- Verifying the main.cf file
- Verifying LLT, GAB, and cluster operation
- Accessing the VCS documentation

About verifying the VCS installation

After successful installation, you can inspect the contents of the key configuration files that you have installed and modified during the process. These files reflect the configuration based on the information you supplied.

Verifying 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.

/etc/llthosts

The file `llthosts(4)` is a database, containing one entry per system, that 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 entries that resemble:

```
0 sxvcs001
1 sxvcs003
```

/etc/llttab

The file `llttab(1M)` contains 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 entries that resemble:

```
set-node sxvcs003
set-cluster 8
link e1000g2 /dev/e1000g:2 - ether - -
link e1000g3 /dev/e1000g:3 - 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, beginning with the `link` command, identify the two network cards used by the LLT protocol.

Refer to the `llttab(4)` manual page for details about how the LLT configuration may be modified. The manual page describes the ordering of the directives in the `llttab` file.

/etc/gabtab

After you install VCS, the file /etc/gabtab contains a `gabconfig(1)` command that configures the GAB driver for use.

The file /etc/gabtab contains a line that resembles:

```
/sbin/gabconfig -c -n2
```

where the `-c` option configures the driver for use and `-n`N 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: Symantec recommends not using the `-c -x` option for `/sbin/gabconfig`. The Gigabit Ethernet controller does not support the use of `-c -x`.

Verifying the main.cf file

The installation process creates the VCS configuration file /etc/VRTSvcs/conf/config/main.cf.

See "Example main.cf, for clusters without the GCO option" on page 132.

See "Example main.cf, for clusters with the GCO option" on page 134.

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 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 and includes the following:

  ■ The ClusterService service group includes the IP, NIC, and VRTSWebApp resources.

  ■ If you configured Cluster Management Console to manage this cluster locally, the main.cf includes the VRTSWebApp resource that includes AppName = cmc attribute.

  ■ If you configured Cluster Connector so that Cluster Management Console can centrally manage this cluster, the main.cf includes the CMC service group.
  The CMC service group includes the ClusterConnectorConfig and Process resources.

  ■ 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 installed VCS with the Global Cluster Option, the ClusterService service group contains an Application resource, wac (wide-area connector), whose 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* and review the chapter on configuration concepts for descriptions and examples of main.cf and types.cf files for Solaris x64 Platform Edition systems.

Example main.cf, for clusters without the GCO option

The following main.cf is for a secure cluster that is managed locally by Cluster Management Console.

```
include "types.cf"
include "ClusterConnectorConfigType.cf"

cluster vcs_cluster1 (
        UserNames = { "CMC_HAD@HA_SERVICES@ifrt01.cdc.veritas.com" =
Im,"CMC_HAD@HA_SERVICES@sxvcs003" = Im }
        ClusterAddress = "10.198.90.223"
        Administrators = {
"CMC_HAD@HA_SERVICES@ifrt01.cdc.veritas.com",
                "CMC_HAD@HA_SERVICES@sxvcs003" }
        SecureClus = 1
        )

system IFRT01 (
        )

system sxvcs003 (
        )

group CMC (
        SystemList = { IFRT01 = 0, sxvcs003 = 1 }
        AutoStartList = { IFRT01, sxvcs003 }
        OnlineRetryLimit = 3
        OnlineRetryInterval = 120
        )

        ClusterConnectorConfig CMC_ClusterConfig (
                MSAddress = "sxvcs003.cdc.veritas.com"
                MSPort = 14145
                ClusterId = 1169779551
                ClusterType = vcs
               Logging = "/opt/VRTScmccc/conf/cc_logging.properties"
                )

        Process CMC_ClusterConnector (
                PathName = "/bin/sh"
               Arguments = "/opt/VRTScmccc/bin/cluster_connector.sh"
                )

        CMC_ClusterConnector requires CMC_ClusterConfig
```

```
                    // resource dependency tree
                    //
                    //      group CMC
                    //      {
                    //      Process CMC_ClusterConnector
                    //          {
                    //          ClusterConnectorConfig CMC_ClusterConfig
                    //          }
                    //      }


         group ClusterService (
                 SystemList = { IFRT01 = 0, sxvcs003 = 1 }
                 AutoStartList = { IFRT01, sxvcs003 }
                 OnlineRetryLimit = 3
                 OnlineRetryInterval = 120
                 )

             IP webip (
                      Device = e1000g0
                        Address = "10.198.90.223"
                        NetMask = "255.255.248.0"
                        )

             NIC csgnic (
                      Device = e1000g0
                      )

             VRTSWebApp VCSweb (
                      Critical = 0
                      AppName = cmc
                      InstallDir = "/opt/VRTSweb/VERITAS"
                      TimeForOnline = 5
                      RestartLimit = 3
                      )

              VCSweb requires webip
             webip requires csgnic

             // resource dependency tree
             //
             //      group ClusterService
             //      {
             //      VRTSWebApp VCSweb
             //          {
             //          IP webip
             //              {
             //              NIC csgnic
             //              }
             //          }

         group VxSS (
```

```
                    SystemList = { IFRT01 = 0, sxvcs003 = 1 }
                    Parallel = 1
                    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
                    //      }
```

Example main.cf, for clusters with the GCO option

If you installed VCS with the Global Cluster option, note that the ClusterService
group also contains the Application resource, wac, required to control the
cluster in a Global Cluster environment.

```
.
.
group ClusterService (
    SystemList = { north = 0, south = 1 }
    AutoStartList = { north, south }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
    )

    Application wac (
        StartProgram = "/opt/VRTSvcs/bin/wacstart"
        StopProgram = "/opt/VRTSvcs/bin/wacstop"
        MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
        RestartLimit = 3
        )
.
.
```

Verifying LLT, GAB, and cluster operation

Before attempting to verify the operation of LLT, GAB, or the cluster, you must:

■   Log in to any node in the cluster as superuser.

■   Place the VCS command directory in your PATH variable:

    # **export PATH=$PATH:/usr/sbin:/sbin:/opt/VRTS/bin:/opt/VRTSvcs/
bin**

Note: If you are using 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.

Verifying LLT

Use the lltstat command to verify that links are active for LLT. This command 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.

Using lltstat -n

In the following example, lltstat -n is typed on each node in the cluster:

Node 1

    # **lltstat -n**

Output resembles:

```
LLT node information:
    Node            State      Links
    *0 north        OPEN       2
     1 south        OPEN       2
```

Node 2

    # **lltstat -n**

Output resembles:

```
LLT node information:
    Node            State      Links
     0 north        OPEN       2
    *1 south        OPEN       2
```

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

Using lltstat -nvv

With LLT configured correctly, the output of `lltstat -n` shows all the nodes in the cluster and two links for each node. If the output shows otherwise, you can use the verbose option of `lltstat`.

For example, type `lltstat -nvv | more` on a node to view additional information about LLT. In the following example, `lltstat -nvv | more` is typed on node north in a two-node cluster:

```
# lltstat -nvv | more
```

Output resembles:

```
Node        State       Link    Status      Address
*0 north    OPEN
                        e1000g:0 UP    08:00:20:93:0E:34
                        e1000g:1 UP    08:00:20:93:0E:34
 1 south    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
                        /dev/e1000g:1  DOWN
```

Note that the output lists 32 nodes. It reports on the two nodes in the cluster, north and south, plus non-existent nodes. For each correctly configured node, the information should show a state of OPEN, a status for each link of UP, and an address for each link. However, the output in the example shows that for the node south the private network connection is possibly broken or the information in the /etc/llttab file may be incorrect.

To obtain information about the ports open for LLT, type lltstat -p on any node. In the following example, lltstat -p is typed on one node in a two-node cluster:

Node 1

```
# lltstat -p
```

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

To verify that GAB is operating, type the following command on each node:

```
# /sbin/gabconfig -a
```

If GAB is operating, the following GAB port membership information is returned:

```
GAB Port Memberships
=================================
Port a gen a36e0003 membership 01
Port h gen fd570002 membership 01
```

Port a indicates that GAB is communicating, gen a36e0003 is a randomly generated number, and membership 01 indicates that nodes 0 and 1 are connected.

Port h indicates that VCS is started, gen fd570002 is a randomly generated number, and membership 01 indicates that nodes 0 and 1 are both running VCS.

If GAB is not operating, 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
```

For more information on GAB, refer to the *Veritas Cluster Server User's Guide.*

Verifying the cluster

To verify that the cluster is operating, type the following command:

```
# hastatus -summary
```

The output resembles:

```
-- SYSTEM STATE
-- System            State            Frozen

A   north            RUNNING          0
A   south            RUNNING          0

-- GROUP STATE
-- Group          System     Probed AutoDisabled   State

B  ClusterService north       Y      N             ONLINE
```

```
      B  ClusterService  south      Y        N              OFFLINE
```

Note the system state. If the value is RUNNING, VCS is successfully installed and running. The group state lists the ClusterService group, which is ONLINE on north and OFFLINE on south. Refer to the `hastatus(1M)` manual page. In the *Veritas Cluster Server User's Guide*, look for a description of system states and the transitions between them.

hasys -display

On one of the nodes, use the `hasys(1M)` command:

   # **/opt/VRTSvcs/bin/hasys -display**

On each node, the output should be similar. For more information on the `hasys -display` command, refer to the `hasys(1M)` manual page. Also refer to the *Veritas Cluster Server User's Guide* for information about administering VCS from the command-line.

The example shows the output when the `hasys -display` command is run on the node north; the list continues with similar information for south (not shown) and any other nodes in the cluster:

```
#System  Attribute          Value

north    AgentsStopped      0

north    AvailableCapacity  100

north    CPUBinding         BindTo None CPUNumber 0

north    CPUUsage           0

north    CPUUsageMonitoring Enabled 0 ActionThreshold 0
                            ActionTimeLimit 0 Action NONE
                            NotifyThreshold 0 NotifyTimeLimit 0

north    Capacity           100

north    ConfigBlockCount   130

north    ConfigCheckSum     46688

north    ConfigDiskState    CURRENT

north    ConfigFile         /etc/VRTSvcs/conf/config

north    ConfigInfoCnt      0

north    ConfigModDate      Fri May 26 17:22:48 2006

north    ConnectorState     Down
```

```
north    CurrentLimits

north    DiskHbStatus

north    DynamicLoad        0

north    EngineRestarted    0

north    EngineVersion      5.0.00.0

north    Frozen             0

north    GUIIPAddr

north    LLTNodeId          0

north    LicenseType        DEMO

north    Limits

north    LinkHbStatus       e1000g:0 UP e1000g:1 UP

north    LoadTimeCounter    0

north    LoadTimeThreshold  600

north    LoadWarningLevel   80

north    NoAutoDisable      0

north    NodeId             0

north    OnGrpCnt           1

north    ShutdownTimeout    120

north    SourceFile         ./main.cf

north    SysInfo            Solaris:thoropt70,Generic_118844-
                            27,5.10,i86pc

north    SysName            north

north    SysState           RUNNING

north    SystemLocation

north    SystemOwner

north    TFrozen            0

north    TRSE               0

north    UpDownState        Up

north    UserInt            0
```

```
north    UserStr

north    VCSFeatures        DR

north    VCSMode            VCS
```

Accessing the VCS documentation

If you had chosen to install the optional package VRTSvcsdc, then the directory /opt/VRTS/docs contains the documentation for VCS in Portable Document Format (PDF). The directory contains the following documents:

- vcs_users.pdf, *Veritas Cluster Server User's Guide*

- vcs_bundled_agents.pdf, *Veritas Cluster Server Bundled Agents Reference Guide*

- vcs_agent_dev.pdf, *Veritas Cluster Server Agent Developer's Guide*

Upgrading to VCS 5.0

This chapter includes the following topics:

- About upgrading to VCS 5.0
- Upgrading VCS using installvcs program
- Upgrading VCS in a secure environment
- Upgrading to the VCS 5.0 Java Console

About upgrading to VCS 5.0

VCS 5.0 supports the following upgrade paths:

From	To
VCS 4.1	VCS 5.0
VCS 4.1 MP1	VCS 5.0

Upgrading VCS using installvcs program

If you are currently running a VCS cluster, including a VCS global cluster, you can run the installer to upgrade to VCS 5.0. The program detects the current configuration and prompts you to indicate whether you want to upgrade. While the installer is not able to verify the validity of the existing configuration, it is able to run extensive upgrade requirement checks before proceeding.

See "About the VCS installation program" on page 46.

If you want to upgrade from an earlier version of VCS to VCS 5.0 and use the Global Cluster option, you must first upgrade to standard VCS 5.0. After adding a license for the Global Cluster option, you can run the gcoconfig wizard. Note

that the Global Cluster option requires a VCS HA/DR license. See the *Veritas Cluster Server User's Guide* for instructions.

Upgrading VCS to 5.0

The upgrade example demonstrates how to upgrade from VCS 4.1 and VCS 4.1 MP1 to VCS 5.0 on nodes IFRT01 and sxvcs003 using installvcs program.

Upgrade tasks include:

- "Removing deprecated resource types" on page 144
- "Installing VRTSperl in non-global zones" on page 145
- "Starting the upgrade" on page 145
- "Checking upgrade requirements" on page 146
- "Removing VCS packages from previous versions and installing VCS 5.0 packages" on page 147
- "Starting VCS" on page 148
- "Completing the upgrade" on page 149
- "Upgrading VCS in a secure environment" on page 150

Removing deprecated resource types

With VCS 5.0, 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 are:

- `Disk`
- `DiskReservation`
- `NFSLock`
- `ServiceGroupHB`
- `CampusCluster`

If you are using the resource type `ServiceGroupHB`, Symantec recommends the use of I/O fencing.

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`.

    a   Remove the resource of the deprecated resource types.
        You must modify the resource dependencies to ensure that the
        configuration works properly.

    b   Save the `main.cf`.

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.

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 ensures that HAD reads the configuration from the local node and
updates it on the remaining nodes.

Installing VRTSperl in non-global zones

To upgrade from VCS 4.1 MP1 to 5.0 in non-global zones, you must install the
VRTSperl package manually. This involves the following steps:

1   Check if any non-global zones exist before the upgrade.

2   Check if VRTSperl is installed in each non-global zone. If not, install
    VRTSperl manually in the non-global zone of each system.

3   Upgrade to VCS 5.0 using installvcs.

Starting the upgrade

Start the installvcs program to automatically upgrade the cluster nodes.

To start the upgrade

1   Log in as superuser.

2   Mount the software disc.
    See "Mounting the product disc" on page 39.

3   Ensure that HAD is running.

4   On the node where the disk was mounted, browse to the `cluster_server`
    directory and start the installvcs program:

    `# ./installvcs`

    See "Starting the software installation" on page 50.

5   After viewing a copyright notice, review the report after the program
    examines the configuration files and discovers the existing cluster
    configuration:

    ```
    VCS configuration files exist on this system with the following
    information:

        Cluster Name: test
        Cluster ID Number: 100
        Systems: IFRT01 sxvcs003
        Service Groups:

    No checks have been made to ensure the validity of the
    configuration.
    VCS version 4.1 is installed on this system
    ```

6   To upgrade to VCS 5.0, press **Enter**.

    ```
    Do you want to upgrade to version 5.0 on these systems using the
    current configuration? [y,n,q,?] (y) y
    ```

Checking upgrade requirements

The installer program verifies that you systems meet the requirements for
upgrading.

To check upgrade requirements

1   Review the output as the program verifies that the upgrade can proceed on
    the systems.

    The installvcs program checks the operating system level and system-to-
    system communication. It also creates logs for the VCS update.

    ```
    installvcs requires that ssh commands used between systems
    execute without prompting for passwords or confirmations.  If
    installvcs hangs or asks for a login password or hangs, stop
    installvcs and run it again with the ssh configured for
    passwordless logins, or configure rsh and use the -rsh option.

    Checking rsh communication with IFRT01................SunOS 5.10
    Checking VCS installation on IFRT01..........................4.1
    Checking architecture on IFRT01.............................i386
    Checking VCS installation on sxvcs003...................... 4.1
    Checking architecture on sxvcs003........................ i386
    ```

2   Review the output as the installer checks the licenses that are currently in
    place on the cluster nodes. The installer also prompts to add additional
    license keys:

```
Checking system licensing

Installing licensing packages
Permanent VCS license registered on IFRT01

Do you want to enter another license key for IFRT01? [y,n,q](n)
n
Permanent VCS license registered on sxvcs003

Do you want to enter another license key for sxvcs003?
[y,n,q](n)
```

Removing VCS packages from previous versions and installing VCS 5.0 packages

The installer program is ready to remove the packages from VCS version 4.1 and install the 5.0 packages on each node after verifying the requirements.

To remove VCS packages from previous versions and add 5.0 packages

1   Review the output as the installer checks for existing packages.

```
Checking installed packages on IFRT01
Checking installed packages on sxvcs003
```

2   Specify the packages to be installed on the cluster systems.

```
Additional packages are typically installed to simplify future
upgrades.
1)  All Veritas Cluster Server packages - 524 MB required
2)  Storage Foundation Enterprise HA packages - 1071 MB required
Select the packages to be installed on all systems? [1-2,q] (2)
1
```

3   The installer lists the packages that will be installed or upgraded. Press
    Enter at the prompt.

```
installvcs will upgrade or install the following VCS packages:
VRTSperl     Veritas Perl 5.8.8 Redistribution
VRTSvlic     Veritas Licensing
VRTSicsco    Symantec Common Infrastructure
VRTSpbx      Symantec Private Branch Exchange
VRTSsmf      Symantec Service Management Framework
VRTSat       Symantec Product Authentication Service
VRTSspt      Veritas Software Support Tools
SYMClma      Symantec License Inventory Agent
VRTSllt      Veritas Low Latency Transport
VRTSgab    Veritas Global Atomic Broadcast
VRTSvxfen    Veritas I/O Fencing
VRTSvcs      Veritas Cluster Server
VRTSvcsmg    Veritas Cluster Server Message Catalogs
VRTSvcsag    Veritas Cluster Server Bundled Agents
VRTSvcsmn    Veritas Cluster Server Man Pages
VRTSvcsdc    Veritas Cluster Server Documentation
```

```
VRTSjre      Veritas Java Runtime Environment Redistribution
VRTSjre15    Veritas Java Runtime Environment Redistribution
VRTScutil    Veritas Cluster Utilities
VRTScscw     Veritas Cluster Server Configuration Wizards
VRTSweb      Veritas Java Web Server
VRTScscm     Veritas Cluster Server Cluster Manager
VRTScssim    Veritas Cluster Server Simulator
VRTSacclib   Veritas ACC Library
VRTScmcs     Veritas Cluster Management Console for single
cluster environments
VRTScmccc    Veritas Cluster Management Console cluster
connector
Press [Enter] to continue:
```

4   The installer is now ready to upgrade VCS. Press **Enter** at the prompt.

```
installvcs is now ready to upgrade VCS.
All VCS processes that are currently running must be stopped.

Are you sure you want to upgrade VCS? [y,n,q] (y) y
Updating gabtab on IFRT01...................................
Done

installvcs must now make configuration updates and stop the
cluster before upgrading VCS packages

Are you ready to begin the Veritas Cluster Server upgrade at
this time? [y,n,q] (y)
```

5   View the output as the program backs up the types.cf and main.cf and freezes the service group configured in the cluster.

```
Backing up types.cf and main.cf........................... Done
Updating types.cf file ................................... Done
Updating gabtab on sxvcs003.............................. Done
```

6   The program attempts to stop VCS and uninstall the packages. Progress indicators show the status of these tasks.

7   The program now attempts to install VCS 5.0 packages. A progress indicator shows the status of the task.

Starting VCS

After the installvcs program upgrades the nodes, you can proceed to start VCS using the installvcs program.

To start VCS

1   The program prompts you to start VCS.

```
Installation completed successfully on all systems
Do you want to start Veritas Cluster Server processes now?
[y,n,q] (y)
```

■ If you have an NFS resource configured in the cluster, enter **n** and do not start VCS here. Edit the `main.cf` file to include details about the NFSRestart agent.

Refer to the *Veritas Cluster Server Bundled Agents Reference Guide* for details about editing the `main.cf` file. After editing this file, attempt to start VCS.

■ If you have configured resources for Oracle, Sybase, or DB2 databases, enter **n** and do not start VCS here.

Refer to the respective Enterprise agent documentation for upgrade information. After upgrading the agents, attempt to start VCS.

■ If you have none of the above resources configured on the node, enter **y** to start VCS.

```
Starting VCS: 0%
Starting VCS: 100%
```

2 The program attempts to unfreeze the service group and online the ClusterService group.

```
Unfreezing group VCS........................................ Done
Onlining ClusterService Group on IFRT01.................... Done
```

Completing the upgrade

After starting the cluster server processes, the upgrade is complete.

To complete upgrade

1 Note the locations of the summary and log files that the program creates.

```
Upgrade log files and summary file are saved at:

        /opt/VRTS/install/logs/installvcs-unique string/
```

2 Other files specific to the installer program are also created in the `/opt/VRTS/install/logs/installvcs-unique string/` directory.

3 Verify that the cluster is operating properly after the upgrade.
See "Verifying the VCS installation" on page 129.

Note: You must reboot your systems if the upgrade is from VCS 4.1 on Solaris 10, or if GAB or LLT cannot be stopped or unloaded successfully during the VCS 5.0 upgrade. Use the following procedure to unfreeze and online service groups as needed.

To unfreeze and online service groups after a system reboot

1 After VCS has started on all systems, verify all resources have been probed. On any system, type:

    # **hastatus -summary**

2 Unfreeze all service groups. On any system, type:

    # **haconf -makerw**
    # **hagrp -unfreeze service_group -persistent**
    # **haconf -dump -makero**

    Where service_group is the name of the service group.

3 Online the groups, if necessary. On any system type:

    # **hagrp -online service_ group -sys system**

    Where system is the system name.

Upgrading VCS in a secure environment

In a secure environment, run the installvcs program on each node to upgrade a cluster to VCS 5.0. On the first node, the program updates the configuration and stops the cluster before upgrading the system. On the other nodes, it uninstalls the previous version and installs VCS 5.0. After the last node is upgraded and started, the upgrade is complete.

Using the halogin command for native OS accounts with VCS

VCS has deprecated the AllowNativeCliUsers attribute. To use native OS accounts with VCS, use the halogin command. Once you run halogin, VCS encrypts and stores your VCS credentials in the your home directory for a specific time period. After running halogin, you do not need to authenticate yourself every time you run a VCS command. In secure clusters, the command also sets up a trust relationship and retrieves a certificate from an authentication broker.

Follow the instructions below if you used the AllowNativeCliUsers attribute.

To set up VCS authentication for clusters running in secure mode

Ensure that each native user running VCS commands must have a home directory on the system from which the user runs VCS commands.

1 Create VCS user accounts for all users and assign privileges to these users.

2 If you run VCS commands from a remote host, set these environment variables:

- ■ VCS_HOST—Name of the VCS node on which you run commands. You may specify the virtual IP address associated with the cluster.
- ■ VCS_DOMAIN—Name of the VxSS domain to which the user belongs.
- ■ VCS_DOMAINTYPE—Type of VxSS domain: `unixpwd`, `nt`, `nis`, `nisplus`, or `vx`.

3 Run the `halogin` command:

halogin *vcsusername password*

Perform steps 2 and 3 for each VCS user.

To set up VCS authentication for clusters not running in secure mode

Ensure that each native user running VCS commands must have a home directory on the system from which the user runs VCS commands.

1 Create VCS user accounts for all users and assign privileges to these users.

2 For each VCS user, run the `halogin` command:

halogin *vcsusername password*

Upgrading to the VCS 5.0 Java Console

When you upgrade to VCS release 5.0, you must also upgrade the Java Console (GUI) that you installed on a separate system. VCS 5.0 supports earlier versions of the Java Console, but earlier VCS versions cannot support later versions of Java Console.

The VCS 5.0 Java Console requires JRE version 1.5. If necessary, you can add it when you add the Java Console package.

Use one of the following applicable procedures:

To upgrade Java console on Solaris x64 Platform Edition

Log in as superuser on the node where you intend to install the package.

Remove the GUI from the previous installation.

pkgrm VRTScscm

Create a directory for installation of the Java Console.

mkdir /tmp/install

Insert the software disc with the VCS software into a drive connected to the system.

The Solaris volume-management software automatically mounts the disc as /cdrom/cdrom0. If it does not, type the command:

cd /cdrom/cdrom0

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

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

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

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

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

`cd /tmp/install`

`gunzip VRTScscm.tag.gz`

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

Extract the compressed file from the tar file.

`tar -xvf VRTScscm.tar`

Install the software:

`pkgadd -d . VRTScscm`

Answer `Yes` if prompted.

To upgrade Java console on Windows Systems

Remove the Java-based cluster manager from previous installations.

From the Control Panel, double-click **Add/Remove Programs**.

Select **Symantec Cluster Manager**.

Click **Add/Remove**.

Follow the uninstall wizard instructions.

Add the new Java-based Cluster Manager.

See "Installing the Java Console on a Windows system" on page 63.

Upgrading to the VCS 5.0 Java Console

When you upgrade to VCS release 5.0, you must also upgrade the Java Console (GUI) that you installed on a separate system. VCS 5.0 supports earlier versions of the Java Console, but earlier VCS versions cannot support later versions of Java Console.

Note: The VCS 5.0 Java Console requires JRE version 1.5. If necessary, you can add it when you add the Java Console package.

Use one of the following applicable procedures:

To upgrade Java console on Solaris x64 Platform Edition

1   Log in as superuser on the node where you intend to install the package.

2   Remove the GUI from the previous installation.

    # `pkgrm VRTScscm`

3   Create a directory for installation of the Java Console.

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

4   Insert the software disc with the VCS software into a drive connected to the system.

The Solaris volume-management software automatically mounts the disc as /cdrom/cdrom0. If it does not, type the command:

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

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

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

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

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

7   Go to the temporary directory and unzip the compressed package file.

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

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

8   Extract the compressed file from the tar file.

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

9   Install the software:

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

10 Answer **Yes** if prompted.

To upgrade Java console on Windows Systems

1   Remove the Java-based cluster manager from previous installations.

   ■   From the Control Panel, double-click **Add/Remove Programs**.

   ■   Select **Symantec Cluster Manager**.

   ■   Click **Add/Remove**.

   ■   Follow the uninstall wizard instructions.

2   Add the new Java-based Cluster Manager.
    See "Installing the Java Console on a Windows system" on page 63.

Adding and removing cluster nodes

This chapter contains the following topics:

■ About adding and removing nodes

■ Adding a node to a cluster

■ Removing a node from a cluster

About adding and removing nodes

After installing VCS and creating a cluster, you can add and remove nodes from the cluster. You can create clusters 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 "Preparing to install and configure VCS" on page 19.

Table 8-1 specifies the tasks involved in adding a cluster. The example demonstrates how to add a node east to already existing nodes, north and south.

Table 8-1   Tasks involved in adding a node to a cluster

Task	Reference
Set up the hardware.	"Setting up the hardware" on page 156
Install the software manually and add a license key.	"Installing the VCS software manually" on page 157

Table 8-1 Tasks involved in adding a node to a cluster

Task	Reference
Configure LLT and GAB.	"Configuring LLT and GAB" on page 157
Add the node to the existing cluster.	"Adding the node to the existing cluster" on page 158
Start VCS and verify the cluster.	"Starting VCS and verifying the cluster" on page 159

Setting up the hardware

Before configuring a new system to an existing cluster, you must physically add the system to the cluster.

Figure 8-1 Adding a node to a three-node cluster using two independent hubs
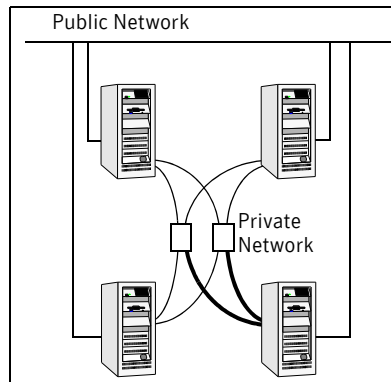
To set up the hardware

1   Connect the VCS private Ethernet controllers.

   ■   If you are expanding from a two-node cluster, you need to use independent hubs for the private network connections, replacing crossover cables if they are used.

   ■   If you already use independent hubs, connect the two Ethernet controllers on the new node to the independent hubs.

   Figure 8-1 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

Install the VCS 5.0 packages manually and install the license key.

See "Installing VCS software manually" on page 85.

See "Adding a license key" on page 87.

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 are adding east to a cluster consisting of north and south:

    ■   If the file on one of the existing nodes resembles:
        ```
        0 north
        1 south
        ```
    ■   Update the file for all nodes, including the new one, resembling:
        ```
        0 north
        1 south
        2 east
        ```

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.

    See "/etc/llttab" on page 130.

    The following example describes a system where node east is the new node on cluster number 2:

    ```
    set-node east
    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
        ```

then the file on the new node should be the same, although it is recommended to 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
```

then, 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
```

See "/etc/gabtab" on page 130.

The -n flag indicates to VCS the number of nodes required to 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 (north and south) 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 east**

3 Enter the following command:

    `# haconf -dump`

4 Copy the main.cf file from an existing node to your new node:

    `# rcp /etc/VRTSvcs/conf/config/main.cf east:/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 8-2 specifies the tasks 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 8-2   Tasks involved in removing a node

Task	Reference
■ Back up the configuration file. ■ Check the status of the nodes and the service groups.	"Verify the status of nodes and service groups" on page 160

Table 8-2   Tasks involved in removing a node

Task	Reference
■ Switch or remove any VCS service groups on the node leaving the cluster. ■ Delete the node from VCS configuration.	"Deleting the leaving node from VCS configuration" on page 161
Modify the llthosts and gabtab files to reflect the change.	"Modifying configuration files on each remaining node" on page 162
On the node leaving 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.	"Unloading LLT and GAB and removing VCS on the leaving node" on page 163

Verify 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 leaving node. Service group grp4 runs only on node C. Service groups grp1 and grp2 do not run on node C.

Deleting the leaving node from VCS configuration

Before removing a node from the cluster, you must remove or switch from the leaving node the service groups on which other service groups depend.

To remove or switch service groups from the leaving node

1   Switch failover service groups from the leaving 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 leaving node; for example, grp4 runs only on the leaving node.

    # **hagrp -dep**

3   If the service group on the leaving node requires other service groups, that is, if it is a parent to service groups on other nodes, then 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 leaving node:

    # **hastop -sys C**

5   Check the status again. The state of the leaving node should be EXITED. Also, any service groups set up for failover should be 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 leaving node from the SystemList of service groups grp3 and grp4.

```
# hagrp -modify grp3 SystemList -delete C
# hagrp -modify grp4 SystemList -delete C
```

7   For service groups that run only on the leaving node, delete the resources
    from the group before deleting the group.

```
# hagrp -resources grp4
    processx_grp4
    processy_grp4
# hares -delete processx_grp4
# hares -delete processy_grp4
```

8   Delete the service group configured to run on the leaving 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, although Symantec recommends using the -n*N* option,
    where *N* is the number of cluster systems.
    If the command has the form /sbin/gabconfig -c -n*N*, where *N* is the
    number of cluster systems, then make sure that *N* is not greater than the
    actual number of nodes in the cluster, or 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 leaving node.

For example, change:

```
0 A
1 B
2 C
```

to:

```
0 A
1 B
```

Unloading LLT and GAB and removing VCS on the leaving node

Perform the tasks on the node leaving 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 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 VRTScscm
# pkgrm VRTSvcsw
```

```
# pkgrm VRTSweb
# pkgrm VRTScscw
# pkgrm VRTScssim
# pkgrm VRTScutil
# pkgrm VRTSjre
# pkgrm VRTSvcsdc
# pkgrm VRTSvcsmn
# pkgrm VRTSvcsag
# pkgrm VRTSvcsmg
# pkgrm VRTSvcs
# pkgrm VRTSvxfen
# pkgrm VRTSgab
# pkgrm VRTSllt
# pkgrm VRTSat
# pkgrm VRTSperl
# pkgrm VRTSvlic
```

6   Remove the LLT and GAB configuration files.

```
# rm /etc/llttab
# rm /etc/gabtab
# rm /etc/llthosts
```

Installing VCS on a single node

This chapter contains the following topics:

About installing VCS on a single node

You can install VCS 5.0 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 166.

See "Creating a single-node cluster manually" on page 167.

Creating a single-node cluster using the installer program

Table 9-3 specifies the tasks involved in installing VCS on a single node using the installer program.

Table 9-3   Tasks to create a single-node cluster using the installer

Task	Reference
Prepare for installation.	"Preparing for a single node installation" on page 166
Install the VCS software on the system using the installer.	"Starting the installer for the single node cluster" on page 166

Preparing for a single node installation

You can use the installer program to install a cluster on a single system for two purposes:

■   To prepare the single node cluster to join a larger cluster

■   To prepare the single node cluster to be a standalone single node cluster

When you prepare it to join a larger cluster, install it with LLT and GAB. For a standalone cluster, you do not need to enable LLT and GAB.

For more information about LLT and GAB:

See "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 50.

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 standalone 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 51.

Creating a single-node cluster manually

Table 9-4 specifies the tasks involved in installing VCS on a single node.

Table 9-4   Tasks to create a single-node cluster manually

Task	Reference
Set the PATH variable	"Setting the PATH variable" on page 167
Install the VCS software manually and add a license key	"Installing VCS software manually" on page 85
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.	"Renaming the LLT and GAB startup files" on page 168
Create and modify the VCS configuration files.	"Configuring VCS" on page 168
Start VCS and verify single-node operation.	"Verifying single-node operation" on page 169

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
  ```

Installing the VCS software manually

Install the VCS 5.0 packages manually and install the license key.

See "Installing VCS software manually" on page 85.

See "Adding a license key" on page 87.

Renaming the LLT and GAB startup files

You may need the LLT and GAB startup files if you need 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

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"
```

To edit the main.cf file

1 Log in as superuser, and move to the directory containing the configuration file:

   # **cd /etc/VRTSvcs/conf/config**

2 Using vi, or another text editor, edit the main.cf file, defining your cluster name and system names. Refer to the following example.

3 Save and close the file.

Refer to the *Veritas Cluster Server User's Guide* for a full description of the main.cf file, how to edit it and verify it.

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 a file named types.cf. 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.

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(1M) 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

Table 9-5 specifies the activities involved in adding nodes 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 and the node that is to join Node A to form a multiple-node cluster as Node B.

Table 9-5   Tasks to add a node to a single-node cluster

Task	Reference
Set up Node B to be compatible with Node A	"Setting up a node to join the single-node cluster" on page 170
■ 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	"Installing and configuring Ethernet cards for private network" on page 171
Connect both nodes to shared storage	"Configuring the shared storage" on page 172
■ Bring up VCS on Node A ■ Edit the configuration file	"Bringing up the existing node" on page 172
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.	"Installing the VCS software manually" on page 172
Edit the configuration files on Node B	"See" on page 172
Start LLT and GAB on Node B	"Starting LLT and GAB" on page 175
■ Start LLT and GAB on Node A ■ Restart VCS on Node A ■ Modify service groups for two nodes	"Reconfiguring VCS on the existing node" on page 176
■ Start VCS on Node B ■ Verify the two-node cluster	"Verifying configuration on both nodes" on page 176

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:

   ■   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 removing the node from the cluster, remove the VCS packages and configuration files.
       See "Removing a node from a cluster" on page 159.

   ■   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 94.

   ■   If you renamed the LLT and GAB startup files, remove them.
       See "Renaming the LLT and GAB startup files" on page 168.

2   If necessary, install VxVM and VxFS.
    See "Installing VxVM, VxFS if necessary" on page 171.

Installing VxVM, VxFS if necessary

If VxVM with the cluster option or VxFS with the cluster option is installed on the existing node in the cluster, then the same versions must also be installed on the new node.

Refer to the appropriate documentation for VxVM and VxFS to verify the versions of the installed products and make sure the same version is running on all nodes that are to use any 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, two separate NICs on each node should be used, such that the failure of one NIC does not restrict heartbeating between the nodes.

See "Setting up the private network" on page 32.

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.

4   Install the ethernet card on Node B.

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 34.

Bringing up the existing node

1   Reboot each system.

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 listed 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

Install the VCS 5.0 packages manually and install the license key.

See "Installing VCS software manually" on page 85.

See "Adding a license key" on page 87.

See

Creating configuration files

1 Create the file /etc/llttab that lists both the nodes.
See "Setting up /etc/llttab" on page 89.

2 Create the file /etc/llthosts. Set up /etc/llthosts for a two-node cluster.
See to "Setting up /etc/llthosts" on page 88.

3 Create the file /etc/gabtab.
See "Configuring group membership and atomic broadcast (GAB)" on page 91.

Configuring LLT and GAB

VCS uses LLT and GAB to replace the functions of TCP/IP for VCS private network communications. LLT and GAB provide the performance and reliability required by VCS for these and other functions.

LLT and GAB must be configured as described in the following sections.

Configuring low latency transport (LLT)

To configure LLT, set up two files: /etc/llthosts and /etc/llttab on each node in the cluster.

Setting up /etc/llthosts

The file llthosts(4) is a database, containing one entry per 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 entries that resemble:

```
0 north
1 south
```

Setting up /etc/llttab

The /etc/llttab 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 llttab file in /opt/VRTSllt.

See "LLT directives" on page 174.

Using vi or another editor, create the file /etc/llttab that contains entries that resemble:

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

The first line must identify the system on which the file exists. In the example above, the value for `set-node` could be north, 0, or the file name /etc/nodename, provided the file contains the name of the system (north in this example). 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/llttab.

LLT directives

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

Table 9-6   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 listed 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 `llttab` to set optional static MAC addresses. The second argument to link is the device name of the network interface. Its format is *device_name*:*device_instance_number*. The remaining four arguments to link are defaults; these arguments should be modified only in advanced configurations. There should be one link directive for each network interface. LLT uses an unregistered Ethernet SAP of 0xCAFE. If the SAP is unacceptable, refer to the `llttab(4)` manual page for information on how to customize SAP. Note that IP addresses do not need to be assigned to the network device; LLT does not use IP addresses.
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 and, in addition to enabling VCS communication, broadcasts heartbeats to monitor each network connection.

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

Additional considerations for LLT

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. For example, if IP is configured on one public interface and LLT on another, and 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, configure the system to assign unique MAC addresses by setting the eeprom(1M) parameter local-mac-address? to true.

Configuring group membership and atomic broadcast (GAB)

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 a simple /etc/gabtab file:

```
/sbin/gabconfig -c -nN
```

Where the `-c` option configures the driver for use and `-n`*N* 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.

Starting LLT and GAB

On the new node, start LLT and GAB.

To start LLT and GAB

1   Start LLT on Node B.

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

2   Start GAB on Node B.

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

Reconfiguring VCS on the existing node

1 On Node A, create the files /etc/llttab, /etc/llthosts, and /etc/gabtab. Use the files created on Node B as a guide, customizing the /etc/llttab for Node A.

2 Start LLT on Node A.

`/etc/rc2.d/S70llt start`

3 Start GAB on Node A.

`/etc/rc2.d/S92gab 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 listed, add the node.

  # **`hagrp -modify` *`group`* `SystemList -add sysB 1`**

Verifying configuration on both nodes

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`

Advanced topics related to installing VCS

This appendix contains the following topics:

- Reconciling major/minor numbers for NFS shared disks
- LLT over UDP
- Minimal downtime upgrade

Reconciling major/minor numbers for NFS shared disks

Your configuration may include disks on the shared bus that support NFS. File systems exported by NFS can be configured 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 the drivers required by the devices (such as a Solaris partition or a VxVM volume), and minor number 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 on which a file system is mounted for export via 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 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 minor numbers that do not match on disk partitions

1   Complete the following steps. 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 used by the SCSI driver:
    # **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 associated with the device name that matches the name for Node A
    displayed in step 2, is "1."

4   Compare instance numbers the device in the output on each node.

    ■   If the instance number from one node is not used on the other (that is,
        it does not appear in the output of step 3), edit /etc/path_to_inst
        to make the second node's instance number similar to that of the first
        node.

■ If the instance numbers are being used on both nodes, edit /etc/
path_to_inst on both nodes. Change the instance number associated
with the device name to an unused number greater than the highest
number used by other devices. For example, the output of step b shows
the instance numbers used by all devices (from 0 to 29), so 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 VxVM volumes used by cluster systems.

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 16:41
/dev/vx/dsk/shareddg/vol3
```
On Node B, the output may resemble:
```
brw-------   1 root   root   36,43000 Mar 22 16:41
/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 used by Veritas
Volume Manager (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`), and 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 for instructions on reconciling the
    Veritas Volume Manager minor numbers, with specific reference to the
    `reminor` option.

    Node where the vxio driver number have been changed require rebooting.

LLT over UDP

VCS 5.0 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 when:

■   LLT must be used over WANs

■   When hardware, such as blade servers, do not support LLT over Ethernet

Performance considerations

Because LLT over UDP is slower that LLT over Ethernet, LLT over UDP should only be used when the hardware configuration makes it necessary.

Configuring LLT over UDP

Following is a checklist for configuring LLT over UDP. Examples are provided in the sections that follow.

■   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 182.
     See the examples in the following sections.

■   Make sure that each NIC has an IP address 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 non-well known UDP port.
     See "Selecting UDP ports" on page 184.

■   Set the broadcast address correctly for direct-attached (non-routed) links.

■   For 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 186.

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

Table A-7 describes the fields of the `link` command shown in the /etc/llttab file examples.

See "Sample configuration: Direct-attached links" on page 185.

See "Sample configuration: Links crossing IP routers" on page 186.

Note that some of these fields differ from the command for standard LLT links.

Table A-7   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 184.
`<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 having broadcasts enabled, 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. Table A-8 describes the fields of the `set-addr` command.

Table A-8   Field description for set-addr command in /etc/llttab

Field	Description
`<node-id>`	The ID of the cluster node; for example, 0.
`<link tag-name>`	The string used by LLT to identify the link; for example link1, link2,....
`<address>`	IP address assigned to the link for the peer node.

Selecting UDP ports

When selecting a UDP port, select an available 16-bit integer from the range described below.

■ Use available ports (that is, ports that are not in use)] in the private range 49152 to 65535

■ Do not use:

■ 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; UDP ports 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 LLT on subnets

You need to make sure to properly configure the netmask and broadcast address when nodes reside on different subnets.

Configuring the netmask

If you have nodes on different subnets, set the netmask so that the nodes can access the subnets in use.

For example:

- 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

If you have 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

The following illustration depicts a typical configuration of direct-attached links employing LLT over UDP.

The configuration represented by the following /etc/lltab file for Node 0 has directly attached crossover links or links connected through a hub or switch. These links do not cross routers.

Because LLT broadcasts requests peer nodes to discover their addresses, the addresses of peer nodes do not need to be specified in the /etc/lltab file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the /etc/lltab file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link <tag-name> <device> <node-range> <link-type> <udp port>
<MTU> <IP-address> <bcast-address>
link link1 /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 would resemble:

```
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

The following illustration depicts a typical configuration of links crossing an IP router employing LLT over UDP. The illustration shows just two nodes of a four-node cluster.

The configuration represented by the following `/etc/llttab` file for Node 1
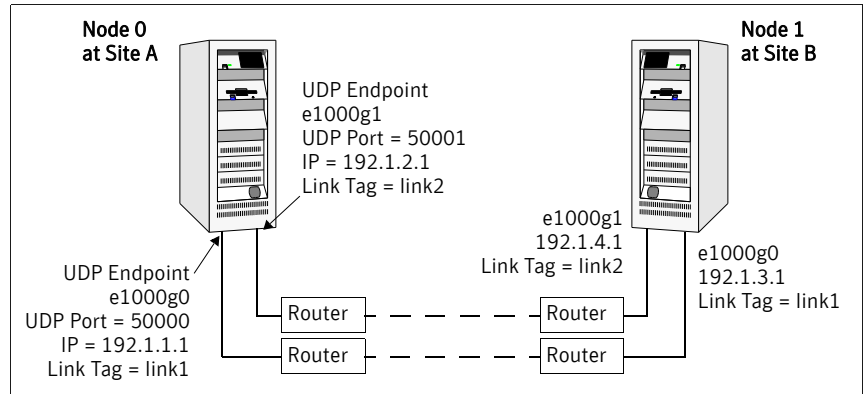
has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the broadcast address does not need to be set in the in the `link` command of the `/etc/llttab` file.

```
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 would resemble:

```
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
```

Minimal downtime upgrade

Use a minimal downtime upgrade to upgrade VCS. This procedure minimizes downtime for the cluster that you want to upgrade. In situations where you can fail over all your service groups to running nodes, downtime equals the time that it takes to offline and online the service groups. In situations where you have a service group that you cannot fail over to a running node during upgrade, downtime for that service group equals the time that it takes to perform an upgrade and reboot the node.

Supported upgrades

Use this procedure to upgrade from VCS 4.0 or 4.1.

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 in order to minimize the downtime for any particular service group.
- Some rough guidelines are:
  - 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 tampering 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.

■　node01 runs sg2.

■　node02 runs sg1.

■　node03 runs sg4.

■　node04 runs sg3.

In your system list, you have each service group failing over to one other node.

■　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 A-2　　　　Four nodes, four service groups, and their failover paths
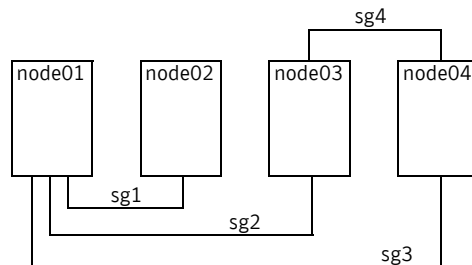
Minimal downtime example overview

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, it is best if sg3 is 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. Reboot node01 and node04 when the upgrade completes. They then rejoin the cluster and you can balance the load on machines by switching service groups.

Performing the minimal downtime example upgrade

This upgrade uses four nodes with four service groups—note that in this scenario the service groups cannot stay online for part of the upgrade. Remember to not add, remove, or change resources or service groups on any nodes during the upgrade as these changes are likely to get lost after the upgrade.

To establish running service groups

1   Establish where your service groups are online.

```
# hagrp -state
#Group          Attribute        System         Value
 sg1            State            node01         |OFFLINE|
 sg1            State            node02         |ONLINE|
 sg2            State            node01         |OFFLINE|
 sg2            State            node03         |ONLINE|
 sg3            State            node01         |OFFLINE|
 sg3            State            node04         |ONLINE|
 sg4            State            node03         |ONLINE|
 sg4            State            node04         |OFFLINE|
```

2   Switch the service groups from all the nodes that you are first upgrading (node02 and node03) to the remaining nodes (node01 and node04).

```
# hagrp -switch sg1 -to node01
# hagrp -switch sg2 -to node01
# hagrp -switch sg4 -to node04
```

3   Verify that your service groups are offline on the nodes targeted for upgrade.

```
# hagrp -state
#Group          Attribute          System         Value
 sg1             State              node01        |ONLINE|
 sg1             State              node02        |OFFLINE|
 sg2             State              node01        |ONLINE|
 sg2             State              node03        |OFFLINE|
 sg3             State              node01        |OFFLINE|
 sg3             State              node04        |ONLINE|
 sg4             State              node03        |OFFLINE|
 sg4             State              node04        |ONLINE|
```

During the next procedure, do not perform any configuration tasks. Do not start any modules.

To perform the minimum downtime upgrade on target nodes

1   On the target nodes, start the 5.0 installer for VCS.

2   Select the VCS installation.

3   Answer **n** when the installer asks:
```
Do you want to upgrade to version 5.0 on these systems using
the current configuration? [y,n,q,?] (y) n
```

4   Answer with the names of the nodes that you want to upgrade:
```
Enter the system names separated by spaces on which to install
VCS: node02 node03
```

5   Select either option **1** or **2** when the installer asks:
```
Select the packages to be installed on all systems? 2
```

6   Answer **n** when the installer completes and asks:
```
Do you want to start Veritas Cluster Server processes now?
[y,n,q] (y) n
```

To edit the configuration and prepare for upgrade node01 and node04

1   When HAD is down on node02 and node03, you see this message:
```
Shutdown completed successfully on all systems.
```

2   After you see the above message, you can make the VCS configuration writable on node01 or node04. Note that you need to make the configuration writable because the installer froze the service groups during the upgrade.
```
# haconf -makerw
```

3   Unfreeze all service groups.
```
# hagrp -unfreeze sg1 -persistent
# hagrp -unfreeze sg2 -persistent
# hagrp -unfreeze sg3 -persistent
# hagrp -unfreeze sg4 -persistent
```

4   Dump the configuration and make it read-only.
```
# haconf -dump -makero
```

To edit the configuration on node02 and node03

1   Open the main.cf file, and delete the `Frozen = 1` line for each service
    group as appropriate.

2   Save and close the file.

3   Reboot node02 and node03.

4   Wait for GAB to come up. In the console's output, look for a line that reads:
    ```
    Starting GAB is done.
    ```

To upgrade and restart your clusters

1   On node01 and node04, take the service groups offline.
    ```
    # hagrp -offline sg1 -sys node01
    # hagrp -offline sg2 -sys node01
    # hagrp -offline sg3 -sys node04
    # hagrp -offline sg4 -sys node04
    ```

2   On node01 and node04, perform the upgrade.
    See "To perform the minimum downtime upgrade on target nodes" on
    page 191.

3   When HAD is down on node01 and node04, you see this message:
    ```
    Shutdown completed successfully on all systems.
    ```

4   Start vxfenconfig on node02 and node03.
    ```
    # vxfenconfig -c
    ```

5   Start your cluster on node02 and node03.
    ```
    # hastart
    ```

6   After the upgrade completes, reboot node01 and node04.

After you have rebooted the nodes, all four nodes now run the latest version of
VCS.

In this example, you achieved minimal downtime because your service groups
were down only from the point when you took them offline on node01 and
node04, to the time VCS brought them online on node02 or node03 as
appropriate.

Configuring the Symantec License Inventory Agent

This appendix includes the following topics:

- About the Symantec License Inventory Manager
- When the Symantec License Inventory Agent is installed
- When the server and access points are installed
- What you can do with the agent after it is installed
- How to remove the agent
- How to order the Symantec License Inventory Manager license and media kit

The Symantec License Inventory Manager installation disc is available separately. For information on how to order the full product, see "How to order the Symantec License Inventory Manager license and media kit" on page 197. The installation media provides online documentation with details on all topics discussed in this appendix.

Read the following Technical Support TechNote for the latest information on updates, patches, and software issues regarding this product:

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

You can also download the *Symantec License Inventory Agent 4.1 Release Notes*, from this website.

About the Symantec License Inventory Manager

The Symantec License Inventory Manager (license inventory manager) is an enterprise asset management tracking tool that inventories Symantec Information Availability products in your network and consolidates critical information on the deployment of these products to facilitate license management and compliance tracking. Using the information provided by the license inventory manager, you can:

■ Determine all the Symantec software products and licenses being used in your enterprise

■ Achieve easier license self-compliance management

■ Know your Enterprise License Agreement deployment status

■ Reduce administrative overhead for managing license compliance

■ Renew support and maintenance based on the licenses you have deployed

■ Gain more control over your Symantec software usage

■ Manage department chargebacks based on actual software usage

■ Use more flexible licensing and pricing models

■ Exploit detailed deployment data to perform return on investment analyses for purchased software

The license inventory manager is a three-tiered system that consists of a server tier, access point tier, and an agent tier. The server tier is the Symantec License Inventory Server, which consolidates and stores information that it gathers from the agents and access points.

The optional access point tier includes Symantec License Inventory Access Points and serves as a consolidation layer between the agents and server.

The agent tier includes Symantec License Inventory Agents, which are deployed on individual hosts in a network. Each agent gathers product information on the supported Symantec products that are installed on the agent's host, then sends the information to an access point or the server.

When the Symantec License Inventory Agent is installed

The Symantec product installer installs or upgrades the agent on the host with the Symantec product. The agent is installed in the following directory:

`/opt/SYMClma`

The agent is installed with a default configuration that minimizes its impact on a running system. The minimum configuration prevents remote communication with the agent to keep its data and interfaces secure.

When the server and access points are installed

The server and access points are not installed automatically. If you want to use the Symantec License Inventory Manager, you must manually install the server and, optionally, the access points. After you install the server and access points, the agents can gather information and you can create inventory reports.

You can install the server and access points from the Symantec License Inventory Manager installation disc.

What you can do with the agent after it is installed

If you are already participating in a Symantec sales program that requires the use of the agent, or if you want to order and deploy the Symantec License Inventory Manager, you can use the agent to track Symantec products on the systems on which it was installed. To use the agent, however, you must manually configure it to enable remote communication between the agent and its server or access point.

Complete instructions for reconfiguring the agent are provided in the *Symantec License Inventory Manager 4.1 Release Notes.* You can download this document from the following website:

http://support.veritas.com/docs/282183

How to remove the agent

If you do not want to use the Symantec License Inventory Manager, you can remove the agent using the operating system package removal commands to remove the agent packages, which include SYMClma and VRTSsmf.

The server and access point also use the VRTSsmf package. If the server or access point is installed on this host with the agent, you can remove the SYMClma package, but not the VRTSsmf package. If neither the server nor the access point is installed on this host, you can remove both the SYMClma and VRTSsmf packages.

If you remove both packages, remove the SYMClma package first.

Table B-1 lists the commands required to remove these packages on the supported platforms.

Table B-1   Package removal commands required to remove the agent

Platform	Package removal command
AIX	`installp -u SYMClma` `installp -u VRTSsmf`
HP-UX	`swremove SYMClma` `swremove VRTSsmf`
Linux	`rpm evv SYMClma` `rpm evv VRTSsmf`
Solaris	`pkgrm SYMClma` `pkgrm VRTSsmf`

Later, you can reinstall the agent with the Symantec License Inventory Manager installation disc. This disc is available in the Symantec License Inventory Manager kit.

How to order the Symantec License Inventory Manager license and media kit

To order a Symantec License Inventory Manager license and media kit, contact your Symantec sales representative.

The installation media provides online documentation for the Symantec License Inventory Manager. You can contact your sales representative to order printed copies of the documentation. The documents you can order include:

- *Symantec License Inventory Manager Installation and Configuration Guide*

- *Symantec License Inventory Manager Administrator's Guide*

- *Symantec License Inventory Manager User's Guide*

Index

W