

Veritas Storage Foundation[™] Cluster File System Installation Guide

Linux

5.0

Veritas Storage Foundation Cluster File System Installation Guide

Copyright © 2010 Symantec Corporation. All rights reserved.

Product version: 5.0
Document version: 5.0.1

Symantec, the Symantec logo, Veritas, and Veritas Storage Foundation Cluster File System are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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

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

The Licensed Software and Documentation are deemed to be “commercial computer software” and “commercial computer software documentation” as defined in FAR Sections 12.212 and DFARS Section 227.7202.

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 Veritas 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 Storage Foundation Cluster File System is a licensed product. See the *Veritas Storage Foundation Cluster File System 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	Installing and configuring the product	
	Hardware overview	8
	Shared storage	9
	Fibre channel switch	9
	Cluster platforms	9
	Software components	10
	Preinstallation	13
	Release Notes	13
	Product licensing	13
	Setting PATH and MANPATH environment variables	15
	Secure communication	15
	Veritas Enterprise Administrator	15
	Prerequisites	16
	Installing the product	17
	Using the log files	20
	Installation log file	20
	Response file	21
	Summary file	21
	Installing language packages	21
	Verifying the configuration files	21
	Low Latency Transport configuration files	22
	Checking Low Latency Transport operation	23
	Group Membership and Atomic Broadcast configuration files	25
	Checking Group Membership and Atomic Broadcast operation	25
	Checking cluster operation	26
	Verifying agent configuration	28
	Synchronizing time on Cluster File Systems	28
	Configuring VCS	28
	main.cf file	29
	SFCFS HA Only	30
	VCS application failover services	30
Chapter 2	Upgrading the product	
	Preparing to upgrade the product	32
	Planning the upgrade	32

Upgrade paths	32
Overview of procedures	34
Phased upgrade	34
Full upgrade	34
Ensuring the file systems are clean (full only)	35
Performing the upgrade (phased or full)	36
Shutting down VCS (phased only)	38
Updating the configuration and confirm startup (phased or full)	39
Upgrading remaining nodes (phased only)	40
Making the file systems clean	41

Chapter 3 Adding and removing a node

Adding a node to a cluster	44
Configuring SFCFS and CVM agents on the new node	46
Removing a node from a cluster	49

Chapter 4 Uninstalling the product

Appendix A Troubleshooting and recovery

Installation issues	55
Incorrect permissions for root on remote system	55
Inaccessible system	56
Storage Foundation Cluster File System problems	56
Unmount failures	56
Mount failures	56
Command failures	57
Performance issues	58
High availability issues	58

Installing and configuring the product

This chapter describes how to install the Veritas Storage Foundation Cluster File System (SFCFS). SFCFS requires several Veritas software packages to configure a cluster and to provide messaging services. These packages include the Veritas Cluster Server (VCS) to monitor systems and application services, Veritas Low Latency Transport (LLT) and Veritas Group Membership and Atomic Broadcast (GAB) for messaging and cluster membership, the Veritas Volume Manager (VxVM) to create the shared volumes necessary for cluster file systems, and the Veritas File System (VxFS) package.

Topics covered in this chapter include:


- [Hardware overview](#)
- [Software components](#)
- [Preinstallation](#)
- [Prerequisites](#)
- [Installing the product](#)
- [Using the log files](#)
- [Verifying the configuration files](#)
- [Verifying agent configuration](#)
- [Synchronizing time on Cluster File Systems](#)
- [Configuring VCS](#)
- [VCS application failover services](#)

Hardware overview

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

The figure below shows the configuration of a cluster file system on a Fibre Channel fabric with a disk array.

Figure 1-1 Four Node SFCFS Cluster Built on Fibre Channel Fabric



Shared storage

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

Fibre channel switch

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

Cluster platforms

There are several hardware platforms that can function as nodes in a cluster file system cluster.

See the *Storage Foundation Cluster File System Release Notes*.

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

Software components

Storage Foundation for Cluster File System is the name of the Veritas Cluster File System product and its supporting software rpms. Storage Foundation Cluster File System HA (SFCFS HA) provides support for application failover functionality of Veritas Cluster Server (VCS) in addition to other SFCFS features.

rpms installed with SFCFS and SFCFS HA

The software rpms listed below are required for implementing cluster file system functionality. They are available on the software disc in the depot directory and are installed on each node in the cluster using the `installer` installation script.

The following table shows the rpm name and contents for each rpm:

rpms	Contents
VRTSllt	Veritas Low Latency Transport
VRTSgab	Veritas Group Membership and Atomic Broadcast
VRTSvxfen	Veritas I/O Fencing
VRTSvcs	Veritas Cluster Server
VRTSvxfs-common	Veritas File System Common
VRTSvxfs-platform	Veritas File System Platform Specific
VRTSvxvm-common	Veritas Volume Manager Common
VRTSvxvm-platform	Veritas Volume Manager Platform
VRTScavf	Veritas Cluster Server Agents for Cluster File System
VRTSglm	Veritas Group Lock Manager
VRTSdsa	Veritas Datacenter Storage Agent
VRTSfsman	Veritas File System Manual Pages
VRTSfssdk	Veritas File System Software Developer Kit
VRTSvcsmsg	Veritas Cluster Server Message Catalogs
VRTSacclib	Veritas ACC Library
VRTSvcsag	Veritas Cluster Server Bundled Agents
VRTSvcsdr	Veritas Cluster Server Disk Reservation Modules and Utilities
VRTScscw	Veritas Cluster Server Configuration Wizards

rpms	Contents
VRTScutil	Veritas Cluster Utilities
VRTSfspro	Veritas File System Management Services Provider
VRTSvdid	Veritas Device Identification API
VRTSddlpr	Veritas Device Discovery Layer Services Provider
VRTSvmpro	Veritas Volume Manager Management Services Provider
VRTSvsvc	Veritas Volume Server and Client Provider
VRTSdcli	Veritas Distributed Command Line Interface
VRTSalloc	Veritas Volume Manager Intelligent Storage Provisioning
VRTSvrpro	Veritas Volume Replicator Client Extension and Provider for Veritas Enterprise Administrator
VRTSvcsvr	Veritas Cluster Server Agents for Veritas Volume Replicator
VRTSvrw	Veritas Volume Replicator Web Console

Additional rpms for SFCFS and SFCFS HA

rpms	Contents
VRTSperl	Veritas Perl 5.8.8 Redistribution
VRTSvlic	Veritas License Utilities
VRTSicisco	Symantec Infrastructure Core Services Common
VRTSspbx	Symantec Private Branch Exchange
VRTSsmf	Symantec Service Management Framework
VRTSatClient	Symantec Product Authentication Service Client
VRTSatServer	Symantec Product Authentication Service
VRTSobc33	Veritas Enterprise Administrator Core Service
VRTSob	Veritas Enterprise Administrator Service
VRTSobgui	Veritas Enterprise Administrator
VRTSccg	Veritas Enterprise Administrator Central Control Grid
VRTSmh	Veritas Centralized Management for Storage Foundation - Managed Host
VRTSaa	Veritas Enterprise Administrator Action Agent
VRTSspt	Veritas Software Support Tools
SYMClma	Symantec License Inventory Agent
VRTSjre	Veritas Java Runtime Environment Redistribution
VRTSjre15	Veritas Java Runtime Environment Redistribution
VRTSweb	Veritas Java Web Server

Optional rpms for SFCFS and SFCFS HA

rpms	Contents
VRTSfdoc	Veritas File System Documentation
VRTSfsmnd	Veritas File System Software Developer Kit Manual Pages
VRTScfdoc	Veritas Cluster File System Documentation
VRTSvmdoc	Veritas Volume Manager Documentation
VRTSvmman	Veritas Volume Manager Manual Pages

VRTS1vmconv	Veritas Linux LVM to VxVM Converter
VRTSvcsmn	Veritas Cluster Server Man Pages
VRTScscm	Veritas Cluster Server Cluster Manager
VRTScssim	Veritas Cluster Server Simulator
VRTScmcs	Veritas Cluster Management Console (Single Cluster Mode)
VRTScmccc	Veritas Cluster Management Console Cluster Connector
VRTSvrdoc	Veritas Volume Replicator Documentation

Patches that supersede the ones in this list. To verify that you have the latest patches, go to the Veritas support website to view the following TechNote:

<http://support.veritas.com/docs/281993>

Note: Most of the above patches are available in the Feature11i bundle. The Feature11i bundle and the EnableVxFS bundle are available from HP software download site <http://h20293.www2.hp.com/>. Search for *Veritas 5.0* on this site and follow the instructions.

Preinstallation

Release Notes

Read the *Release Notes* for all products included with this product. Portable Document Format (.pdf) versions of the *Release Notes* are included on the software disc in the `storage_foundation_cluster_file_system/release_notes` directory and on the documentation disc that came with your software.

Because product *Release Notes* are not installed by any packages, it is recommended that you copy them from the disc to the `/opt/VRTS/docs` directory on your system so that they are available for future reference.

Product licensing

Product keys are available on the License Key certificate. The certificate also includes the number of product licenses purchased. A single key enables product installation on the number and the type of systems for which you purchased the license. A key may enable the operation of more products than specified on the

certificate, but you are legally limited to the number of product licenses purchased.

The `VRTSvlic` package executes Veritas product licensing. The following commands and their manual pages are available on the system after `VRTSvlic` is installed:

```
vxlicinst  installs a license key for a Veritas product
vxlicrep   displays currently installed licenses
vxlictest  retrieves features and their descriptions encoded in a license
           key
```

If you encounter problems while licensing your product, visit the Veritas licensing support website at

<http://www.veritas.com/buy/vLicense/vLicenseHome.jhtml>.

Veritas products are installed under the `/opt` directory on the specified host systems. Verify that the directory `/opt` exists and has write permissions for `root` before starting the installation procedure. Do not make `/opt` a VxFS file system.

Setting PATH and MANPATH environment variables

The software and online manual pages for the packages comprising SFCFS are installed in several different directories. However, there are symbolic links to all commands in the `/opt/VRTS/bin` directory, and symbolic links to all manual pages in `/opt/VRTS/man`. To make all SFCFS commands and manual pages accessible when you do the installation, add `/opt/VRTS/bin` to your `PATH` and `/opt/VRTS/man` to your `MANPATH` environment variables. Command line examples in this guide assume these environment variables are set.

Secure communication

Establishing a secure communication between nodes is required to install and configure a cluster. The `ssh` command provides a remote communication and a greater level of security than the `rsh` command.

See the *Veritas Storage Foundation and High Availability Solutions Getting Started Guide*.

Veritas Enterprise Administrator

The Veritas Enterprise Administrator (VEA) client can be installed and run on any machine that supports the Java Runtime Environment.

VEA is required to access the graphical user interface (GUI) for Veritas Storage Foundation. You can use the GUI to administer disks, volumes, file systems, and database functionality on local or remote machines.

One of the following packages needs to be installed and running on the client:

- Veritas Enterprise Administrator (`VRTSobgui`)
- Veritas Enterprise Administrator for Windows
(`windows/VRTSobgui.msi`)

This is the client package for Windows.

Check the *Veritas Storage Foundation Release Notes* for any patch information before you install VEA.

Requirements

The following are system recommendations for the GUI:

OS	Requirements
Windows XP, NT, ME, 2000, or 98	300MHz Pentium with at least 256MB of memory

Prerequisites

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

Provide the following information when installing the SFCFS:

- The cluster name, beginning with a letter (a-z, A-Z).
- A unique ID from 0-65535 for the cluster. Within the public subnet, a new cluster using a duplicate cluster ID can cause existing clusters to fail.
- The Storage Foundation Cluster File System is also supported without I/O fencing enabled. However, without I/O fencing enabled, split brain scenarios can result in data corruption.
- The host names of the cluster nodes.
- The device names of the network interface cards (NICs) used for the private networks among nodes.
- Establishing communication between nodes is required to install Veritas software from a remote system, or to install and configure a cluster. The node from which the installation utility is run must have permissions to run `rsh` (remote shell) or `ssh` (secure shell) utilities as root on all cluster nodes or remote systems.
See [Veritas Storage Foundation and High Availability Solutions Getting Started Guide](#).
- Symantec recommends configuring the cluster with I/O fencing enabled. I/O fencing requires shared devices to support SCSI-3 Persistent Reservations (PR). Enabling I/O fencing prevents data corruption caused by a split brain scenario.
The Storage Foundation Cluster File System is supported without I/O fencing enabled. However, without I/O fencing enabled, split brain scenarios can result in data corruption.

Installing the product

The product installer is the recommended method to license and install the product. The installer also enables you to configure the product, verify preinstallation requirements, and view the product's description.

At most points during an installation, you can type **b** ("back") to return to a previous section of the installation procedure. The back feature of the installation scripts is context-sensitive, so it returns to the beginning of a grouped section of questions. If an installation procedure hangs, use **Control-c** to stop and exit the program. There is a short delay before the script exits.

The following sample procedure is based on the installation of a Veritas Storage Foundation Cluster File System HA cluster with two nodes: "system01" and "system02." If you are installing on standalone systems only, some steps are unnecessary, and these are indicated. Default responses are enclosed by parentheses. Press **Return** to accept defaults.

Note: If you have obtained a Veritas product from an electronic download site, the single product download files do not contain the `installer` installation script, so you must use the product installation script to install the product. For example, if you download Veritas Cluster File System, use the `installsfcfs` script instead of the `installer` script.

To install the product

- 1 Log in as superuser.
- 2 Insert the appropriate media disc into your system's DVD-ROM drive connected to your system.
- 3 After inserting the software disc, type the following command:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```
- 4 Change directory to `/mnt/cdrom`:

```
# cd /mnt/cdrom
```
- 5 Run the `installer` command to install SFCFS. The `installer` script uses `ssh` to communicate with remote nodes as default:

```
# ./installer
```

If you want to use `rsh` you must specify on the command line:

```
# ./installer -rsh
```
- 6 From the Installation menu, choose the `x` option for Install and select **6** the Veritas Storage Foundation Cluster File System. You are prompted to continue with the installation.
- 7 Enter **y** to continue with the installation.

Do you wish to continue with the installation? [y,n,q] (y)

8 Enter one or more system names to install SFCFS.

Enter the system names separated by spaces on which to install SFCFS: **system01 system02**

9 Enter the license key for system01.

Enter a SFCFS license key for system01?

10 Enter **y** to accept another license key or enter **n** to proceed.

Do you want to enter another license key for system02?
[y,n,q] (n)

11 Enter **1**, **2** or **3** to be installed on all systems.

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

12 Press **Return** to continue.

Press [Return] to continue:

13 Enter **y** for configuring SFCFS.

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

14 Enter **y** or **n** for configuring I/O Fencing in enable mode.

Will you be configuring I/O Fencing in enabled mode?
[y,n,q,?] (y)

See *VCS Administration Guide*.

15 Answer the information about the cluster:

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

Discovering NICs on system01 discovered eth0
inet6 eth1 inet6 sit0

Enter the NIC for the first private heartbeat link on system01: [b,?]

Would you like to configure a second private heartbeat link?
[y,n,q,b,?] (y)

Enter the NIC for the second private heartbeat link on system01: [b,?]

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

.
. .
. .

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

16 Enter **y** or **n** to configure SFCFS to use Veritas Security Services.

Would you like to configure SFCFS to use Veritas Security Services? [y,n,q] (n)

17 Enter **y** or **n** to set the username and password. the Cluster Monitor.

Do you want to set the username and/or password for the Admin user (default username = 'admin', password= 'password')?
[y,n,q] (n)

- 18 Enter **y** or **n** to add another user to the cluster.
Do you want to add another user to the cluster? [y,n,q] (y)
- 19 Enter **y** if the information is correct.
Is this information correct? [y,n,q] (y)
- 20 Enter **y** or **n** to configure SMTP notification.
Do you want to configure SMTP notification? [y,n,q] (y)
- 21 Enter **y** or **n** to configure SNMP notification.
Do you want to configure SNMP notification? [y,n,q] (y)
- 22 Enter **y** to set up a default disk group or **n** to continue.
Do you want to set up a default disk group for each system?
[y,n,q,?] (y)
- 23 Enter **y** to accept the fully qualified domain name.
Is the fully qualified hostname of system "system01" =
system01.veritas.com"? [y,n,q]
- 24 Enter **y** to accept the fully qualified domain name.
Is the fully qualified hostname of sytem "system02" =
system02.veritas.com"? [y,n,q]

25 Enter **y** or **n** to configure Storage Foundation Management Server Management.

```
Enable Storage Foundation management Server Mangement?
[y,n,q] (y)
.
.
Starting SFCFS: 100%

Configuring CFS agents for SFCFSHA:

Confirming VxVM configuration daemons startup....All
systems Enabled
Starting CFS agents.....Started
You declined to set up the name of the default disk group
for system01.
You declined to set up the name of the default disk group for
system02.

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

    /opt/VRTS/install/logs/installer-30ibjC

See the Storage Foundation Cluster File System
Administrator's Guide for more information.
```

When installer installs software, some software may be applied rather than committed. It is the responsibility of the system administrator to commit the software, which can be performed later with the `-c` option of the `installp` command.

Using the log files

After product installation, the installation scripts create three text files in the `/opt/VRTS/install/logs` directory. Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.

Installation log file

The installation log file contains commands executed during the procedure, the output, and any errors generated by the commands. This file is for debugging installation problems and can be used by Veritas Technical Services.

Response file

The response file contains the configuration information entered during the procedure. The response file can be used for future installations by invoking an installation script with the `responsefile` option. The response file passes arguments to the script to automate the installation. This file can be edited to automate installation of additional systems.

Summary file

The summary file contains output of the Veritas product installation scripts. This file shows which products were installed, where the log and response files are for a particular installation, and the messages displayed at the end of installation.

Installing language packages

To install language packages you must mount the language package disc and run the installation script `install_lp`.

To install language packages

- 1 Verify that VEA is not running. To display VEA status, type:

```
# /opt/VRTSob/bin/vxsvcctl status
```

If VEA is running, type the following command to stop it:

```
# /opt/VRTSob/bin/vxsvcctl stop
```

- 2 Insert the language disc and run the `install_lp` script.

```
# cd /mnt/cdrom  
# ./install_lp
```

- 3 Restart VEA:

```
# /opt/VRTSob/bin/vxsvcctl start
```

Verifying the configuration files

You can inspect the contents of the configuration files that were installed and modified after a successful installation process. These files reflect the configuration based on the information you supplied.

To verify the configuration files

- 1 Log in as superuser to any system in the cluster.
- 2 Set up your environment `PATH` variable.

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

Low Latency Transport configuration files

The following files are required by the VCS communication services for Low Latency Transport (LLT).

/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 system in the cluster.

For example, the file `/etc/llthosts` contains entries that resemble:

```
0 system01
1 system02
```

/etc/llttab

The file `llttab(4)` contains information that is derived during installation and used by the utility `lltconfig(1M)`. After installation, this file lists the network links that correspond to the specific system.

For example, the file `/etc/llttab` contains entries that resemble:

```
set-node system01
set-cluster 100
link lan1 lan:1 - ether - -
link lan2 lan:2 - ether - -
```

The first line identifies the local system name. 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.

See the `llttab(4)` manual page.

The manual page describes the ordering of the directives in the `llttab` file.

Checking Low Latency Transport operation

Use the `lltstat` command to verify that links are active for LLT. This command returns information about the links for LLT for the system on which it is typed. See the `lltstat(1M)` manual page.

In the following example, `lltstat -n` is typed on each system in the cluster.

To check LLT operation

- 1 Log into system01.

```
# lltstat -n
```

Output resembles:

```
LLT node information:
Node                               State          Links
* 0 system01                       OPEN           2
  1 system02                       OPEN           2
```

- 2 Log into system02.

```
# lltstat -n
```

Output resembles:

```
LLT node information:
Node                               State          Links
  0 system01                       OPEN           2
* 1 system02                       OPEN           2
```

Note: Each system has two links and that each system is in the `OPEN` state. An asterisk (*) denotes the system on which the command is typed.

With LLT configured correctly, the output of `lltstat -n` shows all of the systems in the cluster and two links for each system. If the output shows otherwise, you can use the verbose option of `lltstat`. For example, type `lltstat -nvv | more` on a system to view additional information about LLT. In the following example, `lltstat -nvv | more` is typed on a system in a two-node cluster.

- 3 Log into system01.

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

Output resembles:

```
Node      State      Link      Status      Address
*0 system01 OPEN      lan1       UP           08:00:20:93:0E:34
          lan2       UP           08:00:20:93:0E:34
  1 system02 OPEN      lan1       UP           08:00:20:8F:D1:F2
          lan2       DOWN        08:00:20:8F:D1:F2
  2          CONNWAIT
          lan1       DOWN
          lan2       DOWN
```

```
.  
. .  
31          CONNWAIT  
                lan1          DOWN  
                lan2          DOWN
```

Note: The output lists 32 nodes. It reports on the two cluster nodes, system01 and system02, plus non-existent nodes. For each correctly configured system, the information shows a state of OPEN, a status for each link of UP, and an address for each link. However, in the example above, the output shows that for node system02, the private network may have failed, or the information in `/etc/llttab` may be incorrect.

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

4 Log into system01.

```
# lltstat -p
```

Output resembles:

```
LLT port information:  
Port  Usage  Cookie  
0     gab    0x0  
opens:      0 1 3 4 5 6 7 8 9 10 11 12 13...  
connects:   0 1
```

Note: The two systems with node ID's 0 and 1 are connected.

See [“/etc/llthosts”](#) on page 22.

Group Membership and Atomic Broadcast configuration files

The following files are required by the VCS communication services for Group Membership and Atomic Broadcast (GAB).

/etc/gabtab

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

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

```
/sbin/gabconfig -c -n N
```

where the `-c` option configures the driver for use and `-n N` specifies that the cluster will not be formed until at least `N` systems are ready to form the cluster. `N` is the number of systems in the cluster.

Checking Group Membership and Atomic Broadcast operation

This section describes how to check GAB operation.

To check GAB operation

- Enter the following command on each node in the cluster.

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

If GAB is operational, the following output displays with GAB port membership information:

```
GAB Port Memberships
```

```
=====
```

```
Port a gen 1bbf01 membership 01
Port b gen 1bbf06 membership 01
Port f gen 1bbf0f membership 01
Port h gen 1bbf03 membership 01
Port v gen 1bbf0b membership 01
Port w gen 1bbf0d membership 01
```

If GAB is not operational, the following output display with no GAB port membership information:

```
GAB Port Memberships
```

```
=====
```

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

Checking cluster operation

This section describes how to check cluster operation.

To check cluster operation

- 1 Enter the following command on any system:

```
# hastatus -summary
```

The output for an SFCFS HA installation resembles:

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

A system01                              RUNNING                            0
A system02                              RUNNING                            0

-- GROUP STATE
-- Group      System      Probed AutoDisabled  State

B cvm        system01  Y      N      ONLINE
B cvm        system02  Y      N      OFFLINE
```

Note: If the State value is running, VCS is successfully installed and running on that node. The group state lists the cvm group, which is online on system01 and offline on system02.

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

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

- 2 Enter the following command on any systems:

```
# hasys -display
```

The example on the next page shows the output of `system01`. The list continues with similar information for `system02` (not shown) and any other systems in the cluster. On each system, the output should be similar.

For more information on the `hasys -display` command, see the `hasys(1M)` manual page. Also refer to the chapter in the *Veritas Cluster Server User's Guide*, "Administering VCS From the Command Line."

#System	Attribute	Value
system01	AgentsStopped	0
system01	AvailableCapacity	1
system01	Capacity	1
system01	ConfigBlockCount	54
system01	ConfigChecksum	29776

#System	Attribute	Value
system01	ConfigDiskState	CURRENT
system01	ConfigFile	/etc/VRTSvcs/conf/config
system01	ConfigInfoCnt	0
system01	ConfigModDate	Tues June 25 23:00:00 2006
system01	CurrentLimits	
system01	DiskHbStatus	
system01	DynamicLoad	0
system01	Frozen	0
system01	GUIIPAddr	
system01	LLTNodeId	0
system01	Limits	
system01	LoadTimeCounter	1890
system01	LoadTimeThreshold	600
system01	LoadWarningLevel	80
system01	MajorVersion	2
system01	MinorVersion	0
system01	NodeId	0
system01	OnGrpCnt	1
system01	ShutdownTimeout	60
system01	SourceFile	./main.cf
system01	SysName	system01
system01	SysState	RUNNING
system01	SystemLocation	
system01	SystemOwner	
system01	TFrozen	0
system01	TRSE	0
system01	UpDownState	Up
system01	UserInt	0

#System	Attribute	Value
system01	UserStr	

Verifying agent configuration

This section describes how to verify the agent configuration.

To verify the agent configuration

- Enter the cluster status command from any node in the cluster:

```
# cfscluster status
```

Output resembles:

```
Node           : system01
Cluster Manager : running
CVM state      : running
No mount point registered with cluster configuration

Node           : system02
Cluster Manager : running
CVM state      : running
No mount point registered with cluster configuration
```

Synchronizing time on Cluster File Systems

SFCFS requires that the system clocks on all nodes are synchronized using some external component such as the Network Time Protocol (NTP) daemon. If the nodes are not in sync, timestamps for change (`ctime`) and modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

Configuring VCS

Configuring VCS means conveying to the VCS engine the definitions of the cluster, service groups, resources, and resource dependencies. VCS uses two configuration files in a default configuration:

- The `main.cf` file defines the entire cluster.
- The `types.cf` file defines the resource types.

By default, both files reside in the directory `/etc/VRTSvcs/conf/config`. Additional files similar to `types.cf` may be present if agents have been added, such as `Oracleypes.cf`.

In a VCS cluster, the first system to be brought online reads the configuration file and creates an internal (in-memory) representation of the configuration.

Systems brought online after the first system derive their information from systems running in the cluster. You must stop the cluster while you are modifying the files from the command line. Changes made by editing the configuration files take effect when the cluster is restarted. The node on which the changes were made should be the first node to be brought back online.

main.cf file

The VCS configuration file `main.cf` is created during the installation procedure. After installation, the `main.cf` file contains the base definitions of the cluster and its nodes. Additionally, the file `types.cf` listed in the include statement defines the bundled agents for VCS resources.

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

A typical VCS configuration file for SFCFS file resembles:

```
include "types.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"

cluster cfs_cluster (
    UserNames = { admin = bIJbIDiFJeJhRJdIG }
    Administrators = { admin }
    HacliUserLevel = COMMANDROOT
)

system system01 (
)

system system02 (
)

group cvm (
    SystemList = { system01 = 0, system02 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { system01, system02 }
)

CFSfsckd vxfsckd (
)

CVMCluster cvm_clus (
    CVMClustName = cfscluster
    CVMNodeId = { system01 = 0, system02 = 1 }
    CVMTransport = gab
    CVMTimeout = 200
)

CVMVxconfigd cvm_vxconfigd (
```

```
Critical = 0  
CVMVxconfigdArgs = { syslog }  
)
```

```
cvm_clus requires cvm_vxconfigd  
vxfsckd requires cvm_clus
```

```
// resource dependency tree  
//  
// group cvm  
// {  
// CFSfsckd vxfsckd  
// {  
// CVMCluster  
// cvm_clus  
// {  
// CVMVxconfigd  
// cvm_vxconfigd  
// }  
// }  
// }
```

SFCFS HA Only

If you configured VCS Cluster Manager (Web Console), a service group, “ClusterService,” was created that includes IP, Process, and Notifier resources. These resources were configured according to information you provided during the installation procedure. A resource dependency was also created.

VCS application failover services

If you installed SFCFS HA, you can begin implementing the application monitoring failover services provided by the Veritas Cluster Server. Information about setting up VCS services is beyond the scope of this document.

See the *Veritas Cluster Server* documentation.

Upgrading the product

If you are running an earlier release of Veritas Storage Foundation Cluster File System, you can upgrade your product using the procedures described in this chapter.

Topics covered in this chapter include:

- [Preparing to upgrade the product](#)
- [Overview of procedures](#)
- [Ensuring the file systems are clean \(full only\)](#)
- [Performing the upgrade \(phased or full\)](#)
- [Making the file systems clean](#)

Note: If your system is running an older version of either Red Hat Enterprise Linux or SUSE Linux Enterprise Server, you must upgrade it before attempting to install the Veritas Storage Foundation Cluster File System software. Consult the *Red Hat* or *SUSE* documentation for more information on upgrading your system.

Preparing to upgrade the product

This section prepares you for the Veritas Storage Foundation Cluster File System upgrade.

Planning the upgrade

Complete the following tasks in advance of upgrading:

- Review the *Veritas Storage Foundation Cluster File System Release Notes* for any late-breaking information on upgrading your system.
- Be sure that the administrator doing the upgrade has root access and a working knowledge of system administration.
- Schedule sufficient outage time for the upgrade.
- Make sure you have upgraded all the file systems to disk layout Version 6, before you upgrade SFCFS to 5.0. Disk layout Version 7 is the recommended version for SFCFS 5.0.

See `vxupgrade(1M)`, `vxfsconvert(1M)`, and `fsadm(1M)`.

See the *Veritas File System Administrator's Guide*.

- Verify all the file systems are working fine and data is intact.
 See `cfsmount(1M)`.

Upgrade paths

The upgrade paths for Veritas Storage Foundation Cluster File System on RHEL4 are:

RHEL4 Update	From	Upgrade To	Tasks
Update 1 or Update 2	Storage Foundation Cluster File System 4.1 GA or any 4.1 MP release	Storage Foundation Cluster File System 5.0	<ul style="list-style-type: none"> ■ Upgrade to RHEL4 Update 3 ■ Proceed to “Overview of procedures” on page 34.
Update 3	Storage Foundation Cluster File System 4.1 GA or any 4.1 MP release	Storage Foundation Cluster File System 5.0	<ul style="list-style-type: none"> ■ Proceed to “Overview of procedures” on page 34.

The upgrade paths for Veritas Storage Foundation Cluster File System on SLES9 are:

SLES9 Service Pack	From	Upgrade To	Tasks
SP 1 or SP 2	Storage Foundation Cluster File System 4.1 GA	Storage Foundation Cluster File System 5.0	<ul style="list-style-type: none"> ■ Upgrade Storage Foundation Cluster File System to latest MP release of 4.1. ^a ■ Upgrade to SLES9 SP3. ■ Proceed to “Overview of procedures” on page 34.
SP 1 or SP 2	Any Storage Foundation Cluster File System 4.1 MP release	Storage Foundation Cluster File System 5.0	<ul style="list-style-type: none"> ■ Upgrade to SLES9 SP3. ■ Proceed to “Overview of procedures” on page 34.
SP 3	Any Storage Foundation Cluster File System 4.1 MP release	Storage Foundation Cluster File System 5.0	<ul style="list-style-type: none"> ■ Proceed to “Overview of procedures” on page 34.

a..4.1 MP2 is the latest MP release for Storage Foundation Cluster File System at the time of this document.

Overview of procedures

There are two ways to upgrade cluster nodes to the latest version of Storage Foundation Cluster File System: phased and full.

The upgrade procedures apply to both the phased and full upgrade procedures unless otherwise noted. Occasionally, steps differ between the two procedures. Screen output is also common between both procedures unless otherwise noted.

Note: Both procedures automatically uninstall the previous version of the software.

Phased upgrade

Note: A phased upgrade should not be performed from one of the nodes in the cluster.

A phased upgrade minimizes downtime by upgrading portions of the cluster, one at a time. Although the entire cluster is offline for a shorter period than a full upgrade, this method requires command-line interaction and some manual configuration.

The stages of the phased upgrade procedure are:

- 1 Select two or more nodes to upgrade.
- 2 Install the new version.
- 3 Shut down VCS on remaining non-upgraded nodes and ensuring the file systems are clean.
- 4 Reboot the upgraded nodes.
- 5 Install the new version on each remaining node and reboot them.

Full upgrade

A full upgrade upgrades the product on the entire cluster and the cluster remains offline for the duration of the procedure. Minimal command-line interaction and some manual configuration are required.

The stages of the full upgrade procedure are:

- 1 Ensuring the file systems are clean.
- 2 Install the new version.

Ensuring the file systems are clean (full only)

Before upgrading to SFCFS 5.0, ensure that the file systems are clean. To ensure that the logs have been replayed and the file systems are marked clean:

To ensure the file systems are clean

- 1 Log in as superuser onto any node in the cluster.
- 2 Offline the group on each node of the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFMount resource.

Note: Repeat [step 2](#) for each SFCFS service group.

- 3 Check and repair each file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

Note: Repeat [step 3](#) for each file system.

Performing the upgrade (phased or full)

This section describes how to upgrade for phased or full.

If you are performing a phased upgrade, select one or more nodes to upgrade.

To perform the upgrade

- 1 Log in as superuser.
- 2 Insert the appropriate media disc per your distribution and architecture into your system's DVD-ROM drive.
- 3 If volume management software is running on your system, the software disc automatically mounts as `/mnt/cdrom`.

If volume management software is not available to mount the disc, you must mount it manually, enter:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```

- 4 Change to the top-level directory on the disc:

```
# cd /mnt/cdrom
```
- 5 Verify there are no VxFS file systems mounted on the nodes being upgraded:

```
# mount -t vxfs
```

- a If any VxFS file systems are mounted, offline the group on each node of the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFSSMount resource.

- b Repeat [step a](#) for each SFCFS service group.
- 6 If performing a phased upgrade, start the installation procedure from the node selected in the cluster. In the following example the phased upgrade is performed on one node in a four-node cluster.

Note: Ensure that the HAD daemon of VCS is running on all nodes selected for the upgrade.

```
# ./installsfdfs system01 system02
```

- a Press **Return**.

If performing a full upgrade, start the installation from any node in the cluster.

```
# ./installsfcfs
```

a Press **y** to upgrade the cluster configuration.

- 7 Press **Return** to begin installing infrastructure packages.
- 8 Press **Return** to begin license verification.
- 9 Press **Return** to begin the uninstall. The uninstall utility checks the system's uninstallation requirements and identifies packages, patches and dependencies.

Note: The output can be extensive and exceed 100 lines, depending on your configuration.

- 10 If VCS is running you are prompted to upgrade and reset the VCS password. To ensure security, passwords are not written to the install log.

```
installsfcfs must now make configuration updates and stop
the cluster before upgrading VCS packages.
```

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

- 11 Enter **y**.
- 12 At the prompt, enter your new password.
- 13 Reenter your new password.
- 14 Output shows information that Cluster Server must be started on a running system. Enter **y** to continue.
- 15 Press **Return** to begin removing the previous packages and installing the new.
- 16 Press **Return** again for summary information about logs and reboots.

Note: Do not remove the log files until the Veritas products are working properly on your system. Technical Support will need these log files for debugging purposes.

- 17 If performing a phased upgrade, proceed to [“Shutting down VCS \(phased only\)”](#) on page 38.
 If performing a full upgrade, proceed to [“Updating the configuration and confirm startup \(phased or full\)”](#) on page 39.

Shutting down VCS (phased only)

Shutdown VCS on remaining nodes that are not being upgraded to preventing them from rejoining the cluster.

To shut down the cluster

- 1 Separate the nodes that are not being upgraded from those that are.

- 2 Check to see if there are frozen CVM and SFCFS groups, enter

```
# /opt/VRTSvcs/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

- a Make configuration read/write:

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

- b Unfreeze the group:

```
# /opt/VRTSvcs/bin/hagrps -unfreeze group -persistent
```

- c Save the configuration:

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

- 3 Offline the group on each of the remaining nodes of the cluster:

```
# hagrps -offline group -sys system01
```

```
# hagrps -offline group -sys system02
```

```
# hagrps -offline group -sys system03
```

```
# hagrps -offline group -sys system04
```

where *group* is the VCS service group that has the CVMVolDg and CFSSMount resource.

Note: Repeat [step 2](#) for each SFCFS service group.

- 4 Check and repair each file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

Note: Repeat [step 3](#) for each file system.

- 5 On each node that is not being upgraded, shutdown VCS, enter the following commands:

```
# /opt/VRTSvcs/bin/hastop -local
```

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

```
# /opt/VRTS/bin/fsclustadm cfsdeinit
```

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

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

- 6 Proceed to “[Updating the configuration and confirm startup \(phased or full\)](#).”

Updating the configuration and confirm startup (phased or full)

Perform the following steps on each upgraded node.

To update the configuration and confirm startup

- 1 Remove the `/etc/VRTSvcs/conf/config/.stale` file, if it exists.

```
# rm -f /etc/VRTSvcs/conf/config/.stale
```
- 2 Reboot the upgraded nodes.

```
# reboot
```
- 3 After the nodes reboot, verify that LLT is running:

```
# lltconfig
LLT is running
```
- 4 Verify GAB is configured:

```
# gabconfig -l | /usr/bin/grep 'Driver.state' | \
/usr/bin/grep Configured
Driver state : Configured
```
- 5 Verify VxVM daemon is started and enabled:

```
# /opt/VRTS/bin/vxdctl mode
mode: enabled
```
- 6 Confirm all upgraded nodes are in a RUNNING state.

```
# /opt/VRTSvcs/bin/hasys -state | /usr/bin/grep RUNNING | \
/usr/bin/wc -l
1
```
- 7 After the configuration is complete, the CVM and SFCFS groups may come up frozen. To find out the frozen CVM and SFCFS groups, enter

```
# /opt/VRTS/bin/hastatus -sum
```

If the groups are frozen, unfreeze CVM and SFCFS groups using the following commands for each group:

 - a Make configuration read/write:

```
# /opt/VRTS/bin/haconf -makerw
```
 - b Unfreeze the group:

```
# /opt/VRTS/bin/hagrps -unfreeze group -persistent
```
 - c Save the configuration:

```
# /opt/VRTS/bin/haconf -dump -makero
```
- 8 If the CVM and SFCFS groups are offline, bring them online on all systems:

```
# /opt/VRTS/bin/hagrps -online group -sys system01
# /opt/VRTS/bin/hagrps -online group -sys system02
```

Note: If the SFCFS service groups do not come online then your file system could be dirty. Proceed to [“Making the file systems clean”](#) on page 41.

- 9 If performing a phased upgrade, proceed to “[Upgrading remaining nodes \(phased only\)](#)” on page 40.
If performing a full upgrade, proceed to [step 7](#) on page 40.

Upgrading remaining nodes (phased only)

To upgrade remaining nodes

- 1 Log in as superuser.
- 2 Insert the appropriate media disc per your distribution and architecture into your system’s DVD-ROM drive.
- 3 Change to the `storage_foundation_cluster_file_system` directory:

```
# cd /mnt/cdrom/storage_foundation_cluster_file_system
```
- 4 Install SFCFS on the remaining nodes. Type:

```
# ./installsfdfs system03 system04
```

When upgrading the remaining nodes, you may be prompted that Cluster Server should have been running during the initial upgrade.
- 5 Press **y** to continue and follow all the prompts.
- 6 `system03system04` Check cluster status. Type:

```
# hastatus -summary
```
- 7 If you are configuring SFCFS for a fenced environment.
See the *Veritas Cluster Server Administrator’s Guide*.
See the *Veritas Storage Foundation Cluster File System Administrator’s Guide*.
- 8 Verify the cluster protocol version, enter the following command that displays the same on a running node:

```
# /opt/VRTS/bin/vxdctl protocolversion
```

If the cluster protocol version is less than 70, then it needs to be upgraded to 70 for SFCFS 5.0.

Note: The cluster protocol version can only be updated on the master node.

One can upgrade the entire cluster using the following command on the master node.

```
# /opt/VRTS/bin/vxdctl upgrade
```

- 9 Type the following command on one upgraded node to enable membership:

```
# gabconfig -xc
```

Making the file systems clean

If you upgrade to SFCFS 5.0 and the file systems are dirty, you have to deport the shared disk group and import it as non-shared. After the import, run `fsck`. `fsck` should succeed. Then deport the disk group and import it back as shared.

To make the file systems clean

- 1 Log in as superuser onto the CVM master node.
- 2 If performing a full upgrade, offline the group on all the nodes in the cluster:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
# hagr -offline group -sys system03
# hagr -offline group -sys system04
```

If performing a phased upgrade, offline the group:

```
# hagr -offline group -sys system01
# hagr -offline group -sys system02
```

where `group` is the VCS service group that has the CVMVolDg and CFSSMount resource.

- 3 Deport the disk group:

```
# vxdg deport diskgroup
```

where `diskgroup` is the SFCFS disk group.

- 4 Import the disk group:

```
# vxdg -C import diskgroup
```

- 5 Start all the volumes in the disk group:

```
# vxvol -g diskgroup startall
```

- 6 Check and repair each file system:

```
# fsck -t vxfs /dev/vx/dsk/diskgroup/volume
```

Note: Repeat [step 3](#) for each file system.

- 7 Deport the disk group:

```
# vxdg deport diskgroup
```

- 8 Import the disk group:

```
# vxdg -s import diskgroup
```

- 9 Start all the volumes in the disk group:

```
# vxvol -g diskgroup startall
```

- 10 If performing a full upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01  
# hares -clear resource -sys system02  
# hares -clear resource -sys system03  
# hares -clear resource -sys system04
```

If performing a phased upgrade, for all the resources that are faulted, run the following command:

```
# hares -clear resource -sys system01  
# hares -clear resource -sys system02
```

- 11 If performing a full upgrade, online the group on all the nodes in the cluster:

```
# hagrps -online group -sys system01  
# hagrps -online group -sys system02  
# hagrps -online group -sys system03  
# hagrps -online group -sys system04
```

If performing a phased upgrade, online the group:

```
# hagrps -online group -sys system01  
# hagrps -online group -sys system02
```

Adding and removing a node

This chapter provides information on how to add a node to an existing cluster and removing a node from a cluster. Topics include:

- [Adding a node to a cluster](#)
- [Configuring SFCFS and CVM agents on the new node](#)
- [Removing a node from a cluster](#)

Adding a node to a cluster

If you want to add a new node to a multi-node cluster, first prepare the new system hardware. Physically connect the new system to the cluster using private networks and attach to any shared storage. Then install the required OS software.

See “[Cluster platforms](#)” on page 9.

To add a node to a cluster

- 1 Log into the new system as superuser.
- 2 After inserting the software disc, type the following command:

```
# mount -o ro /dev/cdrom /mnt/cdrom
```
- 3 Add `/opt/VRTS/bin` to your `PATH` and `/opt/VRTS/man` to your `MANPATH` environment variables
See “[Setting PATH and MANPATH environment variables](#)” on page 15.
- 4 Change to the SFCFS directory.

```
# cd sfcfs
```
- 5 Run the `installsfcfs` script with `-installonly` option to install all the required SFCFS packages on the new node.

```
# ./installsfcfs -installonly
```
- 6 Enter **y** to install SFCFS on these systems.
Do you want to install SFCFS on these systems? [y,n,q] (y)
- 7 Enter the system name of the new node to install SFCFS.
Enter the system names separated by spaces on which to install SFCFS: **system03**
- 8 Enter a license key for system03.
Enter a SFCFS license key for system03:[?]
XXXX-XXXX-XXXX-XXXX-XXXX-X
- 9 Enter **y** or **n** for another license key. You are prompted to press Return to continue.
Do you want to enter another license key for system03?
[y,n,q,?] (n)
- 10 Enter **1** or **2** to be installed on all systems.
Select the rpms to be installed on all systems?
[1-2,q,?] (2)

- 11 Press **Return** to continue.

Press [Return] to continue:

Note: Do not reboot the machine now.

- 12 Create the `/etc/llttab` file the same as it looks on another node in the cluster.

- a Change the `set-node` line to the name of the new node and specify that the LAN ports for the public and private networks are configured the same as on the other cluster nodes:

```
set-node system03
set-cluster 100
link lan2 /dev/lan:2 - ether - -
link lan0 /dev/lan:0 - ether - -
```

- b Copy `/etc/llthosts` from one other node of the cluster and add a new line to it with the name of this node.

- c Copy the same `llthosts` file to all nodes of the cluster.
See “[/etc/llttab](#)” on page 22.

- 13 Create `/etc/gabtab` file the same as it is on another node in the cluster.
See “[/etc/gabtab](#)” on page 25.

For example,

```
/sbin/gabconfig -c -n 2
```

There is no need to reboot the other nodes, just update the `/etc/gabtab` file on the other nodes in the cluster.

- 14 Reboot the system that has been added.

```
# /sbin/shutdown -r now
```

- 15 Start VxVM on the system that has been added.

```
# vxinstall
```

- 16 After starting VxVM, proceed to “[Configuring SFCFS and CVM agents on the new node.](#)”

Configuring SFCFS and CVM agents on the new node

You must configure the SFCFS and CVM agents, after rebooting the new system.

To configure SFCFS and CVM agents on the new node

- 1 Verify if `/etc/VRTSvcs/conf/config/.state` file is present, before starting VCS.
If `/etc/VRTSvcs/conf/config/.state` file is not present, enter:

```
# touch /etc/VRTSvcs/conf/config/.state
```
- 2 Start the VCS server and `vxfen` on `system03`.
 - a Start LLT and GAB on the new node:

```
# /etc/init.d/llt start  
# /etc/init.d/gab start
```
 - b For starting `vxfen` in the disable mode, run the following commands on `system03`:

```
# echo vxfen_mode=disabled > /etc/vxfenmode  
# /etc/init.d/vxfen start
```
 - c For starting `vxfen` in the enabled mode:
 - Copy the following files from one of the existing cluster nodes to `system03`:

```
/etc/vxfenmode  
/etc/vxfendg
```
 - Run the following command:

```
# /etc/init.d/vxfen start
```
- 3 On the new node, verify that the GAB port memberships are a and b. Run the following command:

```
# /sbin/gabconfig -a
```
- 4 Determine the CVM master node:

```
# vxdctl -c mode
```
- 5 Make a backup copy of the `main.cf` file. Enter the following commands:

```
# cd /etc/VRTSvcs/conf/config  
# cp main.cf main.cf.2node
```
- 6 Open the VCS configuration for writing and add the new node. For example:

```
# haconf -makerw  
# hasys -add system03
```
- 7 Add the new node to the CVM system list and specify a failover priority of zero:

```
# hagrps -modify cvm SystemList -add system03 X
```

where `X` is one more than the index of the last system in System list of CVM service group in `/etc/VRTSvcs/conf/config/main.cf`.

- 8 Add the new node to the CVM AutoStartList:


```
# hagrpl -modify cvm AutoStartList -add system03
```
- 9 Node ID can be obtained from `CVMNodeId` of `/etc/VRTSvcs/conf/config/main.cf`. Add the new node, `system03`, and its node ID, #, to the `cvm_clust` resource:


```
# hares -modify cvm_clus CVMNodeId -add system03 2
```
- 10 Write the new VCS configuration to disk:


```
# haconf -dump -makero
```
- 11 Verify the syntax of `main.cf` file:


```
# hacf -verify .
```
- 12 To enable the existing cluster to recognize the new node, execute on all the nodes in the cluster:


```
# /etc/vx/bin/vxclustadm -m vcs -t gab reinit
# /etc/vx/bin/vxclustadm nidmap
```
- 13 Start CVM on the newly added node.
 - Determine the node ID:


```
# cat /etc/llthosts
```
 - Verify that this host ID is seen by the GAB module:


```
# gabconfig -a
```
 - Start the VCS engine.
 - If on the newly added node ports `f`, `u`, `v`, or `w` are present before `hastart`, then the newly added node must be rebooted to properly start VCS. To properly start VCS:


```
# shutdown -r
```
 - If on the newly added node ports `f`, `u`, `v` or `w` were not present before `hastart`, then use the following command to start VCS:


```
# hastart
```
- 14 Check the system status to see whether the new node is online:


```
# hastatus -sum
-- SYSTEM STATE
-- System          State          Frozen
A      system01   RUNNING       0
A      system02   RUNNING       0
A      system03   RUNNING       0

-- GROUP STATE
-- Group   System    Probed  AutoDisabled  State
B   cvm    system01  Y          N             ONLINE
B   cvm    system02  Y          N             ONLINE
B   cvm    system03  Y          N             ONLINE
```
- 15 Add shared disk groups to the cluster configuration:


```
# cfsdgadm add cfsdg system03=sw
```

16 Create a `/mnt` on `system03` and run the following commands:

```
# cfsmntadm modify /mnt add system03=rw
```

Refer to `cfsmntadm` man page for more details.

17 Use `cfsmount` command to cluster mount `/mnt` back on all the nodes:

```
# cfsmount /mnt
```

Removing a node from a cluster

This section describes how to remove a node from a cluster. As in previous examples, the following removes the system `system03` from a three-node cluster. The procedure can be done from any node remaining in the cluster or from a remote host.

To remove a node from a cluster

- 1 Log in as superuser on a node other than `system03`.
- 2 Use `cfsumount` command to unmount the file system `/mnt` on all the nodes:

```
# cfsumount /mnt system03
```
- 3 Stop all the cluster components:

```
# cfscluster stop -f system03
```
- 4 Open the VCS configuration for writing:

```
# haconf -makerw
```
- 5 Remove `system03` from the system list attribute of the CVM and SFCFS service groups:

```
# hagr -modify service_group SystemList -delete system03  
# hagr -modify cvm SystemListf -delete system03
```

where `service_group` is the command that displays the service groups by `hagr -dep cvm`.

If an error message similar to the following is displayed by either of the above commands:

```
VCS:10456:Configuration must be ReadWrite. ('hagr  
-modify ...  
-delete(0x10f)', Sysstate=RUNNING, Channel=IPM, Flags=0x0)
```

Repeat [step 4](#) and the command that failed in [step 5](#).

- 6 Write the new VCS configuration to disk:

```
# haconf -dump -makero
```
- 7 Edit `/etc/llthosts` on the remaining nodes of the cluster, and remove the entry corresponding to the node being removed.
- 8 Edit `/etc/gabtab` on the remaining nodes of the cluster and edit the `gabconfig` command to reflect the correct and new number of nodes in the cluster.
- 9 Login to `system03` and remove the following files:

```
# rm /etc/vxfenmode  
# rm /etc/llthosts  
# rm /etc/llttab  
# rm /etc/gabtab
```
- 10 If fencing was enabled on the cluster, run the following commands:

```
# rm /etc/vxfentab
```

```
# rm /etc/vxfendg
```

11 Reboot system03:

```
# /sbin/shutdown -r
```

12 Change to the `install` directory:

```
# cd /opt/VRTS/install
```

13 From the `scripts` directory, run the `uninstallsfcfs` script and remove SFCFS on system03:

```
# ./uninstallsfcfs
```

If you do not want to remove the Veritas Cluster Server software, enter **n** when prompted to uninstall VCS.

See the *Veritas Cluster Server Installation Guide*, on the software disc.

Uninstalling the product

If you need to uninstall SFCFS software. Use the `uninstallsfcfs` script.

To uninstall SFCFS HA

- 1 Log in as superuser.

Note: Do not use the `hastop -force` command to stop VCS.

- 2 Change directory to `/opt/VRTS/install`:

```
# cd /opt/VRTS/install
```

- 3 Run the `uninstallsfcfs` command to uninstall SFCFS. The `uninstallsfcfs` script uses `ssh` to communicate with remote nodes as default:

```
# ./uninstallsfcfs
```

If you want to use `rsh` you must specify on the command line:

```
# ./uninstallsfcfs -rsh
```

- 4 Enter the system names to uninstall SFCFS.

```
Enter the system names separated by spaces on which to  
uninstall SFCFS: system01 system02
```

- 5 Enter **y** to uninstall SFCFS.

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


Troubleshooting and recovery

Installation issues

If you encounter any issues installing SFCFS, refer to the following paragraphs for typical problems and their solutions.

Incorrect permissions for root on remote system

The permissions are inappropriate. Make sure you have remote root access permission on each system to which you are installing.

```
Checking communication with system01 ..... FAILED
Remote remsh/rcp permissions not available on: system01
Correct permissions and continue
Continue? [Y/N] :
```

Suggested solution: You need to set up the systems to allow remote access using `ssh` or `rsh`.

See the *Veritas Storage Foundation and High Availability Solutions Getting Started Guide*.

Note: Remove remote shell permissions after completing the SFCFS installation and configuration.

Inaccessible system

The system you specified is not accessible. This could be for a variety of reasons such as, the system name was entered incorrectly or the system is not available over the network.

```
Checking communication with system01 ..... FAILED
System not accessible : system01
```

Suggested solution: Verify that you entered the system name correctly; use the `ping(1M)` command to verify the accessibility of the host.

Storage Foundation Cluster File System problems

If there is a device failure or controller failure to a device, the file system may become disabled cluster-wide. To address the problem, unmount file system on all the nodes, then run a full `fsck`. When the file system check completes, mount all nodes again.

Unmount failures

The `umount` command can fail if a reference is being held by an NFS server. Unshare the mount point and try the unmount again.

Mount failures

Mounting a file system can fail for the following reasons:

- The file system is not using disk layout Version 6 or 7.
- The mount options do not match the options of already mounted nodes.
- A cluster file system is mounted by default with the `qio` option enabled if the node has a Quick I/O for Databases license installed, even if the `qio` mount option was not explicitly specified. If the Quick I/O license is not installed, a cluster file system is mounted *without* the `qio` option enabled. So if some nodes in the cluster have a Quick I/O license installed and others do not, a cluster mount can succeed on some nodes and fail on others due to different mount options. To avoid this situation, ensure that Quick I/O licensing is uniformly applied, or be careful to mount the cluster file system with the `qio/noqio` option appropriately specified on each node of the cluster.
See the `mount(1M)` manual page.
- A shared CVM volume was not specified.
- The device is still mounted as a local file system somewhere on the cluster. Unmount the device.

- The `fsck` or `mkfs` command is being run on the same volume from another node, or the volume is mounted in non-cluster mode from another node.
- The `vxfsckd` daemon is not running. This typically happens only if the `CFSfsckd` agent was not started correctly.
- If `mount` fails with an error message:


```
vxfs mount: cannot open mnttab
/etc/mnttab is missing or you do not have root privileges.
```
- If `mount` fails with an error message:


```
vxfs mount: device already mounted, ...
```

 The device is in use by `mount`, `mkfs` or `fsck` on the same node. This error cannot be generated from another node in the cluster.
- If this error message displays:


```
mount: slow
```

 The node may be in the process of joining the cluster.
- If you try to mount a file system that is already mounted without `-o cluster` option (that is, not in shared mode) on another cluster node,


```
# mount -t vxfs /dev/vx/dsk/share/vol01 /vol01
```

 The following error message displays:


```
vxfs mount: /dev/vx/dsk/share/vol01 is already mounted,
/vol01 is busy, allowable number of mount points exceeded,
or cluster reservation failed for the volume
```

Command failures

- Manual pages not accessible with the `man` command. Set the `MANPATH` environment variable as listed under “[Setting PATH and MANPATH environment variables](#)” on page 15.
- The `mount`, `fsck`, and `mkfs` utilities reserve a shared volume. They fail on volumes that are in use. Be careful when accessing shared volumes with other utilities such as `dd`, it is possible for these commands to destroy data on the disk.
- Running some commands, such as `vxupgrade -n 7/vol02`, can generate the following error message:


```
vxfs vxupgrade: ERROR: not primary in a cluster file
system
```

 This means that you can run this command only on the primary, that is, the system that mounted this file system first.

Performance issues

Quick I/O File system performance is adversely affected if a cluster file system is mounted with the `qio` option enabled and Quick I/O is licensed, but the file system is *not* used for Quick I/O files. Because `qio` is enabled by default, if you do not intend to use a shared file system for Quick I/O, explicitly specify the `noqio` option when mounting.

High availability issues

Network partition/jeopardy

Network partition (or *split brain*) is a condition where a network failure can be misinterpreted as a failure of one or more nodes in a cluster. If one system in the cluster incorrectly assumes that another system failed, it may restart applications already running on the other system, thereby corrupting data. CFS tries to prevent this by having redundant heartbeat links.

At least one link must be active to maintain the integrity of the cluster. If all the links go down, after the last network link is broken, the node can no longer communicate with other nodes in the cluster. Thus the cluster is in one of two possible states. Either the last network link is broken (called a network partition condition), or the last network link is okay, but the node crashed, in which case it is not a network partition problem. It is not possible to identify whether it is the first or second state, so a kernel message is issued to indicate that a network partition may exist and there is a possibility of data corruption.

Jeopardy is a condition where a node in the cluster has a problem connecting to other nodes. In this situation, the link or disk heartbeat may be down, so a *jeopardy* warning may be displayed. Specifically, this message appears when a node has only one remaining link to the cluster and that link is a network link. This is considered a critical event because the node may lose its only remaining connection to the network.

Caution: Do not remove the communication links while shared storage is still connected.

Low memory

Under heavy loads, software that manages heartbeat communication links may not be able to allocate kernel memory. If this occurs, a node halts to avoid any chance of network partitioning. Reduce the load on the node if this happens frequently.

A similar situation may occur if the values in the `/etc/llttab` files on all cluster nodes are not correct or identical.

