

Veritas Storage Foundation[™] Cluster File System Installation Guide

HP-UX

5.0

Veritas Storage Foundation Cluster File System Installation Guide

Copyright © 2006 Symantec Corporation. All rights reserved.

SFCFS 5.0

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.

HP-UX is a registered trademark of Hewlett-Packard Development Company, L.P.

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	10
	Shared storage	11
	Fibre channel switch	11
	Cluster platforms	11
	Software components	12
	Required HP-UX patches	15
	Preinstallation	17
	Release Notes	17
	Product licensing	17
	Setting PATH and MANPATH environment variables	18
	Secure communication	18
	Veritas Enterprise Administrator	19
	Prerequisites	19
	Installing the product	20
	Configuring the Components	23
	Using the log files	25
	Installation log file	25
	Response file	25
	Summary file	25
	Verifying the configuration files	26
	Low Latency Transport configuration files	26
	Checking Low Latency Transport operation	27
	Group Membership and Atomic Broadcast configuration files	29
	Checking Group Membership and Atomic Broadcast operation	29
	Checking cluster operation	30
	Verifying agent configuration	32
	Synchronizing time on Cluster File Systems	32
	Configuring VCS	32
	main.cf file	33
	SFCFS HA Only	34
	VCS application failover services	34

Chapter 2	Upgrading the product	
	Preparing to upgrade the product	36
	Planning the upgrade	36
	Upgrade paths	36
	Upgrade Overview	38
	Phased upgrade	38
	Full upgrade	38
	Upgrading from 3.5 to 5.0	39
	Phased upgrade	39
	Full upgrade	42
	Upgrading from 4.1 to 5.0	45
	Phased upgrade	45
	Full upgrade	48
	Upgrading the disk layout versions	50
Chapter 3	Adding and removing a node	
	Adding a node to a cluster	54
	Configuring SFCFS and CVM agents on the new node	56
	Removing a node from a cluster	57
Chapter 4	Uninstalling the product	
Appendix A	Troubleshooting and recovery	
	Installation issues	63
	Incorrect permissions for root on remote system	63
	Resource temporarily unavailable	64
	Inaccessible system	64
	Storage Foundation Cluster File System problems	64
	Unmount failures	65
	Mount failures	65
	Command failures	66
	Performance issues	66
	High availability issues	67

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:

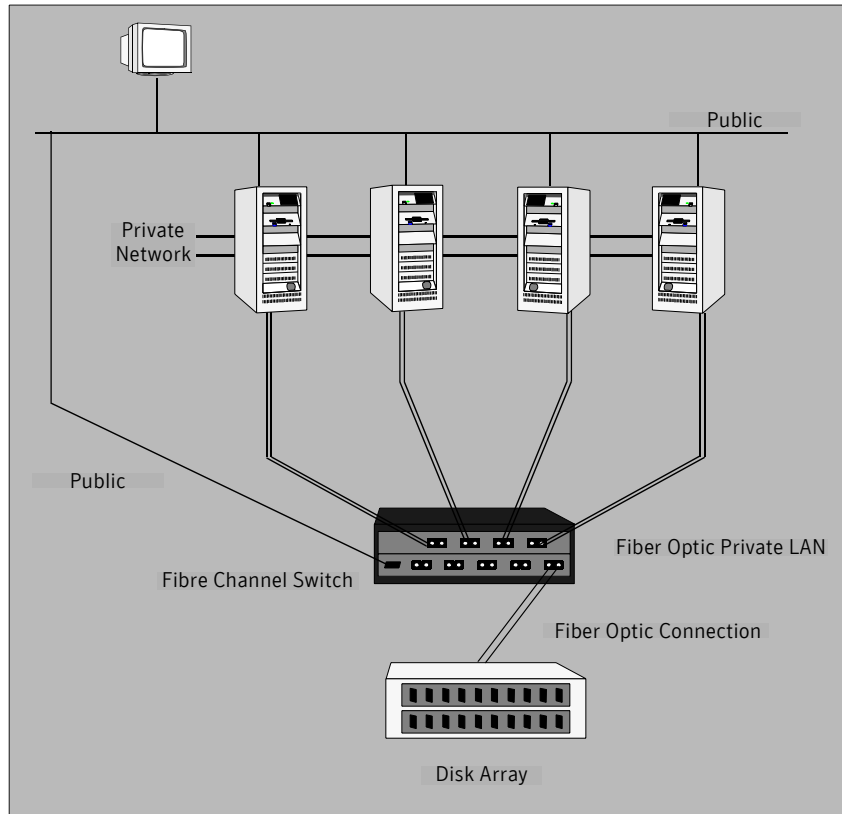
- [Hardware overview](#)
- [Software components](#)
- [Preinstallation](#)
- [Prerequisites](#)
- [Installing the product](#)
- [Using the log files](#)
- [Verifying the configuration files](#)
- [Verifying agent configuration](#)
- [Synchronizing time on Cluster File Systems](#)
- [Configuring VCS](#)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

Install the HP-UX 11i 64-bit operating system with the September 2004 HP-UX 11i Version 2.0 or later on each node and install a Fibre Channel host bus adapter to allow connection to the Fibre Channel switch.

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 packages. 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.

Packages installed with SFCFS only

The software packages 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 package name and contents for each package:

Package	Contents
VRTSperl	Veritas Perl 5.8.8 Redistribution
VRTSvlic	Veritas Licensing
VRTSicisco	Symantec Common Infrastructure
VRTSspbx	Symantec Private Branch Exchange
VRTSsmf	Symantec Service Management Framework
VRTSsat	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
VRTSl1t	Veritas Low Latency Transport
VRTSgab	Veritas Group Membership and Atomic Broadcast
VRTSvxfen	Veritas I/O Fencing

Package	Contents
VRTSvcs	Veritas Cluster Server
VRTSaclib	Veritas ACC Library
VRTSvcsag	Veritas Cluster Server Bundled Agents
VRTSvcsmsg	Veritas Cluster Server Message Catalogs
VRTSjre	Veritas Java Runtime Environment Redistribution
VRTSjre15	Veritas Java Runtime Environment Redistribution
VRTScutil	Veritas Cluster Utilities
VRTSweb	Veritas Java Web Server
VRTScscw	Veritas Cluster Server Configuration Wizards
VRTSvxvm	Veritas Volume Manager Binaries
VRTSdsa	Veritas Datacenter Storage Agent
VRTSfspro	Veritas File System Management Services Provider
VRTSvmpro	Veritas Volume Manager Management Services Provider
VRTSdcli	Veritas Distributed Command Line Interface
VRTSalloc	Veritas Volume Manager Intelligent Storage Provisioning
VRTSvdid	Veritas Device Identification API
VRTSvsvc	Veritas Volume Server and Client Provider
VRTSddlpr	Veritas Device Discovery Layer Services Provider
VRTSvrpro	Veritas Volume Replicator Client Extension and Provider for Veritas Enterprise Administrator
VRTSvcsvr	Veritas Cluster Server Agents for VVR
VRTSvrw	Veritas Volume Replicator Web Console
VRTSvxfs	Veritas File System
VRTSfsman	Veritas File System Manual Pages
VRTSfssdk	Veritas File System Software Developer Kit
VRTSg1m	Veritas Group Lock Manager
VRTScavf	Veritas Cluster Server Agents for Cluster File System

Optional packages for SFCFS and SFCFS HA

Packages	Contents
VRTScfsdc	Veritas Cluster File System Documentation
VRTScmccc	Veritas Cluster Management Console Cluster Connector
VRTScmcs	Veritas Cluster Management Console (Single Cluster Mode)
VRTScscm	Veritas Cluster Server Cluster Manager
VRTScssim	Veritas Cluster Server Simulator
VRTSfdoc	Veritas File System Documentation
VRTSvcscdc	Veritas Cluster Server Documentation
VRTSvcsmn	Veritas Cluster Server Man Pages
VRTSvmdoc	Veritas Volume Manager Documentation
VRTSvrdoc	Veritas Volume Replicator Documentation
VRTSfsmnd	Veritas File System SDK Manuals

Note: SFCFS 5.0 operates only on HP-UX 11i 64-bit operating system with the September 2004 HP-UX 11i Version 2.0 or later. All cluster nodes must be running this OS version.

For cluster file system functionality to work reliably on HP-UX, you must have the required HP-UX patches installed.
See [“Required HP-UX patches”](#) on page 15.

Required HP-UX patches

HP-UX required patches include the following:

HP-UX Patch ID	Description
PHCO_32385	Enables <code>fscat(1M)</code> .
PHCO_32387	Enables <code>getext(1M)</code> .
PHCO_32388	Enables <code>setext(1M)</code> .
PHCO_32389	Enables <code>vxdump(1M)</code> .
PHCO_32390	Enables <code>vxrestore(1M)</code> .
PHCO_32391	Enables <code>vxfsstat(1M)</code> .
PHCO_32392	Enables <code>vxtunefs(1M)</code> .
PHCO_32393	Enables <code>vxupgrade(1M)</code> .
PHCO_32488	Enables LIBC for VxFS 4.1 and later file systems.
PHCO_32523	Enhancement to <code>quota(1)</code> for supporting large uids.
PHCO_32524	Enhancement to <code>edquota</code> for supporting large uids.
PHCO_32551	Enhancement to <code>quotaon/quotaoff</code> for supporting large uids.
PHCO_32552	Enhancement to <code>repquota</code> for supporting large uids.
PHCO_32596	Enables <code>df(1M)</code> .
PHCO_32608	Enables <code>bdf(1M)</code> .
PHCO_32609	Enables <code>fstyp(1M)</code> .
PHCO_32610	Enables <code>mount(1M)</code> .
PHCO_32611	Fix <code>fs_wrapper</code> to accept vxfs from subtype.
PHCO_33238	<code>swapon(1M)</code> cumulative patch.
PHCO_34036	LVM commands patch.
PHCO_34208	SAM cumulative patch.
PHCO_34191	Cumulative <code>libc</code> patch.
PHSS_32674	Obam patch (backend for the SAM patch).
PHKL_31500	Sept04 Base Patch
PHKL_32272	Changes to fix intermittent failures in <code>getacl/setacl</code> .

HP-UX Patch ID	Description
PHKL_32430	Changes to separate vxfs symbols from libdebug.a, so that symbols of VxFS 4.1 and later are easily available in q4/p4.
PHKL_32431	Changes to disallow mounting of a file system on a vnode having VNOMOUNT set. Enhancements for supporting quotas on large uids.
PHKL_33312	LVM Cumulative Patch.
PHKL_34010	Cumulative VM Patch.

In addition to the above patches the `EnableVXFS` bundle needs to be installed before installing the SFCFS 5.0. This bundle is a HP bundle and contains enhancements to various commands to understand the new disk layout Version 6 and later. The `EnableVXFS` bundle contains the following patches:

Patch ID	Description
FSLibEnh	Enhancement to LIBC libraries to understand VxFS disk layout Version 6 and later.
DiskQuota-Enh	Enhancements to various quota related commands to support large uids.
FSCmdsEnh	Enhancements to the <code>mount</code> command to support VxFS 5.0.

Note: Install all the latest required HP-UX patches *before* you install SFCFS. You can use the `swlist` command to determine whether the correct update and patches are installed. The installation procedure terminates if the correct patches are not found. Make sure that `EnableVXFS` bundle has revision B.11.23.04 or later after installing the latest patches.

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/281875>

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.

Also, you can get the patches from Hewlett-Packard's Patch Database offered under the Maintenance and Support section of the HP Services & Support - IT Resource Center. HP's Patch Database provides fast, accurate searches for the latest recommended and superseded patches available for Veritas File System or Veritas Volume Manager.

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.

To prevent conflicts with VxFS manual pages previously installed with JFS/OnLineJFS 3.5, the VxFS 5.0 manual pages are installed in the `/opt/VRTS/vxfs5.0/man` directory. The `/opt/VRTS/vxfs5.0/man` directory is automatically added to `/etc/MANPATH` when the VxFS 5.0 package is installed. Make sure that the `/opt/VRTS/man` directory or the `/opt/VRTS/vxfs5.0/man` directory goes before `/usr/share/man` in your `MANPATH` environment variable so that the latest version of the VxFS manual pages display.

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)
This is the client package for HP-UX.
- 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
HP-UX	512MB of memory
Windows XP, NT, ME, 2000, or 98	300MHz Pentium with at least 256MB of memory

Alternatively, remove the file `/sbin/init.d/vmsa-server`.

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 `installsfdfs` 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 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.
- 4 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom  
# /usr/sbin/mount -F cd9660 /dev/dsk/c3t2d0 /dvdrom
```
- 5 Change directory to `/dvdrom`:

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

```
Do you wish to continue with the installation? [y,n,q] (y)
```
- 9 Enter one or more system names to install SFCFS.

```
Enter the system names separated by spaces on which to install  
SFCFS: system01 system02
```
- 10 Enter the license key for system01.

```
Enter a SFCFS license key for system01?
```
- 11 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)
```
- 12 Enter `1`, `2` or `3` to be installed on all systems.

```
Select the packages to be installed on all systems?  
[1-3,q,?] (2)
```
- 13 Press **Return** to continue.

```
Press [Return] to continue:
```
- 14 Reboot all the nodes on which SFCFS is installed and proceed to [“Configuring the Components.”](#)

Configuring the Components

This sections describes the configuration of SFCFS components.

To configure the components

- 1 Log in as superuser.
- 2 Run the `installer` command to install the SFCFS. For example:

```
# cd /cdrom
# ./installer
```
- 3 From the Installation menu, choose the **C** option for Configuration and select **6** the Veritas Storage Foundation Cluster File System. You are prompted to enter one or more system names.

```
Enter the system names separated by spaces on which to
configure SFCFS: system01 system02
```
- 4 Enter one or more system names to configure SFCFS. You are prompted to enter Return to continue.

```
Press [Return] to continue:
```
- 5 Press Return to continue. You are prompted to configure I/O fencing in enabled mode.

```
Will you be configuring I/O Fencing in enabled mode? [y,n,q]
(y)
```
- 6 Enter **y** or **n** to configure I/O fencing in enable mode. You are prompted to configure VCS for SFCFS.

```
Enter the unique cluster name: [?]
Enter the unique Cluster ID number between 0-65535: [b,?]
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
hpslia05: [b,?] (lan0)
Are you sure you want to use lan0 for the second private
heartbeak link? [y,n,q,b,?] (n)
Do you want to configure an additional low priority
heartbeat link? [y,n,q,b,?] (n)
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)
```

- 7 Answer the prompts to configure VCS for SFCFS. You are prompted to configure SFCFS to use Veritas Security Services.
Would you like to configure SFCFS to use Veritas Security Services? [y,n,q] (n)
- 8 Enter **y** or **n** to configure SFCFS to use Veritas Security Services. You are prompted to add Storage Foundation Cluster File System users.
Do you want to set the username and/or password for the Admin user (default username = 'admin', password='password')? [y,n,q] (n)
- 9 Enter **n** if you want to decline. If you enter **y**, you are prompted to change the password. You are prompted to add another user to the cluster.
Do you want to add another user to the cluster? [y,n,q] (y)
- 10 Enter **n** if you want to decline, enter **y** if you want to add another user. You are prompted to verify the Storage Foundation user.
Is this information correct? [y,n,q] (y)
- 11 Enter **y** or **n** to verify if this information is correct. You are prompted to configure the cluster monitor.
Do you want to configure the Cluster Monitor [y,n,q] (n)
- 12 Enter **y** or **n** to configure the cluster monitor. You are prompted to configure SMTP notification.
Do you want to configure SMTP notification? [y,n,q] (y)
- 13 Enter **y** or **n** to configure SMTP notification. You are prompted to configure SMTP notification.
Do you want to configure SNMP notification? [y,n,q] (y)
- 14 Enter **y** or **n** to configure SNMP notification. You are prompted to set up the enclosure-based naming scheme.
Do you want to set up the enclosure-based naming scheme? [y,n,q,?] (n)
- 15 Enter **y** or **n** to set up the enclosure-based naming scheme. You are prompted to set up a VxVM default disk group for each system.
Do you want to set up a default disk group for each system? [y,n,q,?] (y)

Will you specify one disk group name for all eligible systems? [y,n,q,?] (y)

Specify a default disk group for all systems. [?] sharedg
.
.
.
Is this correct? [y,n,q] (y)

- 16 Enter **y** or **n** if the VxVM default disk group information is correct. You are prompted to enable centralized management.
Enable Centralized Management? [y,n,q] (y) n
- 17 Enter **y** or **n** to enable centralized management. You are prompted to verify the fully qualified domain name for system01.
Is the fully qualified host name system01.domain_name?
[y,n,q]
(y)
- 18 Enter **y** or **n** to verify the fully qualified domain name for system01. You are prompted to start SFCFS process now.
Do you want to start Veritas Storage Foundation for Cluster File System process now? [y,n,q] (y)

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.

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/llhosts`

The file `llhosts(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/llhosts` 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 24.

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](#)
- [Upgrade Overview](#)
- [Upgrading from 3.5 to 5.0](#)
- [Upgrading from 4.1 to 5.0](#)
- [Upgrading the disk layout versions](#)

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 that the prerequisite patches required for SFCFS 5.0 are accessible.

Upgrade paths

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

From	Upgrade to	Tasks
Storage Foundation Cluster File System 3.5 (formerly known as, SANPoint Foundation Suite 3.5)	Storage Foundation Cluster File System 5.0	Proceed to “ Upgrading from 3.5 to 5.0 ” on page 39.
Storage Foundation Cluster File System 3.5 Update 1 (formerly known as, SANPoint Foundation Suite 3.5 Update 1)	Storage Foundation Cluster File System 5.0	Proceed to “ Upgrading from 3.5 to 5.0 ” on page 39.
Storage Foundation Cluster File System 3.5 Update 2 (formerly known as, SANPoint Foundation Suite 3.5 Update 2)	Storage Foundation Cluster File System 5.0	Proceed to “ Upgrading from 3.5 to 5.0 ” on page 39.

From	Upgrade to	Tasks
Storage Foundation Cluster File System 3.5 Update 3 (formerly known as, SANPoint Foundation Suite 3.5 Update 3)	Storage Foundation Cluster File System 5.0	Proceed to “Upgrading from 3.5 to 5.0” on page 39.
Storage Foundation Cluster File System 4.1	Storage Foundation Cluster File System 5.0	Proceed to “Upgrading from 4.1 to 5.0” on page 45.
Storage Foundation Cluster File System 4.1 MP1	Storage Foundation Cluster File System 5.0	Proceed to “Upgrading from 4.1 to 5.0” on page 45.

Upgrade Overview

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

Phased upgrade

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. Each phase of the phased upgrade should be performed on more than one node of the cluster. Cluster will be offline only during [step 4](#) and [step 5](#) below for a phased upgrade.

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

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.
- 4 Modify the configuration information in the main.cf file.
- 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 Install the new version on all the nodes.
- 2 Modify the configuration information in the main.cf file.
- 3 Bring up the cluster.

Upgrading from 3.5 to 5.0

SFCFS can be upgraded from 3.5 to 5.0 using phased or full upgrade procedure.

Phased upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where system01 and system02 are initially upgraded and rest of the cluster is brought up later.

To upgrade from 3.5 to 5.0

- 1 Log in as superuser.
- 2 Select one or more nodes to upgrade, say system01 and system02.
- 3 Insert the appropriate software disc into your system's DVD drive.
- 4 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.

- 5 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom  
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```

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

```
# cd /dvdrom
```

- 7 Offline all SFCFS resources on nodes selected in step 2 by running the following commands on one of the cluster nodes.

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

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

- 8 Remove the VRTScavf and VRTSglm 3.5 packages from these nodes and run the following commands on system01 and system02.

```
# hstop -local  
# swremove VRTScavf VRTSglm
```

- 9 Uninstall VCS 3.5 from system01 and system02. Run the following commands from one of the nodes.

See the *Veritas Cluster Server Installation Guide*.

```
# cd /opt/VRTSvcs/install
# ./uninstallvcs
```

Note: Ignore any errors from the `uninstallvcs` script and proceed with the uninstall of VCS. Also run the following command from rest of the nodes in the cluster on which uninstallation is not performed using `./uninstallvcs`:

```
# hstart
```

- 10 Upgrade the operating system from HP-UX 11i Version 1 to HP-UX 11i Version 2 on system01 and system02.

See the *HP-UX Operating System* documentation.

- 11 Install all the prerequisite patches on system01 and system02. See “[Required HP-UX patches](#)” on page 15.

- 12 Install SFCFS 5.0 on system01 and system02 and reboot these nodes. See “[Installing the product](#)” on page 19.

Note: Do not configure SFCFS after reboot.

- 13 Shutdown VCS on remaining non-upgraded nodes.

- a Run the following commands from one of the non-upgraded nodes.

```
# haconf -makerw
# hagr -unfreeze cvm -persistent
# hagr -unfreeze service_group -persistent
# haconf -dump -makero
```

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

- b Run the following commands on all the non-upgraded nodes.

```
# hstop -local
```

- 14 Start `vxfen` on system01 and system02. `vxfen` can be started either in disable or enable mode. For starting `vxfen` in disabled mode, run the following commands

```
# echo vxfen_mode=disabled > /etc/vxfenmode
# /sbin/init.d/vxfen start
```

See the *Veritas Cluster Server Installation Guide* for information regarding starting `vxfen` in enabled mode.

- 15 Change the configuration files by running the following commands on one of the upgraded nodes, say system01.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# hagr -unfreeze cvm -persistent
# hagr -unfreeze service_group -persistent
# /opt/VRTS/bin/hatype -add CVMVxconfigd
# /opt/VRTS/bin/hares -add cvm_vxconfigd CVMVxconfigd cvm
# /opt/VRTS/bin/hares -modify cvm_vxconfigd Enabled 1
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

If you have configured the VCS Cluster Manager (Web Console), complete [step a](#) through [step d](#) to modify the

`/etc/VRTSvcs/conf/config/main.cf` file.

- a Remove VRTSweb:

```
Process VRTSweb (
    PathName = "/opt/VRTSvcs/bin/haweb"
    Arguments = "10.129.96.64 8181"
)
```

- b Replace it with:

```
VRTSWebApp VCSweb (
    Critical = 0
    AppName = vcs
    InstallDir = "/opt/VRTSweb/VERITAS"
    TimeForOnline = 5
)
```

- c Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named `csgnic` and the public NIC device is `hme0`, add:

```
NIC csgnic (
    Device = hme0
```

- d Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMngr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
ntfr requires csgnic
webip requires csgnic
```

- 16 Verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file by running the following commands on system01:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

- 17 Start VCS on all the upgraded nodes. Run the following command on system01 and system02.

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

- 18 Configure SFCFS on system01 and system02. See “[Using the log files](#)”.

Note: VCS configuration files are not changed during this configuration.

- 19 Upgrade file systems to proper disk layout version as mentioned in “[Upgrading the disk layout versions](#)” on page 50.
- 20 Verify that all the file systems are working properly and data is intact. See `cfsmount(1M)`.
- 21 Run [step 8](#) through [step 12](#), [step 14](#), [step 17](#), [step 18](#) and [step 20](#) on rest of the nodes to be upgraded.

Full upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where all nodes are simultaneously upgraded from 3.5 to 5.0.

- 1 Log in as superuser.
- 2 Insert the appropriate software disc into your system's DVD drive.
- 3 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.
- 4 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom  
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```
- 5 Change to the top-level directory on the disc:

```
# cd /dvdrom
```
- 6 Offline all SFCFS resources on all nodes by running the following commands on one of the cluster nodes.

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

where *service_group* is the command that displays the service groups by `hagr -dep cvm`.
- 7 Remove the VRTScavf and VRTSglm 3.5 packages from these nodes and run the following command on all the systems.

```
# hst -local  
# swremove VRTScavf VRTSglm
```


- 8 Uninstall VCS 3.5 from all the nodes. Run the following commands from one of the nodes.

See the *Veritas Cluster Server Installation Guide*.

```
# cd /opt/VRTSvcs/install
# ./uninstallvcs
```

Note: Ignore any errors from the `uninstallvcs` script and proceed with the uninstall of VCS.

- 9 Upgrade the operating system from HP-UX 11i Version 1 to HP-UX 11i Version 2 on all the nodes.
See the *HP-UX Operating System* documentation.
- 10 Install all the prerequisite patches on all the nodes. See section “[Required HP-UX patches](#)” on page 15.
- 11 Install SFCFS 5.0 and reboot all the nodes.
See “[Installing the product](#)” on page 19.

Note: Do not configure SFCFS after reboot.

- 12 Start `vxfen` on all the nodes. `vxfen` can be started either in disable or enable mode. For starting `vxfen` in disabled mode, run the following commands:

```
# cat vxfen_mode=disabled > /etc/vxfenmode
# /sbin/init.d/vxfen start
```

See the *Veritas Cluster Server Installation Guide* for information regarding starting `vxfen` in enabled mode.

- 13 Change the configuration files by running the following commands from one of the nodes.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# /opt/VRTS/bin/hatype -add CVMVxconfigd
# /opt/VRTS/bin/hares -add cvm_vxconfigd CVMVxconfigd cvm
# /opt/VRTS/bin/hares -modify cvm_vxconfigd Enabled 1
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

If you have configured the VCS Cluster Manager (Web Console), complete [step a](#) through [step d](#) to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

- a Remove VRTSweb:

```
Process VRTSweb (
    PathName = "/opt/VRTSvcs/bin/haweb"
    Arguments = "10.129.96.64 8181"
)
```

- b Replace it with:

```
VRTSWebApp VCSweb (
    Critical = 0
    AppName = vcs
    InstallDir = "/opt/VRTSweb/VERITAS"
    TimeForOnline = 5
)
```

- c Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named `csgnic` and the public NIC device is `hme0`, add:

```
NIC csgnic (
    Device = hme0
```

- d Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMngr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
ntfr requires csgnic
webip requires csgnic
```

- 14 Verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file by running the following command on `system01`:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

- 15 Run the following command on all the nodes to start VCS.

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

- 16 Configure SFCFS on system01 and system02. See “[Using the log files](#)”.

Note: VCS configuration files are not changed during this configuration.

- 17 Upgrade file systems to proper disk layout version as mentioned in “[Upgrading the disk layout versions](#)” on page 50.
- 18 Verify that all the file systems are working properly and data is intact. See `cfsmount(1M)`.

Upgrading from 4.1 to 5.0

SFCFS can be upgraded from 4.1 to 5.0 using phased or full upgrade procedure.

Phased upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where system01 and system02 are initially upgraded and rest of the cluster is brought up later.

To upgrade from 4.1 to 5.0

- 1 Log in as superuser.
- 2 Select one or more nodes to upgrade, say system01 and system02.
- 3 Insert the appropriate software disc into your system's DVD drive.
- 4 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.
- 5 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom  
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```

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

```
# cd /dvdrom
```
- 7 Install all the prerequisite patches on system01 and system02. See “[Required HP-UX patches](#)” on page 15.

If this step is being performed in any phase other than the first phase of the upgrade, the `llhosts` need to be replaced to prevent `vxfsen` from starting after reboot by running the following commands.

- a Move `/etc/llhosts` to `/etc/llhosts.bak` on all the nodes to be upgraded.

```
# mv /etc/llthosts /etc/llthosts.bak
```

- b Install all the prerequisite patches and reboot the machines.
- c Move /etc/llthosts to /etc/llthosts.bak on all the nodes to be upgraded.

```
# mv /etc/llthosts.bak /etc/llthosts
```

- 8 Offline all SFCFS resources on nodes selected in step 2 by running the following commands on one of the cluster nodes.

```
# hagrpf -offline service_group -sys system01
```

```
# hagrpf -offline service_group -sys system02
```

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

- 9 Install SFCFS 5.0 on *system01* and *system02* and reboot these nodes. See “[Installing the product](#)” on page 19.

Note: Do not configure SFCFS after reboot.

- 10 Shutdown VCS on remaining non-upgraded nodes.

- a Run the following commands from one of the non-upgraded nodes.

```
# haconf -makerw
```

```
# hagrpf -unfreeze cvm -persistent
```

```
# hagrpf -unfreeze service_group -persistent
```

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

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

- b Run the following commands on all the non-upgraded nodes.

```
# hastop -local
```

```
# vxfenconfig -U
```

- 11 Start `vxfen` on *system01* and *system02*. `vxfen` can be started either in disable or enable mode. For starting `vxfen` in disabled mode, run the following commands:

```
# echo vxfen_mode=disabled > /etc/vxfenmode
```

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

See the *Veritas Cluster Server Installation Guide* for information regarding starting `vxfen` in enabled mode.

- 12 Change the configuration files by running the following commands on one of the upgraded nodes. For example, system01.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# hagr -unfreeze cvm -persistent
# hagr -unfreeze service_group -persistent
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

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

- 13 If you have configured the VCS Cluster Manager (Web Console), complete step a through step d to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

- a Remove VRTSweb:

```
Process VRTSweb (
    PathName = "/opt/VRTSvcs/bin/haweb"
    Arguments = "10.129.96.64 8181"
)
```

- b Replace it with:

```
VRTSWebApp VCSweb (
    Critical =0
    AppName = vcs
    InstallDir = "/opt/VRTSweb/VERITAS"
    TimeForOnline = 5
)
```

- c Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named `csgnic` and the public NIC device is `hme0`, add:

```
NIC csgnic (
    Device = hme0
```

- d Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMngr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
ntfr requires csgnic
webip requires csgnic
```

- 14 Verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file by running the following commands on system01:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

- 15 Start VCS on all the upgraded nodes. Run the following command on system01 and system02.

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

- 16 Configure SFCFS on system01 and system02. See “[Using the log files](#)”.

Note: VCS configuration files are not changed during this configuration.

- 17 Upgrade file systems to proper disk layout version as mentioned in “[Upgrading the disk layout versions](#)” on page 50.
- 18 Verify that all the file systems are working properly and data is intact. See `cfsmount(1M)`.
- 19 Run [step 7](#), [step 9](#), and [step 16](#) on rest of the nodes to be upgraded.

Full upgrade

Following procedure assumes a 4 node cluster system01, system02, system03, system04 where all nodes are simultaneously upgraded from 4.1 to 5.0.

- 1 Log in as superuser.
- 2 Insert the appropriate software disc into your system's DVD drive.
- 3 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.
- 4 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom  
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```
- 5 Change to the top-level directory on the disc:

```
# cd /dvdrom
```
- 6 Install all the prerequisite patches on all the nodes. See “[Required HP-UX patches](#)” on page 15.
- 7 Offline all SFCFS resources on all nodes by running the following commands on one of the cluster nodes.

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

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

- 8 Install SFCFS 5.0 and reboot all the nodes.
 See “[Installing the product](#)” on page 19.

Note: Do not configure SFCFS after reboot.

- 9 Start `vxfen` on all the nodes. `vxfen` can be started either in disable or enable mode. For starting `vxfen` in disabled mode, run the following commands:

```
# cat vxfen_mode=disabled > /etc/vxfenmode
# /sbin/init.d/vxfen start
```

See the *Veritas Cluster Server Installation Guide* for information regarding starting `vxfen` in enabled mode.

- 10 Change the configuration files by running the following commands from one of the nodes.

```
# /opt/VRTS/bin/hastart
# /opt/VRTS/bin/haconf -makerw
# /opt/VRTS/bin/hares -delete qlogckd
# /opt/VRTS/bin/haconf -dump -makero
# /opt/VRTS/bin/hastop -all -force
```

- 11 If you have configured the VCS Cluster Manager (Web Console), complete [step a](#) through [step d](#) to modify the `/etc/VRTSvcs/conf/config/main.cf` file.

- a Remove VRTSweb:

```
Process VRTSweb (
    PathName = "/opt/VRTSvcs/bin/haweb"
    Arguments = "10.129.96.64 8181"
)
```

- b Replace it with:

```
VRTSWebApp VCSweb (
    Critical = 0
    AppName = vcs
    InstallDir = "/opt/VRTSweb/VERITAS"
    TimeForOnline = 5
)
```

- c Add the NIC resource in the ClusterService group. For example, where the name of the NIC resource is named `csgnic` and the public NIC device is `hme0`, add:

```
NIC csgnic (
    Device = hme0
```

- d Add new dependencies for the new resources in the ClusterService group. For example, using the names of the VRTSWebApp, NotifierMgr, IP, and NIC resources, enter lines that resemble:

```
VCSweb requires webip
ntfr requires csgnic
webip requires csgnic
```

- 12 Verify the syntax of the `/etc/VRTSvcs/conf/config/main.cf` file by running the following commands on system01:

```
# cd /etc/VRTSvcs/conf/config
# /opt/VRTS/bin/hacf -verify .
```

- 13 Run the following command on all the nodes to start VCS.

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

- 14 Configure SFCFS on all the nodes. See [“Using the log files”](#).

Note: VCS configuration files are not changed during this configuration.

- 15 Upgrade file systems to proper disk layout version as mentioned in [“Upgrading the disk layout versions”](#) on page 50.

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

Upgrading the disk layout versions

SFCFS 5.0 supports disk layouts Versions 4, 5, 6 and 7 for locally mounted file systems and disk layouts Versions 6 and 7 for cluster mounted file systems. If you have cluster mounted file systems with disk layout versions lower than 6, then after upgrading to SFCFS 5.0, perform the following additional steps to prepare the file system for being mounted on all nodes of the cluster:

- 1 Select one of the nodes of the cluster and `mount` the file system locally on this node. For example, mount it without the `-ocluster` option. Enter,

```
# mount -F vxfs block_device_path /mnt1
```

- 2 Current disk layout version on a file system can be found using

```
# fstyp -v char_device_path | grep version | \
awk '{print $2}'
```

- 3 On the node selected in [step 1](#), incrementally upgrade the disk layout of this file system to layout Version 6 or layout Version 7. For example, if you had a cluster mounted file system of disk layout Version 4 while running with SFCFS 3.5 on HP-UX 11i Version 1, after upgrading to SFCFS 5.0, you would need to upgrade the disk layout to version 6 or version 7 incrementally as follows:

```
# vxupgrade -n 5 /mnt1
# vxupgrade -n 6 /mnt1
# vxupgrade -n 7 /mnt1
```


- 4 On the node selected in [step 1](#), after the disk layout has been successfully upgraded, `umount` the file system.
`umount /mnt1`
- 5 This file system can be mounted on all nodes of the cluster using `cfsmount`.

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. Install all the prerequisite patches mentioned in “[Required HP-UX patches](#)” on page 15.

See “[Cluster platforms](#)” on page 11.

To add a node to a cluster

- 1 Log into the new system as superuser.
- 2 Determine the block device file for the DVD drive:

```
# ioscan -fnC disk
```

Make a note of the device file as it applies to your system.
- 3 Run the following commands to start PFS (Portable File System):

```
# nohup pfs_mountd &  
# nohup pfsd &
```
- 4 Create a directory in which to mount the software disc and mount the disc using the appropriate drive name. For example:

```
# mkdir -p /dvdrom  
# /usr/sbin/mount -F cdfs /dev/dsk/c3t2d0 /dvdrom
```
- 5 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 18.
- 6 Change to the SFCFS directory.

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

```
# ./installsfcfs -installonly
```
- 8 Enter **y** to install SFCFS on these systems.
Do you want to install SFCFS on these systems? [y,n,q] (y)
- 9 Enter the system name of the new node to install SFCFS.
Enter the system names separated by spaces on which to install SFCFS: **system03**
- 10 Enter a license key for system03.
Enter a SFCFS license key for system03:[?]
XXXX-XXXX-XXXX-XXXX-XXXX-X

- 11 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)
```

- 12 Enter **1** or **2** to be installed on all systems.

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

- 13 Press **Return** to continue.

```
Press [Return] to continue:
```

Note: Do not reboot the machine now.

- 14 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 25.

- 15 Create `/etc/gabtab` file the same as it is on another node in the cluster. See “[/etc/gabtab](#)” on page 28.

For example,

```
/sbin/gabconfig -c -n 3
```

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

- 16 Reboot the system that has been added.

```
# /usr/sbin/shutdown -R -y 0
```

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

```
# vxinstall
```

- 18 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 Start the VCS server and `vxfen` on `system03`.
 - a Use `hastart` on `system03` for starting the VCS server.
 - b For starting `vxfen` in the disable mode, run the following commands on `system03`:

```
# echo vxfen_mode=disabled > /etc/vxfenmode
# /sbin/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:

```
# /sbin/init.d/vxfen start
```
- 2 Check that there are no service groups dependent on CVM, such as SFCFS, that are still online:

```
# hagr -dep cvm
```
- 3 If there are any dependencies, take them offline, then take the CVM service group offline:

```
# hagr -offline cvm -sys system01
# hagr -offline cvm -sys system02
```
- 4 Open the VCS configuration for writing:

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

```
# hagr -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`.
- 6 Add the new node to the CVM AutoStartList:

```
# hagr -modify cvm AutoStartList system01 system02 system03
```
- 7 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
```
- 8 Write the new VCS configuration to disk:

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

- 9 Put the CVM resources back online, in the following order:

```
# hagr -online cvm -sys system01
# hagr -online cvm -sys system02
# hagr -online cvm -sys system03
```

- 10 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
```

- 11 Add shared disk groups to the cluster configuration:

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

- 12 Create a /mnt on system03 and run the following commands:

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

Refer to `cfsmntadm` man page for more details.

- 13 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 Stop all the cluster components:

```
# cfscluster stop -f system03
```

- 3 Open the VCS configuration for writing:

```
# haconf -makerw
```

- 4 Remove `system03` from the `systemList` 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 3](#) and the command that failed in [step 4](#).

- 5 Write the new VCS configuration to disk:

```
# haconf -dump -makero
```
- 6 Edit `/etc/llhosts` on the remaining nodes of the cluster, and remove the entry corresponding to the node being removed.
- 7 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.
- 8 Login to `system03` and remove the following files:

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

```
# rm /etc/vxfentab
# rm /etc/vxfendg
```
- 10 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 `-nN` option, where `N` is the number of cluster systems. If the command has the form `/sbin/gabconfig -c -nN`, where `N` is the number of cluster systems, then make sure that `N` is not greater than the actual number of nodes in the cluster, or GAB does not automatically seed.
Modify `/etc/llhosts` file on each remaining nodes to remove the entry of the leaving node.
- 11 Change to the `install` directory:

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


- 12 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)
```
- 6 Reboot the systems on which SFCFS is uninstalled after successful uninstillation

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.

Resource temporarily unavailable

If the installation fails with the following error message on the console:

```
fork() failed: Resource temporarily unavailable
```

The value of `nkthread` tunable parameter may not be large enough. The `nkthread` tunable requires a minimum value of 600 on all systems in the cluster. To determine the current value of `nkthread`, enter:

```
# kctune -q nkthread
```

If necessary, you can change the value of `nkthread` using the SAM (System Administration Manager) interface, or by running the `kctune` command. If you change the value of `nkthread`, the kernel must be rebuilt for the new value to take effect. It is easier to change the value using SAM because there is an option to process the new kernel immediately. See the `kctune(1M)` and `sam(1M)` manual pages for more information on tuning kernel parameters.

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.

If a system cannot access the software source depot, either `swagentd` is not running on the target system or the `swlist` command cannot see the source depot.

```
Correct /etc/{hosts, nsswitch.conf} and continue from here
Continue? [Y/N] :
```

Suggested solutions: check that `swagentd` is running. Check whether there is an entry for the target system in `/etc/hosts`. If there is no entry, then ensure the `hosts` file is not the primary lookup for the “hosts” entry.

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 `vxfscsd` daemon is not running. This typically happens only if the `CFSfscsd` agent was not started correctly.
- If `mount` fails with an error message:

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

```
vxfscsd 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 -F 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 18.
- 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.