Veritas™ Cluster Server Installation Guide

Linux for IBM System p

5.0

11472871

symantec™

Veritas Cluster Server Installation Guide

Veritas Cluster Server 5.0

Symantec Corporation
20330 Stevens Creek Blvd.
Cupertino, CA 95014
www.symantec.com

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.

Linux is a registered trademark of Linus Torvalds.

Licensing and registration

Veritas Cluster Server is a licensed product. See the *Veritas Cluster Server Installation Guide* for license installation instructions.

Technical support

For technical assistance, visit http://support.veritas.com and select phone or email support. 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.

Contents

Chapter 1

Introducing VCS

This chapter contains the following topics:

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

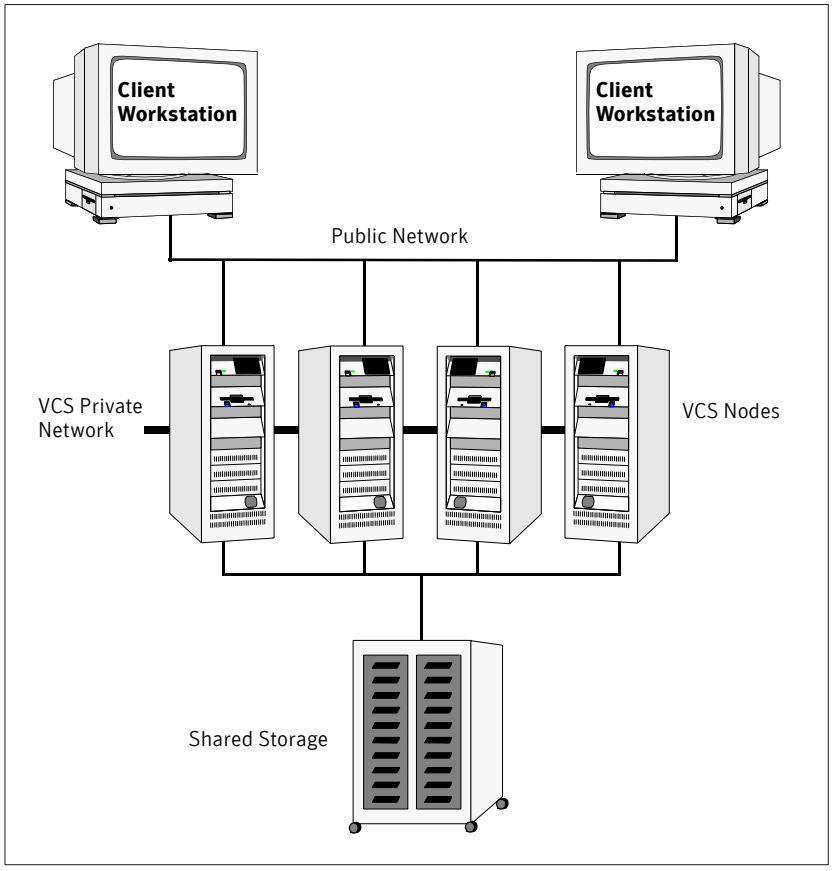

Multiple nodes

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

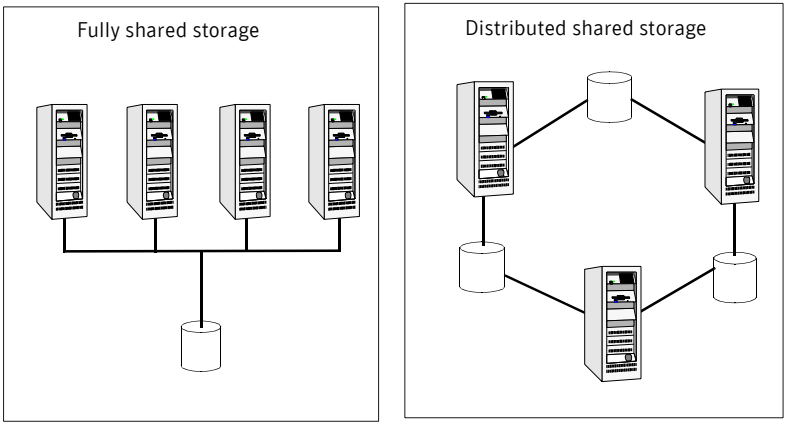

LLT and GAB

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/lltab, 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 117.

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 Linux PPC 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

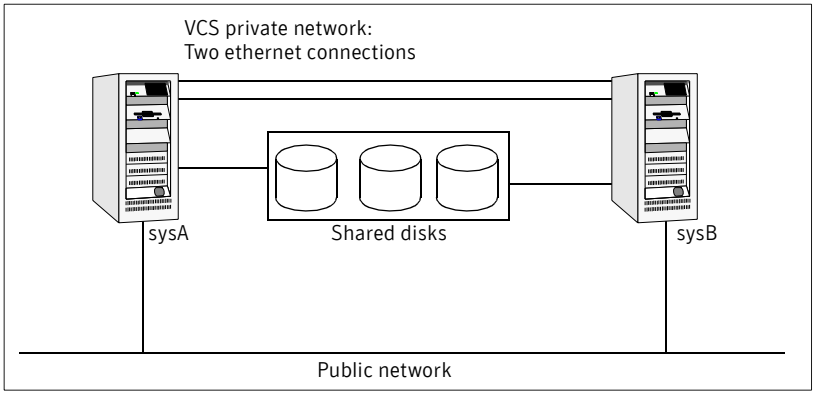

Preexisting network partitions

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.

Chapter 2

Preparing to install and configure VCS

This chapter contains the following topics:

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/285834.

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

http://support.veritas.com/docs/286819.

About VCS 5.0 features

To configure the optional features of the VCS components, make sure to install all RPMs 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
- Veritas Cluster Management Console
- SMTP email notification for VCS events
- SNMP trap notification for VCS events
- 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 56.

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

Figure 2-4 Secure VCS cluster configuration flowchart

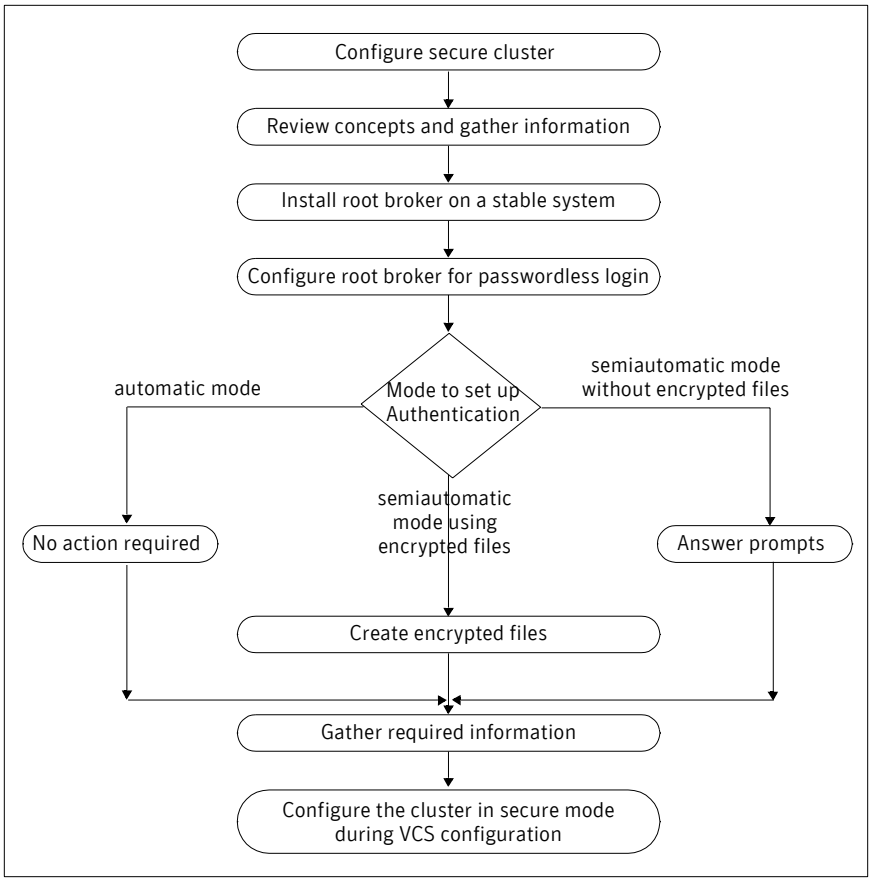

If you decide to enable Authentication Service, the root broker administrator must perform the following preparatory tasks:

- 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 rook broker and authentication brokers are in sync.

Veritas Cluster Management Console

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

You can configure Cluster Management Console to manage single clusters, multiple clusters, or both. To configure Cluster Management Console, do the following:

- To manage multiple clusters, you must have Cluster Management Console management server setup. If the prerequisite is met, then you can manage multiple clusters using direct connection or cluster connector.
  See Veritas Cluster Management Console documentation.
  Depending on the connection mode you would use, do the following:

Direct connection	If you can use direct connection to communicate with the management server, the clusters require no further preparation during VCS installation and configuration. After configuring VCS, you can start the Cluster Management Console from the management server and configure the management server to connect to clusters using direct connection.
Cluster connector	If a firewall separates the management server and cluster nodes, you need to install a component called the cluster connector on each cluster node. The cluster connector enables communication with clusters through firewalls and provides buffering for cluster data. If the console goes offline and then comes back online, it can retrieve data collected during the offline period from the cluster connector buffer. You can configure cluster connector during or after the VCS installation and configuration. Clusters using cluster connector connect to the management server automatically. See "Installing and configuring VCS 5.0" on page 49. See "Establishing cluster communication with the management server" on page 66.

- To manage a single cluster, you must choose the option to install the Cluster Management Console during VCS installation and configuration.
  See "Installing and configuring VCS 5.0" on page 49.

See the *Veritas Cluster Server User's Guide*.

Supported browsers for the Cluster Management Console

Veritas Cluster Management Console is supported on the following browsers:

- Microsoft Internet Explorer 6.0 or later
- Firefox 1.5 or later

Veritas Cluster Management requires the Macromedia Flash Plugin v8.0.

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.

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 85.

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 25
- "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

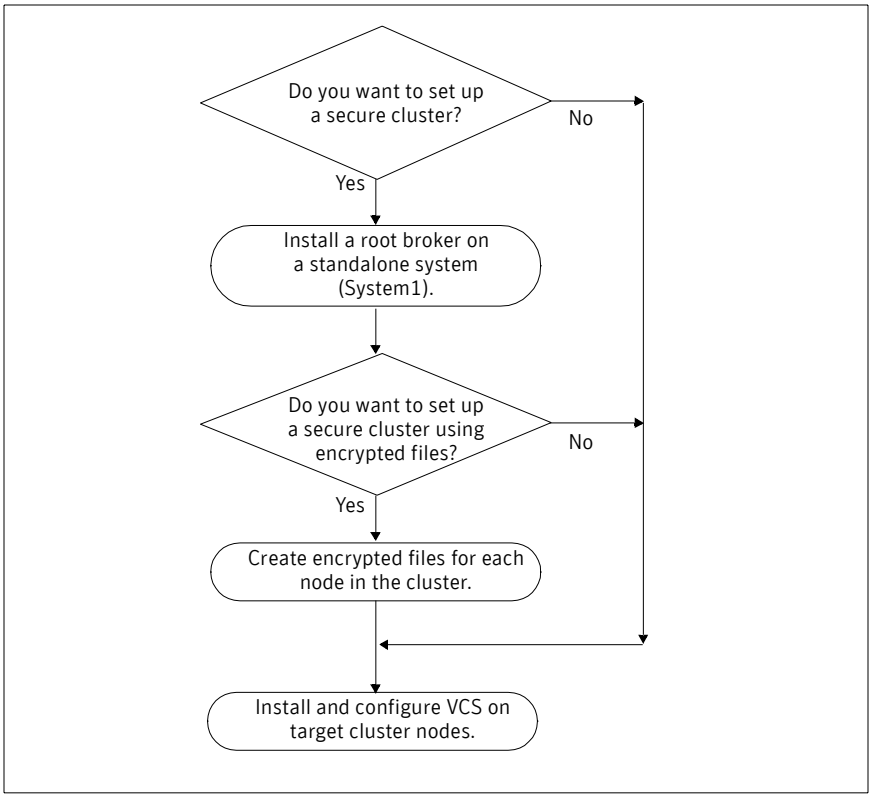

Complete the following preparatory tasks based on the VCS features you want to configure:

- "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	From 1 to 32 Linux PPC systems running the supported Linux PPC operating system version. See "Supported operating systems" on page 26.
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 85.
Disk space	See Table 2-2, "Disk space requirements and totals."
Network Interface Cards (NICs)	In addition to the built-in public NIC, VCS requires at least one more NIC per system. Symantec recommends two additional NICs.
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.

Required disk space

Confirm that your system has enough free disk space to install VCS. The following table shows the approximate disk space usage by directory for the Veritas Cluster Server RPMs.

Table 2-2 Disk space requirements and totals

Packages	/	/opt	/usr	/var	Totals
Required	5 MB	143 MB	25 MB	169 MB	342 MB

Table 2-2 Disk space requirements and totals

Packages	/	/opt	/usr	/var	Totals
Optional	3 MB	39 MB	1 MB	7 MB	50 MB
Required and optional total	8 MB	202 MB	26 MB	176 MB	412 MB

Note: 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.

Supported operating systems

Run VCS 5.0 on these operating systems at the suggested patch levels. Within a cluster, all nodes must use the same operating system version and patch level.

Table 2-3 lists the architectures and operating systems on which VCS 5.0 operates. Symantec supports only those kernel binaries distributed by Red Hat and SUSE.

Table 2-3 Supported operating systems and architectures

Operating System	Kernel	Architecture
Red Hat Enterprise Linux 4 (RHEL 4) Update 3	2.6.9-34.EL 2.6.9-34.smp 2.6.9-34.hugemem	Linux ppc64
SUSE Linux Enterprise Server 9 (SLES 9) with SP3	2.6.5-7.244 2.6.5-7.244-smp 2.6.5-7.244-bigsmp	Linux ppc64

Symantec products will operate on subsequent kernel and patch releases provided the operating systems maintain kernel ABI (application binary interface) compatibility.

Information about the latest supported Red Hat erratas and updates and SUSE service packs is available in the following TechNote. Read this TechNote *before* installing any Symantec product.

http://entsupport.symantec.com/docs/285834.

Supported software

- ext2, ext3, reiserfs, NFS, and bind on LVM2, Veritas Volume Manager (VxVM) 5.0, and raw disks.
- Veritas Volume Manager 5.0 with Veritas File System 5.0

Note: If you plan to set up VCS I/O fencing in a shared storage environment, Symantec recommends using VxVM versions 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 18.

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 RPMs on the system.
   The installvcs program lists the RPMs 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 49.

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 18.

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 49.

Performing pre-installation tasks

Table 2-4 lists the tasks you must perform before proceeding to install VCS.

Table 2-4 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
Configure SuSE network interfaces	"Configuring SuSE network interfaces" on page 33
Set up shared storage for I/O fencing (optional)	"Setting up shared storage" on page 35
Enable communication between systems.	"Enabling communication between systems" on page 38
Review basic instructions to optimize LLT media speeds.	"Optimizing LLT media speed settings on private NICs" on page 38
Review guidelines to help you set the LLT interconnects.	"Guidelines for setting the media speed of the LLT interconnects" on page 38
Set up `ssh` on cluster systems.	"Setting up ssh on cluster systems" on page 39
Obtain license keys.	"Obtaining VCS license keys" on page 40
Mount the product disc	"Mounting the product disc" on page 41

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
```

  If you use the `man` command to access manual pages, set LC_ALL to "C" in your shell for correct page display.

```
# export LC_ALL=C
```

  See incident 82099 on the Red Hat support web site for more information.

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 (NICs).
2. Connect the private NICs 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

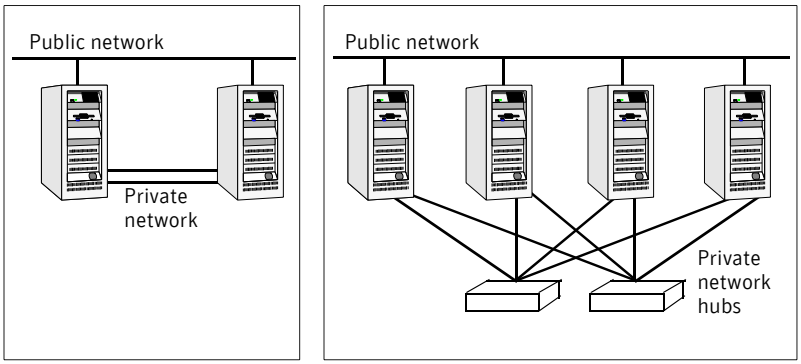

4 Test network connections by temporarily assigning network addresses and use `telnet` or `ping` to verify communications.
LLT uses its own protocol, and does not use TCP/IP. Therefore, to ensure the private network connections are used only for LLT communication and not for TCP/IP traffic, unplumb and 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 47.

Using network switches

You can use network switches instead of hubs.

Configuring SuSE network interfaces

You must perform network configuration on SUSE. In rare cases where RedHat does not automatically configure the network interfaces, RedHat users may also have to perform the network configuration.

For VCS to function properly:

- VCS must be able to find the same network interface names across reboots.
- VCS must have network interfaces up before LLT starts to run.

Symantec suggests the following steps for configuring network interfaces on SUSE.

To configure persistent interface names for network devices

1 Navigate to the hotplug file in the /etc/sysconfig directory:

```
# cd /etc/sysconfig
```

2 Open the hotplug file in an editor.

3 Set HOTPLUG_PCI_QUEUE_NIC_EVENTS to yes:

```
HOTPLUG_PCI_QUEUE_NIC_EVENTS=yes
```

4 Make sure that the interface name to MAC address mapping remains same across the reboots. Symantec recommends adding the PERSISTENT_NAME entries to the configuration files for all the network interfaces (including the network interfaces that are not used).

- Run the command:

  ```
  ifconfig -a
  ```

- For each ethernet interface displayed, do the following:
  - If a file named /etc/sysconfig/network/ifcfg-eth-id-*mac*, where *mac* is the hardware address of that interface, does not exist, then do the following:

    Create the file.

    If a file exists for the same network interface with the name /etc/sysconfig/network/ifcfg-*ethX*, then copy the contents of that file into the newly created file. The variable *ethX* represents the interface name.
  - Add the following line at the end of the file /etc/sysconfig/network/ifcfg-eth-id-*mac*.

    ```
    PERSISTENT_NAME=ethX
    ```

    where *ethX* is the interface name.

Note: The MAC address in the ifcfg-eth-id-*mac* file can be in uppercase or lowercase. SUSE, and therefore the Veritas product installer, ignores the file with lowercase MAC address if the file with uppercase MAC address is present.

For example:

```
# ifconfig -a
eth0    Link encap:Ethernet  HWaddr 00:02:B3:DB:38:FE
        inet addr:10.212.99.30  Bcast:10.212.99.255
        Mask:255.255.254.0
        inet6 addr: fe80::202:b3ff:fedb:38fe/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:453500 errors:0 dropped:0 overruns:0 frame:0
        TX packets:8131 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:35401016 (33.7 Mb)  TX bytes:999899 (976.4 Kb)
        Base address:0xdce0 Memory:fcf20000-fcf40000
```

If a file named etc/sysconfig/network/ifcfg-eth-id-00:02:B3:DB:38:FE does not exist, do the following:

- Create the file.
- If the file /etc/sysconfig/network/ifcfg-eth0 exists, then copy the contents of this file into etc/sysconfig/network/ifcfg-eth-id-00:02:B3:DB:38:FE.

Add the following to the end of the file named etc/sysconfig/network/ifcfg-eth-id-00:02:B3:DB:38:FE,

```
PERSISTENT_NAME=eth0
```

Perform the above steps for all interfaces displayed by `ifconfig -a` command.

To configure interfaces to be up before starting LLT

1. For each network interface that you want LLT to use, find its MAC address by running the `ifconfig` command:

   ```
   # ifconfig eth0
   eth0    Link encap:Ethernet  HWaddr 00:0C:0D:08:C4:32
   ```

   Where **eth0** is the sample network interface name. The output displays 00:0C:0D:08:C4:32 as the interface's MAC address.

2. Navigate to the config file in the /etc/sysconfig/network directory:

   ```
   # cd /etc/sysconfig/network
   ```

3. Open the config file in an editor.

4. Append the string eth-id-*macaddress* to the MANDATORY_DEVICES list in the config file. Separate each address with a space, for example:

   ```
   MANDATORY_DEVICES="eth-id-00:0C:0D:08:C4:31 eth-id-
   00:0C:0D:08:C4:32"
   ```

Note: You must not reboot the system between configuring the persistent interface names and configuring the interfaces to be up before starting LLT.

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 85.

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

Setting up shared storage: SCSI

To set up shared storage

1. Connect the disk to the first cluster system.
2. Power on the disk.
3. Connect a terminator to the other port of the disk.
4. Boot the system. The disk is detected while the system boots.
5. Press CTRL+A to bring up the SCSI BIOS settings for that disk.
   - Set Host adapter SCSI ID = 7, or to an appropriate value for your configuration.
   - Set Host Adapter BIOS in Advanced Configuration Options to Disabled.
6. Format the shared disk and create required partitions on it.
   - Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is /dev/sdc.
     Identify whether the shared disk is sdc, sdb, and so on.
   - Type the following command:
     # **fdisk /dev/*shareddiskname***
     For example, if your shared disk is sdc, type:
     # **fdisk /dev/sdc**
   - Create disk groups and volumes using Volume Manager utilities.
   - To apply a file system on the volumes, type:
     # **mkfs -t *fs-type* /dev/vx/dsk/*disk-group*/*volume***
     For example, enter the following command if the name of the disk group is dg, the name of the volume is vol, and the file system type is vxfs.
     # **mkfs -t vxfs /dev/vx/dsk/dg/vol01**
7. Power off the disk.
8. Remove the terminator from the disk and connect the disk to the other cluster system.
9. Power on the disk.
10. Boot the second system. The system can now detect the disk.
11. Press Ctrl+A to bring up the SCSI BIOS settings for the disk.
    - Set Host adapter SCSI ID = 6, or to an appropriate value for your configuration. Note that the SCSI ID should be different from the one configured on the first cluster system.
    - Set Host Adapter BIOS in Advanced Configuration Options to Disabled.

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

Setting up shared storage: Fibre Channel

To set up shared storage for fibre channel

1 Connect the fibre channel disk to a cluster system.

2 Boot the system and change the settings of the fibre channel. Perform the tasks for all QLogic adapters in the system:

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

3 Verify that the system is detecting the fibre channel disks properly.

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

- Identify your shared disk name. If you have two internal SCSI hard disks, your shared disk is /dev/sdc.
  Identify whether the shared disk is sdc, sdb, and so on.
- Type the following command:
  `# fdisk /dev/shareddiskname`
  For example, if your shared disk is sdc, type:
  `# fdisk /dev/sdc`
- Create disk groups and volumes using Volume Manager utilities.
- To apply a file system on the volumes, type:
  `# mkfs -t fs-type /dev/vx/dsk/disk-group/volume`
  For example, enter the following command if the name of the disk group is dg, the name of the volume is vol, and the file system type is vxfs.
  `# mkfs -t vxfs /dev/vx/dsk/dg/vol01`

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

6 Power off this cluster system.

7 Connect the same disks to the next cluster system.

8 Turn on the power for the second system.

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

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 70.

Warning: The `rsh` and `ssh` commands to the remote systems, where VCS is to be installed, must not print any extraneous characters.

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 do not use Auto_Negotiation, you have to set it to the same speed on all nodes for all NICs used by LLT.

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`.

The Remote Shell (`rsh`) is disabled by default to provide better security. Use `ssh` for remote command execution.

Configuring ssh

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

To configure ssh

1. Log on to the system from which you want to install VCS.
2. Generate a DSA key pair on this system by running the following command:
   ```
   # ssh-keygen -t dsa
   ```
3. Accept the default location of ~/.ssh/id_dsa.
4. When prompted, enter a passphrase and confirm it.
5. Change the permissions of the .ssh directory by typing:
   ```
   # chmod 755 ~/.ssh
   ```
6. The file ~/.ssh/id_dsa.pub contains a line beginning with ssh_dss and ending with the name of the system on which it was created. Copy this line to the /root/.ssh/authorized_keys2 file on all systems where VCS is to be installed.
   If the local system is part of the cluster, make sure to edit the authorized_keys2 file on that system.
7. Run the following commands on the system from which the installation is taking place:
   ```
   # exec /usr/bin/ssh-agent $SHELL
   # ssh-add
   ```
   This step is shell-specific and is valid for the duration the shell is alive.
8. When prompted, enter your DSA passphrase.

You are ready to install VCS on several systems by running the installvcs program on any one of them or on an independent system outside the cluster.

To avoid running the ssh-agent on each shell, run the X-Window system and configure it so that you will not be prompted for the passphrase. Refer to the Red Hat documentation for more information.

9 To verify that you can connect to the systems on which VCS is to be installed, type:

```
# ssh -x -l root north ls
# ssh -x -l root south ifconfig
```

The commands should execute on the remote system without having to enter 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.

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 with the VCS software into a drive connected to the system.
   The disc is automatically mounted.
3. If the disc does not automatically mount, then enter:
   `# mount -o ro /dev/cdrom /mnt/cdrom`
4. Navigate to the location of the RPMs.
   Depending on the OS distribution and architecture, type the command:

RHEL 4	`# cd /mnt/cdrom/rhel4_ppc64/cluster_server`
SLES 9	`# cd /mnt/cdrom/sles9_ppc64/cluster_server`

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 RPMs 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 See "Obtaining VCS license keys" on page 40.

To decide whether to install: ■ the required VCS RPMs ■ all the VCS RPMs	Install only the required RPMs if you do not want to configure any optional components or features. The default option is to install all RPMs. See "Optional VCS RPMs" on page 45.

- 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 eth0. Example: **`eth1, eth2`**

- 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 18. ■ 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: **`eth0`**
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. See "Veritas Cluster Management Console" on page 21. 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 I/O fencing:

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

Optional VCS RPMs

The optional VCS RPMs 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

Chapter 3

Installing and configuring VCS

This chapter contains the following topics:

About installing and configuring VCS

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

Veritas product installer	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 RPMs 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, or Symantec Product Authentication Services 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 17.

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

See “About the uninstallvcs program” on page 81.

Optional features of the installvcs program

Table 3-5 specifies the optional actions that the installvcs program can perform.

Table 3-5 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 51.
Install VCS RPMs without configuring VCS.	See “Installing VCS using installonly option” on page 70.
Configure or reconfigure VCS when VCS RPMs are already installed.	See “Configuring VCS using configure option” on page 70.

Table 3-5 installvcs optional features

Optional action	Reference
Perform secure installations using values stored in a configuration file.	See "Performing VCS installation in a secure environment" on page 70.
Perform automated installations using values stored in a configuration file.	See "Performing automated installations" on page 72.

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 RPMs 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 82.

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 65.

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 RPMs 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

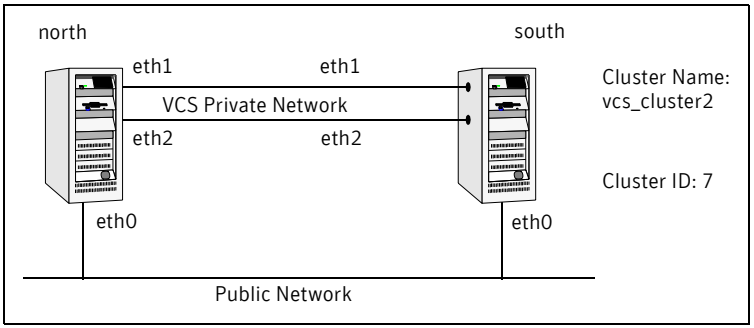

Overview of tasks

Table 3-6 lists the installation and configuration tasks.

Table 3-6 Installation and configuration tasks

Task	Reference
Start the installation process and choose the installation	■ "Checking the systems for installation" on page 51 (optional) ■ "Starting the software installation" on page 51 ■ "Specifying systems for installation" on page 52 ■ "Licensing VCS" on page 53 ■ "Choosing VCS RPMs" on page 53 ■ "Choosing to install VCS RPMs or configure VCS" on page 54
Configure the cluster and optional features	■ "Configuring the cluster" on page 54 ■ "Configuring the cluster in secure mode" on page 56 (optional) ■ "Adding VCS users" on page 57 (optional) ■ "Configuring cluster connector" on page 58 (optional) ■ "Configuring the Cluster Management Console" on page 59 (optional) ■ "Configuring SMTP email notification" on page 60 (optional) ■ "Configuring SNMP trap notification" on page 61 (optional)

Table 3-6 Installation and configuration tasks

Task	Reference
Install the RPMs and create configuration files	■ "Installing the VCS RPMs" on page 62 ■ "Creating VCS configuration files" on page 63
Start VCS and its components	■ "Starting VCS" on page 64 ■ "Completing the installation" on page 64
Perform the post-installation tasks	■ "Copying the installation guide to each node" on page 64 ■ "Setting up I/O fencing" on page 65 ■ "Installing the VCS Java Console" on page 65 ■ "Establishing cluster communication with the management server" on page 66 ■ "Installing cluster connector" on page 67
Verify the cluster	■ "Verifying the cluster after installation" on page 70

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/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, RPMs, 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 79.

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 144.
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

- Checks whether VCS is installed

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 77.

To license VCS

1. Review the output as the utility checks system licensing and installs the licensing RPM.
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 RPMs

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

To install VCS RPMs

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

2 Choose the VCS RPMs to be installed.

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

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

1	Installs only the required VCS RPMs.
2	Installs all the VCS RPMs. 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 RPMs.
3	Installs all the VCS and SF HA RPMs. (default option) If you already installed the SF HA RPMs, the installer does not list this option.

3 View the list of RPMs that the installer would install on each node.
If the current version of a RPM is on a system, the installer removes it from the RPM installation list for the system.

Choosing to install VCS RPMs 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 54.
- Install RPMs on the systems and leave the cluster configuration steps for later.

To install VCS RPMs 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 RPMs.

3 Configure the cluster later.
See "Configuring VCS using configure option" on page 70.

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 eth0.)

```
Enter the NIC for the first private heartbeat NIC on north:
[b,?] eth1
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,?] eth2
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 18.

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 57.

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 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 58.

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 56.

See "Configuring cluster connector" on page 58.

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.

See "Veritas Cluster Management Console" on page 21.

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**.
  See "Configuring the Cluster Management Console" on page 59.

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 the Cluster Management Console

If you want to locally manage this cluster, then you must configure the Cluster Management Console. Note that this cluster can also be a part of the clusters that are centrally managed by the management server.

See "Veritas Cluster Management Console" on page 21.

To configure the Cluster Management Console

1 Review the required information to configure the Cluster Management Console.

2 Choose whether to configure the Cluster Management Console or not. Do one of the following:

- To configure the Cluster Management Console on the systems, press **Enter**.

  ```
  Do you want to configure the Cluster Management Console
  [y,n,q](y)
  ```

- To skip configuring the Cluster Management Console and advance to configuring SMTP, enter **n**.

  See "Configuring SMTP email notification" on page 60.

3 Confirm whether you want to use the discovered public NIC on the first system. Do one of the following:

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

```
Active NIC devices discovered on north: eth0
Enter the NIC for Cluster Management Console to use on north:
[b,?](eth0)
```

4 Confirm whether you want to use the same public NIC on all nodes. Do one of the following:

- If all nodes use the same public NIC, enter **y**.
- If unique NICs are used, enter **n** and enter a NIC for each node.

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

5 Enter the virtual IP address for the Cluster Management Console.

```
Enter the Virtual IP address for Cluster Management Console:
[b,?] 10.10.12.1
```

6 Confirm the default netmask or enter another one:

```
Enter the netmask for IP 10.10.12.1: [b,?] (255.255.240.0)
```

7 Verify and confirm the Cluster Management Console information.

```
Cluster Management Console verification:

    NIC: eth0
    IP: 10.10.12.1
    Netmask: 255.255.240.0

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

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 61.

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 installing VCS RPMs, enter **n**.

See "Installing the VCS RPMs" on page 62.

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)
```

Installing the VCS RPMs

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

See "Creating VCS configuration files" on page 63.

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 RPMs.

Creating VCS configuration files

After installing the RPMs, 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.

Verifying the NIC configuration

The installer verifies on all the nodes if all NICs have PERSISTENT_NAME set correctly.

If the persistent interface names are not configured correctly for the network devices, the installer displays the following warnings:

```
Verifying that all NICs have PERSISTENT_NAME set correctly on
north:

For VCS to run correctly, the names of the NIC cards must be boot
persistent.

CPI WARNING V-9-122-1021
 No PERSISTENT_NAME set for NIC with MAC address
 00:11:43:33:17:28 (present name eth0), though config file
 exists!
CPI WARNING V-9-122-1022
 No config file for NIC with MAC address 00:11:43:33:17:29
 (present name eth1) found!
CPI WARNING V-9-122-1022
 No config file for NIC with MAC address 00:04:23:ac:25:1f
 (present name eth3) found!

PERSISTENT_NAME is not set for all the NICs.
You need to set them manually before the next reboot.
```

Set the PERSISTENT_NAME for all the NICs.

See “Configuring SuSE network interfaces” on page 33.

Warning: If the installer finds the network interface name to be different from the name in the configuration file, then the installer exits.

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 3-7 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 3-7 File description

File	Description
summary file	■ Lists RPMs 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 73.

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 85.

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 or Windows 2000 Professional system, or Linux PPC system. 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.

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*.

Hardware requirements for the Java Console

The following items are the minimum hardware requirements for the Java Console:

- Pentium II 300 megahertz
- 256 megabytes of RAM
- 800x600 display resolution
- 8-bit color depth of the monitor
- Graphics card capable of 2D images

Note: Symantec recommends using the Pentium III, 400MHz, 256MB RAM, and 800x600 display resolution as minimum requirements.

The version of the Java™ 2 Runtime Environment (JRE) requires 32 megabytes of RAM and is supported on Intel Pentium platforms running the Linux kernel v 2.2.12 and glibc v2.1.2-11 (or later). Symantec recommends using 48 megabytes

of RAM, 16-bit color mode, and KDE and KWM window managers used in conjunction with displays set to local hosts.

Installing the Java Console on Linux PPC

To install Java console on Linux PPC

1. Insert the VCS software disc into a drive on the system.
   The software automatically mounts the disc on /mnt/cdrom.
2. If the disc does not get automatically mounted, then enter:
   `# mount -o ro /dev/cdrom /mnt/cdrom`
3. Navigate to the folder containing the RPMs.
   `# cd /mnt/cdrom/dist_arch/cluster_server/rpms`
   Where *dist* is the Linux distribution, rhel4 or sles9 and *arch* is the architecture, ppc64.
4. Install the RPM using `rpm -i` command.
   `# rpm -i VRTScscm-5.0.10.0-MP1_GENERIC.noarch.rpm`

Installing the Java Console on a Windows workstation

You can install the VCS Java Console (Cluster Manager) on a Windows NT workstation or a Windows 2000 Professional Workstation 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.

See “Veritas Cluster Management Console” on page 21.

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 67.
- 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

Cluster connector is a process agent, which you must use if a firewall exists between the management server and any clusters. You can install cluster connector on Linux PPC clusters. You can also a batch feature to install cluster connector on Linux PPC clusters.

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.

For Linux PPC, one of the following two conditions must be true:

- 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 Linux PPC 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.
   # **cd /mnt/cdrom/*dist_arch*/cluster_management_console**
   Where *dist* is rhel4 or sles9 and *arch* is ppc64.
4. Run the installcmc program.
   # **./installcmc -rsh**
5. Proceed to install the cluster connector.
   The installer program 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 117.

Installing VCS using installonly option

In certain situations, users may choose to install the VCS RPMs 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 RPMs 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 54.

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 73.

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 51.
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 49.
   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.

9 On the next system:

- Mount the product disc.
  See "Mounting the product disc" on page 41.
- 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 51.

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 /mnt/cdrom/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 3-8, "Response file variables."

```
#
# installvcs configuration values:
#
$CPI::CFG{AT_ROOTDOMAIN}="root\@east.symantecexample.com";
$CPI::CFG{CMC_CC_CONFIGURED}=1;
$CPI::CFG{CMC_CLUSTERID}{east}=1146235600;
$CPI::CFG{CMC_MSADDR}{east}="mgmtserver1";
$CPI::CFG{CMC_MSADDR}{west}="mgmtserver1";
$CPI::CFG{CMC_MS_ROOT_HASH}="758a33dbd6fae751630058ace3dedb54e5
62fe98";
$CPI::CFG{CMC_SERVICE_PASSWORD}="U2FsdGVkX18vE5tn0hTSWwodThACc+
rX";
$CPI::CFG{ENCRYPTED}="U2FsdGVkX1+k2DHKVcnW7b6vrVghdh+zW4G0WFj5I
JA=";
$CPI::CFG{KEYS}{east}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{KEYS}{west}=[ qw(XXXX-XXXX-XXXX-XXXX-XXXX-XXX) ];
$CPI::CFG{OBC_IGNOREWARNINGS}=0;
$CPI::CFG{OBC_MODE}="STANDALONE";
$CPI::CFG{OPT}{INSTALL}=1;
$CPI::CFG{OPT}{NOEXTRAPKGS}=1;
$CPI::CFG{OPT}{RSH}=1;
$CPI::CFG{SYSTEMS}=[ qw(east west) ];
$CPI::CFG{UPI}="VCS";
$CPI::CFG{VCS_ALLOWCOMMS}="Y";
$CPI::CFG{VCS_CLUSTERID}=13221;
$CPI::CFG{VCS_CLUSTERNAME}="vcs_cluster3";
$CPI::CFG{VCS_CSGNETMASK}="255.255.240.0";
$CPI::CFG{VCS_CSGNIC}{ALL}="eth0";
$CPI::CFG{VCS_CSGVIP}="10.10.12.1";
$CPI::CFG{VCS_LLTLINK1}{east}="eth1";
$CPI::CFG{VCS_LLTLINK1}{west}="eth1";
$CPI::CFG{VCS_LLTLINK2}{east}="eth2";
$CPI::CFG{VCS_LLTLINK2}{west}="eth2";
$CPI::CFG{VCS_SMTPRECP}=[ qw(earnie@symantecexample.com) ];
$CPI::CFG{VCS_SMTPRSEV}=[ qw(SevereError) ];
$CPI::CFG{VCS_SMTPSERVER}="smtp.symantecexample.com";
$CPI::CFG{VCS_SNMPCONS}=[ qw(neptune) ];
$CPI::CFG{VCS_SNMPCSEV}=[ qw(SevereError) ];
$CPI::CFG{VCS_SNMPPORT}=162;
```

Response file variable definitions

Table 3-8 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), and the SNMP trap notification (SNMPPORT, SNMPCONS, and SNMPCSEV).

Table 3-8 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 RPMs must be installed. Configuration can be performed at a later time using the `-configure` option.
$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.

Table 3-8 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{AT_ROOTDOMAIN}	List	Opt'l	Defines the name of the system where the root broker is installed.
$CPI::CFG{OPT}{PKGPATH}	Scalar	Opt'l	Defines a location, typically an NFS mount, from which all remote systems can install product depots. The location must be accessible from all target systems.
$CPI::CFG{OPT}{TMPPATH}	Scalar	Opt'l	Defines the location where a working directory is created to store temporary files and depots 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}{*RPM*}	List	Opt'l	Instructs the installation to not install the optional RPMs designated in the list.
$CPI::CFG{DONOTREMOVE}{*RPM*}	List	Opt'l	Instructs the uninstallation to not remove the optional RPMs 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.
$CPI::CFG{OPT_LOGPATH}	Scalar	Opt'l	Mentions the location where the log files are to be copied. The default location is /opt/VRTS/install/logs.

Table 3-8 Response file variables

Variable	List/ Scalar	Opt'l/ Req'd	Description
$CPI::CFG{CONFIGURE}	Scalar	Opt'l	Performs configuration if the RPMs are already installed using the `-installonly` option.
$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 Management Console.
$CPI::CFG{VCS_CSGNETMASK}	Scalar	Opt'l	Defines the Netmask of the virtual IP address to be used by the Cluster Management 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 3-8 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_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
$CPI::CFG{OPT}{UNINSTALL}	Scalar	Opt'l	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                     = 1249
License Type                      = PERMANENT
OEM ID                            = 478

Features :=
  Platform                        = Linux
  Version                         = 5.0
  Tier                            = 0
  Reserved                        = 0

  Mode                            = VCS
```

Updating product licenses using vxlicinst

You can use the `vxlicinst` command to add the VCS license key on each node. If you have VCS already installed and configured and you are using a demo license, you can replace the demo license.

See "Replacing a VCS demo license with a permanent license" on page 78.

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 3-9 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 3-9 installvcs options

Option and Syntax	Description
`-configure`	Configure VCS after using `-installonly` option to install VCS. See "Configuring VCS using configure option" on page 70.
`-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 RPMs on systems without configuring VCS. See "Installing VCS using installonly option" on page 70.
`-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.

Table 3-9 installvcs options

Option and Syntax	Description
`-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 RPMs 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 RPMs.
`-nolic`	Install product RPMs on systems without licensing or configuration. License-based features or variants are not installed when using this option.
`-nooptionalpkgs`	Specifies that the optional product RPMs such as man pages and documentation need not be installed.
`-nostart`	Bypass starting VCS after completing installation and configuration.
`-pkgpath` *`pkg_path`*	Specifies that *pkg_path* contains all RPMs 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 51.
`-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.

Table 3-9 installvcs options

Option and Syntax	Description
`-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 70. See "Performing automated installations" on page 72.
`-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 18.
`-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 RPMs 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 RPM 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.
- Before removing VCS from fewer than all nodes in a 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 131.

Uninstalling VCS 5.0

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

See "Sample VCS installation and configuration output" on page 295.

Removing VCS 5.0 RPMs

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 RPMs to be uninstalled.
5. If RPMs, such as enterprise agents, are found to be dependent on a VCS RPM, the uninstaller prompt you on whether you want them removed. Enter **y** to remove the designated RPMs.
6. Review the uninstaller report after the verification.
7. Press Enter to uninstall the VCS RPMs.
   ```
   Are you sure you want to uninstall VCS rpms? [y,n,q] (y)
   ```
8. Review the output as the uninstaller stops processes, unloads kernel modules, and removes the RPMs.
9. Note the location of summary and log files that the uninstaller creates after removing all the RPMs.

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.

Uninstalling the Cluster Management Console cluster connector

Perform the following procedure to remove the cluster connector from Linux PPC systems.

Uninstalling cluster connector from Linux PPC systems

Use this procedure to remove the Cluster Management Console cluster connector from each cluster.

On Linux PPC systems, the default installer option is -ssh. If you are performing a remote uninstallation and ssh is not enabled, run the installer program with the -rsh option. Otherwise, the installer generates an error during the uninstallation.

To uninstall cluster connector from Linux PPC systems

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.
   ```
   # cd /mnt/cdrom/dist_arch/cluster_management_console
   ```
   Where *dist* is rhel4 or sles9 and *arch* is ppc64.
4. Run the installcmc program.
   ```
   # ./uninstallcmc -rsh
   ```
5. If prompted, select to uninstall the cluster connector. Otherwise, proceed to step 7.
   The installer program presents a message stating that it will uninstall cluster connector.
6. The uninstall program prompts you for the name of at least one node in the cluster.
   ```
   Enter one system name from each cluster separated by spaces from
   which to uninstall CMC: sysA
   ```
   Based on this, it determines the nodes from which to uninstall and perform the necessary checks.

   Note: If you get an error message similar to this:
   ```
   Checking ssh communication with sysA Enter passphrase for key
   '/.ssh/id_dsa'
   ```
   You must return and set up ssh.
7. Enter **y** to verify that the information up to this point is correct.
   ```
   Is this information correct? [y,n,q] (y)
   ```
   The installer program performs an initial system check of the cluster nodes and checks for installed packages on the cluster nodes. If these checks are satisfactory, the installer program lists the packages to be uninstalled.
8. Enter **y** to verify that you want to uninstall cluster connector.
   ```
   Are you sure you want to uninstall CMC? [y,n,q] (y)
   ```
   The installer program lists package dependencies and uninstallation progress percentages. If the uninstallation is successful, the installer program displays this message followed by the location of the uninstallation logs:
   ```
   Uninstall completed successfully
   ```

Chapter 4

Setting up I/O fencing

This chapter contains the following topics:

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 95

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 RPM 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 35.

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:

```
# fdisk -l
```

- Test the shared disks using the vxfentsthdw script.
  See "Testing the shared disks for SCSI-3" on page 89.

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 92.

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 38.
  After completing the testing process, remove permissions for communication and restore public network connections.
  See "Removing permissions for communication" on page 97.
- 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 89.

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/sdr path on node A and the /dev/sdt path on node B.

From node A, enter:

```
# vxfenadm -i /dev/sdr
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
Vendor id       : EMC
Product id      : SYMMETRIX
Revision        : 5567
Serial Number   : 42031000a
```

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

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

```
# vxfenadm -i /dev/sdc
SCSI ID=>Host: 2 Channel: 0 Id: 0 Lun: E
Vendor id       : HITACHI
Product id      : OPEN-3
Revision        : 0117
Serial Number   : 0401EB6F0002
```

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

Testing the disks using vxfentsthdw script

This procedure uses the /dev/sdr 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/sdr is ready to be configured for I/O Fencing on node
north
```

See "Adding or removing coordinator disks" on page 113.

To test disks using vxfentsthdw script

1. Make sure system-to-system communication is functioning properly.
   See "Enabling communication between systems" on page 38.
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/sdx
/dev/sdr
Enter the disk name to be checked for SCSI-3 PGR on node
south in the format: /dev/sdx
Make sure it's the same disk as seen by nodes north and south
/dev/sdr
```

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/sdr 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:

```
# fdisk -l
```

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 sdr 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 89.
- 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 93.
- 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 /dev/sdz, /dev/sdaa, and sdab.

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 sdz
```

2 Add the other two disks to the disk group:

```
# vxdg -g vxfencoorddg adddisk sdaa
# vxdg -g vxfencoorddg adddisk sdab
```

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 95.

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

- Raw disk

```
/dev/sdz
/dev/sdy
/dev/sdu
```

- DMP disk

```
/dev/vx/rdmp/sdz
/dev/vx/rdmp/sdy
/dev/vx/rdmp/sdu
```

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 113.

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 4-10 describes the methods the utility provides to test storage devices.

Table 4-10 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 4-10 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/sdg, /dev/sdh, and /dev/sdi.
- From the node north, the disks are /dev/sdx, /dev/sdy, and /dev/sdz.

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/sdg 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/sdx 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 113.

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 91.
See "Setting up coordinator disk groups" on page 92.

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 89.

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/sdz south /dev/sdy
north /dev/sdu south /dev/sdw
```

Where the first disk is listed in the first line and is seen by north as /dev/sdz and by south as /dev/sdy. The other disk, in the second line, is seen as /dev/sdu from north and /dev/sdw 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 /etc/sysconfig/vxfen 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 4-11 describes tunable parameters for the VXFEN driver.

Table 4-11 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.
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

Table 4-11 VXFEN tunable parameters

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
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

See "Configuring the VXFEN parameters" on page 103.

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. Unload the module.
   ```
   # /etc/init.d/vxfen stop
   ```
3. Edit the /etc/sysconfig/vxfen file.
   For example, change the entry from:
   ```
   vxfen_min_delay=0
   ```
   to:
   ```
   vxfen_min_delay=30
   ```
4. Start the VXFEN module.
   ```
   # /etc/init.d/vxfen start
   ```
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 4-12 describes how I/O fencing works to prevent data corruption in different failure event scenarios. For each event, corrective operator actions are indicated.

Table 4-12 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.
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.

Table 4-12 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
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 4-12 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 111.

Table 4-12 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 111.

Table 4-12 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.

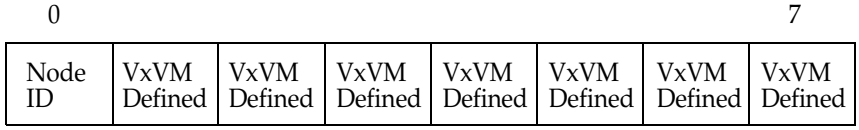

The keys currently assigned to disks can be displayed by using the `vxfenadm` command.

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

```
# vxfenadm -g /dev/sdy
Reading SCSI Registration Keys...
Device Name: /dev/sdy
SCSI ID => Host:  3 Channel:  0 Id:  9 Lun:  0
PRGeneration: -84
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/sdu
```

2 Read the existing keys:

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

The output from this command displays the key:

```
Device Name: /dev/sdu
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/sdu
```

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/sdu
```

7 Remove the temporary key assigned in step 5.

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

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 113.
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 93.
7. Test the recreated disk group for SCSI-3 persistent reservations compliance.
   See "Testing the coordinator disk group using vxfentsthdw -c" on page 99.
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
```

Chapter 5

Verifying the VCS installation

This chapter contains the following topics:

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/lltab 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   north
1   south
```

/etc/llttab

The file `lltab(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 north
set-cluster 2
link eth1 eth1 - ether - -
link eth2 eth2 - ether - -
```

If you use MAC address for the network interface, the file /etc/llttab contains entries that resemble:

```
set-node north
set-cluster 2
link eth1 eth-00:04:23:AC:12:C4 - ether - -
link eth2 eth-00:04:23:AC:12:C5 - ether - -
```

The first line identifies the system. The second line identifies the cluster (that is, the cluster ID you entered during installation). The next two lines, beginning with the `link` command, identify the two network cards used by the LLT protocol.

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

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

Verifying the main.cf file

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

See "Example main.cf for VCS clusters" on page 120.

See "Example main.cf for a centrally managed cluster using Cluster Management Console" on page 122.

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.

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 Linux PPC systems.

Example main.cf for VCS clusters

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

```
include "types.cf"

cluster vcs_cluster2 (
    UserNames = { admin = cDRpdxPmHpzS, smith = dKLhKJkHLh }
    ClusterAddress = "10.10.12.1"
    Administrators = { admin, smith }
    CounterInterval = 5
    SecureClus = 1
)

    system north (
        )

    system south (
        )

    group ClusterService (
        SystemList = { north = 0, south = 1 )
        UserStrGlobal = "LocalCluster@https://10.182.2.76:8443;"
        AutoStartList = { north, south )
        OnlineRetryLimit = 3
        OnlineRetryInterval = 120
        )

    IP webip (
        Device = eth0
        Address = "10.10.12.1"
        NetMask = "255.255.240.0"
        )

    NIC csgnic (
        Device = eth0
        )

    NotifierMngr ntfr (
        SnmpConsoles = { "saturn" = Error, "jupiter" = SevereError }
        SmtpServer = "smtp.example.com"
        SmtpRecipients =  { "ozzie@example.com" = Warning,
                            "harriet@example.com" = Error }
        )

    VRTSWebApp VCSweb (
```

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

    VCSweb requires webip
    ntfr requires csgnic
    webip requires csgnic

// resource dependency tree
    //
    // group ClusterService
    // {
    //      VRTSWebApp VCSweb
    //      {
    //          IP webip
    //          {
    //              NIC csgnic
    //          }
    //      }
    //      NotifierMngr ntfr
    //      {
    //          NIC csgnic
    //      }
    // }
group VxSS (
    SystemList = { north = 0, south = 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 a centrally managed cluster using Cluster Management Console

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

cluster vcs_cluster2 (
    UserNames = { "admin" = hqrJqlQnrMrrPzrLqo }
    Administrators = { "admin" }
    ClusterAddress = "10.10.12.1"
    CounterInterval = 5
    )

system north (
    )

system south (
    )

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

    IP webip (
        Device = eth0
        Address = "10.10.12.1"
        NetMask = "255.255.240.0"
        )

    NIC csgnic (
        Device = eth0
        )

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

    VCSweb requires webip
    webip requires csgnic

group CMC (
    SystemList = { north, south }
    AutoStartList = { north, south }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
```

```
)

ClusterConnectorConfig CMC_ClusterConfig (
    MSAddress = "mgmtserver1.symantecexample.com"
    MSPort = 14145
    ClusterId = "1145613636"
    ClusterType = "vcs"
    ClusterPort = 14141
    VCSLoggingLevel = "TAG_A"
    Logging = "/opt/VRTScmccc/conf/cc_logging.properties"
    ClusterConnectorVersion = "5.0.1000.0"
)

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

CMC_ClusterConnector requires CMC_ClusterConfig
```

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:/sbin:/usr/sbin:/opt/VRTS/bin
```

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
                        eth1 UP         08:00:20:93:0E:34
                        eth2 UP         08:00:20:93:0E:34
1 south     OPEN
                        eth1 UP         08:00:20:8F:D1:F2
                        eth2 DOWN
2               CONNWAIT
                        eth1    DOWN
                        eth2    DOWN
3           CONNWAIT
                        eth1    DOWN
                        eth2    DOWN
.
.
.

31          CONNWAIT
                        eth1     DOWN
                        eth2     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    CPUUsage           0
north    CPUUsageMonitoring Enabled 0 ActionThreshold 0
                            ActionTimeLimit 0 Action NONE
                            NotifyThreshold 0 NotifyTimeLimit 0
north    Capacity           100
north    ConfigBlockCount   142
north    ConfigCheckSum     4085
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       eth1 UP eth2 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            Linux:north,#1 SMP Mon Dec 12 18:32:25
                            UTC 2005,2.6.5-7.244-pseries64,ppc64

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 Veritas Cluster Management Console

The VCS web-based Cluster Management Console enables you to monitor the cluster from any workstation on the public network. Supported browsers are Netscape Navigator 4.0 or later, or Internet Explorer 4.0 or later.

When VCS starts running in the cluster and the ClusterService Group comes up, the Web Console server starts.

To access the Web Console

1 From the browser, navigate to the Web Console by entering:

http://*hostname*:8443/cmc

or

http://*hostname*:8181/cmc

Where hostname is the system name or IP address.

For example:

http://10.10.12.1:8443/cmc

The IP address is the “Cluster virtual IP address” configured into the ClusterService Group.

2 On the Login screen, enter a valid user name and password. By default, the administrator of a new installation can log in as “admin” and use “password” as a password. For security, change your password at your earliest convenience.

3 Click Login to enter the Cluster Summary view.

Accessing the VCS documentation

If you had chosen to install the optional RPM 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*

Chapter 6

Adding and removing cluster nodes

This chapter contains the following topics:

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 17.

Table 6-13 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 6-13 Tasks involved in adding a node to a cluster

Task	Reference
Set up the hardware.	"Setting up the hardware" on page 132

Table 6-13 Tasks involved in adding a node to a cluster

Task	Reference
Install the software manually.	"Preparing for a manual installation" "Installing VCS RPMs for a manual installation"
Add a license key.	"Adding a license key" on page 135
Configure LLT and GAB.	"Configuring LLT and GAB" on page 135
Add the node to the existing cluster.	"Adding the node to the existing cluster" on page 137
Start VCS and verify the cluster.	"Starting VCS and verifying the cluster" on page 137

Setting up the hardware

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

Figure 6-8 Adding a node to a three-node cluster using two independent hubs

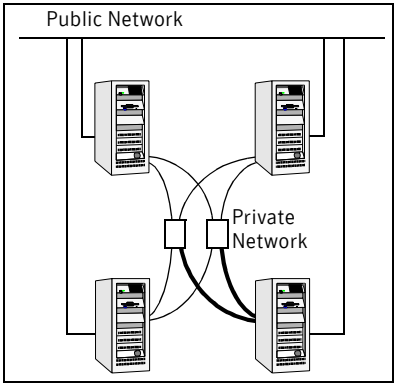

To set up the hardware

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 6-8 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.

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 41.

To prepare for installation

- Depending on the OS distribution, replace the dist in the command with rhel4 or sles9. Replace the arch in the command with ppc64. Type the command:

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

Installing VCS RPMs for a manual installation

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

When selecting the optional RPMs, note:

- Symantec recommends that you install the RPMs for VCS manual pages (VRTSvcsmn) and VCS documentation (VRTSvcsdc). Install the documentation RPM on nodes where you want access to the documentation.
- The I/O fencing RPM (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 85.
- The VCS configuration wizard (VRTScscw) RPM 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 RPMs.

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

To install VCS RPMs on a node

1 Install the required VCS RPMs in the order shown. Do not install any RPMs already installed on the system. Pay special attention to operating system distribution and architecture.

- RHEL4/ppc64, required RPMs:

```
# rpm -i VRTSatClient-4.3.28.0-0.ppc.rpm
# rpm -i VRTSatServer-4.3.28.0-0.ppc.rpm
# rpm -i VRTSicsco-1.3.20.0-0.ppc.rpm
# rpm -i VRTSpbx-1.3.19.0-0.ppc.rpm
# rpm -i VRTSperl-5.0.2.0-linux.ppc64.rpm
# rpm -i VRTSvlic-3.02.18.0-0.ppc64.rpm
# rpm -i VRTSllt-5.0.10.0-MP1_RHEL4.ppc64.rpm
# rpm -i VRTSgab-5.0.10.0-MP1_RHEL4.ppc64.rpm
# rpm -i VRTSvxfen-5.0.10.0-MP1_RHEL4.ppc64.rpm
# rpm -i VRTSvcs-5.0.10.0-MP1_RHEL4.ppc64.rpm
# rpm -i VRTSvcsmg-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTSacclib-5.0.00.0-0.ppc64.rpm
# rpm -i VRTSvcsag-5.0.10.0-MP1_RHEL4.ppc64.rpm
# rpm -i VRTSvcsdr-5.0.10.0-MP1_RHEL4.ppc64.rpm
# rpm -i VRTSjre-1.4-GA1.ppc64.rpm
# rpm -i VRTSjre15-1.5-GA1.ppc64.rpm
# rpm -i VRTScscw-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTSweb-5.0.1-GA1_GENERIC.ppc64.rpm
# rpm -i VRTScutil-5.0-MP1_GENERIC.noarch.rpm
```

- SLES9/ppc64, required RPMs:

```
# rpm -i VRTSacclib-5.0.00.0-0.ppc64.rpm
# rpm -i VRTSatClient-4.3.28.0-0.ppc.rpm
# rpm -i VRTSatServer-4.3.28.0-0.ppc.rpm
# rpm -i VRTScscw-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTScutil-5.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTSgab-5.0.10.0-MP1_SLES9.ppc64.rpm
# rpm -i VRTSicsco-1.3.20.0-0.ppc.rpm
# rpm -i VRTSjre-1.4-GA1.ppc64.rpm
# rpm -i VRTSjre15-1.5-GA1.ppc64.rpm
# rpm -i VRTSllt-5.0.10.0-MP1_SLES9.ppc64.rpm
# rpm -i VRTSpbx-1.3.19.0-0.ppc.rpm
# rpm -i VRTSperl-5.0.2.0-linux.ppc64.rpm
# rpm -i VRTSvcs-5.0.10.0-MP1_SLES9.ppc64.rpm
# rpm -i VRTSvcsag-5.0.10.0-MP1_SLES9.ppc64.rpm
# rpm -i VRTSvcsdr-5.0.10.0-MP1_SLES9.ppc64.rpm
# rpm -i VRTSvcsmg-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTSvlic-3.02.18.0-0.ppc64.rpm
# rpm -i VRTSvxfen-5.0.10.0-MP1_SLES9.ppc64.rpm
# rpm -i VRTSweb-5.0.1-GA1_GENERIC.ppc64.rpm
```

2 Install the optional RPMs, in the order shown. Omit those that you do not want to install.

- RHEL4/ppc64, optional RPMs:

```
# rpm -i VRTScmccc-5.0.00.00-GA_RHEL4.ppc64.rpm
```

```
# rpm -i VRTScmcs-5.0.00.00-GA_RHEL4.ppc64.rpm
# rpm -i VRTScscm-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTScssim-5.0.10.0-MP1_RHEL4.ppc64.rpm
# rpm -i VRTSvcsdc-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTSvcsmn-5.0.10.0-MP1_GENERIC.noarch.rpm
```

- SLES9/ppc64, optional RPMs:

```
# rpm -i VRTScmccc-5.0.00.00-GA_SLES9.ppc64.rpm
# rpm -i VRTScmcs-5.0.00.00-GA_SLES9.ppc64.rpm
# rpm -i VRTScscm-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTScssim-5.0.10.0-MP1_SLES9.ppc64.rpm
# rpm -i VRTSvcsdc-5.0.10.0-MP1_GENERIC.noarch.rpm
# rpm -i VRTSvcsmn-5.0.10.0-MP1_GENERIC.noarch.rpm
```

Adding a license key

After you have installed all RPMs 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.

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 118.
The following example describes a system where node east is the new node on cluster number 2:

```
set-node east
set-cluster 2
link eth1 eth1 - ether - -
link eth2 eth2 - 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 118.
  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 125.

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 6-14 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 6-14 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 138
■ 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 139
Modify the llthosts and gabtab files to reflect the change.	"Modifying configuration files on each remaining node" on page 141
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 RPMs.	"Unloading LLT and GAB and removing VCS on the leaving node" on page 141

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 stop LLT and GAB and remove VCS

1. Stop GAB and LLT:

   ```
   # /etc/init.d/gab stop
   # /etc/init.d/llt stop
   ```

2. To determine the RPMs to remove, enter:

   ```
   # rpm -qa | grep VRTS
   ```

3. To permanently remove the VCS RPMs from the system, use the `rpm -e` command. Start by removing the following RPMs, which may have been optionally installed, in the order shown:

   ```
   # rpm -e VRTScmccc
   # rpm -e VRTScmcs
   # rpm -e VRTScssim
   # rpm -e VRTScscm
   ```

```
# rpm -e VRTSvcsdc
# rpm -e VRTSvcsmn
# rpm -e VRTScutil
# rpm -e VRTSweb
# rpm -e VRTScscw
# rpm -e VRTSjre15
# rpm -e VRTSjre
# rpm -e VRTSvcsdr
# rpm -e VRTSvcsag
# rpm -e VRTSacclib
# rpm -e VRTSvcsmg
# rpm -e VRTSvcs
# rpm -e VRTSvxfen
# rpm -e VRTSgab
# rpm -e VRTSllt
# rpm -e VRTSvlic
# rpm -e VRTSperl
# rpm -e VRTSpbx
# rpm -e VRTSicsco
# rpm -e VRTSatServer
# rpm -e VRTSatClient
```

4 Remove the LLT and GAB configuration files.

```
# rm /etc/lltab
# rm /etc/gabtab
# rm /etc/llthosts
```

Chapter 7

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 144.

See "Creating a single-node cluster manually" on page 145.

Creating a single-node cluster using the installer program

Table 7-15 specifies the tasks involved in installing VCS on a single node using the installer program.

Table 7-15 Tasks to create a single-node cluster using the installer

Task	Reference
Prepare for installation.	"Preparing for a single node installation" on page 144
Install the VCS software on the system using the installer.	"Starting the installer for the single node cluster" on page 144

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 13.

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 51.

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 53.

Creating a single-node cluster manually

Table 7-16 specifies the tasks involved in installing VCS on a single node.

Table 7-16 Tasks to create a single-node cluster manually

Task	Reference
Set the PATH variable	"Setting the PATH variable" on page 145
Install the VCS software manually and adding a license key	"Installing the VCS software manually" on page 146
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 146
Modify the VCS startup file for single-node operation.	"Modifying the startup files" on page 146
Create and modify the VCS configuration files.	"Configuring VCS" on page 146
Start VCS and verify single-node operation.	"Verifying single-node operation" on page 147

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 RPMs manually and install the license key.

See "Preparing for a manual installation" on page 133.

See "Installing VCS RPMs for a manual installation" on page 133.

See "Adding a license key" on page 135.

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/init.d/llt /etc/init.d/llt.old
# mv /etc/init.d/gab /etc/init.d/gab.old
```

Modifying the startup files

Modify the VCS startup file /etc/sysconfig/vcs to include the `-onenode` option as follows:

Change the line:

```
ONENODE=no
```

To:

```
ONENODE=yes
```

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 7-17 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 7-17 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 149
■ 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 150
Connect both nodes to shared storage	"Configuring the shared storage" on page 150
■ Bring up VCS on Node A ■ Edit the configuration file ■ Edit the startup scripts	"Bringing up the existing node" on page 150
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"
Edit the configuration files on Node B	"Configuring LLT and GAB" on page 151
Start LLT and GAB on Node B	"Starting LLT and GAB" on page 154

Table 7-17 Tasks to add a node to a single-node cluster

Task	Reference
■ 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 154
■ Start VCS on Node B ■ Verify the two-node cluster	"Verifying configuration on both nodes" on page 155

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 RPMs and configuration files.
  See "Removing a node from a cluster" on page 138.
- If the node you plan to add as Node B is also currently a single VCS node, uninstall VCS.
  See "Uninstalling VCS 5.0" on page 82
- If you renamed the LLT and GAB startup files, remove them.
  See "Renaming the LLT and GAB startup files" on page 146.

2 If necessary, install VxVM and VxFS.
See "Installing VxVM, VxFS if necessary" on page 149.

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 35.

See "Setting up shared storage: Fibre Channel" on page 37.

Bringing up the existing node

1. Restart Node A.
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 Edit the VCS system configuration file /etc/sysconfig/vcs, and remove the "-onenode" option.

Change the line:

```
ONENODE=yes
```

To:

```
ONENODE=no
```

9 Rename the GAB and LLT startup files so they can be used.

```
# mv /etc/init.d/gab.old /etc/init.d/gab
# mv /etc/init.d/llt.old /etc/init.d/llt
```

Installing the VCS software manually

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

See "Preparing for a manual installation" on page 133.

See "Installing VCS RPMs for a manual installation" on page 133.

See "Adding a license key" on page 135.

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/lltab 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 152.

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

```
set-node north
set-cluster 2
link eth1 eth1 - ether - -
link eth2 eth2 - 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 `lltab`(4) manual page.

Table 7-18 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 7-18 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 `lltab` to set optional static MAC addresses. The second argument to link is the device name of the network interface. Its format is *device_name*:*device_instance_number*. The remaining four arguments to link are defaults; these arguments should be modified only in advanced configurations. There should be one link directive for each network interface. LLT uses an unregistered Ethernet SAP of 0xCAFE. If the SAP is unacceptable, refer to the `lltab(4)` manual page for information on how to customize SAP. Note that IP addresses do not need to be assigned to the network device; LLT does not use IP addresses.
`set-cluster`	Assigns a unique cluster number. Use this directive when more than one cluster is configured on the same physical network connection. LLT uses a default cluster number of zero.
`link-lowpri`	Use this directive in place of `link` for public network interfaces. This directive prevents VCS communication on the public network until the network is the last link, and reduces the rate of heartbeat broadcasts. Note that LLT distributes network traffic evenly across all available network connections and, in addition to enabling VCS communication, broadcasts heartbeats to monitor each network connection.

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

Additional considerations for LLT

You must attach each network interface that is configured for LLT to a separate and distinct physical network.

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/init.d/llt start
   ```
2. Start GAB on Node B.
   ```
   # /etc/init.d/gab start
   ```

Reconfiguring VCS on the existing node

1. On Node A, create the files /etc/lltab, /etc/llthosts, and /etc/gabtab. Use the files created on Node B as a guide, customizing the /etc/lltab for Node A.
2. Start LLT on Node A.
   ```
   # /etc/init.d/llt start
   ```
3. Start GAB on Node A.
   ```
   # /etc/init.d/gab start
   ```
4. Check the membership of the cluster.
   ```
   # gabconfig -a
   ```
5. Start VCS on Node A.
   ```
   # hastart
   ```
6. Make the VCS configuration writable.
   ```
   # haconf -makerw
   ```
7. Add Node B to the cluster.
   ```
   # hasys -add sysB
   ```
8. Add Node B to the system list of each service group.
   - List the service groups.
     ```
     # hagrp -list
     ```
   - For each service group 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**

Appendix A

Advanced topics related to installing VCS

This appendix contains the following topics:

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/lltttab explicitly depending on the subnet for each link.
  See "Broadcast address in the /etc/lltttab file" on page 158.
  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/lltttab 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 160.
- 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/lltttab file.
  See "Sample configuration: Links crossing IP routers" on page 162.

Broadcast address in the /etc/lltttab 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 udp - udp  50000  -  10.20.30.1  10.20.30.255
link link2 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.

```
# ifconfig
eth2 Link encap:Ethernet HWaddr 00:04:23:AC:2B:E4
 inet addr:10.20.30.1 Bcast:10.20.30.255 Mask:255.255.255.0
eth3 Link encap:Ethernet HWaddr 00:04:23:AC:2B:E5
 inet addr:10.20.31.1 Bcast:10.20.31.255 Mask:255.255.255.0
```

The link command in the /etc/lltttab file

Table A-19 describes the fields of the `link` command shown in the /etc/lltttab file examples.

See "Sample configuration: Direct-attached links" on page 161.

See “Sample configuration: Links crossing IP routers” on page 162.

Note that some of these fields differ from the command for standard LLT links.

Table A-19 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 udp. A place holder string. On other unix platforms like Solaris or HP, this entry points to a device file (for example, /dev/udp). Linux does not have devices for protocols. So this field is ignored.
`<node-range>`	Nodes using the link. “`-`” indicates *all* cluster nodes are to be configured for this link.
`<link-type>`	Type of link; must be “`udp`” for LLT over UDP.
`<udp-port>`	Unique UDP port in the range of 49152-65535 for the link. See “Selecting UDP ports” on page 160.
`<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-20 describes the fields of the `set-addr` command.

See “Sample configuration: Links crossing IP routers” on page 162.

Table A-20 Field description for set-addr command in /etc/llttab

Field	Description
`<node-id>`	The ID of the cluster node; for example, 0.

Table A-20 Field description for set-addr command in /etc/llttab

Field	Description
`<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 -au | more
Active Internet connections (servers and established)
Proto Recv-Q Send-Q Local Address Foreign Address         State
udp        0      0 *:32768
*:*
udp        0      0 *:956
*:*
udp        0      0 *:tftp
*:*
udp        0      0 *:sunrpc
*:*
udp        0      0 *:ipp
*:*
```

Look in the UDP section of the output; 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 udp - udp 50000 - 192.168.30.1 192.168.30.255
link link2 udp - udp 50001 - 192.168.31.1 192.168.31.255
```

Sample configuration: Direct-attached links

The following illustration depicts a typical configuration of direct-attached links employing LLT over UDP.

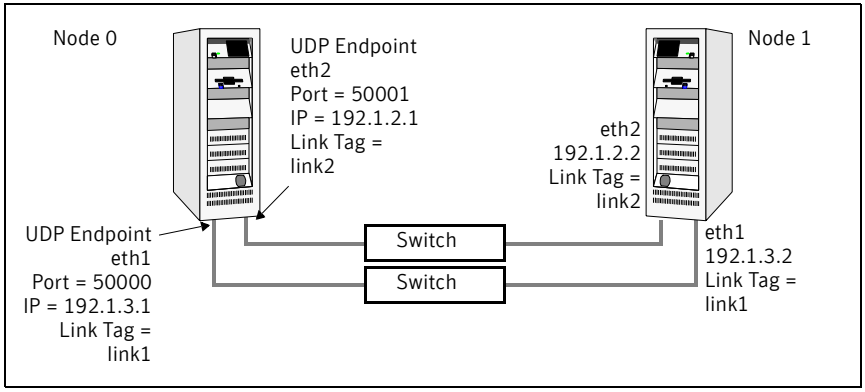

The configuration represented by the following /etc/llttab file for Node 0 has directly attached crossover links or links connected through a hub or switch. These links do not cross routers.

Because LLT broadcasts requests peer nodes to discover their addresses, the addresses of peer nodes do not need to be specified in the /etc/llttab file using the `set-addr` command. For direct attached links, you do need to set the broadcast address of the links in the /etc/llttab file. Verify that the IP addresses and broadcast addresses are set correctly by using the `ifconfig -a` command.

```
set-node Node0
set-cluster 1
#configure Links
#link <tag-name> <device> <node-range> <link-type> <udp port>
<MTU> <IP-address> <bcast-address>
link link1 udp - udp 50000 - 192.1.2.1 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.1 192.1.3.255
```

The file for Node 1 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 udp - udp 50000 - 192.1.2.2 192.1.2.255
link link2 udp - udp 50001 - 192.1.3.2 192.1.3.255
```

Sample configuration: Links crossing IP routers

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.

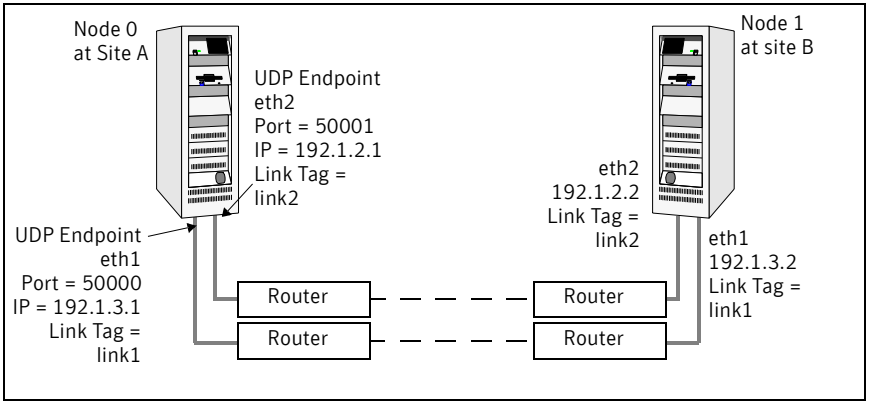

The configuration represented by the following `/etc/llttab` file for Node 1 has links crossing IP routers. Notice that IP addresses are shown for each link on each peer node. In this configuration broadcasts are disabled. Hence, the

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 udp - udp 50000 - 192.1.3.1 -
link link2 udp - udp 50001 - 192.1.4.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr <node-id> <link tag-name> <address>
set-addr        0 link1 192.1.1.1
set-addr        0 link2 192.1.2.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

The `/etc/llttab` file on Node 0 would resemble:

```
set-node Node0
set-cluster 1

link link1 udp - udp 50000 - 192.1.1.1 -
link link2 udp - udp 50001 - 192.1.2.1 -

#set address of each link for all peer nodes in the cluster
#format: set-addr <node-id> <link tag-name> <address>
set-addr        1 link1 192.1.3.1
set-addr        1 link2 192.1.4.1
set-addr        2 link1 192.1.5.2
set-addr        2 link2 192.1.6.2
set-addr        3 link1 192.1.7.3
set-addr        3 link2 192.1.8.3

#disable LLT broadcasts
set-bcasthb     0
set-arp         0
```

Setting up a trust relationship between two authentication brokers

This procedure is a general prerequisite to add secure direct connection clusters to a management server or a peer management server.

To set up the trust relationship

1 Identify which two systems with authentication brokers are to participate in the trust relationship.

To set up a peer management server, these systems are:

- The local management server host
- The remote management server host

For adding secure direct connection clusters, these systems are:

- The system that hosts the authentication broker used by the cluster (can be a cluster node in the cluster)
- The local management server host

2 On one of the two systems, run the following command (the dashes preceding the command options are double dashes):

```
# vssat setuptrust --broker systemname:2821 --securitylevel low
```

Where systemname is the DNS server name or the IP address of the second (other) system.

3 On the second (other) system, obtain a command prompt and run the following command:

```
vssat setuptrust --broker systemname:2821 --securitylevel low
Where systemname is the DNS server name or the IP address of the
first system.
```

If you are adding secure direct connection clusters, you are finished with this procedure.

If you are adding a peer management server, you must continue with the following steps:

4 On both systems, obtain a command prompt and run the following command (the dashes in steps 4–6 are single dashes):

```
# gaserver -list
```

This command returns the ID, system name, and address of the system that runs the command. Record these data for both systems.

5 On the first system, run the following command:

```
# gaserver -add ID systemname address [port]
```

Where ID, systemname, and address are the ID, system name, and IP address of the second (other) system, and port is the default Cluster Management Console port, 14145.

6 On the second (other) system, obtain a command prompt and run the following command:

```
# gaserver -add ID systemname address [port]
```

Where ID, systemname, and address are the ID, system name, and IP address of the first system, and port is the default Cluster Management Console port, 14145.

7 To verify the peer management server, run the gaserver -list command on each management server host. If both command runs return data for both systems, the peer management server is configured successfully.

Index

M

N

O

P

R

S

T

U

V

W