

Veritas Storage Foundation™ Cluster File System Administrator's Guide

HP-UX 11i v3

5.1 Service Pack 1



Veritas Storage Foundation™ Cluster File System Administrator's Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Product version: 5.1 SP1

Document version: 5.1SP1.0

Legal Notice

Copyright © 2011 Symantec Corporation. All rights reserved.

Symantec, the Symantec logo, Veritas, Veritas Storage Foundation, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate 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.

THE 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 as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043
<http://www.symantec.com>

Technical Support

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

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

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

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

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

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

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

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

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

- Product release level

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

Licensing and registration

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

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

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

Support agreement resources

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

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

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec Web site.

<https://sort.symantec.com/documents>

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

docs@symantec.com

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Technical Support	4	
Chapter 1	Technical overview of Storage Foundation Cluster File System	13
	Storage Foundation Cluster File System architecture	13
	About the symmetric architecture	13
	About Storage Foundation Cluster File System primary/secondary failover	14
	About single-host file system semantics using Group Lock Manager	14
	About Veritas File System features supported in cluster file systems	14
	Veritas File System features in cluster file systems	15
	Veritas File System features not in cluster file systems	16
	Storage Foundation Cluster File System benefits and applications	17
	How Storage Foundation Cluster File System works	17
	When to use Storage Foundation Cluster File System	18
Chapter 2	Storage Foundation Cluster File System architecture	21
	About Veritas Cluster Server architecture	22
	About Storage Foundation Cluster File System and the Group Lock Manager	23
	Storage Foundation Cluster File System namespace	23
	About asymmetric mounts	24
	Primary and secondary	25
	Determining or moving primaryship	25
	Synchronize time on Cluster File Systems	25
	File system tunables	26
	Setting the number of parallel fsck threads	26
	About Storage Checkpoints	26
	Storage Foundation Cluster File System backup strategies	27
	Parallel I/O	29
	I/O error handling policy	29

Recovering for I/O failures	30
Single network link and reliability	30
Configuring low priority a link	30
Split-brain and jeopardy handling	31
Jeopardy state	31
Jeopardy handling	32
About I/O fencing	32
About I/O fencing for Storage Foundation Cluster File System	
HA in virtual machines that do not support SCSI-3 PR	32
About preventing data corruption with I/O fencing	33
About I/O fencing components	34
About I/O fencing configuration files	38
How disk-based I/O fencing works in different event	
scenarios	40
About server-based I/O fencing	44
About Veritas Volume Manager cluster functionality	47
Shared disk groups overview	48
Storage Foundation Cluster File System and Veritas Volume Manager	
cluster functionality agents	53
Veritas Volume Manager cluster functionality	54

Chapter 3

Administering Storage Foundation Cluster File	
System and its components	55
About Storage Foundation Cluster File System administration	55
Administering CFS	56
Using cfsmount to mount CFS file systems	56
Resizing CFS file systems	57
Verifying the status of CFS file systems	57
Verifying CFS port	57
CFS agents and AMF support	58
CFS agent log files	58
Storage Foundation Cluster File System commands	
mount, fsclustadm, and fsadm commands	59
Time synchronization for Cluster File Systems	60
Growing a Storage Foundation Cluster File System	60
The /etc/fstab file	61
When the Storage Foundation Cluster File System primary	
fails	61
Storage Checkpoints on Storage Foundation Cluster File	
System	61
Snapshots on Storage Foundation Cluster File System	61
Administering VCS	63

Configuring VCS to start Oracle with a specified Pfile	64
Verifying VCS configuration	64
Starting and stopping VCS	64
Configuring destination-based load balancing for LLT	64
Administering CVM	65
Listing all the CVM shared disks	65
Establishing CVM cluster membership manually	65
Changing the CVM master manually	65
Importing a shared disk group manually	68
Deporting a shared disk group manually	68
Starting shared volumes manually	68
Evaluating the state of CVM ports	69
Verifying if CVM is running in an SFCFS cluster	69
Verifying CVM membership state	69
Verifying the state of CVM shared disk groups	70
Verifying the activation mode	70
CVM log files	70
Administering ODM	71
Verifying the ODM port	71
Starting ODM	71
Stopping ODM	71
Administering I/O Fencing	72
About administering I/O fencing	72
About the vxfcntlshdw utility	73
About the vxfenadm utility	80
About the vxfcntlpre utility	85
About the vxfcntlswap utility	88
About administering the coordination point server	97
Enabling or disabling the preferred fencing policy	118
About VXFEN tunable parameters	120
Administering SFCFS global clusters	122
Setting up a disaster recovery fire drill	122
Configuring the fire drill service group using the Fire Drill Setup wizard	123
Verifying a successful fire drill	125
Scheduling a fire drill	125

Chapter 4

Using Veritas Extension for Oracle Disk

Manager	127
About Oracle Disk Manager	127
How Oracle Disk Manager improves database performance	129

About Oracle Disk Manager and Storage Foundation Cluster File System	130
About Oracle Disk Manager and Oracle Managed Files	131
How Oracle Disk Manager works with Oracle Managed Files	131
Setting up Veritas Extension for Oracle Disk Manager	133
Preparing existing database storage for Oracle Disk Manager	134
Converting Quick I/O files to Oracle Disk Manager files	134
Verifying that Oracle Disk Manager is configured	135
Disabling the Oracle Disk Manager feature	137
Using Cached ODM	138
Enabling Cached ODM for file systems	139
Modifying Cached ODM settings for individual files	139
Adding Cached ODM settings via the cachemap	140
Making the caching settings persistent across mounts	141
Chapter 5 Clustered NFS	143
About Clustered NFS	143
Requirements	143
Understanding how Clustered NFS works	144
Basic design	144
Internal Clustered NFS functionality	144
cfsshare manual page	146
Configure and unconfigure Clustered NFS	146
Configure Clustered NFS	146
Unconfigure Clustered NFS	148
Administering Clustered NFS	148
Displaying the NFS shared CFS file systems	148
Sharing a CFS file system previously added to VCS	148
Unsharing the previous shared CFS file system	149
Adding an NFS shared CFS file system to VCS	150
Deleting the NFS shared CFS file system from VCS	150
Adding a Virtual IP address to VCS	150
Deleting a Virtual IP address from VCS	151
Adding an IPv6 Virtual IP address to VCS	151
Deleting an IPv6 Virtual IP address from VCS	151
Changing the share options associated with an NFS share	151
Sharing a file system checkpoint	152
Samples for configuring a Clustered NFS	152
Sample main.cf file	156
How to mount an NFS-exported file system on the NFS clients	160
Debugging Clustered NFS	160

Chapter 6	Common Internet File System	161
	About Common Internet File System	161
	Requirements	162
	Understanding how Samba works	162
	Configuring Clustered NFS and Common Internet File System on CFS	162
	cfsshare manual page	162
	Configuring Common Internet File System in user mode	163
	Configuring Common Internet File System in domain mode	164
	Configuring Common Internet File System in ads mode	166
	Administering Common Internet File System	168
	Sharing a CFS file system previously added to VCS	170
	Unsharing the previous shared CFS file system	171
	Sample main.cf	171
	Debugging Common Internet File System	175
Chapter 7	Troubleshooting SFCFS	177
	About troubleshooting SFCFS	177
	Troubleshooting fenced configurations	177
	Example of a pre-existing network partition (split-brain)	178
	Recovering from a pre-existing network partition (split-brain)	178
	Troubleshooting I/O fencing	180
	SCSI reservation errors during bootup	180
	The vxfsentsthdw utility fails when SCSI TEST UNIT READY command fails	180
	Node is unable to join cluster while another node is being ejected	181
	System panics to prevent potential data corruption	181
	How vxfsen driver checks for preexisting split-brain condition	181
	Cluster ID on the I/O fencing key of coordinator disk does not match the local cluster's ID	183
	Clearing keys after split-brain using vxfsenclearpre command	184
	Registered keys are lost on the coordinator disks	184
	Replacing defective disks when the cluster is offline	184
	Troubleshooting server-based I/O fencing	187
	Troubleshooting server-based fencing on the SFCFS cluster nodes	188
	Troubleshooting server-based I/O fencing in mixed mode	192
	Troubleshooting CVM	196

	CVM group is not online after adding a node to the cluster	196
	Shared disk group cannot be imported	196
	Error importing shared disk groups	197
	Unable to start CVM	197
	CVMVolDg not online even though CVMCluster is online	198
Appendix A	Creating a starter database	199
	Creating a database for Oracle 10g or 11g	199
	Creating database tablespace on shared raw VxVM volumes (option 1)	199
	Creating database tablespace on CFS (option 2)	201
Glossary	203
Index	209

Technical overview of Storage Foundation Cluster File System

This chapter includes the following topics:

- [Storage Foundation Cluster File System architecture](#)
- [About Veritas File System features supported in cluster file systems](#)
- [Storage Foundation Cluster File System benefits and applications](#)

Storage Foundation Cluster File System architecture

The Veritas Storage Foundation Cluster File System (SFCFS) allows clustered servers to mount and use a file system simultaneously as if all applications using the file system were running on the same server. The Veritas Volume Manager cluster functionality (CVM) makes logical volumes and raw device applications accessible throughout a cluster.

This section includes the following topics:

- About the symmetric architecture
- About Storage Foundation Cluster File System primary/secondary failover
- About single-host file system semantics using Group Lock Manager

About the symmetric architecture

SFCFS uses a symmetric architecture in which all nodes in the cluster can simultaneously function as metadata servers. SFCFS still has some remnants of

the old master/slave or primary/secondary concept. The first server to mount each cluster file system becomes its primary; all other nodes in the cluster become secondaries. Applications access the user data in files directly from the server on which they are running. Each SFCFS node has its own intent log. File system operations, such as allocating or deleting files, can originate from any node in the cluster.

About Storage Foundation Cluster File System primary/secondary failover

If the server on which the SFCFS primary is running fails, the remaining cluster nodes elect a new primary. The new primary reads the intent log of the old primary and completes any metadata updates that were in process at the time of the failure.

If a server on which an SFCFS secondary is running fails, the primary reads the intent log of the failed secondary and completes any metadata updates that were in process at the time of the failure.

About single-host file system semantics using Group Lock Manager

SFCFS uses the Veritas Group Lock Manager (GLM) to reproduce UNIX single-host file system semantics in clusters. This is most important in write behavior. UNIX file systems make writes appear to be atomic. This means that when an application writes a stream of data to a file, any subsequent application that reads from the same area of the file retrieves the new data, even if it has been cached by the file system and not yet written to disk. Applications can never retrieve stale data, or partial results from a previous write.

To reproduce single-host write semantics, system caches must be kept coherent and each must instantly reflect any updates to cached data, regardless of the cluster node from which they originate. GLM locks a file so that no other node in the cluster can update it simultaneously, or read it before the update is complete.

About Veritas File System features supported in cluster file systems

The Veritas Storage Foundation Cluster File System is based on the Veritas File System (VxFS).

Most of the major features of VxFS local file systems are available on cluster file systems, including the following features:

- Extent-based space management that maps files up to a terabyte in size

- Fast recovery from system crashes using the intent log to track recent file system metadata updates
- Online administration that allows file systems to be extended and defragmented while they are in use

Every VxFS manual page has a section on Storage Foundation Cluster File System Issues with information on whether the command functions on a cluster-mounted file system and indicates any difference in behavior from local mounted file systems.

Veritas File System features in cluster file systems

[Table 1-1](#) describes the VxFS supported features and commands for SFCFS.

Table 1-1 Veritas File System features in cluster file systems

Features	Description
FileSnap	A FileSnap is a space-optimized copy of a file in the same name space, stored in the same file system. See the <i>Veritas Storage Foundation Advanced Features Administrator's Guide</i> .
Fast file system recovery	VxFS provides fast recovery of a file system from system failure. See the <i>Veritas File System Administrator's Guide</i> .
Quick I/O for databases	The Quick I/O for Databases feature, using clustered Oracle Disk Manager (ODM), is supported on SFCFS. See the <i>Veritas Storage Foundation Advanced Features Administrator's Guide</i> .
Storage Checkpoints	Storage Checkpoints are supported on cluster file systems, but are licensed only with other Veritas products.
Snapshots	Snapshots are supported on cluster file systems.
Quotas	Quotas are supported on cluster file systems.
NFS mounts	You export the NFS file systems from the Cluster. You can NFS export CFS file systems in a distributed highly available way.
Nested Mounts	You can use a directory on a cluster mounted file system as a mount point for a local file system or another cluster file system.
Freeze and thaw	Synchronizing operations, which require freezing and thawing file systems, are done on a cluster-wide basis.

Table 1-1 Veritas File System features in cluster file systems (*continued*)

Features	Description
Memory mapping	Shared memory mapping established by the <code>mmap()</code> function is supported on SFCFS. See the <code>mmap(2)</code> manual page.
Disk layout versions	Currently, only the Version 5, 6, 7, and 8 disk layouts can be created, and only the Version 4, 5, 6, 7, and 8 disk layouts can be mounted. Use the <code>fstyp -v special_device</code> command to ascertain the disk layout version of a VxFS file system. Use the <code>vxupgrade</code> command to update the older disk layout version.
Locking	Advisory file and record locking are supported on SFCFS. For the <code>F_GETLK</code> command, if there is a process holding a conflicting lock, the <code>l_pid</code> field returns the process ID of the process holding the conflicting lock. The nodeid-to-node name translation can be done by examining the <code>/etc/llthosts</code> file or with the <code>fsclustadm</code> command. Mandatory locking, and deadlock detection supported by traditional <code>fcntl</code> locks, are not supported on SFCFS. See the <code>fcntl(2)</code> manual page.

Veritas File System features not in cluster file systems

[Table 1-2](#) describes functionality as not supported and may not be expressly prevented from operating on cluster file systems, but the actual behavior is indeterminate.

It is not advisable to use unsupported functionality on SFCFS, or to alternate mounting file systems with these options as local and cluster mounts.

Table 1-2 Veritas File System features not in cluster file systems

Unsupported features	Comments
qlog	Quick log is not supported.
Swap files	Swap files are not supported on cluster mounted file system.
mknod	The <code>mknod</code> command cannot be used to create devices on a cluster mounted file system.
Cache advisories	Cache advisories are set with the <code>mount</code> command on individual file systems, but are not propagated to other nodes of a cluster.

Table 1-2 Veritas File System features not in cluster file systems (*continued*)

Unsupported features	Comments
Cached Quick I/O	This Quick I/O for Databases feature that caches data in the file system cache is not supported.
Commands that depend on file access times	File access times may appear different across nodes because the <code>atime</code> file attribute is not closely synchronized in a cluster file system. So utilities that depend on checking access times may not function reliably.

Storage Foundation Cluster File System benefits and applications

This section describes the SFCFS benefits and applications.

This section includes the following topics:

- How Storage Foundation Cluster File System works
- When to use Storage Foundation Cluster File System

How Storage Foundation Cluster File System works

SFCFS simplifies or eliminates system administration tasks that result from the following hardware limitations:

- The SFCFS single file system image administrative model simplifies administration by enabling the execution of all file system management commands from any node.
- You can create and manage terabyte-sized volumes, so partitioning file systems to fit within disk limitations is usually not necessary.
- SFCFS can support file systems with up to 256 terabyte in size, so only extremely large data farms must be partitioned because of file system addressing limitations.
- Because all servers in a cluster have access to SFCFS cluster-shareable file systems, keeping data consistent across multiple servers is automatic. All cluster nodes have access to the same data, and all data is accessible by all servers using single server file system semantics.
- Because all files can be accessed by all servers, applications can be allocated to servers to balance load or meet other operational requirements. Similarly,

failover becomes more flexible because it is not constrained by data accessibility.

- Because each SFCFS file system can be on any node in the cluster, the file system recovery portion of failover time in an n -node cluster can be reduced by a factor of n by distributing the file systems uniformly across cluster nodes.
- Enterprise RAID subsystems can be used more effectively because all of their capacity can be mounted by all servers, and allocated by using administrative operations instead of hardware reconfigurations.
- Larger volumes with wider striping improve application I/O load balancing. Not only is the I/O load of each server spread across storage resources, but with SFCFS shared file systems, the loads of all servers are balanced against each other.
- Extending clusters by adding servers is easier because each new server's storage configuration does not need to be set up—new servers simply adopt the cluster-wide volume and file system configuration.
- The clusterized Oracle Disk Manager (ODM) feature that makes file-based databases perform as well as raw partition-based databases is available to applications running in a cluster.

When to use Storage Foundation Cluster File System

You should use SFCFS for any application that requires the sharing of files, such as for home directories and boot server files, Web pages, and for cluster-ready applications. SFCFS is also applicable when you want highly available standby data, in predominantly read-only environments where you just need to access data, or when you do not want to rely on NFS for file sharing.

Almost all applications can benefit from SFCFS. Applications that are not “cluster-aware” can operate on and access data from anywhere in a cluster. If multiple cluster applications running on different servers are accessing data in a cluster file system, overall system I/O performance improves due to the load balancing effect of having one cluster file system on a separate underlying volume. This is automatic; no tuning or other administrative action is required.

Many applications consist of multiple concurrent threads of execution that could run on different servers if they had a way to coordinate their data accesses. SFCFS provides this coordination. Such applications can be made cluster-aware allowing their instances to co-operate to balance client and data access load, and thereby scale beyond the capacity of any single server. In such applications, SFCFS provides shared data access, enabling application-level load balancing across cluster nodes.

SFCFS provides the following features:

- For single-host applications that must be continuously available, SFCFS can reduce application failover time because it provides an already-running file system environment in which an application can restart after a server failure.
- For parallel applications, such as distributed database management systems and Web servers, SFCFS provides shared data to all application instances concurrently. SFCFS also allows these applications to grow by the addition of servers, and improves their availability by enabling them to redistribute load in the event of server failure simply by reassigning network addresses.
- For workflow applications, such as video production, in which very large files are passed from station to station, the SFCFS eliminates time consuming and error prone data copying by making files available at all stations.
- For backup, the SFCFS can reduce the impact on operations by running on a separate server, accessing data in cluster-shareable file systems.

The following are examples of applications and how they might work with SFCFS:

- **Using Storage Foundation Cluster File System on file servers**
Two or more servers connected in a cluster configuration (that is, connected to the same clients and the same storage) serve separate file systems. If one of the servers fails, the other recognizes the failure, recovers, assumes the primaryship, and begins responding to clients using the failed server's IP addresses.
- **Using Storage Foundation Cluster File System on web servers**
Web servers are particularly suitable to shared clustering because their application is typically read-only. Moreover, with a client load balancing front end, a Web server cluster's capacity can be expanded by adding a server and another copy of the site. A SFCFS-based cluster greatly simplifies scaling and administration for this type of application.

Storage Foundation Cluster File System architecture

This chapter includes the following topics:

- [About Veritas Cluster Server architecture](#)
- [About Storage Foundation Cluster File System and the Group Lock Manager](#)
- [Storage Foundation Cluster File System namespace](#)
- [About asymmetric mounts](#)
- [Primary and secondary](#)
- [Determining or moving primaryship](#)
- [Synchronize time on Cluster File Systems](#)
- [File system tunables](#)
- [Setting the number of parallel fsck threads](#)
- [About Storage Checkpoints](#)
- [Storage Foundation Cluster File System backup strategies](#)
- [Parallel I/O](#)
- [I/O error handling policy](#)
- [Recovering for I/O failures](#)
- [Single network link and reliability](#)
- [Split-brain and jeopardy handling](#)

- [About I/O fencing](#)
- [About Veritas Volume Manager cluster functionality](#)
- [Storage Foundation Cluster File System and Veritas Volume Manager cluster functionality agents](#)
- [Veritas Volume Manager cluster functionality](#)

About Veritas Cluster Server architecture

The Group Membership and Atomic Broadcast (GAB) and Low Latency Transport (LLT) are VCS-specific protocols implemented directly on Ethernet data link. They run on redundant data links that connect the nodes in a cluster. VCS requires redundant cluster communication links to avoid single points of failure.

GAB provides membership and messaging for the cluster and its applications. GAB membership also provides orderly startup and shutdown of a cluster. The `/etc/gabtab` file is used to configure GAB. This file contains the `gabconfig` command run by GAB on startup. For example, the `-n <number>` option of the command specifies the number of nodes in the cluster. GAB is configured automatically when you run the SFCFS installation script, but you may have to reconfigure GAB when adding nodes to a cluster.

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

LLT provides kernel-to-kernel communications and monitors network communications. The `LLT/etc/llthosts` and `/etc/llttab` files are configured to set system IDs within a cluster, set cluster IDs for multiple clusters, and tune network parameters such as heartbeat frequency. LLT is implemented so that cluster membership changes are reflected quickly, which in turn enables fast responses.

As with GAB, LLT is configured automatically when you run the VCS installation script. The `/etc/llttab` and `/etc/llthosts` files contain information you provide during installation. You may also have to reconfigure LLT when adding nodes to a cluster.

See the `llttab(4)` and the `llthosts(4)` manual pages.

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

Each component in SFCFS registers with a GAB membership port. The port membership identifies nodes that have formed a cluster for the individual components.

[Table 2-1](#) describes the port memberships.

Table 2-1 Port memberships

Port	Description
port a	heartbeat membership
port b	I/O fencing membership
port d	Oracle Disk Manager (ODM) membership
port f	Cluster File system membership
port h	Veritas Cluster Server communication between GAB and High Availability Daemon (HAD)
port u	Cluster Volume Manager (CVM) port for redirecting commands from CVM slaves to CVM master
port v	Cluster Volume Manager membership
port w	Cluster Volume Manager daemons on different nodes communicate with one another using this port, but receive cluster membership information through GAB (port v)

About Storage Foundation Cluster File System and the Group Lock Manager

SFCFS uses the Veritas Group Lock Manager (GLM) to reproduce UNIX single-host file system semantics in clusters. UNIX file systems make writes appear atomic. This means when an application writes a stream of data to a file, a subsequent application reading from the same area of the file retrieves the new data, even if it has been cached by the file system and not yet written to disk. Applications cannot retrieve stale data or partial results from a previous write.

To reproduce single-host write semantics, the file system must keep system caches coherent, and each must instantly reflect updates to cached data, regardless of the node from which the updates originate.

Storage Foundation Cluster File System namespace

The mount point name must remain the same for all nodes mounting the same cluster file system. This is required for the VCS mount agents (online, offline, and monitoring) to work correctly.

About asymmetric mounts

A VxFS file system mounted with the `mount -o cluster` option is a cluster, or shared mount, as opposed to a non-shared or local mount. A file system mounted in shared mode must be on a VxVM shared volume in a cluster environment. A local mount cannot be remounted in shared mode and a shared mount cannot be remounted in local mode when you use the `mount -o remount` option. A single clustered file system can be mounted with different read/writes options on different nodes. These are called asymmetric mounts.

Asymmetric mounts allow shared file systems to be mounted with different read/write capabilities. For example, one node in the cluster can mount read/write, while other nodes mount read-only.

When a primary mounts "ro", this means that neither this node nor any other node is allowed to write the file system. Secondaries can only mount "ro", if the primary mounts "ro". Otherwise, the primary mounts either "rw" or "ro,crw", and the secondaries have the same choice.

You can specify the cluster read-write (`crw`) option when you first mount the file system, or the options can be altered when doing a remount (`mount -o remount`).

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

[Figure 2-1](#) describes the first column showing the mode in which the primary is mounted:

Figure 2-1 Primary and secondary mounts

		Secondary		
		ro	rw	ro, crw
Primary	ro	X		
	rw		X	X
	ro, crw		X	X

The check marks indicate the mode secondary mounts can use for a given mode of the primary.

Mounting the primary with only the `-o cluster,ro` option prevents the secondaries from mounting in a different mode; that is, read-write.

Note: `rw` implies read-write capability throughout the cluster.

Primary and secondary

A file system cluster consists of one primary, and up to 63 secondaries. The primary-secondary terminology applies to one file system, not to a specific node (or hardware platform). You can have the same cluster node be primary for one shared file system, while at the same time it is secondary for another shared file system. Such distribution of file system primaryship to balance the load on a cluster is a recommended administrative policy.

See “[Distribute the load on a cluster](#)” on page 61.

For CVM, a single cluster node is the master for all shared disk groups and shared volumes in the cluster.

See “[About Veritas Volume Manager cluster functionality](#)” on page 47.

Determining or moving primaryship

The first node of a cluster file system to mount is called the primary node. Other nodes are called secondary nodes. If a primary node fails, an internal election process determines which of the secondaries becomes the primary file system.

To determine primaryship

- To determine primaryship, type the following command:

```
# fsclustadm -v showprimary mount_point
```

To give primaryship to a node

- To give primaryship to a node, type the following command on the node:

```
# fsclustadm -v setprimary mount_point
```

Synchronize 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 inode (`ctime`) and data modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

File system tunables

Using the `tunefstab` file updates the tunable parameters at the time of mount. The file system `tunefstab` parameters are set to be identical on all nodes by propagating the parameters to each cluster node. When the file system is mounted on the node, the `tunefstab` parameters of the primary node are used. Symantec recommends that this file be identical on each node.

Setting the number of parallel fsck threads

This section describes how to set the number of parallel fsck threads.

The number of parallel fsck threads that could be active during recovery was set to 4. For example, if a node failed over 12 file systems, log replay for the 12 file systems will not complete at the same time. The number was set to 4 since parallel replay of a large number of file systems would put memory pressure on systems with less memory. However, on larger systems the restriction of 4 parallel processes replaying is not necessary.

This value gets tuned in accordance with available physical memory in the system.

To set the number of parallel fsck threads

- ◆ On all nodes in the cluster, edit the `/opt/VRTSvcs/bin/CFSfsckd/CFSfsckd.env` file and set `FSCKD_OPTS="-n N"`.

where *N* is the number of parallel fsck threads desired and value of *N* has to be between 4 and 128.

Note: The default number of parallel fsck threads is determined by the amount of memory on the machine, unless overridden by the user.

About Storage Checkpoints

Veritas File System (VxFS) provides a Storage Checkpoint feature that quickly creates a persistent image of a file system at an exact point in time. Storage Checkpoints significantly reduce I/O overhead by identifying and maintaining only the file system blocks that have changed since the last Storage Checkpoint or backup via a copy-on-write technique.

Storage Checkpoints provide:

- Persistence through reboots and crashes.
- The ability for data to be immediately writeable by preserving the file system metadata, the directory hierarchy, and user data.

Storage Checkpoints are actually data objects that are managed and controlled by the file system. You can create, remove, and rename Storage Checkpoints because they are data objects with associated names.

Unlike a disk-based mirroring technology that requires a separate storage space, Storage Checkpoints minimize the use of disk space by using a Storage Checkpoint within the same free space available to the file system.

After you create a Storage Checkpoint of a mounted file system, you can also continue to create, remove, and update files on the file system without affecting the logical image of the Storage Checkpoint. A Storage Checkpoint preserves not only the name space (directory hierarchy) of the file system, but also the user data as it existed at the moment the file system image was captured.

You can use a Storage checkpoint in many ways. For example, you can use them to:

- Create a stable image of the file system that can be backed up to tape.
- Provide a mounted, on-disk backup of the file system so that end users can restore their own files in the event of accidental deletion. This is especially useful in a home directory, engineering, or email environment.
- Create a copy of an application's binaries before installing a patch to allow for rollback in case of problems.
- Create an on-disk backup of the file system in that can be used in addition to a traditional tape-based backup to provide faster backup and restore capabilities.
- Test new software on a point-in-time image of the primary fileset without jeopardizing the live data in the current primary fileset by mounting the Storage Checkpoints as writable.

Storage Foundation Cluster File System backup strategies

The same backup strategies used for standard VxFS can be used with SFCFS because the APIs and commands for accessing the namespace are the same. File System checkpoints provide an on-disk, point-in-time copy of the file system. Because performance characteristics of a checkpointed file system are better in

certain I/O patterns, they are recommended over file system snapshots (described below) for obtaining a frozen image of the cluster file system.

File System snapshots are another method of a file system on-disk frozen image. The frozen image is non-persistent, in contrast to the checkpoint feature. A snapshot can be accessed as a read-only mounted file system to perform efficient online backups of the file system. Snapshots implement “copy-on-write” semantics that incrementally copy data blocks when they are overwritten on the snapped file system. Snapshots for cluster file systems extend the same copy-on-write mechanism for the I/O originating from any cluster node.

Mounting a snapshot filesystem for backups increases the load on the system because of the resources used to perform copy-on-writes and to read data blocks from the snapshot. In this situation, cluster snapshots can be used to do off-host backups. Off-host backups reduce the load of a backup application from the primary server. Overhead from remote snapshots is small when compared to overall snapshot overhead. Therefore, running a backup application by mounting a snapshot from a relatively less loaded node is beneficial to overall cluster performance.

The following are several characteristics of a cluster snapshot:

- A snapshot for a cluster mounted file system can be mounted on any node in a cluster. The file system can be a primary, secondary, or secondary-only. A stable image of the file system is provided for writes from any node. See the `mount_vxfs` manual page for more information on secondary-only (`seconly`) is a CFS mount option.
- Multiple snapshots of a cluster file system can be mounted on the same or different cluster nodes.
- A snapshot is accessible only on the node mounting the snapshot. The snapshot device cannot be mounted on two nodes simultaneously.
- The device for mounting a snapshot can be a local disk or a shared volume. A shared volume is used exclusively by a snapshot mount and is not usable from other nodes as long as the snapshot is mounted on that device.
- On the node mounting a snapshot, the snapped file system cannot be unmounted while the snapshot is mounted.
- A SFCFS snapshot ceases to exist if it is unmounted or the node mounting the snapshot fails. However, a snapshot is not affected if another node leaves or joins the cluster.
- A snapshot of a read-only mounted file system cannot be taken. It is possible to mount a snapshot of a cluster file system only if the snapped cluster file system is mounted with the `crw` option.

In addition to frozen images of file systems, there are volume-level alternatives available for shared volumes using mirror split and rejoin. Features such as Fast Mirror Resync and Space Optimized snapshot are also available.

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

Parallel I/O

Some distributed applications read and write to the same file concurrently from one or more nodes in the cluster; for example, any distributed application where one thread appends to a file and there are one or more threads reading from various regions in the file. Several high-performance compute (HPC) applications can also benefit from this feature, where concurrent I/O is performed on the same file. Applications do not require any changes to use parallel I/O.

Traditionally, the entire file is locked to perform I/O to a small region. To support parallel I/O, SFCFS locks ranges in a file that correspond to I/O requests. Two I/O requests conflict if at least one is a write request, and the I/O range of the request overlaps the I/O range of the other.

The parallel I/O feature enables I/O to a file by multiple threads concurrently, as long as the requests do not conflict. Threads issuing concurrent I/O requests could be executing on the same node, or on different nodes in the cluster.

An I/O request that requires allocation is not executed concurrently with other I/O requests. Note that when a writer is extending the file and readers are lagging behind, block allocation is not necessarily done for each extending write.

Predetermine the file size and preallocate the file to avoid block allocations during I/O. This improves the concurrency of applications performing parallel I/O to the file. Parallel I/O also avoids unnecessary page cache flushes and invalidations using range locking, without compromising the cache coherency across the cluster.

For applications that update the same file from multiple nodes, the `-nomtime` mount option provides further concurrency. Modification and change times of the file are not synchronized across the cluster, which eliminates the overhead of increased I/O and locking. The timestamp seen for these files from a node may not have the time updates that happened in the last 60 seconds.

I/O error handling policy

I/O errors can occur for several reasons, including failures of Fibre Channel links, host-bus adapters, and disks. SFCFS disables the file system on the node encountering I/O errors. The file system remains available from other nodes.

After the hardware error is fixed (for example, the Fibre Channel link is reestablished), the file system can be force unmounted and the mount resource can be brought online from the disabled node to reinstate the file system.

Recovering for I/O failures

The disabled file system can be restored by a force unmount and the resource will be brought online without rebooting, which also brings the shared disk group resource online.

Note: If the jeopardy condition is not fixed, the nodes are susceptible to leaving the cluster again on subsequent node failure.

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

Single network link and reliability

Certain environments may prefer using a single private link or a public network for connecting nodes in a cluster, despite the loss of redundancy for dealing with network failures. The benefits of this approach include simpler hardware topology and lower costs; however, there is obviously a tradeoff with high availability.

For the above environments, SFCFS provides the option of a single private link, or using the public network as the private link if I/O fencing is present. I/O fencing is used to handle split-brain scenarios. The option for single network is given during installation.

See [“About preventing data corruption with I/O fencing”](#) on page 33.

Configuring low priority a link

LLT can be configured to use a low-priority network link as a backup to normal heartbeat channels. Low-priority links are typically configured on the customer's public or administrative network. This typically results in a completely different network infrastructure than the cluster private interconnect, and reduces the chance of a single point of failure bringing down all links. The low-priority link is not used for cluster membership traffic until it is the only remaining link. In normal operation, the low-priority link carries only heartbeat traffic for cluster membership and link state maintenance. The frequency of heartbeats drops 50 percent to reduce network overhead. When the low-priority link is the only remaining network link, LLT also switches over all cluster status traffic. Following

repair of any configured private link, LLT returns cluster status traffic to the high-priority link.

LLT links can be added or removed while clients are connected. Shutting down GAB or the high-availability daemon, had, is not required.

To add a link

- To add a link, type the following command:

```
# lltconfig -d device -t device_tag
```

where *device_tag* is a tag to identify particular link in subsequent commands, and is displayed by `lltstat(1M)`.

To remove a link

- To remove a link, type the following command:

```
# lltconfig -u device_tag
```

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

Changes take effect immediately and are lost on the next reboot. For changes to span reboots you must also update the `/etc/llttab` file.

Note: LLT clients will not know how things are going until you only have one LLT link left and GAB declares jeopardy

Split-brain and jeopardy handling

A split-brain occurs when the cluster membership view differs among the cluster nodes, increasing the chance of data corruption. With I/O fencing, the potential for data corruption is eliminated. I/O fencing requires disks that support SCSI-3 PGR.

Jeopardy state

In the absence of I/O fencing, SFCFS installation requires two heartbeat links. When a node is down to a single heartbeat connection, SFCFS can no longer discriminate between loss of a system and loss of the final network connection. This state is defined as jeopardy.

SFCFS detects jeopardy and responds to it in ways that prevent data corruption in some split-brain situations. However, data corruption can still occur in other situations:

- All links go down simultaneously.
- A node hangs and is unable to respond to heartbeat messages.

To eliminate the chance of data corruption in these scenarios, I/O fencing is required. With I/O fencing, the jeopardy state does not require special handling by the SFCFS stack.

Jeopardy handling

For installations that do not support SCSI-3 PGR, jeopardy handling prevents some potential split-brain conditions. If any cluster node fails following a jeopardy state notification, all cluster file systems that were mounted on the failed node or nodes are disabled on all remaining nodes. If a leave reconfiguration happens after jeopardy state notification, then the nodes which have received the jeopardy state notification leave the cluster.

About I/O fencing

I/O fencing protects the data on shared disks when nodes in a cluster detect a change in the cluster membership that indicates a split-brain condition.

The fencing operation determines the following:

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

This decision prevents possible data corruption. The installer installs the I/O fencing driver, VRTSvxfen, when you install Storage Foundation Cluster File System HA. To protect data on shared disks, you must configure I/O fencing after you install and configure Storage Foundation Cluster File System HA.

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

See [“About preventing data corruption with I/O fencing”](#) on page 33.

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

About I/O fencing for Storage Foundation Cluster File System HA in virtual machines that do not support SCSI-3 PR

In a traditional I/O fencing implementation, where the coordination points are coordination point servers (CP servers) or coordinator disks, Veritas Clustered

Volume Manager and Veritas I/O fencing modules provide SCSI-3 persistent reservation (SCSI-3 PR) based protection on the data disks. This SCSI-3 PR protection ensures that the I/O operations from the losing node cannot reach a disk that the surviving sub-cluster has already taken over.

In virtualized environments that do not support SCSI-3 PR, Storage Foundation Cluster File System HA attempts to provide reasonable safety for the data disks. Storage Foundation Cluster File System HA requires you to configure non-SCSI3 server-based I/O fencing in such environments. Non-SCSI3 fencing uses CP servers as coordination points with some additional configuration changes to support I/O fencing in such environments.

About preventing data corruption with I/O fencing

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

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

Problems arise when the mechanism that detects the failure breaks down because symptoms appear identical to those of a failed node. For example, if a system in a two-node cluster fails, the system stops sending heartbeats over the private interconnects. The remaining node then takes corrective action. The failure of the private interconnects, instead of the actual nodes, presents identical symptoms and causes each node to determine its peer has departed. This situation typically results in data corruption because both nodes try to take control of data storage in an uncoordinated manner.

In addition to a broken set of private networks, other scenarios can generate this situation. If a system is so busy that it appears to stop responding or "hang," the other nodes could declare it as dead. This declaration may also occur for the nodes that use the hardware that supports a "break" and "resume" function. When a node drops to PROM level with a break and subsequently resumes operations, the other nodes may declare the system dead. They can declare it dead even if the system later returns and begins write operations.

SFCFS uses I/O fencing to remove the risk that is associated with split-brain. I/O fencing allows write access for members of the active cluster. It blocks access to storage from non-members.

About SCSI-3 Persistent Reservations

SCSI-3 Persistent Reservations (SCSI-3 PR) are required for I/O fencing and resolve the issues of using SCSI reservations in a clustered SAN environment. SCSI-3 PR

enables access for multiple nodes to a device and simultaneously blocks access for other nodes.

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

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

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

About I/O fencing operations

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

About I/O fencing components

The shared storage for SFCFS must support SCSI-3 persistent reservations to enable I/O fencing. SFCFS involves two types of shared storage:

- Data disks—Store shared data
See [“About data disks”](#) on page 35.

- Coordination points—Act as a global lock during membership changes
See “[About coordination points](#)” on page 35.

About data disks

Data disks are standard disk devices for data storage and are either physical disks or RAID Logical Units (LUNs).

These disks must support SCSI-3 PR and must be part of standard VxVM or CVM disk groups. CVM is responsible for fencing data disks on a disk group basis. Disks that are added to a disk group and new paths that are discovered for a device are automatically fenced.

About coordination points

Coordination points provide a lock mechanism to determine which nodes get to fence off data drives from other nodes. A node must eject a peer from the coordination points before it can fence the peer from the data drives. Racing for control of the coordination points to fence data disks is the key to understand how fencing prevents split-brain.

Note: Typically, a fencing configuration for a cluster must have three coordination points. Symantec also supports server-based fencing with a single CP server as its only coordination point with a caveat that this CP server becomes a single point of failure.

The coordination points can be disks, servers, or both.

- Coordinator disks

Disks that act as coordination points are called coordinator disks. Coordinator disks are three standard disks or LUNs set aside for I/O fencing during cluster reconfiguration. Coordinator disks do not serve any other storage purpose in the SFCFS configuration.

Dynamic Multi-pathing (DMP) allows coordinator disks to take advantage of the path failover and the dynamic adding and removal capabilities of DMP. On cluster nodes with HP-UX 11i v3, you must use DMP devices or iSCSI devices for I/O fencing. The following changes in HP-UX 11i v3 require you to not use raw devices for I/O fencing:

- Provides native multipathing support
- Does not provide access to individual paths through the device file entries

The metanode interface that HP-UX provides does not meet the SCSI-3 PR requirements for the I/O fencing feature. You can configure coordinator disks to use Veritas Volume Manager Dynamic Multi-pathing (DMP) feature. See the *Veritas Volume Manager Administrator's Guide*.

■ Coordination point servers

The coordination point server (CP server) is a software solution which runs on a remote system or cluster. CP server provides arbitration functionality by allowing the SFCFS cluster nodes to perform the following tasks:

- Self-register to become a member of an active SFCFS cluster (registered with CP server) with access to the data drives
- Check which other nodes are registered as members of this active SFCFS cluster
- Self-unregister from this active SFCFS cluster
- Forcefully unregister other nodes (preempt) as members of this active SFCFS cluster

In short, the CP server functions as another arbitration mechanism that integrates within the existing I/O fencing module.

Note: With the CP server, the fencing arbitration logic still remains on the SFCFS cluster.

Multiple SFCFS clusters running different operating systems can simultaneously access the CP server. TCP/IP based communication is used between the CP server and the SFCFS clusters.

About preferred fencing

The I/O fencing driver uses coordination points to prevent split-brain in a VCS cluster. By default, the fencing driver favors the subcluster with maximum number of nodes during the race for coordination points. With the preferred fencing feature, you can specify how the fencing driver must determine the surviving subcluster.

You can configure the preferred fencing policy using the cluster-level attribute PreferredFencingPolicy as follows:

- Enable system-based preferred fencing policy to give preference to high capacity systems.
- Enable group-based preferred fencing policy to give preference to service groups for high priority applications.

- Disable preferred fencing policy to use the default node count-based race policy.

See [“How preferred fencing works”](#) on page 37.

See [“Enabling or disabling the preferred fencing policy”](#) on page 118.

How preferred fencing works

The I/O fencing driver uses coordination points to prevent split-brain in a VCS cluster. At the time of a network partition, the fencing driver in each subcluster races for the coordination points. The subcluster that grabs the majority of coordination points survives whereas the fencing driver causes a system panic on nodes from all other subclusters. By default, the fencing driver favors the subcluster with maximum number of nodes during the race for coordination points.

This default racing preference does not take into account the application groups that are online on any nodes or the system capacity in any subcluster. For example, consider a two-node cluster where you configured an application on one node and the other node is a standby-node. If there is a network partition and the standby-node wins the race, the node where the application runs panics and VCS has to bring the application online on the standby-node. This behavior causes disruption and takes time for the application to fail over to the surviving node and then to start up again.

The preferred fencing feature lets you specify how the fencing driver must determine the surviving subcluster. The preferred fencing solution makes use of a fencing parameter called node weight. VCS calculates the node weight based on online applications and system capacity details that you provide using specific VCS attributes, and passes to the fencing driver to influence the result of race for coordination points. At the time of a race, the racer node adds up the weights for all nodes in the local subcluster and in the leaving subcluster. If the leaving subcluster has a higher sum (of node weights) then the racer for this subcluster delays the race for the coordination points. Thus, the subcluster that has critical systems or critical applications wins the race.

The preferred fencing feature uses the cluster-level attribute PreferredFencingPolicy that takes the following race policy values:

- Disabled (default): Preferred fencing is disabled.
When the PreferredFencingPolicy attribute value is set as Disabled, VCS sets the count based race policy and resets the value of node weight as 1.
- System: Based on the capacity of the systems in a subcluster.
If one system is more powerful than others in terms of architecture, number of CPUs, or memory, this system is given preference in the fencing race.

When the PreferredFencingPolicy attribute value is set as System, VCS calculates node weight based on the system-level attribute FencingWeight.

- **Group:** Based on the higher priority applications in a subcluster. The fencing driver takes into account the service groups that are online on the nodes in any subcluster. In the event of a network partition, the node with higher priority service groups is given preference in the fencing race. When the PreferredFencingPolicy attribute value is set as Group, VCS calculates node weight based on the group-level attribute Priority for those service groups that are active.

See the *Veritas Cluster Server Administrator's Guide* for more information on the VCS attributes.

See [“Enabling or disabling the preferred fencing policy”](#) on page 118.

About I/O fencing configuration files

[Table 2-2](#) lists the I/O fencing configuration files.

Table 2-2 I/O fencing configuration files

File	Description
/etc/rc.config.d/vxfenconf	<p>This file stores the start and stop environment variables for I/O fencing:</p> <ul style="list-style-type: none"> ■ VXFEN_START—Defines the startup behavior for the I/O fencing module after a system reboot. Valid values include: <ul style="list-style-type: none"> 1—Indicates that I/O fencing is enabled to start up. 0—Indicates that I/O fencing is disabled to start up. ■ VXFEN_STOP—Defines the shutdown behavior for the I/O fencing module during a system shutdown. Valid values include: <ul style="list-style-type: none"> 1—Indicates that I/O fencing is enabled to shut down. 0—Indicates that I/O fencing is disabled to shut down. <p>The installer sets the value of these variables to 1 at the end of Storage Foundation Cluster File System HA configuration.</p>
/etc/vxfendg	<p>This file includes the coordinator disk group information.</p> <p>This file is not applicable for server-based fencing.</p>

Table 2-2 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfenmode	<p>This file contains the following parameters:</p> <ul style="list-style-type: none"> ■ vxfen_mode <ul style="list-style-type: none"> ■ scsi3—For disk-based fencing ■ customized—For server-based fencing ■ disabled—To run the I/O fencing driver but not do any fencing operations. ■ vxfen_mechanism <p>This parameter is applicable only for server-based fencing. Set the value as cps.</p> ■ scsi3_disk_policy <p>You must configure the vxfen module to use DMP devices or iSCSI devices, and set the SCSI-3 disk policy as dmp.</p> ■ security <p>This parameter is applicable only for server-based fencing.</p> <p>1—Indicates that Symantec Product Authentication Service is used for CP server communications. This setting is the default.</p> <p>0—Indicates that communication with the CP server is in non-secure mode.</p> <p>Note: The CP server and the Storage Foundation Cluster File System HA clusters must have the same security setting.</p> ■ List of coordination points <p>This list is required only for server-based fencing configuration.</p> <p>Coordination points in a server-based fencing can include coordinator disks, CP servers, or a mix of both. If you use coordinator disks, you must create a coordinator disk group with the coordinator disk names.</p> <p>Refer to the sample file /etc/vxfen.d/vxfenmode_cps for more information on how to specify the coordination points.</p> ■ single_cp <p>This parameter is applicable only for server-based fencing which uses a single highly available CP server as its coordination point.</p>

Table 2-2 I/O fencing configuration files (*continued*)

File	Description
/etc/vxfentab	<p>When I/O fencing starts, the vxfen startup script creates this /etc/vxfentab file on each node. The startup script uses the contents of the /etc/vxfendg and /etc/vxfenmode files. Any time a system is rebooted, the fencing driver reinitializes the vxfentab file with the current list of all the coordinator points.</p> <p>Note: The /etc/vxfentab file is a generated file; do not modify this file.</p> <p>An example of the /etc/vxfentab file in a disk-based fencing configuration on one node resembles as follows:</p> <pre>/dev/vx/rdmp/c1t1d0 /dev/vx/rdmp/c2t1d0 /dev/vx/rdmp/c3t1d0</pre> <p>For server-based fencing, the /etc/vxfentab file also includes the security settings information.</p> <p>For server-based fencing with single CP server, the /etc/vxfentab file also includes the single_cp settings information.</p>

How disk-based I/O fencing works in different event scenarios

Table 2-3 describes how I/O fencing works to prevent data corruption in different failure event scenarios. For each event, review the corrective operator actions.

Table 2-3 I/O fencing scenarios

Event	Node A: What happens?	Node B: What happens?	Operator action
Both private networks fail.	<p>Node A races for majority of coordinator disks.</p> <p>If Node A wins race for coordinator disks, Node A ejects Node B from the shared disks and continues.</p>	<p>Node B races for majority of coordinator disks.</p> <p>If Node B loses the race for the coordinator disks, Node B panics and removes itself from the cluster.</p>	When Node B is ejected from cluster, repair the private networks before attempting to bring Node B back.

Table 2-3 I/O fencing scenarios (*continued*)

Event	Node A: What happens?	Node B: What happens?	Operator action
Both private networks function again after event above.	Node A continues to work.	Node B has crashed. It cannot start the database since it is unable to write to the data disks.	Restart Node B after private networks are restored.
One private network fails.	Node A prints message about an IOFENCE on the console but continues.	Node B prints message about an IOFENCE on the console but continues.	Repair private network. After network is repaired, both nodes automatically use it.
Node A hangs.	Node A is extremely busy for some reason or is in the kernel debugger. When Node A is no longer hung or in the kernel debugger, any queued writes to the data disks fail because Node A is ejected. When Node A receives message from GAB about being ejected, it panics and removes itself from the cluster.	Node B loses heartbeats with Node A, and races for a majority of coordinator disks. Node B wins race for coordinator disks and ejects Node A from shared data disks.	Verify private networks function and restart Node A.

Table 2-3 I/O fencing scenarios (*continued*)

Event	Node A: What happens?	Node B: What happens?	Operator action
<p>Nodes A and B and private networks lose power. Coordinator and data disks retain power.</p> <p>Power returns to nodes and they restart, but private networks still have no power.</p>	<p>Node A restarts and I/O fencing driver (vxfen) detects Node B is registered with coordinator disks. The driver does not see Node B listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node A from joining the cluster. Node A console displays:</p> <p>Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.</p>	<p>Node B restarts and I/O fencing driver (vxfen) detects Node A is registered with coordinator disks. The driver does not see Node A listed as member of cluster because private networks are down. This causes the I/O fencing device driver to prevent Node B from joining the cluster. Node B console displays:</p> <p>Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.</p>	<p>Resolve preexisting split-brain condition.</p> <p>See “System panics to prevent potential data corruption” on page 181.</p>

Table 2-3 I/O fencing scenarios (*continued*)

Event	Node A: What happens?	Node B: What happens?	Operator action
Node A crashes while Node B is down. Node B comes up and Node A is still down.	Node A is crashed.	<p>Node B restarts and detects Node A is registered with the coordinator disks. The driver does not see Node A listed as member of the cluster. The I/O fencing device driver prints message on console:</p> <p>Potentially a preexisting split brain. Dropping out of the cluster. Refer to the user documentation for steps required to clear preexisting split brain.</p>	<p>Resolve preexisting split-brain condition.</p> <p>See “System panics to prevent potential data corruption” on page 181.</p>
<p>The disk array containing two of the three coordinator disks is powered off.</p> <p>No node leaves the cluster membership</p>	Node A continues to operate as long as no nodes leave the cluster.	Node B continues to operate as long as no nodes leave the cluster.	<p>Power on the failed disk array so that subsequent network partition does not cause cluster shutdown, or replace coordinator disks.</p> <p>See “Replacing I/O fencing coordinator disks when the cluster is online” on page 89.</p>

Table 2-3 I/O fencing scenarios (*continued*)

Event	Node A: What happens?	Node B: What happens?	Operator action
<p>The disk array containing two of the three coordinator disks is powered off.</p> <p>Node B gracefully leaves the cluster and the disk array is still powered off. Graceful leave implies a clean shutdown so that vxfs is properly unconfigured.</p>	<p>Node A continues to operate in the cluster.</p>	<p>Node B has left the cluster.</p>	<p>Power on the failed disk array so that subsequent network partition does not cause cluster shutdown, or replace coordinator disks.</p> <p>See “Replacing I/O fencing coordinator disks when the cluster is online” on page 89.</p>
<p>The disk array containing two of the three coordinator disks is powered off.</p> <p>Node B abruptly crashes or a network partition occurs between node A and node B, and the disk array is still powered off.</p>	<p>Node A races for a majority of coordinator disks. Node A fails because only one of the three coordinator disks is available. Node A panics and removes itself from the cluster.</p>	<p>Node B has left cluster due to crash or network partition.</p>	<p>Power on the failed disk array and restart I/O fencing driver to enable Node A to register with all coordinator disks, or replace coordinator disks.</p> <p>See “Replacing defective disks when the cluster is offline” on page 184.</p>

About server-based I/O fencing

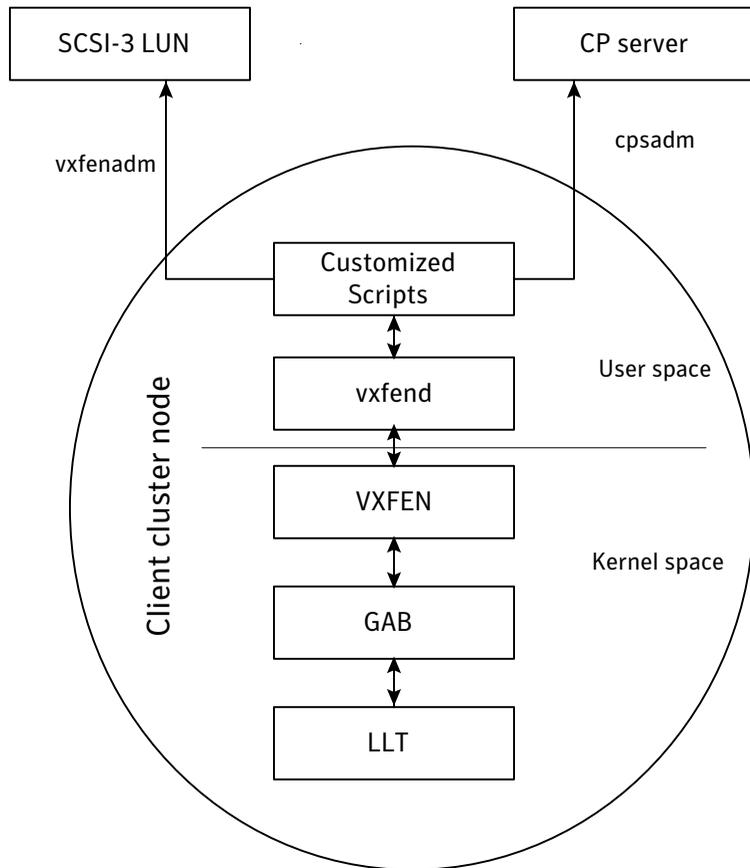
In a disk-based I/O fencing implementation, the vxfs driver handles various SCSI-3 PR based arbitration operations completely within the driver. I/O fencing also provides a framework referred to as customized fencing wherein arbitration operations are implemented in custom scripts. The vxfs driver invokes the custom scripts.

The CP server-based coordination point uses a customized fencing framework. Note that SCSI-3 PR based fencing arbitration can also be enabled using customized fencing framework. This allows the user to specify a combination of SCSI-3 LUNs and CP servers as coordination points using customized fencing. Customized fencing can be enabled by specifying vxfs_mode=customized and vxfs_mechanism=cps in the `/etc/vxfsmode` file.

Moreover, both `/etc/vxfenmode` and `/etc/vxfentab` files contain additional parameter "security" which indicates if communication between CP server and SFCFS cluster nodes is secure.

Figure 2-2 displays a schematic of the customized fencing options.

Figure 2-2 Customized fencing



A user level daemon `vxfend` interacts with the `vxfen` driver, which in turn interacts with `GAB` to get the node membership update. Upon receiving membership updates, `vxfend` invokes various scripts to race for the coordination point and fence off data disks. The `vxfend` daemon manages various fencing agents. The customized fencing scripts are located in the `/opt/VRTSvcs/vxfen/bin/customized/cps` directory.

The scripts that are involved include the following:

- `generate_snapshot.sh`: Retrieves the SCSI ID's of the coordinator disks and/or UUID ID's of the CP servers
For information about the UUID (Universally Unique Identifier), see the *Veritas Cluster Server Administrator's Guide*.
- `join_local_node.sh`: Registers the keys with the coordinator disks or CP servers
- `race_for_coordination_point.sh`: Races to determine a winner after cluster reconfiguration
- `unjoin_local_node.sh`: Removes the keys that are registered in `join_local_node.sh`
- `fence_data_disks.sh`: Fences the data disks from access by the losing nodes.
- `local_info.sh`: Lists local node's configuration parameters and coordination points, which are used by the `vxfen` driver.

About the CP server database

CP server requires a database for storing the registration keys of the SFCFS cluster nodes. CP server uses a SQLite database for its operations. By default, the database is located at `/etc/VRTScps/db`.

For a single node VCS cluster hosting a CP server, the database can be placed on a local file system. For an SFHA cluster hosting a CP server, the database must be placed on a shared file system. The file system must be shared among all nodes that are part of the SFHA cluster.

In an SFHA cluster hosting the CP server, the shared database is protected by setting up SCSI-3 PR based I/O fencing. SCSI-3 PR based I/O fencing protects against split-brain scenarios.

Warning: The CP server database must not be edited directly and should only be accessed using `cpsadm(1M)`. Manipulating the database manually may lead to undesirable results including system panics.

About the CP server user types and privileges

The CP server supports the following user types, each with a different access level privilege:

- CP server administrator (admin)
- CP server operator

Different access level privileges permit the user to issue different commands. If a user is neither a CP server admin nor a CP server operator user, then the user has guest status and can issue limited commands.

The user types and their access level privileges are assigned to individual users during SFCFS cluster configuration for fencing. During the installation process, you are prompted for a user name, password, and access level privilege (CP server admin or CP server operator).

To administer and operate a CP server, there must be at least one CP server admin.

A root user on a CP server is given all the administrator privileges, and these administrator privileges can be used to perform all the CP server specific operations.

About Veritas Volume Manager cluster functionality

Veritas Volume Manager cluster functionality (CVM) allows up to 64 nodes in a cluster to simultaneously access and manage a set of disks under VxVM control (VM disks). The same logical view of the disk configuration and any changes are available on each node. When the cluster functionality is enabled, all cluster nodes can share VxVM objects. Features provided by the base volume manager, such as mirroring, fast mirror resync and dirty region logging are also supported in the cluster environment.

To implement cluster functionality, VxVM works together with the cluster monitor daemon provided by the host operating system or by VCS. The cluster monitor informs VxVM of changes in cluster membership. Each node starts up independently and has its own cluster monitor, plus its own copies of the operating system and CVM. When a node joins a cluster it gains access to shared disks. When a node leaves a cluster, it no longer has access to shared disks. A node joins a cluster when the cluster monitor is started on that node.

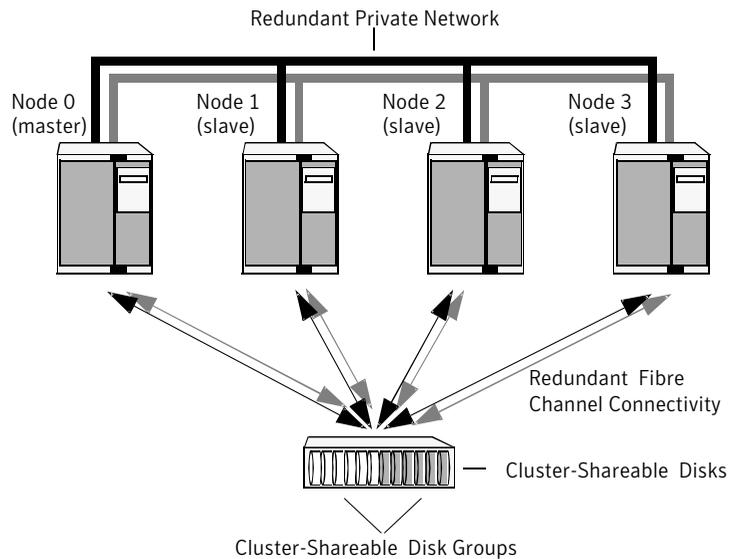
Note: RAID-5 volumes are not supported on a shared disk group.

[Figure 2-3](#) illustrates a simple cluster arrangement consisting of four nodes with similar or identical hardware characteristics (CPUs, RAM and host adapters), and configured with identical software (including the operating system).

The nodes are fully connected by a private network and they are also separately connected to shared external storage (either disk arrays or JBODs: just a bunch of disks) via Fibre Channel. Each node has two independent paths to these disks, which are configured in one or more cluster-shareable disk groups.

The private network allows the nodes to share information about system resources and about each other's state. Using the private network, any node can recognize which nodes are currently active, which are joining or leaving the cluster, and which have failed. The private network requires at least two communication channels to provide redundancy against one of the channels failing. If only one channel were used, its failure would be indistinguishable from node failure—a condition known as network partitioning.

Figure 2-3 Example of a four node cluster



To the cluster monitor, all nodes are the same. VxVM objects configured within shared disk groups can potentially be accessed by all nodes that join the cluster. However, the cluster functionality of VxVM requires one node to act as the master node; all other nodes in the cluster are slave nodes. Any node is capable of being the master node, which is responsible for coordinating certain VxVM activities.

VxVM designates the first node to join a cluster the master node. If the master node leaves the cluster, one of the slave nodes is chosen to be the new master. In the preceding example, node 0 is the master node and nodes 1, 2 and 3 are slave nodes.

Shared disk groups overview

This section provides an overview of shared disk groups.

This section includes the following topics:

- Private and shared disk groups
- Activation modes of shared disk groups
- Connectivity policy of shared disk groups
- Limitations of shared disk groups

Private and shared disk groups

Table 2-4 describes the disk group types.

Table 2-4 Disk group types

Disk group	Description
Private	Belongs to only one node. A private disk group is only imported by one system. Disks in a private disk group may be physically accessible from one or more systems, but import is restricted to one system only. The root disk group is always a private disk group.
Shared	Is shared by all nodes. A shared (or cluster-shareable) disk group is imported by all cluster nodes. Disks in a shared disk group must be physically accessible from all systems that may join the cluster.

In a cluster, most disk groups are shared. Disks in a shared disk group are accessible from all nodes in a cluster, allowing applications on multiple cluster nodes to simultaneously access the same disk. A volume in a shared disk group can be simultaneously accessed by more than one node in the cluster, subject to licensing and disk group activation mode restrictions.

You can use the `vxdg` command to designate a disk group as cluster-shareable. When a disk group is imported as cluster-shareable for one node, each disk header is marked with the cluster ID. As each node subsequently joins the cluster, it recognizes the disk group as being cluster-shareable and imports it. You can also import or deport a shared disk group at any time; the operation takes places in a distributed fashion on all nodes.

Each physical disk is marked with a unique disk ID. When cluster functionality for VxVM starts on the master, it imports all shared disk groups (except for any that have the `noautoimport` attribute set). When a slave tries to join a cluster, the master sends it a list of the disk IDs that it has imported, and the slave checks to see if it can access them all. If the slave cannot access one of the listed disks, it abandons its attempt to join the cluster. If it can access all of the listed disks, it imports the same shared disk groups as the master and joins the cluster. When a node leaves the cluster, it deports all its imported shared disk groups, but they remain imported on the surviving nodes.

Reconfiguring a shared disk group is performed with the co-operation of all nodes. Configuration changes to the disk group happen simultaneously on all nodes and the changes are identical. Such changes are atomic in nature, which means that they either occur simultaneously on all nodes or not at all.

Whether all members of the cluster have simultaneous read and write access to a cluster-shareable disk group depends on its activation mode setting.

The data contained in a cluster-shareable disk group is available as long as at least one node is active in the cluster. The failure of a cluster node does not affect access by the remaining active nodes. Regardless of which node accesses a cluster-shareable disk group, the configuration of the disk group looks the same.

Note: Applications running on each node can access the data on the VM disks simultaneously. VxVM does not protect against simultaneous writes to shared volumes by more than one node. It is assumed that applications control consistency (by using Veritas Storage Foundation Cluster File System or a distributed lock manager, for example).

Activation modes of shared disk groups

A shared disk group must be activated on a node in order for the volumes in the disk group to become accessible for application I/O from that node. The ability of applications to read from or to write to volumes is dictated by the activation mode of a shared disk group. Valid activation modes for a shared disk group are `exclusivewrite`, `readonly`, `sharedread`, `sharedwrite`, and `off` (inactive).

Note: The default activation mode for shared disk groups is `off` (inactive).

Special uses of clusters, such as high availability (HA) applications and off-host backup, can use disk group activation to explicitly control volume access from different nodes in the cluster.

[Table 2-5](#) describes activation modes for shared disk groups.

Table 2-5 Activation modes for shared disk groups

Activation mode	Description
<code>exclusivewrite</code> (ew)	The node has exclusive write access to the disk group. No other node can activate the disk group for write access.

Table 2-5 Activation modes for shared disk groups (*continued*)

Activation mode	Description
<code>readonly (ro)</code>	The node has read access to the disk group and denies write access for all other nodes in the cluster. The node has no write access to the disk group. Attempts to activate a disk group for either of the write modes on other nodes fail.
<code>sharedread (sr)</code>	The node has read access to the disk group. The node has no write access to the disk group, however other nodes can obtain write access.
<code>sharedwrite (sw)</code>	The node has write access to the disk group.
<code>off</code>	The node has neither read nor write access to the disk group. Query operations on the disk group are permitted.

[Table 2-6](#) summarizes the allowed and conflicting activation modes for shared disk groups.

Table 2-6 Allowed and conflicting activation modes

Disk group activated in cluster as...	exclusive-write	readonly	sharedread	sharedwrite
<code>exclusivewrite</code>	Fails	Fails	Succeeds	Fails
<code>readonly</code>	Fails	Succeeds	Succeeds	Fails
<code>sharedread</code>	Succeeds	Succeeds	Succeeds	Succeeds
<code>sharedwrite</code>	Fails	Fails	Succeeds	Succeeds

To share disk groups

- Shared disk groups can be automatically activated in any mode during disk group creation or during manual or auto-import. To control auto-activation of shared disk groups, the `/etc/default/vxdg` file must be created.

The `/etc/default/vxdg` file must contain the following lines:

```
enable_activation=true
default_activation_mode=activation-mode
```

The `activation-mode` is one of `exclusivewrite`, `readonly`, `sharedread`, `sharedwrite`, or `off`.

When a shared disk group is created or imported, it is activated in the specified mode. When a node joins the cluster, all shared disk groups accessible from the node are activated in the specified mode.

The activation mode of a disk group controls volume I/O from different nodes in the cluster. It is not possible to activate a disk group on a given node if it is activated in a conflicting mode on another node in the cluster. When enabling activation using the defaults file, it is recommended that this file be made identical on all nodes in the cluster. Otherwise, the results of activation are unpredictable.

Note: If the `/etc/default/vxdg` file is edited while the `vxconfigd` daemon is already running, the `vxconfigd` process must be restarted for the changes in the defaults file to take effect.

If the default activation mode is anything other than `off`, an activation following a cluster join, or a disk group creation or import can fail if another node in the cluster has activated the disk group in a conflicting mode.

To display the activation mode for a shared disk group, use the `vxdg list` command.

You can also use the `vxdg` command to change the activation mode on a shared disk group.

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

Connectivity policy of shared disk groups

The nodes in a cluster must always agree on the status of a disk. In particular, if one node cannot write to a given disk, all nodes must stop accessing that disk before the results of the write operation are returned to the caller. Therefore, if a node cannot contact a disk, it should contact another node to check on the disk's status. If the disk fails, no node can access it and the nodes can agree to detach the disk. If the disk does not fail, but rather the access paths from some of the nodes fail, the nodes cannot agree on the status of the disk.

[Table 2-7](#) describes the policies for resolving this type of discrepancy.

Table 2-7 Policies

Policy	Description
Global	The detach occurs cluster-wide (globally) if any node in the cluster reports a disk failure. This is the default policy.

Table 2-7 Policies (*continued*)

Policy	Description
Local	In the event of disks failing, the failures are confined to the particular nodes that saw the failure. However, this policy is not highly available because it fails the node even if one of the mirrors is available. Note that an attempt is made to communicate with all nodes in the cluster to ascertain the disks' usability. If all nodes report a problem with the disks, a cluster-wide detach occurs.

Limitations of shared disk groups

The cluster functionality of VxVM does not support RAID-5 volumes, or task monitoring for cluster-shareable disk groups. These features can, however, be used in private disk groups that are attached to specific nodes of a cluster. Online relayout is supported provided that it does not involve RAID-5 volumes.

The root disk group cannot be made cluster-shareable. It must be private.

Only raw device access may be performed via the cluster functionality of VxVM. It does not support shared access to file systems in shared volumes unless the appropriate software, such as Veritas Storage Foundation Cluster File System, is installed and configured.

If a shared disk group contains unsupported objects, deport it and then re-import the disk group as private on one of the cluster nodes. Reorganize the volumes into layouts that are supported for shared disk groups, and then deport and re-import the disk group as shared.

Storage Foundation Cluster File System and Veritas Volume Manager cluster functionality agents

Agents are VCS processes that manage predefined resource types. SFCFS and CVM require agents to interact with VCS. Agents bring resources online, take resources offline, monitor resources, and report any state changes to VCS. VCS bundled agents are part of VCS and are installed when VCS is installed. The SFCFS and CVM agents are add-on resources to VCS specifically for the Veritas File System and Veritas Volume Manager.

See the *Veritas Storage Foundation Cluster File System Installation Guide*.

Veritas Volume Manager cluster functionality

The Veritas Volume Manager cluster functionality (CVM) makes logical volumes accessible throughout a cluster. CVM enables multiple hosts to concurrently access the logical volumes under its control. A VxVM cluster comprises nodes sharing a set of devices. The nodes are connected across a network. If one node fails, other nodes can access the devices. The VxVM cluster feature presents the same logical view of the device configurations, including changes, on all nodes. You configure CVM shared storage after VCS sets up a cluster configuration.

Administering Storage Foundation Cluster File System and its components

This chapter includes the following topics:

- [About Storage Foundation Cluster File System administration](#)
- [Administering CFS](#)
- [Administering VCS](#)
- [Administering CVM](#)
- [Administering ODM](#)
- [Administering I/O Fencing](#)
- [Administering SFCFS global clusters](#)

About Storage Foundation Cluster File System administration

The Veritas Storage Foundation Cluster File System is a shared file system that enables multiple hosts to mount and perform file operations concurrently on the same file. To operate in a cluster configuration, SFCFS requires the integrated set of Veritas products included in the Veritas Storage Foundation Cluster File System.

To configure a cluster, SFCFS requires the Veritas Cluster Server (VCS). VCS supplies two major components integral to SFCFS. The LLT package provides node-to-node communications and monitors network communications. The GAB

package provides cluster state, configuration, and membership service, and monitors the heartbeat links between systems to ensure that they are active. There are several other packages supplied by VCS that provide application failover support when installing SFCFS HA.

See the *Veritas Storage Foundation Cluster File System Installation Guide*.

SFCFS also requires the cluster functionality (CVM) of the Veritas Volume Manager (VxVM) to create the shared volumes necessary for mounting cluster file systems.

For more information on these products, refer to the *Veritas Volume Manager* and *Veritas Cluster Server* documentation.

Administering CFS

This section describes some of the major aspects of cluster file system administration.

This section provides instructions for the following CFS administration tasks:

- Using `cfsmount` to mount CFS file systems
See “[Using cfsmount to mount CFS file systems](#)” on page 56.
- Resizing CFS file systems
See “[Resizing CFS file systems](#)” on page 57.
- Verifying the status of CFS file systems
See “[Verifying the status of CFS file systems](#)” on page 57.
- Verifying CFS port
See “[Verifying CFS port](#)” on page 57.

If you encounter issues while administering CFS, refer to the troubleshooting section for assistance.

Using `cfsmount` to mount CFS file systems

To mount a CFS file system using `cfsmount`:

```
# cfsmount /oradata1
Mounting...
[/dev/vx/dsk/oradatadg/oradatavol]
mounted successfully at /oradata1 on galaxy
[/dev/vx/dsk/oradatadg/oradatavol]
mounted successfully at /oradata1 on nebula
```

Resizing CFS file systems

If you see a message on the console indicating that a CFS file system is full, you may want to resize the file system. The `vxresize` command lets you resize a CFS file system. It extends the file system and the underlying volume.

See the `vxresize (1M)` manual page for information on various options.

The following command resizes an Oracle data CFS file system (the Oracle data volume is CFS mounted):

```
# vxresize -g oradatadg oradatavol +2G
```

Verifying the status of CFS file systems

Run the "cfscluster status" command to see the status of the nodes and their mount points:

```
# cfscluster status
```

```
Node           : nebula
Cluster Manager : not-running
CVM state      : not-running
MOUNT POINT    SHARED VOLUME  DISK GROUP      STATUS
/ocrvote       ocrvotevol     galaxy_ocr       NOT MOUNTED
/oracle         ora_vol        galaxy_ora       NOT MOUNTED
/crshome        ora_crs_vol    galaxy_crs       NOT MOUNTED
/oradata1      ora_data1_vol  galaxy_data1     NOT MOUNTED
/arch           archivol       galaxy_data1     NOT MOUNTED

Node           : galaxy
Cluster Manager : running
CVM state      : running
MOUNT POINT    SHARED VOLUME  DISK GROUP      STATUS
/ocrvote       ocrvotevol     galaxy_ocr       MOUNTED
/oracle         ora_vol        galaxy_ora       MOUNTED
/crshome        ora_crs_vol    galaxy_crs       MOUNTED
/oradata1      ora_data1_vol  galaxy_data1     MOUNTED
/arch           archivol       galaxy_data1     MOUNTED
```

Verifying CFS port

CFS uses port 'f' for communication between nodes. The CFS port state can be verified as follows:

```
# gabconfig -a | grep "Port f"
```

CFS agents and AMF support

The CFS agents (CFMount and CFSfsckd) are AMF-aware. In this release the CFS agents use the V51 framework.

CFS agent log files

You can use the CFS agent log files that are located in the directory `/var/VRTSvcs/log` to debug CFS issues.

```
# cd /var/VRTSvcs/log
# ls
CFMount_A.log
CFSfsckd_A.log
engine_A.log
```

The agent framework information is located in the `engine_A.log` file while the agent entry point information is located in the `CFMount_A.log` and `CFSfsckd_A.log` files.

Storage Foundation Cluster File System commands

[Table 3-1](#) describes the SFCFS commands.

Table 3-1 SFCFS commands

Commands	Description
<code>cfcluster</code>	Cluster configuration command See the <code>cfcluster(1M)</code> manual page for more information.
<code>cfmntadm</code>	Adds, deletes, modifies, and sets policy on cluster mounted file systems See the <code>cfmntadm(1M)</code> manual page for more information.
<code>cfsdgadm</code>	adds or deletes shared disk groups to and from a cluster configuration See the <code>cfsdgadm(1M)</code> manual page for more information.
<code>cfsmount</code>	mounts a cluster file system on a shared volume See the <code>cfsmount(1M)</code> manual page for more information.
<code>cfsumount</code>	unmounts a cluster file system on a shared volume See the <code>cfsumount(1M)</code> manual page for more information.

Table 3-1 SFCFS commands (*continued*)

Commands	Description
<code>cfsshare</code>	Clustered NFS (CNFS) and Common Internet File System (CIFS) configuration command See the <code>cfsshare(1M)</code> manual page for more information.

mount, fsclustadm, and fsadm commands

The `mount` and `fsclustadm` commands are important for configuring cluster file systems.

mount

The `mount` command with the `-o cluster` option lets you access shared file systems.

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

fsclustadm

The `fsclustadm` command reports various attributes of a cluster file system. Using `fsclustadm` you can show and set the primary node in a cluster, translate node IDs to host names and vice versa, list all nodes that currently have a cluster mount of the specified file system mount point, and determine whether a mount is a local or cluster mount. The `fsclustadm` command operates from any node in a cluster on which the file system is mounted, and can control the location of the primary for a specified mount point.

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

fsadm

The `fsadm` command can be invoked from the primary or secondary node.

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

Run commands safely in a cluster environment

Any UNIX command that can write to a raw device must be used carefully in a shared environment to prevent data from being corrupted. For shared VxVM volumes, SFCFS provides protection by reserving the volumes in a cluster to prevent VxFS commands, such as `fsck` and `mkfs`, from inadvertently damaging a mounted file system from another node in a cluster. However, commands such

as `dd` execute without any reservation, and can damage a file system mounted from another node. Before running this kind of command on a file system, be sure the file system is not mounted on a cluster. You can run the `mount` command to see if a file system is a shared or local mount.

Time synchronization for 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 creation (`ctime`) and modification (`mtime`) may not be consistent with the sequence in which operations actually happened.

Growing a Storage Foundation Cluster File System

There is a master node for CVM as well as a primary for SFCFS. When growing a file system, you grow the volume and slaves from the CVM master, and then grow the file system from any SFCFS node. The CVM master and the SFCFS node can be different nodes.

To determine the primary file system in a cluster

- To determine the primary file system in a cluster, type the following command:

```
# fsclustadm -v showprimary mount_point
```

To determine that the current node is the master CVM node

- To determine if the current node is the master CVM node, type the following command:

```
# vxctl -c mode
```

To actually increase the size of the file system

- 1 On the master CVM node, type the following command:

```
# vxassist -g shared_disk_group growto volume_name newlength
```

- 2 On any SFCFS node, type the following command:

```
# fsadm -F vxfs -b newsize -r device_name mount_point
```

The /etc/fstab file

In the `/etc/fstab` file, do not specify any cluster file systems to mount-at-boot because mounts initiated from `fstab` occur before cluster configuration begins. For cluster mounts, use the VCS configuration file to determine which file systems to enable following a reboot.

When the Storage Foundation Cluster File System primary fails

If the server on which the SFCFS primary is running fails, the remaining cluster nodes elect a new primary. The new primary reads the file system intent log and completes any metadata updates that were in process at the time of the failure. Application I/O from other nodes may block during this process and cause a delay. When the file system is again consistent, application processing resumes.

Because nodes using a cluster file system in secondary node do not update file system metadata directly, failure of a secondary node does not require metadata repair. SFCFS recovery from secondary node failure is therefore faster than from primary node failure.

Distribute the load on a cluster

Distributing the workload in a cluster provides performance and failover advantages.

For example, if you have eight file systems and four nodes, designating two file systems per node as the primary would be beneficial. Primaryship is determined by which node first mounts the file system. You can also use the `fsclustadm` to designate a SFCFS primary. The `fsclustadm setprimary` command can also define the order in which primaryship is assumed if the current primary fails. After setup, the policy is in effect as long as one or more nodes in the cluster have the file system mounted.

Storage Checkpoints on Storage Foundation Cluster File System

See the *Veritas Storage Foundation Advanced Features Administrator's Guide* for how to create and maintain Storage Checkpoints. The creation of Storage Checkpoints works the same on CFS file systems as they do on local mount file systems.

Snapshots on Storage Foundation Cluster File System

A snapshot provides a consistent point-in-time image of a VxFS file system. A snapshot can be accessed as a read-only mounted file system to perform efficient

online backups of the file system. Snapshots implement copy-on-write semantics that incrementally copy data blocks when they are overwritten on the snapped file system.

See the *Veritas Storage Foundation Advanced Features Administrator's Guide*.

Snapshots for cluster file systems extend the same copy-on-write mechanism for the I/O originating from any node in the cluster.

Cluster snapshot characteristics

A cluster snapshot has the following characteristics:

- A snapshot for a cluster mounted file system can be mounted on any node in a cluster. The file system can be a primary, secondary, or secondary-only. A stable image of the file system is provided for writes from any node.
- Multiple snapshots of a cluster file system can be mounted on the same or a different node in a cluster.
- A snapshot is accessible only on the node mounting a snapshot. The snapshot device cannot be mounted on two different nodes simultaneously.
- The device for mounting a snapshot can be a local disk or a shared volume. A shared volume is used exclusively by a snapshot mount and is not usable from other nodes in a cluster as long as the snapshot is active on that device.
- On the node mounting a snapshot, the snapped file system cannot be unmounted while the snapshot is mounted.
- A SFCFS snapshot ceases to exist if it is unmounted or the node mounting the snapshot fails. A snapshot, however, is not affected if any other node leaves or joins the cluster.
- A snapshot of a read-only mounted file system cannot be taken. It is possible to mount snapshot of a cluster file system only if the snapped cluster file system is mounted with the `crw` option.

Performance considerations

Mounting a snapshot file system for backup increases the load on the system because of the resources used to perform copy-on-writes and to read data blocks from the snapshot. In this situation, cluster snapshots can be used to do off-host backups. Off-host backups reduce the load of a backup application from the primary server. Overhead from remote snapshots is small when compared to overall snapshot overhead. Therefore, running a backup application by mounting a snapshot from a relatively less loaded node is beneficial to overall cluster performance.

Creating a snapshot on a Storage Foundation Cluster File System

To create and mount a snapshot on a two-node cluster using SFCFS administrative interface commands.

To create a snapshot on a cluster file system

- 1 To create a VxFS file system on a shared VxVM volume, type the following command:

```
# mkfs -F vxfs /dev/vx/rdisk/cfsdg/vol1

version 7 layout
15728640 sectors, 15728640 blocks of size 1024, log size 16384 blocks
largefiles supported
```

- 2 To mount the file system on all nodes, type the following commands:

```
# cfsmntadm add cfsdg vol1 /mnt1 all=cluster
# cfsmount /mnt1
```

The `cfsmntadm` command adds an entry to the cluster manager configuration, then the `cfsmount` command mounts the file system on all nodes.

- 3 Add the snapshot on a previously created volume (`snapvol` in this example) to the cluster manager configuration. For example:

```
# cfsmntadm add snapshot cfsdg snapvol /mnt1 /mnt1snap system01=ro
```

The snapshot of a cluster file system is accessible only on the node where it is created; the snapshot file system itself cannot be cluster mounted.

- 4 Create and locally mount the snapshot file system on `system01`, type the following command:

```
# cfsmount /mnt1snap
```

- 5 A snapped file system cannot be unmounted until all of its snapshots are unmounted. Unmount and destroy the snapshot before trying to unmount the snapped cluster file system, type the following command:

```
# cfsunmount /mnt1snap
```

Administering VCS

This section provides instructions for the following VCS administration tasks:

- Viewing available Veritas devices and drivers
- Verifying VCS configuration
See “[Verifying VCS configuration](#)” on page 64.
- Starting and stopping VCS

If you encounter issues while administering VCS, refer to the troubleshooting section for assistance.

Configuring VCS to start Oracle with a specified Pfile

If you want to configure VCS such that Oracle starts with a specified Pfile, modify the `main.cf` file for the Oracle group as follows:

Verifying VCS configuration

To verify the VCS configuration:

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

Starting and stopping VCS

To start VCS on each node:

```
# hastart
```

To stop VCS on each node:

```
# hastop -local
```

You can also use the command `hastop -all`; however, make sure that you wait for port 'h' to close before restarting VCS.

Configuring destination-based load balancing for LLT

Destination-based load balancing for LLT is turned off by default. Symantec recommends destination-based load balancing when the cluster setup has more than two nodes and more active LLT ports.

To configure destination-based load balancing for LLT

- ◆ Run the following command to configure destination-based load balancing:

```
lltconfig -F linkburst: 0
```

Administering CVM

Listing all the CVM shared disks

You can use the following command to list all the CVM shared disks:

```
# vxdisk -o alldgs list |grep shared
```

Establishing CVM cluster membership manually

In most cases you do not have to start CVM manually; it normally starts when VCS is started.

Run the following command to start CVM manually:

```
# vxclustadm -m vcs -t gab startnode
```

```
vxclustadm: initialization completed
```

Note that `vxclustadm` reads `main.cf` for cluster configuration information and is therefore not dependent upon VCS to be running. You do not need to run the `vxclustadm startnode` command as normally the `hastart` (VCS start) command starts CVM automatically.

To verify whether CVM is started properly:

```
# vxclustadm nidmap
```

Name	CVM Nid	CM Nid	State
galaxy	0	0	Joined: Master
nebula	1	1	Joined: Slave

Changing the CVM master manually

You can change the CVM master manually from one node in the cluster to another node, while the cluster is online. CVM migrates the master node, and reconfigures the cluster.

Symantec recommends that you switch the master when the cluster is not handling VxVM configuration changes or cluster reconfiguration operations. In most cases, CVM aborts the operation to change the master, if CVM detects that any configuration changes are occurring in the VxVM or the cluster. After the master change operation starts reconfiguring the cluster, other commands that require configuration changes will fail.

See [“Errors during CVM master switching”](#) on page 67.

To change the master online, a cluster protocol version 100 or later is required.

To change the CVM master manually

- 1 To view the current master, use one of the following commands:

```
# vxclustadm nidmap
Name                CVM Nid    CM Nid    State
galaxy              0          0         Joined: Slave
nebula              1          1         Joined: Master

# vxdctl -c mode
mode: enabled: cluster active - MASTER
master: nebula
```

In this example, the CVM master is nebula.

- 2 From any node on the cluster, run the following command to change the CVM master:

```
# vxclustadm setmaster nodename
```

where *nodename* specifies the name of the new CVM master.

The following example shows changing the master on a cluster from nebula to galaxy:

```
# vxclustadm setmaster galaxy
```

- 3 To monitor the master switching, use the following command:

```
# vxclustadm -v nodestate
state: cluster member
      nodeId=0
      masterId=0
      neighborId=1
      members[0]=0xf
      joiners[0]=0x0
      leavers[0]=0x0
      members[1]=0x0
      joiners[1]=0x0
      leavers[1]=0x0
      reconfig_seqnum=0x9f9767
      vxfen=off
state: master switching in progress
reconfig: vxconfigd in join
```

In this example, the state indicates that master is being changed.

- 4 To verify whether the master has successfully changed, use one of the following commands:

```
# vxclustadm nidmap
Name           CVM Nid   CM Nid   State
galaxy         0         0        Joined: Master
nebula         1         1        Joined: Slave

# vxctl -c mode
mode: enabled: cluster active - MASTER
master: galaxy
```

Errors during CVM master switching

Symantec recommends that you switch the master when the cluster is not handling VxVM or cluster configuration changes.

In most cases, CVM aborts the operation to change the master, if CVM detects any configuration changes in progress. CVM logs the reason for the failure into the system logs. In some cases, the failure is displayed in the `vxclustadm setmaster` output as follows:

```
# vxclustadm setmaster galaxy
VxVM vxclustadm ERROR V-5-1-0 Master switching, a reconfiguration or
```

```
a transaction is in progress.  
Try again
```

In some cases, if the master switching operation is interrupted by another reconfiguration operation, the master change fails. In this case, the existing master remains the master of the cluster. After the reconfiguration is complete, reissue the `vxclustadm setmaster` command to change the master.

If the master switching operation has started the reconfiguration, any command that initiates a configuration change fails with the following error:

```
Node processing a master-switch request. Retry operation.
```

If you see this message, retry the command after the master switching has completed.

Importing a shared disk group manually

You can use the following command to manually import a shared disk group:

```
# vxvdg -s import dg_name
```

Deporting a shared disk group manually

You can use the following command to manually deport a shared disk group:

```
# vxvdg deport dg_name
```

Note that the deport of a shared disk group removes the SCSI-3 PGR keys on the disks. It also removes the 'shared' flag on the disks.

Starting shared volumes manually

Following a manual CVM shared disk group import, the volumes in the disk group need to be started manually, as follows:

```
# vxvol -g dg_name startall
```

To verify that the volumes are started, run the following command:

```
# vxprint -htrg dg_name | grep ^v
```

Evaluating the state of CVM ports

CVM kernel (vxio driver) uses port 'v' for kernel messaging and port 'w' for vxconfigd communication between the cluster nodes. The following command displays the state of CVM ports:

```
# gabconfig -a | egrep "Port [vw]"
```

Verifying if CVM is running in an SFCFS cluster

You can use the following options to verify whether CVM is up or not in an SFCFS cluster.

The following output is displayed on a node that is not a member of the cluster:

```
# vxdctl -c mode
mode: enabled: cluster inactive
# vxclustadm -v nodestate
state: out of cluster
```

On the master node, the following output is displayed:

```
# vxdctl -c mode
mode: enabled: cluster active - MASTER
master: galaxy
```

On the slave nodes, the following output is displayed:

```
# vxdctl -c mode
mode: enabled: cluster active - SLAVE
master: nebula
```

The following command lets you view all the CVM nodes at the same time:

```
# vxclustadm nidmap
```

Name	CVM Nid	CM Nid	State
galaxy	0	0	Joined: Master
nebula	1	1	Joined: Slave

Verifying CVM membership state

The state of CVM can be verified as follows:

```
# vxclustadm -v nodestate
state: joining
```

```
nodeId=0
masterId=0
neighborId=1
members=0x3
joiners=0x0
leavers=0x0
reconfig_seqnum=0x72a10b
vxfen=on
```

The state indicates that CVM has completed its kernel level join and is in the middle of vxconfigd level join.

The `vxctl -c mode` command indicates whether a node is a CVM master or CVM slave.

Verifying the state of CVM shared disk groups

You can use the following command to list the shared disk groups currently imported in the SFCFS cluster:

```
# vxkg list |grep shared

oradatadg enabled,shared 1052685125.1485.galaxy
```

Verifying the activation mode

In an SFCFS cluster, the activation of shared disk group should be set to “shared-write” on each of the cluster nodes.

To verify whether the “shared-write” activation is set:

```
# vxkg list dg_name |grep activation

local-activation: shared-write
```

If “shared-write” activation is not set, run the following command:

```
# vxkg -g dg_name set activation=sw
```

CVM log files

The `/var/VRTSvcs/log` directory contains the agent log files.

```
# cd /var/VRTSvcs/log
# ls -l *CVM* engine_A.log
CVMCluster_A.log      # CVM Agent log
CVMVolDg_A.log        # CVM VolDg Agent log
```

```
CVMVxconfigd_A.log # CVM vxconfigd Agent log  
engine_A.log      # VCS log
```

You can use the `cmdlog` file to view the list of CVM commands that have been executed. The file is located at `/var/adm/vx/cmdlog`.

See the *Veritas Volume Manager Administrator's Guide* for more information.

Administering ODM

This section provides instructions for the following ODM administration tasks:

- [Verifying the ODM port](#)
- [Starting ODM](#)

If you encounter issues while administering ODM, refer to the troubleshooting section for assistance.

Verifying the ODM port

It is recommended to enable ODM in SFCFS. Run the following command to verify that ODM is running:

```
# gabconfig -a | grep "Port d"
```

Starting ODM

The following procedure provides instructions for starting ODM.

To start ODM

- ◆ Run the following command:

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

Stopping ODM

The following procedure provides instructions for stopping ODM.

To stop ODM

- 1 On all nodes, issue the following command:

```
# hastop -local
```

- 2 Next, issue the following command:

```
#!/sbin/init.d/odm stop
```

Note: The administrator does not usually need to stop or start ODM. Normally, ODM is stopped during `shutdown -r`. ODM is started while rebooting, going to multi-user mode.

Administering I/O Fencing

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

About administering I/O fencing

The I/O fencing feature provides the following utilities that are available through the VRTSvxfen depot:

<code>vxfentsthdw</code>	Tests hardware for I/O fencing See “About the vxfentsthdw utility” on page 73.
<code>vxfenconfig</code>	Configures and unconfigures I/O fencing Checks the list of coordinator disks used by the vxfen driver.
<code>vxfenadm</code>	Displays information on I/O fencing operations and manages SCSI-3 disk registrations and reservations for I/O fencing See “About the vxfenadm utility” on page 80.
<code>vxfenclearpre</code>	Removes SCSI-3 registrations and reservations from disks See “About the vxfenclearpre utility” on page 85.
<code>vxfenswap</code>	Replaces coordinator points without stopping I/O fencing See “About the vxfenswap utility” on page 88.
<code>vxfendisk</code>	Generates the list of paths of disks in the diskgroup. This utility requires that Veritas Volume Manager is installed and configured.

The I/O fencing commands reside in the `/opt/VRTS/bin` folder. Make sure you added this folder path to the `PATH` environment variable.

Refer to the corresponding manual page for more information on the commands.

About the `vxfcntlsthdw` utility

You can use the `vxfcntlsthdw` utility to verify that shared storage arrays to be used for data support SCSI-3 persistent reservations and I/O fencing. During the I/O fencing configuration, the testing utility is used to test a single disk. The utility has other options that may be more suitable for testing storage devices in other configurations. You also need to test coordinator disk groups.

See *Veritas Storage Foundation Cluster File System Installation Guide* to set up I/O fencing.

The utility, which you can run from one system in the cluster, tests the storage used for data by setting and verifying SCSI-3 registrations on the disk or disks you specify, setting and verifying persistent reservations on the disks, writing data to the disks and reading it, and removing the registrations from the disks.

Refer also to the `vxfcntlsthdw(1M)` manual page.

About general guidelines for using `vxfcntlsthdw` utility

Review the following guidelines to use the `vxfcntlsthdw` utility:

- The utility requires two systems connected to the shared storage.

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

- The two nodes must have `ssh` (default) or `remsh` communication. If you use `remsh`, launch the `vxfcntlsthdw` utility with the `-n` option.
After completing the testing process, you can remove permissions for communication and restore public network connections.
- To ensure both systems are connected to the same disk during the testing, you can use the `vxfenadm -i diskpath` command to verify a disk's serial number. See [“Verifying that the nodes see the same disk”](#) on page 84.
- For disk arrays with many disks, use the `-m` option to sample a few disks before creating a disk group and using the `-g` option to test them all.
- The utility indicates a disk can be used for I/O fencing with a message resembling:

The disk /dev/vx/rmp/clt1d0 is ready to be configured for I/O Fencing on node galaxy

If the utility does not show a message stating a disk is ready, verification has failed.

- If the disk you intend to test has existing SCSI-3 registration keys, the test issues a warning before proceeding.

About the vxfcntlhdw command options

[Table 3-2](#) describes the methods that the utility provides to test storage devices.

Table 3-2 vxfcntlhdw options

vxfcntlhdw option	Description	When to use
-n	Utility uses remsh for communication.	Use when remsh is used for communication.
-r	Non-destructive testing. Testing of the disks for SCSI-3 persistent reservations occurs in a non-destructive way; that is, there is only testing for reads, not writes. May be used with -m, -f, or -g options.	Use during non-destructive testing. See “Performing non-destructive testing on the disks using the -r option” on page 77.
-t	Testing of the return value of SCSI TEST UNIT (TUR) command under SCSI-3 reservations. A warning is printed on failure of TUR testing.	When you want to perform TUR testing.
-d	Use DMP devices. May be used with -c or -g options.	By default, the script picks the DMP paths for disks in the diskgroup. If you want the script to use the raw paths for disks in the diskgroup, use the -w option.
-w	Use raw devices. May be used with -c or -g options.	With the -w option, the script picks the raw paths for disks in the diskgroup. By default, the script uses the -d option to pick up the DMP paths for disks in the disk group.

Table 3-2 vxfsentsthdw options (*continued*)

vxfsentsthdw option	Description	When to use
-c	Utility tests the coordinator disk group prompting for systems and devices, and reporting success or failure.	For testing disks in coordinator disk group. See “ Testing the coordinator disk group using vxfsentsthdw -c option ” on page 75.
-m	Utility runs manually, in interactive mode, prompting for systems and devices, and reporting success or failure. May be used with -r and -t options. -m is the default option.	For testing a few disks or for sampling disks in larger arrays. See “ Testing the shared disks using the vxfsentsthdw -m option ” on page 77.
-f <i>filename</i>	Utility tests system/device combinations listed in a text file. May be used with -r and -t options.	For testing several disks. See “ Testing the shared disks listed in a file using the vxfsentsthdw -f option ” on page 79.
-g <i>disk_group</i>	Utility tests all disk devices in a specified disk group. May be used with -r and -t options.	For testing many disks and arrays of disks. Disk groups may be temporarily created for testing purposes and destroyed (ungrouped) after testing. See “ Testing all the disks in a disk group using the vxfsentsthdw -g option ” on page 79.

Testing the coordinator disk group using vxfsentsthdw -c option

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

The procedure in this section uses the following disks for example:

- From the node galaxy, the disks are /dev/vx/rdmp/c1t1d0, /dev/vx/rdmp/c2t1d0, and /dev/vx/rdmp/c3t1d0.
- From the node nebula, the same disks are seen as /dev/vx/rdmp/c4t1d0, /dev/vx/rdmp/c5t1d0, and /dev/vx/rdmp/c6t1d0.

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

To test the coordinator disk group using `vxfcntlsthaw -c`

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

```
# vxfcntlsthaw -c vxfcntlsthaw
```

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

```
Enter the first node of the cluster: galaxy
```

```
Enter the second node of the cluster: nebula
```

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

```
ALL tests on the disk /dev/vx/rdmp/c1t1d0 have PASSED.
```

```
The disk is now ready to be configured for I/O Fencing on node  
galaxy as a COORDINATOR DISK.
```

```
ALL tests on the disk /dev/vx/rdmp/c4t1d0 have PASSED.
```

```
The disk is now ready to be configured for I/O Fencing on node  
nebula as a COORDINATOR DISK.
```

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

Removing and replacing a failed disk

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

To remove and replace a failed disk

- 1 Use the `vxdiskadm` utility to remove the failed disk from the disk group.
Refer to the *Veritas Volume Manager Administrator's Guide*.
- 2 Add a new disk to the node, initialize it, and add it to the coordinator disk group.
See the *Veritas Storage Foundation Cluster File System Installation Guide* for instructions to initialize disks for I/O fencing and to set up coordinator disk groups.
If necessary, start the disk group.
See the *Veritas Volume Manager Administrator's Guide* for instructions to start the disk group.
- 3 Retest the disk group.
See [“Testing the coordinator disk group using `vxfcntlsthdw -c` option”](#) on page 75.

Performing non-destructive testing on the disks using the `-r` option

You can perform non-destructive testing on the disk devices when you want to preserve the data.

To perform non-destructive testing on disks

- ◆ To test disk devices containing data you want to preserve, you can use the `-r` option with the `-m`, `-f`, or `-g` options.

For example, to use the `-m` option and the `-r` option, you can run the utility as follows:

```
# vxfcntlsthdw -rm
```

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

Testing the shared disks using the `vxfcntlsthdw -m` option

Review the procedure to test the shared disks. By default, the utility uses the `-m` option.

This procedure uses the `/dev/vx/rmdp/c1t1d0` disk in the steps.

If the utility does not show a message stating a disk is ready, verification has failed. Failure of verification can be the result of an improperly configured disk array. It can also be caused by a bad disk.

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

```
The disk /dev/vx/rdmp/c1t1d0 is ready to be configured for  
I/O Fencing on node galaxy
```

Note: For A/P arrays, run the `vxfcntlshdw` command only on active enabled paths.

To test disks using `vxfcntlshdw` script

- 1 Make sure system-to-system communication is functioning properly.
- 2 From one node, start the utility.

```
# vxfcntlshdw [-n]
```

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

```
***** WARNING!!!!!!!!!! *****  
THIS UTILITY WILL DESTROY THE DATA ON THE DISK!!  
  
Do you still want to continue : [y/n] (default: n) y  
Enter the first node of the cluster: galaxy  
Enter the second node of the cluster: nebula
```

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

```
Enter the disk name to be checked for SCSI-3 PGR on node  
galaxy in the format: /dev/vx/rdmp/cxtxdx  
/dev/vx/rdmp/c2t13d0  
Enter the disk name to be checked for SCSI-3 PGR on node  
nebula in the format: /dev/vx/rdmp/cxtxdx  
Make sure it's the same disk as seen by nodes galaxy and nebula  
/dev/vx/rdmp/c2t13d0
```

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

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

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

```
ALL tests on the disk /dev/vx/rdmp/c1t1d0 have PASSED
The disk is now ready to be configured for I/O Fencing on node
galaxy
...
Removing test keys and temporary files, if any ...
.
.
```

- 7 Run the `vxfcntlsthdw` utility for each disk you intend to verify.

Testing the shared disks listed in a file using the `vxfcntlsthdw -f` option

Use the `-f` option to test disks that are listed in a text file. Review the following example procedure.

To test the shared disks listed in a file

- 1 Create a text file `disks_test` to test two disks shared by systems `galaxy` and `nebula` that might resemble:

```
galaxy /dev/vx/rdmp/c2t2d1 nebula /dev/vx/rdmp/c3t2d1
galaxy /dev/vx/rdmp/c2t2d1 nebula /dev/vx/rdmp/c3t2d1
```

where the first disk is listed in the first line and is seen by `galaxy` as `/dev/vx/rdmp/c2t2d1` and by `nebula` as `/dev/vx/rdmp/c3t2d1`. The other disk, in the second line, is seen as `/dev/vx/rdmp/c2t2d2` from `galaxy` and `/dev/vx/rdmp/c3t2d2` from `nebula`. Typically, the list of disks could be extensive.

- 2 To test the disks, enter the following command:

```
# vxfcntlsthdw -f disks_test
```

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

Testing all the disks in a disk group using the `vxfcntlsthdw -g` option

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

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

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

To test all the disks in a diskgroup

- 1 Create a diskgroup for the disks that you want to test.
- 2 Enter the following command to test the diskgroup test_disks_dg:

```
# vxfentsthdw -g test_disks_dg
```

The utility reports the test results one disk at a time.

Testing a disk with existing keys

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

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

```
***** WARNING!!!!!!!!!! *****
```

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

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

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

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

About the vxfenadm utility

Administrators can use the vxfenadm command to troubleshoot and test fencing configurations.

The command's options for use by administrators are as follows:

-s read the keys on a disk and display the keys in numeric, character, and node format

Note: The -g and -G options are deprecated. Use the -s option.

-i read SCSI inquiry information from device

-m	register with disks
-n	make a reservation with disks
-p	remove registrations made by other systems
-r	read reservations
-x	remove registrations

Refer to the `vxfsadm(1m)` manual page for a complete list of the command options.

About the I/O fencing registration key format

The keys that the `vxfs` driver registers on the data disks and the coordinator disks consist of eight bytes. The key format is different for the coordinator disks and data disks.

The key format of the coordinator disks is as follows:

Byte	0	1	2	3	4	5	6	7
Value	V	F	cID 0x	cID 0x	cID 0x	cID 0x	nID 0x	nID 0x

where:

- VF is the unique identifier that carves out a namespace for the keys (consumes two bytes)
- cID 0x is the LLT cluster ID in hexadecimal (consumes four bytes)
- nID 0x is the LLT node ID in hexadecimal (consumes two bytes)

The `vxfs` driver uses this key format in both `scsi3` mode and customized mode of I/O fencing.

The key format of the data disks that are configured as failover disk groups under VCS is as follows:

Byte	0	1	2	3	4	5	6	7
Value	A+nID	V	C	S				

where nID is the LLT node ID

For example: If the node ID is 1, then the first byte has the value as B ('A' + 1 = B).

The key format of the data disks configured as parallel disk groups under CVM is as follows:

Byte	0	1	2	3	4	5	6	7
Value	A+nID	P	G	R	DGcount	DGcount	DGcount	DGcount

where DGcount is the count of disk group in the configuration

Displaying the I/O fencing registration keys

You can display the keys that are currently assigned to the disks using the `vxfenadm` command.

The variables such as `disk_7`, `disk_8`, and `disk_9` in the following procedure represent the disk names in your setup.

To display the I/O fencing registration keys

- 1 To display the key for the disks, run the following command:

```
# vxfenadm -s disk_name
```

For example:

- To display the key for the coordinator disk `/dev/vx/rmdp/c1t1d0` from the system with node ID 1, enter the following command:

```
# vxfenadm -s /dev/vx/rmdp/c1t1d0
key[1]:
  [Numeric Format]: 86,70,68,69,69,68,48,48
  [Character Format]: VFDEED00
* [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
```

The `-s` option of `vxfenadm` displays all eight bytes of a key value in three formats. In the numeric format,

- The first two bytes, represent the identifier VF, contains the ASCII value 86, 70.
- The next four bytes contain the ASCII value of the cluster ID 57069 encoded in hex (0xDEED) which are 68, 69, 69, 68.
- The remaining bytes contain the ASCII value of the node ID 0 (0x00) which are 48, 48. Node ID 1 would be 01 and node ID 10 would be 0A. An asterisk before the Node Format indicates that the `vxfenadm` command is run from the node of a cluster where LLT is configured and running.

- To display the keys on a CVM parallel disk group:

```
# vxfenadm -s /dev/vx/rmdp/disk_7
```

```
Reading SCSI Registration Keys...
```

```

Device Name: /dev/vx/rdmp/disk_7
Total Number Of Keys: 1
key[0]:
  [Numeric Format]: 66,80,71,82,48,48,48,48
  [Character Format]: BPGR0001
  [Node Format]: Cluster ID: unknown Node ID: 1 Node Name: nebula

```

■ To display the keys on a VCS failover disk group:

```
# vxfenadm -s /dev/vx/rdmp/disk_8
```

```
Reading SCSI Registration Keys...
```

```

Device Name: /dev/vx/rdmp/disk_8
Total Number Of Keys: 1
key[0]:
  [Numeric Format]: 65,86,67,83,0,0,0,0
  [Character Format]: AVCS
  [Node Format]: Cluster ID: unknown Node ID: 0 Node Name: galaxy

```

2 To display the keys that are registered in all the disks specified in a disk file:

```
# vxfenadm -s all -f disk_filename
```

For example:

To display all the keys on coordinator disks:

```
# vxfenadm -s all -f /etc/vxfentab
```

```

Device Name: /dev/vx/rdmp/disk_9
Total Number Of Keys: 2
key[0]:
  [Numeric Format]: 86,70,66,69,65,68,48,50
  [Character Format]: VFBEAD02
  [Node Format]: Cluster ID: 48813 Node ID: 2 Node Name: unknown
key[1]:
  [Numeric Format]: 86,70,68,69,69,68,48,48
  [Character Format]: VFDEED00
* [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy

```

You can verify the cluster ID using the `lltstat -C` command, and the node ID using the `lltstat -N` command. For example:

```
# lltstat -C
57069
```

If the disk has keys which do not belong to a specific cluster, then the `vxfenadm` command cannot look up the node name for the node ID and hence prints the node name as unknown. For example:

```
Device Name: /dev/vx/rdmp/disk_7
Total Number Of Keys: 1
key[0]:
  [Numeric Format]: 86,70,45,45,45,45,48,49
  [Character Format]: VF----01
  [Node Format]: Cluster ID: unknown Node ID: 1 Node Name: nebula
```

For disks with arbitrary format of keys, the `vxfenadm` command prints all the fields as unknown. For example:

```
[Numeric Format]: 65,66,67,68,49,50,51,45
[Character Format]: ABCD123-
[Node Format]: Cluster ID: unknown Node ID: unknown
Node Name: unknown
```

Verifying that the nodes see the same disk

To confirm whether a disk (or LUN) supports SCSI-3 persistent reservations, two nodes must simultaneously have access to the same disks. Because a shared disk is likely to have a different name on each node, check the serial number to verify the identity of the disk. Use the `vxfenadm` command with the `-i` option to verify that the same serial number for the LUN is returned on all paths to the LUN.

For example, an EMC array is accessible by the `/dev/vx/rdmp/c2t13d0` path on node A and by the `/dev/vx/rdmp/c2t11d0` path on node B.

To verify that the nodes see the same disks

- 1 Verify the connection of the shared storage for data to two of the nodes on which you installed SFCFS.
- 2 From node A, enter the following command:

```
# vxfsadm -i /dev/vx/rdmp/c2t13d0

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

The same serial number information should appear when you enter the equivalent command on node B using the `/dev/vx/rdmp/c2t11d0` path.

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

```
# vxfsadm -i /dev/vx/rdmp/c2t1d0

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

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

About the `vxfsclearpre` utility

You can use the `vxfsclearpre` utility to remove SCSI-3 registrations and reservations on the disks.

See [“Removing preexisting keys”](#) on page 85.

Removing preexisting keys

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

You can also use this procedure to remove the registration and reservation keys created by another node from a disk.

To clear keys after split-brain

- 1 Stop VCS on all nodes.

```
# hastop -all
```

- 2 Make sure that the port h is closed on all the nodes. Run the following command on each node to verify that the port h is closed:

```
# gabconfig -a
```

Port h must not appear in the output.

- 3 Stop I/O fencing on all nodes. Enter the following command on each node:

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

- 4 If you have any applications that run outside of VCS control that have access to the shared storage, then shut down all other nodes in the cluster that have access to the shared storage. This prevents data corruption.

- 5 Start the vxfcntlpre script:

```
# vxfcntlpre
```

6 Read the script's introduction and warning. Then, you can choose to let the script run.

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

In some cases, informational messages resembling the following may appear on the console of one of the nodes in the cluster when a node is ejected from a disk/LUN. You can ignore these informational messages.

```
<date> <system name> scsi: WARNING: /sbus@3,0/lpfs@0,0/  
sd@0,1(sd91):  
<date> <system name> Error for Command: <undecoded  
cmd 0x5f> Error Level: Informational  
<date> <system name> scsi: Requested Block: 0 Error Block 0  
<date> <system name> scsi: Vendor: <vendor> Serial Number:  
0400759B006E  
<date> <system name> scsi: Sense Key: Unit Attention  
<date> <system name> scsi: ASC: 0x2a (<vendor unique code  
0x2a>), ASCQ: 0x4, FRU: 0x0
```

The script cleans up the disks and displays the following status messages.

```
Cleaning up the coordinator disks...
```

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

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

```
You can retry starting fencing module. In order to  
restart the whole product, you might want to  
reboot the system.
```

7 Start the fencing module.

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

8 Start VCS on all nodes.

```
# hstart
```

About the vxfsnwap utility

The vxfsnwap utility allows you to replace coordinator disks in a cluster that is online. The utility verifies that the serial number of the new disks are identical on all the nodes and the new disks can support I/O fencing.

Refer to the vxfsnwap(1M) manual page.

You can replace the coordinator disks without stopping I/O fencing in the following cases:

- The disk becomes defective or inoperable and you want to switch to a new diskgroup.
See [“Replacing I/O fencing coordinator disks when the cluster is online”](#) on page 89.
See [“Replacing the coordinator diskgroup in a cluster that is online”](#) on page 91.
If you want to replace the coordinator disks when the cluster is offline, you cannot use the vxfsnwap utility. You must manually perform the steps that the utility does to replace the coordinator disks.
See [“Replacing defective disks when the cluster is offline”](#) on page 184.
- You want to switch the disk interface between raw devices and DMP devices.
- The keys that are registered on the coordinator disks are lost.
In such a case, the cluster might panic when a network partition occurs. You can replace the coordinator disks with the same disks using the vxfsnwap command. During the disk replacement, the missing keys register again without any risk of data corruption.
See [“Refreshing lost keys on coordinator disks”](#) on page 95.

You can also use the vxfsnwap utility to migrate between the disk-based and the server-based fencing without incurring application downtime in the Storage Foundation Cluster File System HA cluster.

See [“Migrating from disk-based to server-based fencing in an online cluster”](#) on page 112.

See [“Migrating from server-based to disk-based fencing in an online cluster”](#) on page 114.

If the vxfsnwap operation is unsuccessful, then you can use the `-a cancel` of the vxfsnwap command to manually roll back the changes that the vxfsnwap utility does.

- For disk-based fencing, use the vxfsnwap `-g diskgroup -a cancel` command to cancel the vxfsnwap operation.
You must run this command if a node fails during the process of disk replacement, or if you aborted the disk replacement.

- For server-based fencing, use the `vxfsenwap -a cancel` command to cancel the `vxfsenwap` operation.

Replacing I/O fencing coordinator disks when the cluster is online

Review the procedures to add, remove, or replace one or more coordinator disks in a cluster that is operational.

Warning: The cluster might panic if any node leaves the cluster membership before the `vxfsenwap` script replaces the set of coordinator disks.

To replace a disk in a coordinator diskgroup when the cluster is online

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxfsenadm -d

I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Import the coordinator disk group.

The file `/etc/vxfsendg` includes the name of the disk group (typically, `vxfsencoordg`) that contains the coordinator disks, so use the command:

```
# vxsdg -tfc import `cat /etc/vxfsendg`
```

where:

- t specifies that the disk group is imported only until the node restarts.
- f specifies that the import is to be done forcibly, which is necessary if one or more disks is not accessible.
- C specifies that any import locks are removed.

- 4 If your setup uses VRTSvxvm version 5.1 RP2 (or later) or 5.1 SP1 (or later), then skip to step 5. You need not set `coordinator=off` to add or remove disks. For other VxVM versions, perform this step:

Turn off the coordinator attribute value for the coordinator disk group.

```
# vxdg -g vxfencoordg set coordinator=off
```

- 5 To remove disks from the coordinator disk group, use the VxVM disk administrator utility `vxdiskadm`.
- 6 Perform the following steps to add new disks to the coordinator disk group:
 - Add new disks to the node.
 - Initialize the new disks as VxVM disks.
 - Check the disks for I/O fencing compliance.
 - Add the new disks to the coordinator disk group and set the coordinator attribute value as "on" for the coordinator disk group.

See the *Veritas Storage Foundation Cluster File System Installation Guide* for detailed instructions.

Note that though the disk group content changes, the I/O fencing remains in the same state.

- 7 From one node, start the `vxfenswap` utility. You must specify the diskgroup to the utility.

The utility performs the following tasks:

- Backs up the existing `/etc/vxfentab` file.
 - Creates a test file `/etc/vxfentab.test` for the diskgroup that is modified on each node.
 - Reads the diskgroup you specified in the `vxfenswap` command and adds the diskgroup to the `/etc/vxfentab.test` file on each node.
 - Verifies that the serial number of the new disks are identical on all the nodes. The script terminates if the check fails.
 - Verifies that the new disks can support I/O fencing on each node.
- 8 If the disk verification passes, the utility reports success and asks if you want to commit the new set of coordinator disks.

- 9 Review the message that the utility displays and confirm that you want to commit the new set of coordinator disks. Else skip to step 10.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfentab.test` file to the `/etc/vxfentab` file.

- 10 If you do not want to commit the new set of coordinator disks, answer `n`.
The `vxfsnwap` utility rolls back the disk replacement operation.

Replacing the coordinator diskgroup in a cluster that is online

You can also replace the coordinator diskgroup using the `vxfsnwap` utility. The following example replaces the coordinator disk group `vxfsncoorddg` with a new disk group `vxfsndg`.

To replace the coordinator diskgroup

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxfsnadm -d
```

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Find the name of the current coordinator diskgroup (typically `vxfsncoorddg`) that is in the `/etc/vxfsndg` file.

```
# cat /etc/vxfsndg
vxfsncoorddg
```

- 4 Find the alternative disk groups available to replace the current coordinator diskgroup.

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c4t0d1	auto:cdsdisk	-	(vxfendg)	online
c4t0d2	auto:cdsdisk	-	(vxfendg)	online
c4t0d3	auto:cdsdisk	-	(vxfendg)	online
c4t0d4	auto:cdsdisk	-	(vxfencoorddg)	online
c4t0d5	auto:cdsdisk	-	(vxfencoorddg)	online
c4t0d6	auto:cdsdisk	-	(vxfencoorddg)	online

- 5 Validate the new disk group for I/O fencing compliance. Run the following command:

```
# vxfentsthdw -c vxfendg
```

See [“Testing the coordinator disk group using vxfentsthdw -c option”](#) on page 75.

- 6 If the new disk group is not already deported, run the following command to deport the disk group:

```
# vxdg deport vxfendg
```

- 7 Make sure that the `/etc/vxfenmode` file is updated to specify the correct disk policy.

See the *Veritas Storage Foundation Cluster File System Installation Guide* for more information.

- 8 From any node, start the `vxfenswap` utility. For example, if `vxfendg` is the new diskgroup that you want to use as the coordinator diskgroup:

```
# vxfenswap -g vxfendg [-n]
```

The utility performs the following tasks:

- Backs up the existing `/etc/vxfentab` file.
- Creates a test file `/etc/vxfentab.test` for the diskgroup that is modified on each node.
- Reads the diskgroup you specified in the `vxfenswap` command and adds the diskgroup to the `/etc/vxfentab.test` file on each node.
- Verifies that the serial number of the new disks are identical on all the nodes. The script terminates if the check fails.

- Verifies that the new disk group can support I/O fencing on each node.
- 9 If the disk verification passes, the utility reports success and asks if you want to replace the coordinator disk group.
 - 10 Review the message that the utility displays and confirm that you want to replace the coordinator disk group. Else skip to step 13.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfentab.test` file to the `/etc/vxfentab` file.

The utility also updates the `/etc/vxfendg` file with this new diskgroup.

- 11 Set the coordinator attribute value as "on" for the new coordinator disk group.

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

Set the coordinator attribute value as "off" for the old disk group.

```
# vxdg -g vxfencoordg set coordinator=off
```

- 12 Verify that the coordinator disk group has changed.

```
# cat /etc/vxfendg
vxfendg
```

The swap operation for the coordinator disk group is complete now.

- 13 If you do not want to replace the coordinator disk group, answer n at the prompt.

The `vxfenswap` utility rolls back any changes to the coordinator diskgroup.

Adding disks from a recovered site to the coordinator diskgroup

In a campus cluster environment, consider a case where the primary site goes down and the secondary site comes online with a limited set of disks. When the primary site restores, the primary site's disks are also available to act as coordinator disks. You can use the `vxfenswap` utility to add these disks to the coordinator diskgroup.

To add new disks from a recovered site to the coordinator diskgroup

- 1** Make sure system-to-system communication is functioning properly.
- 2** Make sure that the cluster is online.

```
# vxfenadm -d
```

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3** Verify the name of the coordinator diskgroup.

```
# cat /etc/vxfendg
vxfencoorddg
```

- 4** Run the following command:

```
# vxdisk -o alldgs list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c1t1d0	auto:cdsdisk	-	(vxfencoorddg)	online
c2t1d0	auto	- -	offline	
c3t1d0	auto	- -	offline	

- 5** Verify the number of disks used in the coordinator diskgroup.

```
# vxfenconfig -l
```

```
I/O Fencing Configuration Information:
=====
Count                : 1
Disk List
Disk Name            Major Minor  Serial Number      Policy
/dev/vx/rdmp/c1t1d0      32  48  R450 00013154 0312      dmp
```

- 6 When the primary site comes online, start the `vxfsnwap` utility on any node in the cluster:

```
# vxfsnwap -g vxfsncooordg [-n]
```

- 7 Verify the count of the coordinator disks.

```
# vxfsnconfig -l
I/O Fencing Configuration Information:
=====
Single Disk Flag      : 0
Count                 : 3
Disk List
Disk Name             Major  Minor  Serial Number      Policy
/dev/vx/rdmp/c1t1d0   32   48   R450 00013154 0312    dmp
/dev/vx/rdmp/c2t1d0   32   32   R450 00013154 0313    dmp
/dev/vx/rdmp/c3t1d0   32   16   R450 00013154 0314    dmp
```

Refreshing lost keys on coordinator disks

If the coordinator disks lose the keys that are registered, the cluster might panic when a network partition occurs.

You can use the `vxfsnwap` utility to replace the coordinator disks with the same disks. The `vxfsnwap` utility registers the missing keys during the disk replacement.

To refresh lost keys on coordinator disks

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the cluster is online.

```
# vxfenadm -d
```

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Run the following command to view the coordinator disks that do not have keys:

```
# vxfenadm -s all -f /etc/vxfentab
```

```
Device Name: /dev/vx/rdmp/clt1d0
Total Number of Keys: 0
No keys...
...
```

- 4 On any node, run the following command to start the vxfenswap utility:

```
# vxfenswap -g vxfencoordg [-n]
```

- 5 Verify that the keys are atomically placed on the coordinator disks.

```
# vxfenadm -s all -f /etc/vxfentab
```

```
Device Name: /dev/vx/rdmp/clt1d0
Total Number of Keys: 4
...
```

About administering the coordination point server

This section describes how to perform administrative and maintenance tasks on the coordination point server (CP server).

For more information about the `cpsadm` command and the associated command options, see the `cpsadm(1M)` manual page.

Environment variables associated with the coordination point server

Table 3-3 describes the environment variables that are required for the `cpsadm` command. The `cpsadm` command detects these environment variables and uses their value when communicating with the CP server. They are used to authenticate and authorize the user.

Note: The environment variables are not required when the `cpsadm` command is run on the CP server. The environment variables are required when the `cpsadm` command is run on the SFCFS cluster nodes.

Table 3-3 cpsadm command environment variables

Environment variable	Description
CPS_USERNAME	This is the fully qualified username as configured in VxSS (vssat showcred).
CPS_DOMAINTYPE	One of the following values: <ul style="list-style-type: none"> ■ vx ■ unixpwd ■ nis ■ nisplus ■ ldap

The environment variables must be exported directly on the shell before running the `cpsadm` command. For example,

```
# export CPS_USERNAME=
# export CPS_DOMAINTYPE=
```

Additionally, the username and domaintype values are the same as those added onto the CP server. To view these values run the following `cpsadm` command:

```
# cpsadm -s cp_server -a list_users
```

Where *cp_server* is the CP server's virtual IP address or virtual hostname.

Adding and removing SFCFS cluster entries from the CP server database

- To add a SFCFS cluster to the CP server database

Type the following command:

```
# cpsadm -s cp_server -a add_clus -c cluster_name -u uuid
```

- To remove a SFCFS cluster from the CP server database

Type the following command:

```
# cpsadm -s cp_server -a rm_clus -u uuid
```

cp_server The CP server's virtual IP address or virtual hostname.

cluster_name The SFCFS cluster name.

uuid The UUID (Universally Unique ID) of the SFCFS cluster.

Adding and removing a SFCFS cluster node from the CP server database

- To add a SFCFS cluster node from the CP server database

Type the following command:

```
# cpsadm -s cp_server -a add_node -u uuid -n nodeid  
-h host
```

- To remove a SFCFS cluster node from the CP server database

Type the following command:

```
# cpsadm -s cp_server -a rm_node -u uuid -n nodeid
```

cp_server The CP server's virtual IP address or virtual hostname.

uuid The UUID (Universally Unique ID) of the SFCFS cluster.

nodeid The node id of the SFCFS cluster node.

host Hostname

Adding or removing CP server users

- To add a user

Type the following command:

```
# cpsadm -s cp_server -a add_user -e user_name -f user_role
-g domain_type -u uuid
```

- To remove a user

Type the following command:

```
# cpsadm -s cp_server -a rm_user -e user_name -g domain_type
```

<i>cp_server</i>	The CP server's virtual IP address or virtual hostname.
<i>user_name</i>	The user to be added to the CP server configuration.
<i>user_role</i>	The user role, either <code>cps_admin</code> or <code>cps_operator</code> .
<i>domain_type</i>	The domain type, for example <code>vx</code> , <code>unixpwd</code> , <code>nis</code> , etc.
<i>uuid</i>	The UUID (Universally Unique ID) of the SFCFS cluster.

Listing the CP server users

To list the CP server users

Type the following command:

```
# cpsadm -s cp_server -a list_users
```

Listing the nodes in all the SFCFS clusters

To list the nodes in all the SFCFS cluster

Type the following command:

```
# cpsadm -s cp_server -a list_nodes
```

Listing the membership of nodes in the SFCFS cluster

To list the membership of nodes in SFCFS cluster

Type the following command:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

<i>cp_server</i>	The CP server's virtual IP address or virtual hostname.
------------------	---

cluster_name The SFCFS cluster name.

Preempting a node

To preempt a node

Type the following command:

```
# cpsadm -s cp_server -a preempt_node -u uuid -n nodeid  
-v victim_node id
```

cp_server The CP server's virtual IP address or virtual hostname.

uuid The UUID (Universally Unique ID) of the SFCFS cluster.

nodeid The node id of the SFCFS cluster node.

victim_node id The victim node's node id.

Registering and unregistering a node

- To register a node

Type the following command:

```
# cpsadm -s cp_server -a reg_node -u uuid -n nodeid
```

- To unregister a node

Type the following command:

```
# cpsadm -s cp_server -a unreg_node -u uuid -n nodeid
```

cp_server The CP server's virtual IP address or virtual hostname.

uuid The UUID (Universally Unique ID) of the SFCFS cluster.

nodeid The nodeid of the SFCFS cluster node.

Enable and disable access for a user to a SFCFS cluster

- To enable access for a user to a SFCFS cluster

Type the following command:

```
# cpsadm -s cp_server -a add_clus_to_user -e user  
-f user_role -g domain_type -u uuid
```

- To disable access for a user to a SFCFS cluster

Type the following command:

```
# cpsadm -s cp_server -a rm_clus_from_user -e user_name  
-f user_role -g domain_type -u uuid
```

cp_server The CP server's virtual IP address or virtual hostname.

user_name The user name to be added to the CP server.

user_role The user role, either *cps_admin* or *cps_operator*.

domain_type The domain type, for example *vx*, *unixpwd*, *nis*, etc.

uuid The UUID (Universally Unique ID) of the SFCFS cluster

Stopping the CP server

To stop the CP server

Type the following command:

```
# cpsadm -s cp_server -a halt_cps
```

Checking the connectivity of CP servers

To check the connectivity of a CP server

Type the following command:

```
# cpsadm -s cp_server -a ping_cps
```

Taking a CP server database snapshot

To take a CP server database snapshot

Type the following command:

```
# cpsadm -s cp_server -a db_snapshot
```

Replacing coordination points for server-based fencing in an online cluster

Use the following procedure to perform a planned replacement of customized coordination points (CP servers or SCSI-3 disks) without incurring application downtime on an online SFCFS cluster.

Note: If multiple clusters share the same CP server, you must perform this replacement procedure in each cluster.

You can use the `vxfsnwap` utility to replace coordination points when fencing is running in customized mode in an online cluster, with `vxfen_mechanism=cps`. The utility also supports migration from server-based fencing (`vxfen_mode=customized`) to disk-based fencing (`vxfen_mode=scsi3`) and vice versa in an online cluster.

However, if the SFCFS cluster has fencing disabled (`vxfen_mode=disabled`), then you must take the cluster offline to configure disk-based or server-based fencing.

See [“Deployment and migration scenarios for CP server”](#) on page 106.

You can cancel the coordination point replacement operation at any time using the `vxfsnwap -a cancel` command.

See [“About the vxfsnwap utility”](#) on page 88.

To replace coordination points for an online cluster

- 1 Ensure that the SFCFS cluster nodes and users have been added to the new CP server(s). Run the following commands:

```
# cpsadm -s cpserver -a list_nodes  
# cpsadm -s cpserver -a list_users
```

If the SFCFS cluster nodes are not present here, prepare the new CP server(s) for use by the SFCFS cluster.

See the Veritas Storage Foundation Cluster File System Installation Guide for instructions.

- 2 Ensure that fencing is running on the cluster using the old set of coordination points and in customized mode.

For example, enter the following command:

```
# vxfenadm -d
```

The command returns:

```
I/O Fencing Cluster Information:  
=====  
Fencing Protocol Version: <version>  
Fencing Mode: Customized  
Cluster Members:  
* 0 (galaxy)  
1 (nebula)  
RFSM State Information:  
node 0 in state 8 (running)  
node 1 in state 8 (running)
```

- 3 Back up the `/etc/vxfenmode` file on each of the client cluster nodes.

- 4 Use a text editor to access `/etc/vxfenmode` and update the values to the new CP server (coordination points).

The values of the `/etc/vxfenmode` file have to be updated on all the nodes in the SFCFS cluster.

Review and if necessary, update the `vxfenmode` parameters for security, the coordination points, and if applicable to your configuration, `vxfendg`.

Refer to the text information within the `vxfenmode` file for additional information about these parameters and their new possible values.

- 5 Run the `vxfenswap` utility from one of the nodes of the cluster.

The `vxfenswap` utility requires secure ssh connection to all the cluster nodes. Use `-n` to use rsh instead of default ssh.

```
# vxfenswap [-n]
```

If validation of coordination points from all the nodes fails, the `vxfenswap` utility rolls back the coordination point replacement operation. Proceed to restore `/etc/vxfenmode` with the backed up file on all the VCS cluster nodes.

You are then prompted to commit the change. Enter `y` for yes.

Confirm the successful execution of the `vxfenswap` utility by checking the coordination points currently used by the `vxfen` driver.

For example, run the following command:

```
# vxfenconfig -l
```

Refreshing registration keys on the coordination points for server-based fencing

Replacing keys on a coordination point (CP server) when the SFCFS cluster is online involves refreshing that coordination point's registrations. You can perform a planned refresh of registrations on a CP server without incurring application downtime on the SFCFS cluster. You must refresh registrations on a CP server if the CP server agent issues an alert on the loss of such registrations on the CP server database.

The following procedure describes how to refresh the coordination point registrations.

To refresh the registration keys on the coordination points for server-based fencing

- 1 Ensure that the SFCFS cluster nodes and users have been added to the new CP server(s). Run the following commands:

```
# cpsadm -s cp_server -a list_nodes  
  
# cpsadm -s cp_server -a list_users
```

If the SFCFS cluster nodes are not present here, prepare the new CP server(s) for use by the SFCFS cluster.

See the Veritas Storage Foundation Cluster File System Installation Guide for instructions.

- 2 Ensure that fencing is running on the cluster in customized mode using the coordination points mentioned in the `/etc/vxfenmode` file.

For example, enter the following command:

```
# vxfenadm -d  
  
=====
```

```
Fencing Protocol Version: 201  
Fencing Mode: CUSTOMIZED  
Cluster Members:  
* 0 (galaxy)  
1 (nebula)  
RFSM State Information:  
node 0 in state 8 (running)  
node 1 in state 8 (running)
```

- 3 List the coordination points currently used by I/O fencing :

```
# vxfenconfig -l
```

4 Run the `vxfsnwap` utility from one of the nodes of the cluster.

The `vxfsnwap` utility requires secure ssh connection to all the cluster nodes. Use `-n` to use rsh instead of default ssh.

For example:

```
# vxfsnwap [-n]
```

The command returns:

```
VERITAS vxfsnwap version <version> <platform>  
The logfile generated for vxfsnwap is  
/var/VRTSvcs/log/vxfen/vxfsnwap.log.  
19156  
Please Wait...  
VXFEN vxfsnconfig NOTICE Driver will use customized fencing  
- mechanism cps  
Validation of coordination points change has succeeded on  
all nodes.  
You may commit the changes now.  
WARNING: This may cause the whole cluster to panic  
if a node leaves membership before the change is complete.
```

5 You are then prompted to commit the change. Enter `y` for yes.

The command returns a confirmation of successful coordination point replacement.

6 Confirm the successful execution of the `vxfsnwap` utility. If CP agent is configured, it should report ONLINE as it succeeds to find the registrations on coordination points. The registrations on the CP server and coordinator disks can be viewed using the `cpsadm` and `vxfsnadm` utilities respectively.

Note that a running online coordination point refreshment operation can be canceled at any time using the command:

```
# vxfsnwap -a cancel
```

Deployment and migration scenarios for CP server

[Table 3-4](#) describes the supported deployment and migration scenarios, and the procedures you must perform on the SFCFS cluster and the CP server.

Table 3-4 CP server deployment and migration scenarios

Scenario	CP server	SFCFS cluster	Action required
Setup of CP server for a SFCFS cluster for the first time	New CP server	New SFCFS cluster using CP server as coordination point	<p>On the designated CP server, perform the following tasks:</p> <ol style="list-style-type: none"> 1 Prepare to configure the new CP server. 2 Configure the new CP server. <p>On the SFCFS cluster nodes, configure server-based I/O fencing.</p> <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p>
Add a new SFCFS cluster to an existing and operational CP server	Existing and operational CP server	New SFCFS cluster	<p>On the SFCFS cluster nodes, configure server-based I/O fencing.</p> <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p>
Replace the coordination point from an existing CP server to a new CP server	New CP server	Existing SFCFS cluster using CP server as coordination point	<p>On the designated CP server, perform the following tasks:</p> <ol style="list-style-type: none"> 1 Prepare to configure the new CP server. 2 Configure the new CP server. 3 Prepare the new CP server for use by the SFCFS cluster. <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p> <p>On a node in the SFCFS cluster, run the <code>vxfsnswap</code> command to move to replace the CP server:</p> <p>See “Replacing coordination points for server-based fencing in an online cluster” on page 101.</p>
Replace the coordination point from an existing CP server to an operational CP server coordination point	Operational CP server	Existing SFCFS cluster using CP server as coordination point	<p>On the designated CP server, prepare to configure the new CP server manually.</p> <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p> <p>On a node in the SFCFS cluster, run the <code>vxfsnswap</code> command to move to replace the CP server:</p> <p>See “Replacing coordination points for server-based fencing in an online cluster” on page 101.</p>

Table 3-4 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFS cluster	Action required
<p>Enabling fencing in a SFCFS cluster with a new CP server coordination point</p>	<p>New CP server</p>	<p>Existing SFCFS cluster with fencing configured in disabled mode</p>	<p>Note: Migrating from fencing in disabled mode to customized mode incurs application downtime on the SFCFS cluster.</p> <p>On the designated CP server, perform the following tasks:</p> <ol style="list-style-type: none"> 1 Prepare to configure the new CP server. 2 Configure the new CP server 3 Prepare the new CP server for use by the SFCFS cluster <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p> <p>On the SFCFS cluster nodes, perform the following:</p> <ol style="list-style-type: none"> 1 Stop all applications, VCS, and fencing on the SFCFS cluster. 2 To stop VCS, use the following command (to be run on all the SFCFS cluster nodes): <pre># hstop -local</pre> 3 Stop fencing using the following command: <pre># /sbin/init.d/vxfen stop</pre> 4 Reconfigure I/O fencing on the SFCFS cluster. <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p>

Table 3-4 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFS cluster	Action required
Enabling fencing in a SFCFS cluster with an operational CP server coordination point	Operational CP server	Existing SFCFS cluster with fencing configured in disabled mode	<p>Note: Migrating from fencing in disabled mode to customized mode incurs application downtime.</p> <p>On the designated CP server, prepare to configure the new CP server.</p> <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for this procedure.</p> <p>On the SFCFS cluster nodes, perform the following tasks:</p> <ol style="list-style-type: none"> 1 Stop all applications, VCS, and fencing on the SFCFS cluster. 2 To stop VCS, use the following command (to be run on all the SFCFS cluster nodes): # hastop -local 3 Stop fencing using the following command: # /sbin/init.d/vxfen stop 4 Reconfigure fencing on the SFCFS cluster. <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p>

Table 3-4 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFS cluster	Action required
<p>Enabling fencing in a SFCFS cluster with a new CP server coordination point</p>	<p>New CP server</p>	<p>Existing SFCFS cluster with fencing configured in scsi3 mode</p>	<p>On the designated CP server, perform the following tasks:</p> <ol style="list-style-type: none"> 1 Prepare to configure the new CP server. 2 Configure the new CP server 3 Prepare the new CP server for use by the SFCFS cluster <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p> <p>Based on whether the cluster is online or offline, perform the following procedures:</p> <p>For a cluster that is online, perform the following task on the SFCFS cluster:</p> <ul style="list-style-type: none"> ◆ Run the <code>vxfsenwrap</code> command to migrate from disk-based fencing to the server-based fencing. See “Migrating from disk-based to server-based fencing in an online cluster” on page 112. <p>For a cluster that is offline, perform the following tasks on the SFCFS cluster:</p> <ol style="list-style-type: none"> 1 Stop all applications, VCS, and fencing on the SFCFS cluster. 2 To stop VCS, use the following command (to be run on all the SFCFS cluster nodes): <pre># hastop -local</pre> 3 Stop fencing using the following command: <pre># /sbin/init.d/vxfen stop</pre> 4 Reconfigure I/O fencing on the SFCFS cluster. <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p>

Table 3-4 CP server deployment and migration scenarios (*continued*)

Scenario	CP server	SFCFS cluster	Action required
Enabling fencing in a SFCFS cluster with an operational CP server coordination point	Operational CP server	Existing SFCFS cluster with fencing configured in disabled mode	<p>On the designated CP server, prepare to configure the new CP server.</p> <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for this procedure.</p> <p>Based on whether the cluster is online or offline, perform the following procedures:</p> <p>For a cluster that is online, perform the following task on the SFCFS cluster:</p> <ul style="list-style-type: none"> ◆ Run the <code>vxfsenwap</code> command to migrate from disk-based fencing to the server-based fencing. <p>See “Migrating from disk-based to server-based fencing in an online cluster” on page 112.</p> <p>For a cluster that is offline, perform the following tasks on the SFCFS cluster:</p> <ol style="list-style-type: none"> 1 Stop all applications, VCS, and fencing on the SFCFS cluster. 2 To stop VCS, use the following command (to be run on all the SFCFS cluster nodes): <ul style="list-style-type: none"> # <code>hastop -local</code> 3 Stop fencing using the following command: <ul style="list-style-type: none"> # <code>/sbin/init.d/vxfen stop</code> 4 Reconfigure fencing on the SFCFS cluster. <p>See the <i>Veritas Storage Foundation Cluster File System Installation Guide</i> for the procedures.</p>
Refreshing registrations of SFCFS cluster nodes on coordination points (CP servers/coordinator disks) without incurring application downtime	Operational CP server	Existing SFCFS cluster using the CP server as coordination point	<p>On the SFCFS cluster run the <code>vxfsenwap</code> command to refresh the keys on the CP server:</p> <p>See “Refreshing registration keys on the coordination points for server-based fencing” on page 104.</p>

Migrating from disk-based to server-based fencing in an online cluster

You can migrate between disk-based fencing and server-based fencing without incurring application downtime in the Storage Foundation Cluster File System HA clusters.

You can migrate from disk-based fencing to server-based fencing in the following cases:

- You want to leverage the benefits of server-based fencing.
- You want to replace faulty coordinator disks with CP servers.

Similarly, you can migrate from server-based fencing to disk-based fencing when you want to perform maintenance tasks on the CP server systems.

See [“Migrating from server-based to disk-based fencing in an online cluster”](#) on page 114.

To migrate from disk-based fencing to server-based fencing

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the Storage Foundation Cluster File System HA cluster is online and uses disk-based fencing.

```
# vxfenadm -d
```

For example, if Storage Foundation Cluster File System HA cluster uses disk-based fencing:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Make sure that you performed the following tasks on the designated CP server:
 - Preparing to configure the new CP server.
 - Configuring the new CP server
 - Preparing the new CP server for use by the SFCFS cluster

See the *Veritas Storage Foundation Cluster File System Installation Guide* for the procedures.

- 4 Create a new `/etc/vxfenmode.test` file on each Storage Foundation Cluster File System HA cluster node with the fencing configuration changes such as the CP server information.

Refer to the sample `vxfenmode` files in the `/etc/vxfen.d` folder.

- 5 From any node in the Storage Foundation Cluster File System HA cluster, start the `vxfenswap` utility:

```
# vxfenswap [-n]
```

- 6 Review the message that the utility displays and confirm whether you want to commit the change.

- If you do not want to commit the new fencing configuration changes, press Enter or answer **n** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) n
```

The `vxfenswap` utility rolls back the migration operation.

- If you want to commit the new fencing configuration changes, answer **y** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfenmode.test` file to the `/etc/vxfenmode` file.

- 7 After the migration is complete, verify the change in the fencing mode.

```
# vxfenadm -d
```

For example, after the migration from disk-based fencing to server-based fencing in the Storage Foundation Cluster File System HA cluster:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 8 Verify the current coordination points that the vxfen driver uses.

```
# vxfenconfig -l
```

Migrating from server-based to disk-based fencing in an online cluster

You can migrate between disk-based fencing and server-based fencing without incurring application downtime in the Storage Foundation Cluster File System HA clusters.

You can migrate from server-based fencing to disk-based fencing when you want to perform maintenance tasks on the CP server systems.

Similarly, you can migrate from disk-based fencing to server-based fencing in the following cases:

- You want to leverage the benefits of server-based fencing.
- You want to replace faulty coordinator disks with CP servers.

See [“Migrating from disk-based to server-based fencing in an online cluster”](#) on page 112.

To migrate from server-based fencing to disk-based fencing

- 1 Make sure system-to-system communication is functioning properly.
- 2 Make sure that the Storage Foundation Cluster File System HA cluster is online.

```
# vxfenadm -d
```

For example, if Storage Foundation Cluster File System HA cluster uses server-based fencing:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

- 3 Make sure that you performed the following preparatory tasks to configure disk-based fencing:
 - Identifying disks to use as coordinator disks
 - Setting up coordinator disk group
 - Creating I/O configuration files

See the *Veritas Storage Foundation Cluster File System Installation Guide* for the procedures.

- 4 Create a new `/etc/vxfenmode.test` file with the fencing configuration changes such as the `scsi3` disk policy information.

Refer to the sample `vxfenmode` files in the `/etc/vxfen.d` folder.

- 5 From any node in the Storage Foundation Cluster File System HA cluster, start the `vxfenswap` utility:

```
# vxfenswap -g diskgroup [-n]
```

- 6 Review the message that the utility displays and confirm whether you want to commit the change.

- If you do not want to commit the new fencing configuration changes, press Enter or answer **n** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) n
```

The `vxfenmode` utility rolls back the migration operation.

- If you want to commit the new fencing configuration changes, answer **y** at the prompt.

```
Do you wish to commit this change? [y/n] (default: n) y
```

If the utility successfully commits, the utility moves the `/etc/vxfenmode.test` file to the `/etc/vxfenmode` file.

7 After the migration is complete, verify the change in the fencing mode.

```
# vxfenadm -d
```

For example, after the migration from server-based fencing to disk-based fencing in the Storage Foundation Cluster File System HA cluster:

```
I/O Fencing Cluster Information:
=====
Fencing Protocol Version: 201
Fencing Mode: SCSI3
Fencing SCSI3 Disk Policy: dmp
Cluster Members:
  * 0 (galaxy)
  1 (nebula)
RFSM State Information:
  node 0 in state 8 (running)
  node 1 in state 8 (running)
```

8 Verify the current coordination points that the `vxfen` driver uses.

```
# vxfenconfig -l
```

Migrating from non-secure to secure setup for CP server and SFCFS cluster communication

The following procedure describes how to migrate from a non-secure to secure set up for the CP server and SFCFS cluster.

To migrate from non-secure to secure setup for CP server and SFCFS cluster

- 1 Stop fencing on all the SFCFS cluster nodes of all the clusters (which are using the CP servers).

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

- 2 Stop all the CP servers using the following command on each CP server:

```
# hagrps -offline CPSSG -any
```

- 3 Ensure that security is configured for communication between CP servers and SFCFS cluster nodes.

- 4 Modify /etc/vxcps.conf on each CP server to set security=1.

- 5 Start CP servers using the following command on all of them:

```
# hagrps -online CPSSG -any
```

- 6 Add the following user for each client node on each CP server:

```
_HA_VCS_hostname@HA_SERVICES@FQHN
```

where, hostname is the client node name without qualification, and FQHN is Fully Qualified Host Name of the client node.

Add the users to the CP server database.

For example, issue the following commands on the CP server (mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-g vx
```

```
User _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com successfully added
```

```
# cpsadm -s mycps1.symantecexample.com -a add_user -e\  
_HA_VCS_nebula@HA_SERVICES@nebula.symantec.com\  
-g vx
```

```
User _HA_VCS_nebula@HA_SERVICES@nebula.symantec.com successfully added
```

7 Authorize the user to administer the cluster.

For example, issue the following command on the CP server (mycps1.symantecexample.com):

```
# cpsadm -s mycps1.symantecexample.com -a\  
add_clus_to_user -c cpcluster\  
-u {f0735332-1dd1-11b2-a3cb-e3709c1c73b9}\  
-e _HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com\  
-f cps_operator -g vx
```

Cluster successfully added to user

```
_HA_VCS_galaxy@HA_SERVICES@galaxy.symantec.com privileges.
```

8 Modify /etc/vxfenmode file on each SFCFS cluster node to set security=1.

9 After modifying the /etc/vxfenmode file, run the vxfen init script on the SFCFS cluster node to start fencing.

For an example:

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

Enabling or disabling the preferred fencing policy

You can enable or disable the preferred fencing feature for your I/O fencing configuration.

You can enable preferred fencing to use system-based race policy or group-based race policy. If you disable preferred fencing, the I/O fencing configuration uses the default count-based race policy.

See [“About preferred fencing”](#) on page 36.

See [“How preferred fencing works”](#) on page 37.

To enable preferred fencing for the I/O fencing configuration

1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

3 To enable system-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as System.

```
# haclus -modify PreferredFencingPolicy System
```

- Set the value of the system-level attribute FencingWeight for each node in the cluster.

For example, in a two-node cluster, where you want to assign galaxy five times more weight compared to nebula, run the following commands:

```
# hasys -modify galaxy FencingWeight 50  
# hasys -modify nebula FencingWeight 10
```

- Save the VCS configuration.

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

4 To enable group-based race policy, perform the following steps:

- Make the VCS configuration writable.

```
# haconf -makerw
```

- Set the value of the cluster-level attribute PreferredFencingPolicy as Group.

```
# haclus -modify PreferredFencingPolicy Group
```

- Set the value of the group-level attribute Priority for each service group. For example, run the following command:

```
# hagrps -modify service_group Priority 1
```

Make sure that you assign a parent service group an equal or lower priority than its child service group. In case the parent and the child service groups are hosted in different subclusters, then the subcluster that hosts the child service group gets higher preference.

- Save the VCS configuration.

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

- 5 To view the fencing node weights that are currently set in the fencing driver, run the following command:

```
# vxfenconfig -a
```

To disable preferred fencing for the I/O fencing configuration

- 1 Make sure that the cluster is running with I/O fencing set up.

```
# vxfenadm -d
```

- 2 Make sure that the cluster-level attribute UseFence has the value set to SCSI3.

```
# haclus -value UseFence
```

- 3 To disable preferred fencing and use the default race policy, set the value of the cluster-level attribute PreferredFencingPolicy as Disabled.

```
# haconf -makerw
```

```
# haclus -modify PreferredFencingPolicy Disabled
```

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

About VXFEN tunable parameters

The section describes the VXFEN tunable parameters and how to reconfigure the VXFEN module.

[Table 3-5](#) describes the tunable parameters for the VXFEN driver.

Table 3-5 VXFEN tunable parameters

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
vxfen_debug_sz	Size of debug log in bytes <ul style="list-style-type: none">■ Values Default: 131072 (128 KB) Minimum: 65536 (64 KB) Maximum: 524288 (512 KB)

Table 3-5 VXFEN tunable parameters (*continued*)

vxfen Parameter	Description and Values: Default, Minimum, and Maximum
vxfen_max_delay	<p>Specifies the maximum number of seconds that the smaller sub-cluster waits before racing with larger sub-clusters for control of the coordinator disks when a split-brain occurs.</p> <p>This value must be greater than the vxfen_min_delay value.</p> <ul style="list-style-type: none"> ■ Values <ul style="list-style-type: none"> Default: 60 Minimum: 1 Maximum: 600
vxfen_min_delay	<p>Specifies the minimum number of seconds that the smaller sub-cluster waits before racing with larger sub-clusters for control of the coordinator disks when a split-brain occurs.</p> <p>This value must be smaller than the vxfen_max_delay value.</p> <ul style="list-style-type: none"> ■ Values <ul style="list-style-type: none"> Default: 1 Minimum: 1 Maximum: 600
vxfen_vxfnd_tmt	<p>Specifies the time in seconds that the I/O fencing driver VxVx waits for the I/O fencing daemon VXFEND to return after completing a given task.</p> <ul style="list-style-type: none"> ■ Values <ul style="list-style-type: none"> Default: 60 Minimum: 10 Maximum: 600

In the event of a network partition, the smaller sub-cluster delays before racing for the coordinator disks. The time delay allows a larger sub-cluster to win the race for the coordinator disks. The vxfen_max_delay and vxfen_min_delay parameters define the delay in seconds.

Configuring the VXFEN module parameters

After adjusting the tunable kernel driver parameters, you must reconfigure the VXFEN module for the parameter changes to take effect.

The following example procedure changes the value of the vxfen_min_delay parameter.

To configure the VxFEN parameters and reconfigure the VxFEN module

- 1 Configure the tunable parameter.

```
# /usr/sbin/kctune tunable=value
```

For example:

```
# /usr/sbin/kctune vxfen_min_delay=100
```

- 2 Stop VCS.

```
# hastop -local
```

- 3 Reboot the node.

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

Administering SFCFS global clusters

This section provides instructions for the following global cluster administration tasks:

- About setting up a fire drill
See [“Setting up a disaster recovery fire drill”](#) on page 122.
- Configuring the fire drill service group using the wizard
See [“Configuring the fire drill service group using the Fire Drill Setup wizard”](#) on page 123.
- Verifying a successful fire drill
See [“Verifying a successful fire drill”](#) on page 125.
- Scheduling a fire drill
See [“Scheduling a fire drill”](#) on page 125.

Setting up a disaster recovery fire drill

The Disaster Recovery Fire Drill procedure tests the fault-readiness of a configuration by mimicking a failover from the primary site to the secondary site. This procedure is done without stopping the application at the primary site and disrupting user access, interrupting the flow of replicated data, or causing the secondary site to need resynchronization.

The initial steps to create a fire drill service group on the secondary site that closely follows the configuration of the original application service group and

contains a point-in-time copy of the production data in the Replicated Volume Group (RVG). Bringing the fire drill service group online on the secondary site demonstrates the ability of the application service group to fail over and come online at the secondary site, should the need arise. Fire drill service groups do not interact with outside clients or with other instances of resources, so they can safely come online even when the application service group is online.

You must conduct a fire drill only at the secondary site; do not bring the fire drill service group online on the node hosting the original application.

Before you perform a fire drill in a disaster recovery setup that uses VVR, perform the following steps:

- Set the value of the ReuseMntPt attribute to 1 for all Mount resources.
- Configure the fire drill service group.
See “[Configuring the fire drill service group using the Fire Drill Setup wizard](#)” on page 123.
- After the fire drill service group is taken offline, reset the value of the ReuseMntPt attribute to 0 for all Mount resources.

VCS also supports HA fire drills to verify a resource can fail over to another node in the cluster.

Note: You can conduct fire drills only on regular VxVM volumes; volume sets (vset) are not supported.

VCS provides hardware replication agents for array-based solutions, such as Hitachi Truecopy, EMC SRDF, and so on . If you are using hardware replication agents to monitor the replicated data clusters, refer to the VCS replication agent documentation for details on setting up and configuring fire drill.

Configuring the fire drill service group using the Fire Drill Setup wizard

Use the Fire Drill Setup Wizard to set up the fire drill configuration.

The wizard performs the following specific tasks:

- Creates a Cache object to store changed blocks during the fire drill, which minimizes disk space and disk spindles required to perform the fire drill.
- Configures a VCS service group that resembles the real application group.

The wizard works only with application groups that contain one disk group. The wizard sets up the first RVG in an application. If the application has more than one RVG, you must create space-optimized snapshots and configure VCS manually, using the first RVG as reference.

You can schedule the fire drill for the service group using the `fdsched` script.

See [“Scheduling a fire drill”](#) on page 125.

Running the fire drill setup wizard

To run the wizard

- 1 Start the RVG Secondary Fire Drill wizard on the VVR secondary site, where the application service group is offline and the replication group is online as a secondary:

```
# /opt/VRTSvcs/bin/fdsetup
```

- 2 Read the information on the Welcome screen and press the **Enter** key.
- 3 The wizard identifies the global service groups. Enter the name of the service group for the fire drill.
- 4 Review the list of volumes in disk group that could be used for a spaceoptimized snapshot. Enter the volumes to be selected for the snapshot. Typically, all volumes used by the application, whether replicated or not, should be prepared, otherwise a snapshot might not succeed.

Press the **Enter** key when prompted.

- 5 Enter the cache size to store writes when the snapshot exists. The size of the cache must be large enough to store the expected number of changed blocks during the fire drill. However, the cache is configured to grow automatically if it fills up. Enter disks on which to create the cache.

Press the **Enter** key when prompted.

- 6 The wizard starts running commands to create the fire drill setup.

Press the **Enter** key when prompted.

The wizard creates the application group with its associated resources. It also creates a fire drill group with resources for the application (Oracle, for example), the CFSSMount, and the RVGSnapshot types.

The application resources in both service groups define the same application, the same database in this example. The wizard sets the FireDrill attribute for the application resource to 1 to prevent the agent from reporting a concurrency violation when the actual application instance and the fire drill service group are online at the same time.

Configuring local attributes in the fire drill service group

The fire drill setup wizard does not recognize localized attribute values for resources. If the application service group has resources with local (per-system) attribute values, you must manually set these attributes after running the wizard.

Verifying a successful fire drill

Bring the fire drill service group online on a node that does not have the application running. Verify that the fire drill service group comes online. This action validates that your disaster recovery solution is configured correctly and the production service group will fail over to the secondary site in the event of an actual failure (disaster) at the primary site.

If the fire drill service group does not come online, review the VCS engine log to troubleshoot the issues so that corrective action can be taken as necessary in the production service group.

You can also view the fire drill log, located at `/tmp/fd-servicegroup.pid`

Remember to take the fire drill offline once its functioning has been validated. Failing to take the fire drill offline could cause failures in your environment. For example, if the application service group were to fail over to the node hosting the fire drill service group, there would be resource conflicts, resulting in both service groups faulting.

Scheduling a fire drill

You can schedule the fire drill for the service group using the `fdsched` script. The `fdsched` script is designed to run only on the lowest numbered node that is currently running in the cluster. The scheduler runs the command `hagrp -online firedrill_group -any` at periodic intervals.

To schedule a fire drill

- 1 Add the file `/opt/VRTSvcs/bin/fdsched` to your crontab.
- 2 To make fire drills highly available, add the `fdsched` file to each node in the cluster.

Using Veritas Extension for Oracle Disk Manager

This chapter includes the following topics:

- [About Oracle Disk Manager](#)
- [About Oracle Disk Manager and Storage Foundation Cluster File System](#)
- [About Oracle Disk Manager and Oracle Managed Files](#)
- [Setting up Veritas Extension for Oracle Disk Manager](#)
- [Preparing existing database storage for Oracle Disk Manager](#)
- [Converting Quick I/O files to Oracle Disk Manager files](#)
- [Verifying that Oracle Disk Manager is configured](#)
- [Disabling the Oracle Disk Manager feature](#)
- [Using Cached ODM](#)

About Oracle Disk Manager

Veritas Extension for Oracle Disk Manager is specifically designed for Oracle10g or later to enhance file management and disk I/O throughput. The features of Oracle Disk Manager are best suited for databases that reside in a file system contained in Veritas File System. Oracle Disk Manager allows Oracle10g or later users to improve database throughput for I/O intensive workloads with special I/O optimization.

Veritas Extension for Oracle Disk Manager supports Oracle Resilvering. With Oracle Resilvering, the storage layer receives information from the Oracle database

as to which regions or blocks of a mirrored datafile to resync after a system crash. Oracle Resilvering avoids overhead from the VxVM DRL, which increases performance.

Oracle Disk Manager reduces administrative overhead by providing enhanced support for Oracle Managed Files. Veritas Extension for Oracle Disk Manager has Quick I/O-like capabilities, but is transparent to the user. Unlike Veritas Quick I/O, files managed using Veritas Extension for Oracle Disk Manager do not require special file naming conventions. The Oracle Disk Manager interface uses regular database files. If you are upgrading to Oracle10g or later, you should convert from Quick I/O to Oracle Disk Manager.

Database administrators can choose the datafile type used with the Oracle product. Historically, choosing between file system files and raw devices was based on manageability and performance. The exception to this is a database intended for use with Oracle Parallel Server, which requires raw devices on most platforms. If performance is not as important as administrative ease, file system files are typically the preferred file type. However, while an application may not have substantial I/O requirements when it is first implemented, I/O requirements may change. If an application becomes dependent upon I/O throughput, converting datafiles from file system to raw devices is often necessary.

Oracle Disk Manager was designed to work with Oracle10g or later to provide both performance and manageability. Oracle Disk Manager provides support for Oracle's file management and I/O calls for database storage on VxFS file systems and on raw volumes or partitions. This feature is provided as a dynamically-loaded shared library with which Oracle binds when it is loaded. The Oracle Disk Manager library works with an Oracle Disk Manager driver that is loaded in the kernel to perform its functions.

If you are upgrading to Oracle10g or later, you should convert from Quick I/O to Oracle Disk Manager.

The benefits of using Oracle Disk Manager are as follows:

- True kernel asynchronous I/O for files and raw devices
- Reduced system call overhead
- Improved file system layout by preallocating contiguous files on a VxFS file system
- Performance on file system files that is equal to raw devices
- Transparent to users
- Contiguous datafile allocation

How Oracle Disk Manager improves database performance

Oracle Disk Manager improves database I/O performance to VxFS file systems by:

- Supporting kernel asynchronous I/O
- Supporting direct I/O and avoiding double buffering
- Avoiding kernel write locks on database files
- Supporting many concurrent I/Os in one system call
- Avoiding duplicate opening of files per Oracle instance
- Allocating contiguous datafiles

About kernel asynchronous I/O support

Asynchronous I/O performs non-blocking system level reads and writes, allowing the system to perform multiple I/O requests simultaneously. Kernel asynchronous I/O is better than library asynchronous I/O because the I/O is queued to the disk device drivers in the kernel, minimizing context switches to accomplish the work.

About direct I/O support and avoiding double buffering

I/O on files using `read()` and `write()` system calls typically results in data being copied twice: once between the user and kernel space, and the other between kernel space and the disk. In contrast, I/O on raw devices is copied directly between user space and disk, saving one level of copying. As with I/O on raw devices, Oracle Disk Manager I/O avoids the extra copying. Oracle Disk Manager bypasses the system cache and accesses the files with the same efficiency as raw devices. Avoiding double buffering reduces the memory overhead on the system. Eliminating the copies from kernel to user address space significantly reduces kernel mode processor utilization, freeing more processor cycles to execute the application code.

About avoiding kernel write locks on database files

When database I/O is performed by way of the `write()` system call, each system call acquires and releases a kernel write lock on the file. This lock prevents simultaneous write operations on the same file. Because database systems usually implement their own locks for managing concurrent access to files, write locks unnecessarily serialize I/O writes. Oracle Disk Manager bypasses file system locking and lets the database server control data access.

About supporting many concurrent I/Os in one system call

When performing asynchronous I/O, an Oracle process may try to issue additional I/O requests while collecting completed I/Os, or it may try to wait for particular I/O requests synchronously, as it can do no other work until the I/O is completed. The Oracle process may also try to issue requests to different files. All this activity can be accomplished with one system call when Oracle uses the Oracle Disk Manager I/O interface. This interface reduces the number of system calls performed to accomplish the same work, reducing the number of user space/kernel space context switches.

About avoiding duplicate file open calls

Oracle Disk Manager allows files to be opened once, providing a “file identifier.” This is called “identifying” the files. The same file identifiers can be used by any other processes in the Oracle instance. The file status is maintained by the Oracle Disk Manager driver in the kernel. The reduction in file open calls reduces processing overhead at process initialization and termination, and it reduces the number of file status structures required in the kernel.

About allocating contiguous datafiles

Oracle Disk Manager can improve performance for queries, such as sort and parallel queries, that use temporary tablespaces. Without Oracle Disk Manager, Oracle does not initialize the datafiles for the temporary tablespaces. Therefore, the datafiles become sparse files and are generally fragmented. Sparse or fragmented files lead to poor query performance. When using Oracle Disk Manager, the datafiles are initialized for the temporary tablespaces and are allocated in a contiguous fashion, so that they are not sparse.

About Oracle Disk Manager and Storage Foundation Cluster File System

Oracle Disk Manager supports access to clustered files in the SFCFS environment. With a Veritas Storage Foundation Cluster File System license, ODM supports SFCFS files in a serially-exclusive mode which allows access to each SFCFS file by one node at a time, but does not allow simultaneous access from multiple nodes.

See the `mount_odm(1)` man page for more information on its cluster support modes.

About Oracle Disk Manager and Oracle Managed Files

Oracle10g or later offers a feature known as Oracle Managed Files (OMF). OMF manages datafile attributes such as file names, file location, storage attributes, and whether or not the file is in use by the database. OMF is only supported for databases that reside in file systems. OMF functionality is greatly enhanced by Oracle Disk Manager.

OMF is a file management feature that:

- Eliminates the task of providing unique file names
- Offers dynamic space management by way of the tablespace auto-extend functionality of Oracle10g or later

The main requirement for OMF is that the database be placed in file system files. There are additional prerequisites imposed upon the file system itself.

OMF should only be used in file systems that reside within striped logical volumes, which support dynamic file system growth. File systems intended for OMF use must also support large, extensible files in order to facilitate tablespace auto-extension. Raw partitions cannot be used for OMF.

By default, OMF datafiles are created with auto-extensibility. This attribute reduces capacity planning associated with maintaining existing databases and implementing new applications. Due to disk fragmentation that occurs as the tablespace grows over time, database administrators have been somewhat cautious when considering auto-extensible tablespaces. Oracle Disk Manager eliminates this concern.

When Oracle Disk Manager is used in conjunction with OMF, special care is given within Veritas Extension for Disk Manager to ensure that contiguous disk space is allocated to datafiles, including space allocated to a tablespace when it is auto-extended. The table and index scan throughput does not decay as the tablespace grows.

How Oracle Disk Manager works with Oracle Managed Files

The following example illustrates the relationship between Oracle Disk Manager and Oracle Managed Files (OMF). The example shows the `init.ora` contents and the command for starting the database instance. To simplify Oracle UNDO management, the new Oracle10g or later `init.ora` parameter `UNDO_MANAGEMENT` is set to `AUTO`. This is known as System-Managed Undo.

Note: Before building an OMF database, you need the appropriate `init.ora` default values. These values control the location of the `SYSTEM` tablespace, online redo logs, and control files after the `CREATE DATABASE` statement is executed.

```
$ cat initPROD.ora
UNDO_MANAGEMENT = AUTO
DB_CREATE_FILE_DEST = '/PROD'
DB_CREATE_ONLINE_LOG_DEST_1 = '/PROD'
db_block_size = 4096
db_name = PROD
$ sqlplus /nolog
SQL> connect / as sysdba
SQL> startup nomount pfile= initPROD.ora
```

The Oracle instance starts.

```
Total System Global Area 93094616 bytes
Fixed Size 279256 bytes
Variable Size 41943040 bytes
Database Buffers 50331648 bytes
Redo Buffers 540672 bytes
```

To implement a layout that places files associated with the `EMP_TABLE` tablespace in a directory separate from the `EMP_INDEX` tablespace, use the `ALTER SYSTEM` statement. This example shows how OMF handles file names and storage clauses and paths. The layout allows you to think of the tablespaces as objects in a file system as opposed to a collection of datafiles. Since OMF uses the Oracle Disk Manager file resize function, the tablespace files are initially created with the default size of 100 MB and ODM increases the size as needed. Use the `MAXSIZE` attribute to limit growth.

The following example shows the commands for creating an OMF database and for creating the `EMP_TABLE` and `EMP_INDEX` tablespaces in their own locale:

```
SQL> create database PROD;
```

Note: The directory must exist for OMF to work, so the `SQL*Plus HOST` command is used to create the directories:

The database is created.

```
SQL> HOST mkdir /PROD/EMP_TABLE;
SQL> ALTER SYSTEM SET DB_CREATE_FILE_DEST = '/PROD/EMP_TABLE';
```

The system is altered.

```
SQL> create tablespace EMP_TABLE DATAFILE AUTOEXTEND ON MAXSIZE \
500M;
```

A tablespace is created.

```
SQL> ALTER SYSTEM SET DB_CREATE_FILE_DEST = '/PROD/EMP_INDEX';
```

The system is altered.

```
SQL> create tablespace EMP_INDEX DATAFILE AUTOEXTEND ON MAXSIZE \
100M;
```

A tablespace is created.

Use the `ls` command to show the newly created database:

```
$ ls -lFR
total 638062
drwxr-xr-x 2 oracle10g dba 96 May  3 15:43 EMP_INDEX/
drwxr-xr-x 2 oracle10g dba 96 May  3 15:43 EMP_TABLE/
-rw-r--r-- 1 oracle10g dba 104858112 May 3 17:28 ora_1_BEhYgc0m.log
-rw-r--r-- 1 oracle10g dba 104858112 May 3 17:27 ora_2_BEhYu4NA.log
-rw-r--r-- 1 oracle10g dba 806912 May 3 15:43 ora_BEahlfUX.ctl
-rw-r--r-- 1 oracle10g dba 10489856 May 3 15:43 ora_sys_undo_BEajPSVq.dbf
-rw-r--r-- 1 oracle10g dba 104861696 May 3 15:4 ora_system_BEaiFE8v.dbf
-rw-r--r-- 1 oracle10g dba 186 May 3 15:03 PROD.ora

./EMP_INDEX:
total 204808
-rw-r--r-- 1 oracle10g dba 104861696 May 3 15:43
ora_emp_inde_BEakGfun.dbf

./EMP_TABLE:
total 204808
-rw-r--r-- 1 oracle10g dba 104861696 May 3 15:43
ora_emp_tabl_BEaklLqK.dbf
```

Setting up Veritas Extension for Oracle Disk Manager

Veritas Extension for Oracle Disk Manager is part of Veritas Storage Foundation for Cluster File System. Veritas Extension for Oracle Disk Manager is enabled once Veritas Storage Foundation for Cluster File System and Oracle10g or later

are installed. The Veritas Extension for Oracle Disk Manager library is linked to the library in the `{ORACLE_HOME}/lib` directory.

If you are performing a local Oracle installation, not on the SFCFS file system, then ODM linking needs to be performed on all nodes in the cluster.

Before setting up Veritas Extension for Oracle Disk Manager, the following conditions must be met:

- | | |
|---------------|---|
| Prerequisites | <ul style="list-style-type: none">■ Veritas Storage Foundation for cluster File System must be installed on your system.■ Oracle10g, or later, must be installed on your system.■ If Cached Quick I/O is available, do not enable Oracle Disk Manager when Cached Quick I/O is enabled for datafiles. |
| Usage Notes | <ul style="list-style-type: none">■ When the Quick I/O feature is available, Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O. Do not turn off the Quick I/O mount option, which is the default.■ Oracle uses default file access methods if Oracle10g or later or Veritas Storage Foundation for Cluster File System is not installed, or VxFS 5.1 is not available in the kernel. |

Preparing existing database storage for Oracle Disk Manager

Non-Quick I/O files in a VxFS file system work with Oracle Disk Manager without any changes. The files are found and identified for Oracle Disk Manager I/O by default. To take full advantage of Oracle Disk Manager datafiles, files should not be fragmented.

If you are using Quick I/O files in a VxFS file system and you want to move to Oracle Disk Manager, convert the Quick I/O files to normal files using the `qio_convertdbfiles -u` command.

You must be running Oracle10g or later to use Oracle Disk Manager.

Converting Quick I/O files to Oracle Disk Manager files

If you plan to run the Veritas product with Oracle10g or later, and you have been using Quick I/O files, Symantec recommends that you convert your Quick I/O files to regular files. This should be done after you upgrade.

Note: If you are running an earlier version of Oracle (Oracle 8.x or lower), you should not convert your Quick I/O files because Oracle Disk Manager is for Oracle10g or later only.

The Oracle Disk Manager uses the Quick I/O driver to perform asynchronous I/O, do not turn off the Quick I/O mount option, which is the default.

To convert Quick I/O files to Oracle Disk Manager files

- 1 As Oracle DBA, run `qio_getdbfiles` to retrieve a list of all datafiles.

```
$ /opt/VRTS/bin/qio_getdbfiles -T ora -a
```

The list is compiled in a file named `mkqio.dat`.

- 2 Shutdown the database.
- 3 As Oracle DBA, run `qio_convertdbfiles` in the directory containing the `mkqio.dat` file. The `qio_convertdbfiles` script converts all Quick I/O files to ODM files.

```
$ /opt/VRTS/bin/qio_convertdbfiles -T ora -u
```

- 4 Restart the database instance.

Verifying that Oracle Disk Manager is configured

Before verifying that Oracle Disk Manager is configured, make sure that the following conditions are met:

- Prerequisites
- `/opt/VRTSodm/lib/libodm.sl` must exist.
 - If you are using Oracle 10g, `$ORACLE_HOME/lib/libodm10.sl` is linked to `/opt/VRTSodm/lib/libodm.sl`.
 - The `VRTSdbed` license must be valid.
 - The `VRTSodm` package must be installed.

To verify that Oracle Disk Manager is configured

- 1 Verify that the ODM feature is included in the license:

```
# /opt/VRTS/bin/vxlicrep | grep ODM
```

The output verifies that ODM is enabled.

Note: Verify that the license key containing the ODM feature is not expired. If the license key has expired, you will not be able to use the ODM feature.

- 2 Check that the `VRTSodm` package is installed:

```
(root@slias19)[/] swlist VRTSodm
# Initializing...
# Contacting target "slias19"...
#
# Target:  slias19:/
#
# VRTSodm          5.0.31.5.%20090322 Veritas Oracle Disk Manager
# VRTSodm.ODM-KRN  5.0.31.5.%20090322 Veritas ODM kernel files
# VRTSodm.ODM-MAN  5.0.31.5.%20090322 Veritas ODM manual pages
# VRTSodm.ODM-RUN  5.0.31.5.%20090322 Veritas ODM commands
```

- 3 Check that `libodm.sl` is present.

```
# ls -lL /opt/VRTSodm/lib/libodm.sl
-rw-r--r-- 1 root sys 14336 Apr 25 18:42
/opt/VRTSodm/lib/libodm.sl
```

To verify that Oracle Disk Manager is running

- 1 Start the Oracle database.
- 2 Check that the instance is using the Oracle Disk Manager function:

```
# cat /dev/odm/stats
# echo $?
0
```

- 3 Verify that the Oracle Disk Manager is loaded:

```
# /usr/sbin/kcmodule -P state odm  
state loaded
```

- 4 In the alert log, verify the Oracle instance is running. The log should contain output similar to the following:

```
Oracle instance running with ODM: Veritas 5.1.100.00 ODM Library,  
Version 2.0
```

Disabling the Oracle Disk Manager feature

Since the Oracle Disk Manager feature uses regular files, you can access these files as regular VxFS files as soon as the feature is disabled.

Note: To convert to VxFS with Quick I/O, disable Oracle Disk Manager using the following procedure, then convert the files to Quick I/O files.

See [“Converting Quick I/O files to Oracle Disk Manager files”](#) on page 134.

Before disabling the Oracle Disk Manager feature, you may want to back up your files.

To disable the Oracle Disk Manager feature in an Oracle instance

- 1 Shut down the database instance.
- 2 Use the `rm` and `ln` commands to remove the link to the Oracle Disk Manager Library.

For Oracle 11g, enter:

```
# rm ${ORACLE_HOME}/lib/libodm11.sl
# ln -s ${ORACLE_HOME}/lib/libodmd11.sl \
${ORACLE_HOME}/lib/libodm11.sl
```

For Oracle 10g, enter:

```
# rm ${ORACLE_HOME}/lib/libodm10.sl
# ln -s ${ORACLE_HOME}/lib/libodmd10.sl \
${ORACLE_HOME}/lib/libodm10.sl

# rm ${ORACLE_HOME}/lib/libodm9.sl
# ln -s ${ORACLE_HOME}/lib/libodmd9.sl \
${ORACLE_HOME}/lib/libodm9.sl
```

For HP-UX IA

For Oracle 11g, enter:

```
# rm ${ORACLE_HOME}/lib/libodm11.so
# ln -s ${ORACLE_HOME}/lib/libodmd11.sl \
${ORACLE_HOME}/lib/libodm11.so
```

For Oracle 10g, enter:

```
# rm ${ORACLE_HOME}/lib/libodm10.so
# ln -s ${ORACLE_HOME}/lib/libodmd10.sl \
${ORACLE_HOME}/lib/libodm10.so

# rm ${ORACLE_HOME}/lib/libodm9.so
# ln -s ${ORACLE_HOME}/lib/libodmd9.sl \
${ORACLE_HOME}/lib/libodm9.so
```

- 3 Restart the database instance.

Using Cached ODM

ODM I/O normally bypasses the file system cache and directly reads from and writes to disk. Cached ODM enables some I/O to use caching and read ahead, which

can improve ODM I/O performance. Cached ODM performs a conditional form of caching that is based on per-I/O hints from Oracle. The hints indicate what Oracle does with the data. ODM uses these hints to perform caching and read ahead for some reads, but ODM avoids caching other reads, even for the same file.

You can enable cached ODM for local mount files and cluster mount files.

See [“Enabling Cached ODM for file systems”](#) on page 139.

Cached ODM can be configured in two ways. The primary configuration method is to turn caching on or off for all I/O on a per-file basis. The secondary configuration method is to adjust the ODM cachemap. The cachemap maps file type and I/O type combinations into caching advisories.

See [“Modifying Cached ODM settings for individual files”](#) on page 139.

See [“Adding Cached ODM settings via the cachemap”](#) on page 140.

Enabling Cached ODM for file systems

Cached ODM is initially disabled on a file system. You enable Cached ODM for a file system by setting the `odm_cache_enable` option of the `vxtunefs` command after the file system is mounted.

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

Note: The `vxtunefs` command enables conditional caching for all of the ODM files on the file system.

To enable Cached ODM for a file system

- 1 Enable Cached ODM on the VxFS file system `/database01`:

```
# vxtunefs -s -o odm_cache_enable=1 /database01
```

- 2 Optionally, you can make this setting persistent across mounts by adding a file system entry in the file `/etc/vx/tunefstab`:

```
/dev/vx/dsk/datadg/database01 odm_cache_enable=1
```

See the `tunefstab(4)` manual page.

Modifying Cached ODM settings for individual files

You can use the `odmadm setcachefile` command to override the cachemap for a specific file so that ODM caches either all or none of the I/O to the file. The caching state can be ON, OFF, or DEF (default). The DEF caching state is conditional

caching, meaning that for each I/O, ODM consults the cachemap and determines whether the specified file type and I/O type combination should be cached. The ON caching state causes the specified file always to be cached, while the OFF caching state causes the specified file never to be cached.

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

Note: The cache advisories operate only if Cached ODM is enabled for the file system. If the `odm_cache_enable` flag is zero, Cached ODM is OFF for all of the files in that file system, even if the individual file cache advisory for a file is ON.

To enable unconditional caching on a file

- ◆ Enable unconditional caching on the file `/mnt1/file1`:

```
# odmadm setcachefile /mnt1/file1=on
```

With this command, ODM caches all reads from `file1`.

To disable caching on a file

- ◆ Disable caching on the file `/mnt1/file1`:

```
# odmadm setcachefile /mnt1/file1=off
```

With this command, ODM does not cache reads from `file1`.

To check on the current cache advisory settings for a file

- ◆ Check the current cache advisory settings of the files `/mnt1/file1` and `/mnt2/file2`:

```
# odmadm getcachefile /mnt1/file1 /mnt2/file2
/mnt1/file1,ON
/mnt2/file2,OFF
```

To reset all files to the default cache advisory

- ◆ Reset all files to the default cache advisory:

```
# odmadm resetcachefiles
```

Adding Cached ODM settings via the cachemap

You can use the `odmadm setcachemap` command to configure the cachemap. The cachemap maps file type and I/O type combinations to caching advisories. ODM uses the cachemap for all files that have the default conditional cache setting.

Such files are those for which caching has not been turned on or off by the `odmadm setcachefile` command.

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

By default, the cachemap is empty, but you can add caching advisories by using the `odmadm setcachemap` command.

To add caching advisories to the cachemap

- ◆ Add a caching advisory to the cachemap:

```
# odmadm setcachemap data/data_read_seq=cache,readahead
```

With this example command, ODM uses caching and readahead for I/O to online log files (`data`) that have the `data_read_seq` I/O type. You can view the valid file type and I/O type values from the output of the `odmadm getcachemap` command.

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

Making the caching settings persistent across mounts

By default, the Cached ODM settings are not persistent across mounts. You can make the settings persistent by creating the `/etc/vx/odmadm` file and listing the caching advisory settings in the file

To make the caching setting persistent across mounts

- ◆ Create the `/etc/vx/odmadm` file to list files and their caching advisories. In the following example of the `/etc/vx/odmadm` file, if you mount the `/dev/vx/dsk/rootdg/vol1` device at `/mnt1`, `odmadm` turns off caching for `/mnt1/oradata/file1`:

```
setcachemap data/read_data_header=cache
setcachemap all/datapump=cache,readahead
device /dev/vx/dsk/rootdg/vol1
setcachefile oradata/file1=off
```


Clustered NFS

This chapter includes the following topics:

- [About Clustered NFS](#)
- [Requirements](#)
- [Understanding how Clustered NFS works](#)
- [cfsshare manual page](#)
- [Configure and unconfigure Clustered NFS](#)
- [Administering Clustered NFS](#)
- [How to mount an NFS-exported file system on the NFS clients](#)
- [Debugging Clustered NFS](#)

About Clustered NFS

In previous releases, the SFCFS stack only allowed an active/passive setup for NFS serving due to the complexity of lock reclamation in an active/active configuration. This new Clustered NFS (CNFS) feature gracefully handles failure of any node and reclaim the locks in such a way as to not accidentally lose any existing lock grants without notification.

This release only supports NFS Version 3.

See [“How to mount an NFS-exported file system on the NFS clients”](#) on page 160.

Requirements

- Prior knowledge of NFS is a prerequisite.

Understanding how Clustered NFS works

This Clustered NFS feature allows the same file system mounted across multiple nodes using CFS to be shared over NFS from any combination of those nodes without any loss of functionality during failover. The failover of NFS lock servers includes all the locks being released by the old node then reclaimed by clients talking to the new node during the grace period.

Basic design

The basic design is to have VCS manage Virtual IP (VIP) resources that can failover between nodes and to add extra code into the steps used to handle these resources to properly handle the NFS level operations. All other involved resources are active on all nodes participating. The lock data, which is saved into a shared area, is managed and used in lock-step with the Virtual IP resources to ensure that all locks are reclaimed properly by clients while preventing any inappropriate locks from being taken at the wrong time. Such interfering locks are prevented through a combination of stopping services and using new features of the lock handling inside the VxFS kernel code.

To communicate with the new code in the kernel, the `fsclustadm` command has been modified to add command line interfaces to the private ioctl calls.

Note: You must have at least one VIP configured for each CNFS server.

Internal Clustered NFS functionality

This section describes the internal functionality of the triggers and actions scripts that are a part of the Clustered NFS solution.

preonline trigger

The `preonline` script copies the lock state files created by the node status monitor (normally called `statd` or `rpc.statd` daemon) during IP failovers and node failures.

The `preonline` script does the following on IP failover or node failure:

- Finds the IP and the node it was last online on.
- Finds the node on which the IP is next going to go online on.
- Calls `/opt/VRTS/bin/fsclustadm frlpause_enable` and `/opt/VRTS/bin/fsclustadm frlock_pause` to ensure that file system does not give out any new locks during the failover.

- Stops lock and status services on all nodes to prevent granting locks.
- Copies all the files from `/locks/sm/lastonline/sm/` to `/locks/sm/nextonline/sm/` directory.
where *locks* is the file system created for storing lock information.
where *lastonline* is the node on which the VIP resource was previous online.
where *nextonline* is the node on which the VIP resource will go online next.
- Calls `/opt/VRTS/bin/fsclustadm frlock_resume` to resume giving out locks.

Note: At the end of the preonline trigger all lock services have stopped on all nodes and no NFS locks can be requested until they are restarted.

postonline trigger

The `postonline` script for each VIP does the following during an IP failover or node failure:

- Starts lock services, triggers reclaim, and grace mode on all nodes.
- The restarting of status monitor scans all lock status files in the state directory and contacts all nodes to reclaim their locks. The state files get deleted after they are processed and reclaim messages sent appropriately.
- The lock server goes into grace mode and only allows clients to recover their locks during the grace period. It does not give out any new locks during the grace period.

postoffline trigger

The `postoffline` script does the following on IP failover:

- Calls `/opt/VRTS/bin/fsclustadm frlpause_disable` to reduce the internal usage counter.
- Each call to `/opt/VRTS/bin/fsclustadm frlpause_enable` needs to be matched with a call to `/opt/VRTS/bin/fsclustadm frlpause_disable` as the kernel keeps an internal counter to track the number of IP addresses active on a system. If there are no active IPs on a system, it will be in disabled mode.

Note: This trigger is only called for an administratively initiated failover. An actual system failure and reboot discards the local state being manipulated in this stage. This is the one trigger called on the node that was previously hosting the VIP, while the others are called on the server taking over.

Actions

- On each node, a `/opt/VRTSvcs/bin/IP/actions/nfscfs` file is installed. This file is used to start and stop the NFS locking daemons on a specified node. The `action` script is used instead of using `rsh`, `ssh` or `hacli` for remote command execution from the triggers.
- On each node, a `/opt/VRTSvcs/bin/ApplicationNone/actions/nfscfsapp` file is installed. This file is used while configuring and unconfiguring the Clustered NFS solution using `cfsshare config` and `cfsshare unconfig` commands.

cfsshare manual page

This Clustered NFS feature adds a new configuration utility called `cfsshare` to the `VRTScavf` package and several scripts that are added into the VCS configuration to manage parallel NFS server resources. The `cfsshare` command modifies VCS resources that were created by other utilities such as `cfsmntadm`.

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

Configure and unconfigure Clustered NFS

This section describes how to configure and unconfigure Clustered NFS.

Configure Clustered NFS

```
cfsshare config -p nfs [-n] shared_disk_group shared_volume mount_point
```

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

The CNFS solution requires a shared file system such as `/locks` that is mounted on all cluster nodes. This file system is not a data file system; it contains the lock state files corresponding to NFS clients holding locks on the CNFS servers.

Note: If CIFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of Clustered NFS.

The local state tracking directory contains a file for each NFS client that has a transaction with the NFS server. The local state tracking directory is:

```
/var/statmon/sm
```

This creates a symlink to `/locks/sm/nodename/sm` on all the cluster nodes. This allows the lock state files for any cluster node to be accessed by other nodes in the cluster, even when the node is down.

The `-n` option can be used if the user does not want `cfsshare` to create the symlink to `/locks/sm/nodename/sm`. If this option is used, then the user needs to manually create a symlink.

The `config` option adds this shared file system to VCS configuration; it creates the corresponding CFSSMount resource in a special parallel service group called `cfsnfssg`. This also creates an NFS resource in the `cfsnfssg` service group. In addition to this, a separate resource of the new type `ApplicationNone` is created to monitor `lockd` and `statd` daemons.

If you run the `cfsshare config -n` option, you need to perform the following procedure:

- 1 On each node, create the following directory inside the locks directory, if it does not already exist:

```
# mkdir -p /locks/sm/nodename/sm
```

- 2 Run the following commands on any one cluster node to set the owner, group, and permissions of `/locks/sm` appropriately:

```
# chown -R daemon:daemon /locks/sm
# chmod -R 755 /locks/sm
```

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m user -l /var/run \
-c /etc/samba/smb.conf -t /usr lockdg vollocks /locks
```

See the Common Internet File System chapter for explanation regarding the various CIFS-related options.

Service group `cfsnfssg_dummy`

As part of CNFS configuration, a service group called `cfsnfssg_dummy` gets created. This service group is mostly offline.

There is a limit on the number of service groups that can be created in VCS. If this limit is reached, then `cfsnfssg_dummy` serves as the service group in which resources get created during `cfsshare unshare` and `cfsshare delete` operations.

See the *Veritas Cluster Server Administrator's Guide* for information about the `GroupLimit` attribute.

Unconfigure Clustered NFS

```
cfsshare unconfig -p nfs
```

This command is used to undo all the steps during the config phase.

Note: If there are any CFS file systems still being shared or any Virtual IP is added the `cfsshare unconfig` command fails.

Administering Clustered NFS

This section describes Clustered NFS scenarios.

See “[Samples for configuring a Clustered NFS](#)” on page 152.

See “[Sample main.cf file](#)” on page 156.

Displaying the NFS shared CFS file systems

```
cfsshare display
```

This command displays the CFS file systems that are currently being NFS shared by the cluster nodes.

Sharing a CFS file system previously added to VCS

```
cfsshare share mount_point [share_options]
```

Note: You can use the `-N` option to specify the NFS share options.

You can also use the `-p nfs` to specify that the protocol to be used is NFS.

See the `cfsshare(1M)` manual page for more information.

Before running this command, the user should have run `cfsmntadm` command to add the shared file system to VCS configuration and the `cfsmount` command to

mount the shared file system at the *mount_point*. Once these commands have been executed, the CFSSMount resource corresponding to the *mount_point* gets created in either a default service group (with a name similar to `vrts_cfsshare_cfs_int_cfsmountnumber`) or in a separate service group, as specified by the user.

The `cfsshare share` command moves the CFSSMount resource corresponding to the *mount_point* and the associated CVMVolDg resource to the `cfsnfssg` service group (that was created using the `config` option). In addition, this command also creates a share resource on top of the CFSSMount resource in the same `cfsnfssg` service group.

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSSMount and CVMVolDg resources in the `cfsnfssg` service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Unsharing the previous shared CFS file system

```
cfsshare unshare mount_point
```

Before running this command, the user is supposed to have run the `cfsshare share` command.

The `cfsshare unshare` command enables the user to stop sharing the file system mounted at the *mount_point*. This command moves the Share, CFSSMount, and CVMVolDg resources corresponding to the *mount_point* from `cfsnfssg` service group to a newly created service group. The Share resource is taken offline and then deleted.

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSSMount and CVMVolDg resources in the newly created service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Running the `cfsmntadm delete` command does not remove the `ActivationMode` attribute. If no volumes or vsets in the disk group are in the VCS config, you must use the `cfsgadm delete` to remove this `ActivationMode` attribute.

Adding an NFS shared CFS file system to VCS

```
cfsshare add shared_disk_group shared_volume mount_point[share_options] \  
node_name=[mount_options]....
```

```
cfsshare add shared_disk_group shared_volume mount_point[share_options] \  
all=[mount_options]
```

```
cfsshare add -p nfs [-N nfs_share_options] shared_disk_group \  
shared_volume mount_point <node_name=[mount_point]>
```

```
cfsshare add -p nfs [-N nfs_share_options] shared_disk_group \  
shared_volume mount_point all=[mount_point]
```

This command adds the CFS file system to VCS configuration in the `cfsnfssg` service group, then mounts the file system at the `mount_point` and NFS shares the CFS file system.

Deleting the NFS shared CFS file system from VCS

```
cfsshare delete mount_point
```

Before running this command, the user is supposed to have run the `cfsshare add` command to create the required resources (Share, CFSMount, and CVMVolDg, if needed) in the `cfsnfssg` service group.

This command unshares the CFS file system mounted at the `mount_point`, unmounts the CFS file system, and removes the CFS file system from VCS configuration.

Adding a Virtual IP address to VCS

```
cfsshare addvip network_interface address netmask
```

This command is used to create a new non-parallel/failover service group that contains a NIC resource for the given network device and an IP resource for the Virtual IP address.

The `cfsshare addvip` command lets you specify only one network interface, that is assumed to be present on all cluster nodes. If you want to specify different network interfaces for different cluster nodes, then you need to run certain VCS commands. For example:

```
# haconf -makerw  
# hares -local vip1 Device  
# hares -modify vip1 Device lan1 -sys system01
```

```
# hares -modify vip1 Device lan2 -sys system02
# haconf -dump -makero
```

where *vip1* is the Virtual IP resource created by the `cfsshare addvip` command.
where *system01* and *system02* are the cluster nodes.

Deleting a Virtual IP address from VCS

```
cfsshare deletevip address
```

This command is used to delete the non-parallel/failover service group corresponding to the Virtual IP address.

Adding an IPv6 Virtual IP address to VCS

```
cfsshare addvip6 network_interface ipv6_address prefixlen
```

This command is used to create a new non-parallel/failover service group which contains a NIC resource for the given network device and an IP resource for the IPv6 Virtual IP address.

Deleting an IPv6 Virtual IP address from VCS

```
cfsshare deletevip6 ipv6_address
```

This command is used to delete the non-parallel/failover service group corresponding to the IPv6 Virtual IP address.

Changing the share options associated with an NFS share

This section describes how to change the share options associated with an NFS share.

To change the share options associated with an NFS share

- 1 On any node in the cluster, run `cfsshare unshare` to unshare the file system:

```
# cfsshare unshare mount_point
```

- 2 On any node in the cluster, run `cfsshare share` to share the file system with the desired share options:

```
# cfsshare share -p nfs mount_point share_options
```

Note: The `cfsshare unshare` operation can affect NFS clients that might have mounted the `mount_point` file system.

Sharing a file system checkpoint

This section describes how to share a file system checkpoint.

To share a file system checkpoint

- 1 To add the checkpoint to the VCS configuration, enter:

```
# cfsmntadm add ckpt ckptname mntpt_of_fs mntpt_of_checkpoint \  
all=cluster,rw
```

where `ckptname` is the checkpoint name.

where `mntpt_of_fs` is the name of the mount point of the file system.

where `mntpt_of_checkpoint` is the mount point for the checkpoint.

- 2 To mount the checkpoint, enter:

```
# cfsmount mntpt_of_checkpoint
```

- 3 Run the `cfsshare share` command to share this checkpoint:

```
# cfsshare share -p nfs mntpt_of_checkpoint
```

Samples for configuring a Clustered NFS

There are two samples for configuring a Clustered NFS.

Sample 1

This sample is intended to use `cfsshare` command to config and control this feature.

To configure a Clustered NFS (Sample 1)

- 1 Configure a VCS configuration for CFS/CVM, enter:

```
# cfscluster config
```

- 2 Configure CNFS components, enter:

```
# cfsshare config -p nfs shared_disk_group shared_volume \  
mount_point
```

For example:

```
# cfsshare config -p nfs cfsdg vollocks /locks
```

- 3 Add and mount the CFS file system to the VCS configuration, enter:

```
# cfsmntadm add shared_disk_group shared_volume mount_point \  
[service_group] all=[mount_options]  
# cfsmount mount_point
```

For example:

```
# cfsmntadm add cfsdg vol1 /mnt1 all=delaylog,largefiles  
# cfsmount /mnt1
```

- 4 Share the CFS file system, enter:

```
# cfsshare share mount_point [share_options]
```

For example:

```
# cfsshare share /mnt1 rw
```

- 5 Run the following command to display the currently exported file systems:

```
# cfsshare display
```

SHARE RESOURCE	MOUNTPOINT	SHARE OPTIONS
share1	/mnt1	rw

- 6 Add the Virtual IP addresses for users to access the shared CFS file systems, enter:

```
# cfsshare addvip network_interface address netmask
```

For example:

```
# cfsshare addvip lan0 10.182.111.161 255.255.240.0
```

- 7 Delete a previously added Virtual IP address from the configuration, enter:

```
# cfsshare deletevip address
```

For example:

```
# cfsshare deletevip 10.182.111.161
```

- 8 Unshare CFS file system, enter:

```
# cfsshare unshare mount_point
```

For example:

```
# cfsshare unshare /mnt1
```

- 9 Unmount and remove the CFS file system from the VCS configuration, enter

```
# cfsumount mount_point
```

```
# cfsmntadm delete mount_point
```

For example:

```
# cfsumount /mnt1
```

```
# cfsmntadm delete /mnt1
```

- 10 Unconfigure NFS shared CFS file systems, enter:

```
# cfsshare unconfig -p nfs
```

Sample 2

This sample is intended to use the `cfsshare` command to add a CFS file system to VCS configuration and mount them. Then share them through NFS, unshare, unmount, and remove the CFS file systems from VCS configuration.

To configure Clustered NFS (Sample 2)

- 1 Configure a VCS configuration for CFS/CVM, enter:

```
# cfscluster config
```

- 2 Configure the CNFS components, enter:

```
# cfsshare config -p nfs shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p nfs cfsdg vollocks /locks
```

- 3 Add and mount the NFS shared CFS file system to the VCS configuration, enter:

```
# cfsshare add shared_disk_group shared_volume mount_point \  
[share_options] all=[mount_options]
```

For example:

```
# cfsshare add cfsdg vol1 /mnt1 all=rw
```

- 4 Add the Virtual IP addresses for users to access the shared CFS file systems, enter:

```
# cfsshare addvip network_interface address netmask
```

For example:

```
# cfsshare addvip lan0 10.182.111.161 255.255.240.0
```

- 5 Delete a previously added Virtual IP address from the configuration, enter:

```
# cfsshare deletevip address
```

For example:

```
# cfsshare deletevip 10.182.111.161
```

- 6 Unshare, unmount, and remove the CFS file system from the VCS configuration, enter:

```
# cfsshare delete mount_point
```

For example:

```
# cfsshare delete /mnt1
```

- 7 Unconfigure CNFS components, enter:

```
# cfsshare unconfig -p nfs
```

Sample main.cf file

This is a sample `main.cf` file.

```
include "OracleASMTypes.cf"
include "types.cf"
include "ApplicationNone.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.cf"

cluster cfs_cluster (
  UserNames = { admin = anoGniNkoJooMwoIn1 }
  Administrators = { admin }
  HacliUserLevel = COMMANDROOT
)

system galaxy (
)

system nebula (
)

group cfsnfssg (
  SystemList = { galaxy = 0, nebula = 1 }
  AutoFailOver = 0
  Parallel = 1
  AutoStartList = { galaxy, nebula }
)
```

```
ApplicationNone app (
  MonitorProgram = "/opt/VRTSvc/bin/ApplicationNone/lockdstatdmon"
)

CFSMount cfsmount2 (
  Critical = 0
  MountPoint = "/fsqamnt2"
  BlockDevice = "/dev/vx/dsk/dg1/vol2"
  NodeList = { galaxy, nebula }
)

CFSMount cfsnfs_locks (
  Critical = 0
  MountPoint = "/locks"
  BlockDevice = "/dev/vx/dsk/dg1/vollocks"
  NodeList = { galaxy, nebula }
)

CVMVolDg cvmvoldg1 (
  Critical = 0
  CVMDiskGroup = dg1
  CVMActivation @galaxy = sw
  CVMActivation @nebula = sw
  CVMVolume = { vol2 }
  CVMVolumeIoTest = { vol2 }
)

NFS nfs (
)

Share share1 (
  PathName = "/fsqamnt2"
)

requires group cvm online local firm
cfsmount2 requires cvmvoldg1
cfsnfs_locks requires cvmvoldg1
share1 requires cfsmount2
share1 requires nfs

// resource dependency tree
```

```
//
// group cfsnfssg
// {
//   ApplicationNone app
//   CFSMount cfsnfs_locks
//     {
//       CVMVolDg cvmvoldg1
//     }
//   Share share1
//     {
//       NFS nfs
//       CFSMount cfsmount2
//         {
//           CVMVolDg cvmvoldg1
//         }
//     }
// }

group cfsnfssg_dummy (
  SystemList = { galaxy = 0, nebula = 1 }
  AutoFailOver = 0
  Parallel = 1
  AutoStartList = { galaxy, nebula }
)

requires group cvm online local firm

// resource dependency tree
//
// group cfsnfssg_dummy
// {
// }

group cvm (
  SystemList = { galaxy = 0, nebula = 1 }
  AutoFailOver = 0
  Parallel = 1
  AutoStartList = { galaxy, nebula }
)
```

```
CFSfsckd vxfsckd (
  ActivationMode @galaxy = { dgl = sw }
  ActivationMode @nebula = { dgl = sw }
)

CVMcluster cvm_clus (
  CVMClustName = cfs_cluster
  CVMNodeId = { galaxy = 0, nebula = 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
//     }
//   }
// }

group vip1 (
  SystemList = { galaxy = 0, nebula = 1 }
  AutoStartList = { galaxy, nebula }
  PreOnline @galaxy = 1
  PreOnline @nebula = 1
)

IP vip1 (
```

```

    Device = lan0
    Address = "10.182.111.161"
    NetMask = "255.255.252.0"
)

NIC nic1 (
    Device = lan0
)

requires group cfsnfssg online local firm
vip1 requires nic1

// resource dependency tree
//
// group vip1
// {
//   IP vip1
//   {
//     NIC nic1
//   }
// }

```

How to mount an NFS-exported file system on the NFS clients

This section describes how to mount an NFS-exported file system on the NFS clients.

To mount an NFS-exported file system on the NFS clients

- ◆ Run the following command:

```
# mount -F nfs -o vers=3 VIP_address:remote_filesystem mount_point
```

Debugging Clustered NFS

The `cfsshare` command logs error messages to the VCS logs in the `/var/VRTSvcs/log` directory. The `fsclustadm frlpause_query` command may be used to display the current local copy of the global state. This may be useful in debugging any issues with this Clustered NFS feature.

Common Internet File System

This chapter includes the following topics:

- [About Common Internet File System](#)
- [Requirements](#)
- [Understanding how Samba works](#)
- [Configuring Clustered NFS and Common Internet File System on CFS](#)
- [cfsshare manual page](#)
- [Configuring Common Internet File System in user mode](#)
- [Configuring Common Internet File System in domain mode](#)
- [Configuring Common Internet File System in ads mode](#)
- [Administering Common Internet File System](#)
- [Debugging Common Internet File System](#)

About Common Internet File System

This Common Internet File System (CIFS) feature lets you share CFS file systems using CIFS protocol that can be accessed by Window clients. Upon node failure or service group failover, the CIFS shares continue to be served by other cluster nodes.

Use the `cfsshare` command to configure your CIFS shares.

See “[cfsshare manual page](#)” on page 162.

Requirements

- CIFS user mode can run with the default Samba packages that are part of the OS.
- CIFS ads mode can run with the default Samba packages that are part of the OS and Kerberos version KRB5CLIENT_E.1.6.2.08.
- CIFS domain mode requires Samba version 3.2 or later.
- Prior knowledge of Samba is a prerequisite.

Understanding how Samba works

Samba is a networking tool that enables a UNIX server to participate in Windows networks. There are two parts in Samba, one being the server which shares out files and printers for other PC's to use, and the other being the client utilities, which allows a UNIX system to access files and printers on other Windows or Samba servers.

Configuring Clustered NFS and Common Internet File System on CFS

You can configure CNFS and CIFS on the same cluster. However, any given CFS file system can be shared using only one of the protocols (NFS or CIFS) at any given time.

cfsshare manual page

This CIFS feature adds a new functionality to the `cfsshare` utility in the VRTScavf package and several scripts that are added into the VCS configuration to manage parallel NFS and CIFS server resources. The `cfsshare` command modifies VCS resources that were created by other utilities such as `cfsmntadm`.

Note: The `cfsshare` command takes `nfs` protocol by default if the `-p` option is not used.

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

```
man -M /opt/VRTS/man cfsshare
```

Configuring Common Internet File System in user mode

This section describes how to configure CIFS in user mode.

In this mode, user authentication happens on the cluster nodes itself.

You must have NIS or some other mechanism configured on the cluster nodes to ensure the same users/groups have the same user/groups IDs on all cluster nodes.

A shared file system needs to be specified during the config operation. This file system is used to store the `smbpasswd` file, which contains the encrypted passwords for users. This way, users for whom passwords have been created on one cluster node, can authenticate themselves against other cluster nodes as well.

You must backup your existing `smb.conf` file and ensure that the `/var/log/samba` file exists on all cluster nodes, before running the `cfsshare config` command.

Note: If CNFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of CIFS.

```
cfsshare config [-n] -p cifs -l samba_lockdir -c \  
samba_config_file -t samba_topdir -m user \  
shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p cifs -m user -l /var/run \  
-c /etc/samba/smb.conf -t /usr lockdg vollocks /locks
```

Note: Check the path of the Samba lock directory and PID directory using the `smbd -b` command. Use the path of the PID directory for the `-l samba_lock` option.

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

If the `-n` option is specified when using the `cfsshare config` command, you must follow the steps to complete the CIFS configuration:

To complete the CIFS configuration when using the `-n` option

- 1 Copy the following lines to your `smb.conf` file:

```
security = user
passwd backend = smbpasswd
smbpasswd file = pvtdir/smbpasswd
```

where *pvt*dir is the private directory of your Samba installation.

- 2 Run the following command to backup your existing `smbpasswd` file:

```
# cp -f pvtdir/smbpasswd pvtdir/smbpasswd.OLD
```

- 3 Create a symlink in *pvt*dir, that points to the `smbpasswd` file created in the locks file system:

```
# ln -sf mntpt/cifs/smbpasswd pvtdir/smbpasswd
```

where *mnt*pt is the mount point.

To unconfigure CIFS run the following command:

```
# cfsshare unconfig -p cifs
```

Note: The unconfigure operation fails if any file systems or VIPs are being shared via CIFS protocol. Use the `cfsshare delete` command to stop sharing the file systems via CIFS. Use the `cfsshare deletevip` command to remove the VIPs.

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m user -l /var/run \  
-c /etc/samba/smb.conf -t /usr lockdg vollocks /locks
```

Configuring Common Internet File System in domain mode

This section describes how to configure CIFS in domain mode.

In this mode, user authentication happens on the NT4-style Domain controllers.

The cluster nodes act as member servers in the domain. You must have performed additional steps to make user and group mapping via winbind work.

A shared file system needs to be specified during the config operation. This file system is used to replicate the `secrets.tdb` file (machine password file) across all cluster nodes. Only one of the cluster nodes joins the domain using the cluster name. Once you have copied this file to all the cluster nodes, the Domain controller sees all cluster nodes as one member server.

The shared file system can also be used to store any `tdb` files that needs to be shared across all cluster nodes. Appropriate symlinks must be created on all cluster nodes.

You must backup your existing `smb.conf` file and ensure that the `/var/log/samba` file exists on all cluster nodes, before running the `cfsshare config` command.

Note: If CNFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of CIFS.

```
cfsshare config [-n] -p cifs -l samba_lockdir -c \  
samba_config_file -t samba_topdir -m domain \  
shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p cifs -m domain -l \  
/var/run -c /etc/samba/smb.conf -t /usr -s sfstest-ad \  
-d SFSTEST-AD2 -u Administrator lockdg vollocks /locks
```

Note: Check the path of the Samba lock directory and PID directory using the `smbd -b` command. Use the path of the PID directory for the `-l samba_lock` option.

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

If the `-n` option is specified when using the `cfsshare config` command, you must follow the steps to complete the CIFS configuration:

To complete the CIFS configuration when using the `-n` option

- 1 Copy the following lines to your `smb.conf` file:

```
security = domain
workgroup = domainname
password server = Domain_Controller_of_the_domain
```

- 2 Run the following command to backup your existing `secrets.tdb` file:

```
# mv -f pvtmdir/secrets.tdb pvtmdir/secrets.tdb.OLD
```

where `pvtmdir` is the private directory of your Samba installation.

- 3 Copy the `secrets.tdb` file created in the locks file system to the private directory of your Samba installation:

```
# cp -f mntpt/cifs/secrets.tdb pvtmdir/secrets.tdb
```

where `mntpt` is the mount point.

To unconfigure CIFS run the following command:

```
# cfsshare unconfig -p cifs
```

The unconfigure operation fails if any file systems are being shared via CIFS protocol.

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m domain -l \  
/var/run -c /etc/samba/smb.conf -t /usr -s sfstest-ad \  
-d SFSTEST-AD2 -u Administrator lockdg vollocks /locks
```

Configuring Common Internet File System in ads mode

This section describes how to configure CIFS in ads mode.

In this mode, user authentication happens on the active directory using Kerberos. The cluster nodes act as member servers in the domain. You must have performed additional steps to make user and group mapping via winbind work.

A shared file system needs to be specified during the config operation. This file system is used to replicate the `secrets.tdb` file (machine password file) across all cluster nodes. Only one of the cluster nodes joins the domain using the cluster name. Once you have copied this file to all the cluster nodes, the Domain controller sees all cluster nodes as one member server.

You must have configured Kerberos on all cluster nodes. The time on all cluster nodes needs to be synced up with the AD server/KDC.

The shared file system can also be used to store any `tdb` file that needs to be shared across all cluster nodes. Appropriate symlinks must be created on all cluster nodes.

You must backup your existing `smb.conf` file and ensure that the `/var/log/samba` file exists on all cluster nodes, before running the `cfsshare config` command.

Note: If CNFS is already configured on the cluster, then specify the same *shared_volume* and *mount_point* for configuration of CIFS.

```
cfsshare config [-n] -p cifs -l samba_lockdir -c \  
samba_config_file -t samba_topdir -m ads \  
shared_disk_group shared_volume mount_point
```

For example:

```
# cfsshare config -p cifs -m ads -l /var/run \  
-c /etc/samba/smb.conf -t /usr -s sfstest-ad -d \  
SFSTEST-AD2 -r SFSTEST-AD2.LOCAL -u Administrator \  
lockdg lockvol /locks
```

Note: Check the path of the Samba lock directory and PID directory using the `smbd -b` command. Use the path of the PID directory for the `-l samba_lock` option.

Note: The `cfsshare config` command fails if the *shared_volume* specified is already registered with VCS. Verify that the *shared_volume* is not registered with VCS by examining the output from the following command:

```
# /opt/VRTS/bin/cfsmntadm display
```

If the `-n` option is specified when using the `cfsshare config` command, you must follow the steps to complete the CIFS configuration:

To complete the CIFS configuration when using the `-n` option

- 1 Copy the following lines to your `smb.conf` file:

```
security = ads
workgroup = domainname
password server = AD_server_of_the_domain
realm = realm_name
```

- 2 Run the following command to backup your existing `secrets.tdb` file:

```
# mv -f pvtmdir/secrets.tdb pvtmdir/secrets.tdb.OLD
```

where `pvtmdir` is the private directory of your Samba installation.

- 3 Copy the `secrets.tdb` file created in the locks file system to the private directory of your Samba installation:

```
# cp -f mntpt/cifs/secrets.tdb pvtmdir/secrets.tdb
```

where `mntpt` is the mount point.

To unconfigure CIFS run the following command:

```
# cfsshare unconfig -p cifs
```

Note: The unconfigure operation fails if any file systems are being shared via CIFS protocol.

You can configure both CNFS and CIFS at the same time by running the `cfsshare config -p all` command.

```
# cfsshare config -p all -m ads -l /var/run \
-c /etc/samba/smb.conf -t /usr -s sfstest-ad -d \
SFSTEST-AD2 -r SFSTEST-AD2.LOCAL -u Administrator \
lockdg lockvol /locks
```

Administering Common Internet File System

To be able to access a CIFS exported file system from a Windows client you must first add a Virtual IP. You must specify this Virtual IP while sharing a file system via CIFS.

Adding a virtual IP:

```
cfsshare addvip device address netmask [networkhosts]
```

For example:

```
# cfsshare addvip lan0 10.182.79.216 \
255.255.240.0 10.182.79.215
```

The `cfsshare addvip` command lets you specify only one network interface, that is assumed to be present on all cluster nodes. If you want to specify different network interfaces for different cluster nodes, then you need to run certain VCS commands. For example:

```
# haconf -makerw
# hares -local vip1 Device
# hares -modify vip1 Device lan1 -sys system01
# hares -modify vip1 Device lan2 -sys system02
# haconf -dump -makero
```

where `vip1` is the Virtual IP resource created by the `cfsshare addvip` command.

where `system01` and `system02` are the cluster nodes.

Adding and sharing a file system via CIFS:

```
cfsshare add -p cifs -v address -n cifs_share_name \
shared_disk_group shared_volume mount_point \
share_options all=[ mount_options ]
```

For example:

```
# cfsshare add -p cifs -v 10.182.79.216 \
-n sh1 sharedg vol1 /mnt1 "readonly=no" all=
```

Note: You can also use the `-c` option to specify the CIFS share options.

See the `cfsshare(1M)` manual page for more information.

Run the following command to display the currently exported file systems:

```
# cfsshare display

SAMBASHARE RESOURCE      MOUNTPPOINT      SAMBASHARE OPTIONS
sambashare1              /mnt1            path=/mnt1;readonly=no
```

Deleting a previous shared file system:

```
cfsshare delete -p cifs mount_point
```

For example:

```
# cfsshare delete -p cifs /mnt1
```

Deleting the VIP added previously:

```
cfsshare deletevip address
```

For example:

```
# cfsshare deletevip 10.182.79.216
```

Sharing a CFS file system previously added to VCS

Use one of the following commands:

```
cfsshare share -p cifs -v address -n cifs share name \  
[-C cifs_share_options] mount_point
```

For example:

```
# cfsshare share -p cifs -v 10.182.79.216 -n sh1 -C readonly=no /mnt1
```

Or

```
cfsshare share -p cifs -v address -n cifs share name \  
mount_point [share_options]
```

For example:

```
# cfsshare share -p cifs -v 10.182.79.216 -n sh1 /mnt1 readonly=no
```

Note: You must have added the *address* using the `cfsshare addvip` command before running the `cfsshare share` command.

See the `cfsshare(1M)` manual page for more information.

Before running this command, you should have run `cfsmntadm` command to add the shared file system to VCS configuration and the `cfsmount` command to mount the shared file system at the *mount_point*. Once these commands have been executed, the CFSMount resource corresponding to the *mount_point* gets created in either a default service group (with a name similar to `vrts_vea_cfs_int_cfsmountnumber`) or in a separate service group, as specified by the user.

The `cfsshare share` command moves the CFSMount resource corresponding to the *mount_point* and the associated CVMVolDg resource to the `cfsnfssg` service group (that was created using the `config` option). This command also creates a

share resource on top of the CFSSMount resource in the same cfsnfssg service group.

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSSMount and CVMVolDg resources in the cfsnfssg service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Unsharing the previous shared CFS file system

```
cfsshare unshare mount_point
```

Before running this command, you should have run the `cfsshare share` command.

The `cfsshare unshare` command enables the user to stop sharing the file system mounted at the `mount_point`. This command moves the Share, CFSSMount, and CVMVolDg resources corresponding to the `mount_point` from cfsnfssg service group to a newly created service group. The SambaShare resource is taken offline and then deleted.

Note: VCS does not have the functionality to move resources across service groups. The `cfsshare` command creates new CFSSMount and CVMVolDg resources in the newly created service group and deletes the corresponding resources from the original service group.

The newly created resource names are different from the original resource names.

Running the `cfsmntadm delete` command does not remove the `ActivationMode` attribute. If no volumes or vsets in the disk group are in the VCS config, you must use the `cfsdgadm delete` to remove this `ActivationMode` attribute.

Sample main.cf

```
include "OracleASMTypes.cf"
include "types.cf"
include "ApplicationNone.cf"
include "CFSTypes.cf"
include "CVMTypes.cf"
include "Db2udbTypes.cf"
include "OracleTypes.cf"
include "SybaseTypes.cf"
```

```
cluster node1node2 (
    HacliUserLevel = COMMANDROOT
)

system node1 (
)

system node2 (
)

group cfsnfssg (
    SystemList = { node1 = 0, node2 = 1 }
    AutoFailOver = 0
    Parallel = 1
    AutoStartList = { node1, node2 }
)

Application Samba_winbind (
    StartProgram = "/opt/VRTSvcs/bin/ApplicationNone/winbindmonitor.sh start"
    StopProgram = "/opt/VRTSvcs/bin/ApplicationNone/winbindmonitor.sh stop"
    PidFiles = { "/var/run/winbindmonitor.pid" }
)

ApplicationNone app (
    MonitorProgram = "/opt/VRTSvcs/bin/ApplicationNone/lockdstatdmon"
)

CFSMount cfsmount2 (
    Critical = 0
    MountPoint = "/mnt1"
    BlockDevice = "/dev/vx/dsk/sharedg/voll1"
    NodeList = { node1, node2 }
)

CFSMount cfsnfs_locks (
    Critical = 0
    MountPoint = "/locks"
    BlockDevice = "/dev/vx/dsk/lockdg/vollocks"
    NodeList = { node1, node2 }
)

CVMVolDg cvmvoldg1 (
```

```
Critical = 0
CVMDiskGroup = lockdg
CVMVolume = { vollocks }
CVMActivation @node1 = sw
CVMActivation @node2 = sw
)

CVMVolDg cvmvoldg2 (
    Critical = 0
    CVMDiskGroup = sharedg
    CVMVolume = { voll }
    CVMActivation @node1 = sw
    CVMActivation @node2 = sw
)

NetBios Samba_netbios (
    SambaServerRes = SambaServerResource
    NetBiosName = node1node2
)

SambaServer SambaServerResource (
    ConfFile = "/etc/samba/smb.conf"
    LockDir = "/var/run"
)

SambaShare sambashare1 (
    SambaServerRes = SambaServerResource
    ShareName = sh1
    ShareOptions =
        "path=/mnt1;msdfsroot=yes;msdfsproxy=\\10.182.79.216\\sh1_dfs\\"
)

requires group cvm online local firm
Samba_winbind requires Samba_netbios
cfsmount2 requires cvmvoldg2
cfsnfs_locks requires cfsnfs_locks
cfsnfs_locks requires cvmvoldg1
sambashare1 requires SambaServerResource
sambashare1 requires cfsmount2

group cfsnfssg_dummy (
    SystemList = { node1 = 0, node2 = 1 }
    AutoFailOver = 0
```

```
Parallel = 1
AutoStartList = { node1, node2 }
)

requires group cvm online local firm

// resource dependency tree
//
// group cfsnfssg_dummy
// {
// }

group cvm (
  SystemList = { node1 = 0, node2 = 1 }
  AutoFailOver = 0
  Parallel = 1
  AutoStartList = { node1, node2 }
)

CFSfsckd vxfsckd (
  ActivationMode @node1 = { lockdg = sw }
  ActivationMode @node2 = { lockdg = sw }
)

CVMcluster cvm_clus (
  CVMClustName = node1node2
  CVMNodeId = { node1 = 1, node2 = 0 }
  CVMTransport = gab
  CVMTimeout = 200
)

CVMvxconfigd cvm_vxconfigd (
  Critical = 0
  CVMVxconfigdArgs = { syslog }
)

cvm_clus requires cvm_vxconfigd
vxfsckd requires cvm_clus

group vip1 (
```

```
SystemList = { node1 = 0, node2 = 1 }
AutoStartList = { node1, node2 }
PreOnline @node1 = 1
PreOnline @node2 = 1
)

IP vip1 (
    Device = lan0
    Address = "10.182.79.216"
    NetMask = "255.255.248.0"
)

NIC nic1 (
    Device = lan0
)

SambaShare sambasharel_dfs (
    SambaServerRes = SambaServerResource
    ShareName = sh1_dfs
    ShareOptions = "path=/mnt1;readonly=no"
)

requires group cfsnfssg online local firm
sambasharel_dfs requires vip1
vip1 requires nic1
```

Debugging Common Internet File System

The `cfsshare` command logs error messages to the VCS logs in the `/var/VRTSvcs/log` directory.

Troubleshooting SFCFS

This chapter includes the following topics:

- [About troubleshooting SFCFS](#)
- [Troubleshooting fenced configurations](#)
- [Troubleshooting I/O fencing](#)
- [Troubleshooting CVM](#)

About troubleshooting SFCFS

SFCFS contains several component products, and as a result can be affected by any issue with component products. The first step in case of trouble should be to identify the source of the problem. It is rare to encounter problems in SFCFS itself; more commonly the problem can be traced to setup issues or problems in component products.

Use the information in this chapter to diagnose the source of problems. Indications may point to SFCFS set up or configuration issues, in which case solutions are provided wherever possible. In cases where indications point to a component product or to Oracle as the source of a problem, it may be necessary to refer to the appropriate documentation to resolve it.

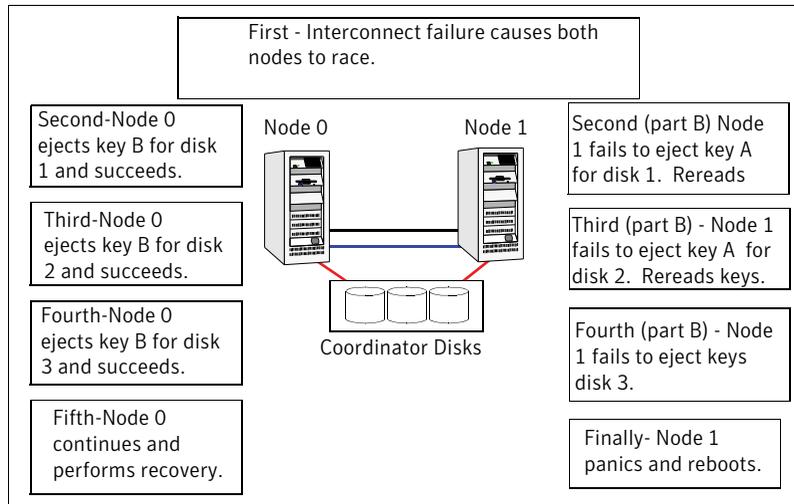
Troubleshooting fenced configurations

The following information describes network partitioning in a fenced environment. See the *Veritas Cluster Server Administrator's Guide*.

Example of a pre-existing network partition (split-brain)

Figure 7-1 shows a two-node cluster in which the severed cluster interconnect poses a potential split-brain condition.

Figure 7-1 Pre-existing network partition (split-brain)



Because the fencing module operates identically on each system, both nodes assume the other is failed, and carry out fencing operations to insure the other node is ejected. The VCS GAB module on each node determines the peer has failed due to loss of heartbeats and passes the membership change to the fencing module.

Each side “races” to gain control of the coordinator disks. Only a registered node can eject the registration of another node, so only one side successfully completes the command on each disk.

The side that successfully ejects the peer from a majority of the coordinator disks wins. The fencing module on the winning side then passes the membership change up to VCS and other higher-level packages registered with the fencing module, allowing VCS to invoke recovery actions. The losing side forces a kernel panic and reboots.

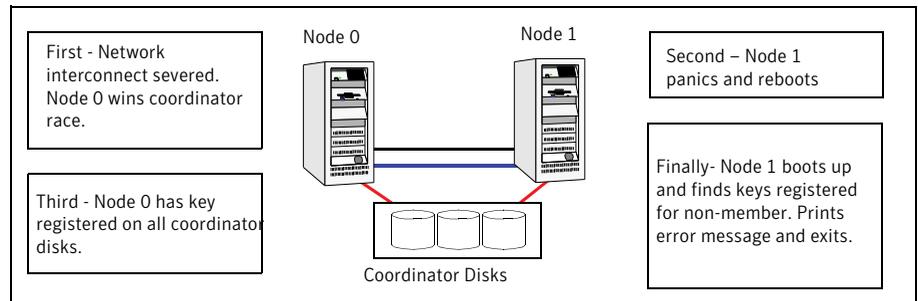
Recovering from a pre-existing network partition (split-brain)

The fencing module `vxfen` prevents a node from starting up after a network partition and subsequent panic and reboot of a node.

Example Scenario I

Figure 7-2 scenario could cause similar symptoms on a two-node cluster with one node shut down for maintenance. During the outage, the private interconnect cables are disconnected.

Figure 7-2 Example scenario I



In example scenario I, the following occurs:

- Node 0 wins a coordinator race following to a network failure.
- Node 1 panics and reboots.
- Node 0 has keys registered on the coordinator disks. When Node 1 boots up, it sees the Node 0 keys, but cannot see Node 0 in the current GAB membership. It senses a potential preexisting split brain and causes the vxfen module to print an error message to the console. The vxfen module prevents fencing from starting, which, in turn, prevents VCS from coming online.
 Suggested solution: Shut down Node 1, reconnect the cables, and restart Node 1.

Example Scenario II

Similar to example scenario I, if private interconnect cables are disconnected in a two-node cluster, Node 1 is fenced out of the cluster, panics, and reboots. If before the private interconnect cables are fixed and Node 1 rejoins the cluster, Node 0 reboots and remote (or just reboots). No node can write to the data disks until the private networks are fixed. This is because GAB membership cannot be formed, therefore the cluster cannot be formed.

Suggested solution: Shut down both nodes, reconnect the cables, restart the nodes.

Example Scenario III

Similar to example scenario II, if private interconnect cables are disconnected in a two-node cluster, Node 1 is fenced out of the cluster, panics, and reboots. If

before the private interconnect cables are fixed and Node 1 rejoins the cluster, Node 0 panics due to hardware failure and cannot come back up, Node 1 cannot rejoin.

Suggested solution: Shut down Node 1, reconnect the cables, restart the node. You must then clear the registration of Node 0 from the coordinator disks.

To fix scenario III

- 1 On Node 1, type the following command:

```
# /opt/VRTSvcS/vxfen/bin/vxfenclearpre
```

- 2 Restart the node.

Troubleshooting I/O fencing

The following sections discuss troubleshooting the I/O fencing problems. Review the symptoms and recommended solutions.

SCSI reservation errors during bootup

When restarting a node of an SFCFS cluster, SCSI reservation errors may be observed such as:

```
date system name kernel: scsi3 (0,0,6) : RESERVATION CONFLICT
```

This message is printed for each disk that is a member of any shared disk group which is protected by SCSI-3 PR I/O fencing. This message may be safely ignored.

The vxfentsthdw utility fails when SCSI TEST UNIT READY command fails

While running the vxfentsthdw utility, you may see a message that resembles as follows:

```
Issuing SCSI TEST UNIT READY to disk reserved by other node  
FAILED.  
Contact the storage provider to have the hardware configuration  
fixed.
```

The disk array does not support returning success for a SCSI TEST UNIT READY command when another host has the disk reserved using SCSI-3 persistent reservations. This happens with the Hitachi Data Systems 99XX arrays if bit 186 of the system mode option is not enabled.

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

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

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

If you see these messages when the new node is booting, the vxfen startup script on the node makes up to five attempts to join the cluster.

To manually join the node to the cluster when I/O fencing attempts fail

- ◆ If the vxfen script fails in the attempts to allow the node to join the cluster, restart vxfen driver with the command:

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

If the command fails, restart the new node.

System panics to prevent potential data corruption

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

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

See [“How vxfen driver checks for preexisting split-brain condition”](#) on page 181.

How vxfen driver checks for preexisting split-brain condition

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

For example, suppose the cluster of system 1 and system 2 is functioning normally when the private network links are broken. Also suppose system 1 is the ejected system. When system 1 restarts before the private network links are restored, its membership configuration does not show system 2; however, when it attempts to register with the coordinator disks, it discovers system 2 is registered with them. Given this conflicting information about system 2, system 1 does not join the cluster and returns an error from vxfenconfig that resembles:

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

```
I/O Fencing Disabled!
```

Note: During the system boot, because the HP-UX rc sequencer redirects the stderr of all rc scripts to the file /etc/rc.log, the error messages will not be printed on the console. It will be logged in the /etc/rc.log file.

Also, the following information is displayed on the console:

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

Note: If syslogd is configured with the `-D` option, then the informational message will not be printed on the console. The messages will be logged in the system buffer. The system buffer can be read with the `dmesg` command.

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

To resolve actual and apparent potential split-brain conditions

- ◆ Depending on the split-brain condition that you encountered, do the following:

- | | | |
|---|---|---|
| Actual potential split-brain condition—system 2 is up and system 1 is ejected | 1 | Determine if system1 is up or not. |
| | 2 | If system 1 is up and running, shut it down and repair the private network links to remove the split-brain condition. |
| | 3 | Restart system 1. |

Apparent potential split-brain condition—system 2 is down and system 1 is ejected

- 1 Physically verify that system 2 is down.
 Verify the systems currently registered with the coordinator disks. Use the following command:

```
# vxfenadm -s all -f /etc/vxfentab
```

 The output of this command identifies the keys registered with the coordinator disks.
- 2 Clear the keys on the coordinator disks as well as the data disks using the `vxfenclearpre` command.
 See [“Clearing keys after split-brain using vxfenclearpre command”](#) on page 184.
- 3 Make any necessary repairs to system 2.
- 4 Restart system 2.

Cluster ID on the I/O fencing key of coordinator disk does not match the local cluster’s ID

If you accidentally assign coordinator disks of a cluster to another cluster, then the fencing driver displays an error message similar to the following when you start I/O fencing:

```
000068 06:37:33 2bdd5845 0 ... 3066 0 VXFEN WARNING V-11-1-56
Coordinator disk has key with cluster id 48813
which does not match local cluster id 57069
```

The warning implies that the local cluster with the cluster ID 57069 has keys. However, the disk also has keys for cluster with ID 48813 which indicates that nodes from the cluster with cluster ID 48813 potentially use the same coordinator disk.

You can run the following commands to verify whether these disks are used by another cluster. Run the following commands on one of the nodes in the local cluster. For example, on galaxy:

```
galaxy> # lltstat -C
57069

galaxy> # cat /etc/vxfentab
/dev/vx/rdmp/disk_7
/dev/vx/rdmp/disk_8
```

```
/dev/vx/rdmp/disk_9

galaxy> # vxfenadm -s /dev/vx/rdmp/disk_7
Reading SCSI Registration Keys...
Device Name: /dev/vx/rdmp/disk_7
Total Number Of Keys: 1
key[0]:
  [Numeric Format]: 86,70,48,49,52,66,48,48
  [Character Format]: VFBEAD00
  [Node Format]: Cluster ID: 48813 Node ID: 0 Node Name: unknown
```

where *disk_7*, *disk_8*, and *disk_9* represent the disk names in your setup.

Recommended action: You must use a unique set of coordinator disks for each cluster. If the other cluster does not use these coordinator disks, then clear the keys using the `vxfcntlpre` command before you use them as coordinator disks in the local cluster.

See [“About the vxfcntlpre utility”](#) on page 85.

Clearing keys after split-brain using vxfcntlpre command

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

See [“About the vxfcntlpre utility”](#) on page 85.

Registered keys are lost on the coordinator disks

If the coordinator disks lose the keys that are registered, the cluster might panic when a cluster reconfiguration occurs.

To refresh the missing keys

- ◆ Use the `vxfenswap` utility to replace the coordinator disks with the same disks. The `vxfenswap` utility registers the missing keys during the disk replacement.

See [“Refreshing lost keys on coordinator disks”](#) on page 95.

Replacing defective disks when the cluster is offline

If the disk becomes defective or inoperable and you want to switch to a new diskgroup in a cluster that is offline, then perform the following procedure.

In a cluster that is online, you can replace the disks using the `vxfenswap` utility.

See “[About the vxfenswap utility](#)” on page 88.

Review the following information to replace coordinator disk in the coordinator disk group, or to destroy a coordinator disk group.

Note the following about the procedure:

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

To replace a disk in the coordinator disk group when the cluster is offline

- 1 Log in as superuser on one of the cluster nodes.
- 2 If VCS is running, shut it down:

```
# hstop -all
```

Make sure that the port `h` is closed on all the nodes. Run the following command to verify that the port `h` is closed:

```
# gabconfig -a
```

- 3 Stop I/O fencing on each node:

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

This removes any registration keys on the disks.

- 4 Import the coordinator disk group. The file `/etc/vxfendg` includes the name of the disk group (typically, `vxfencoorddg`) that contains the coordinator disks, so use the command:

```
# vxdg -tfc import `cat /etc/vxfendg`
```

where:

- t specifies that the disk group is imported only until the node restarts.
- f specifies that the import is to be done forcibly, which is necessary if one or more disks is not accessible.
- C specifies that any import locks are removed.

- 5 To remove disks from the disk group, use the VxVM disk administrator utility, `vxdiskadm`.

You may also destroy the existing coordinator disk group. For example:

- Verify whether the coordinator attribute is set to on.

```
# vxdg list vxfencoorddg | grep flags: | grep coordinator
```

- Destroy the coordinator disk group.

```
# vxdg -o coordinator destroy vxfencoorddg
```

- 6 Add the new disk to the node and initialize it as a VxVM disk.

Then, add the new disk to the vxfencoorddg disk group:

- If you destroyed the disk group in step 5, then create the disk group again and add the new disk to it.

- If the diskgroup already exists, then add the new disk to it.

```
# vxdg -g vxfencoorddg -o coordinator adddisk disk_name
```

- 7 Test the recreated disk group for SCSI-3 persistent reservations compliance.

See [“Testing the coordinator disk group using vxfentsthdw -c option”](#) on page 75.

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

```
# vxdg deport `cat /etc/vxfendg`
```

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

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

- 10 Verify that the I/O fencing module has started and is enabled.

```
# gabconfig -a
```

Make sure that port b membership exists in the output for all nodes in the cluster.

```
# vxfenadm -d
```

Make sure that I/O fencing mode is not disabled in the output.

- 11 If necessary, restart VCS on each node:

```
# hastart
```

Troubleshooting server-based I/O fencing

All CP server operations and messages are logged in the `/var/VRTScps/log` directory in a detailed and easy to read format. The entries are sorted by date and time. The logs can be used for troubleshooting purposes or to review for any possible security issue on the system that hosts the CP server.

The following files contain logs and text files that may be useful in understanding and troubleshooting a CP server:

- `/var/VRTScps/log/cpserver_[ABC].log`
- `/var/VRTSat/vrtsat_broker.txt` (Security related)
- If the `vxcperv` process fails on the CP server, then review the following diagnostic files:
 - `/var/VRTScps/diag/FFDC_CPS_pid_vxcperv.log`
 - `/var/VRTScps/diag/stack_pid_vxcperv.txt`

Note: If the `vxcperv` process fails on the CP server, these files are present in addition to a core file. VCS restarts `vxcperv` process automatically in such situations.

The file `/var/VRTSvcsvcs/log/vxfen/vxfend_[ABC].log` contains logs and text files that may be useful in understanding and troubleshooting fencing-related issues on a SFCFS cluster (client cluster) node.

See [“Troubleshooting issues related to the CP server service group”](#) on page 187.

See [“Checking the connectivity of CP server”](#) on page 190.

See [“Issues during fencing startup on SFCFS cluster nodes set up for server-based fencing”](#) on page 188.

See [“Issues during online migration of coordination points”](#) on page 191.

See [“Troubleshooting server-based I/O fencing in mixed mode”](#) on page 192.

See [“Checking keys on coordination points when `vxfen_mechanism` value is set to `cps`”](#) on page 195.

Troubleshooting issues related to the CP server service group

If you cannot bring up the CPSSG service group after the CP server configuration, perform the following steps:

- Verify that the CPSSG service group and its resources are valid and properly configured in the VCS configuration.

- Check the VCS engine log (`/var/VRTSvcs/log/engine_[ABC].log`) to see if any of the CPSSG service group resources are FAULTED.
- Review the sample dependency graphs to make sure the required resources are configured correctly.

Troubleshooting server-based fencing on the SFCFS cluster nodes

The file `/var/VRTSvcs/log/vxfen/vxfend_[ABC].log` contains logs and text files that may be useful in understanding and troubleshooting fencing-related issues on a SFCFS cluster (client cluster) node.

Issues during fencing startup on SFCFS cluster nodes set up for server-based fencing

Table 7-1 Fencing startup issues on SFCFS cluster (client cluster) nodes

Issue	Description and resolution
cpsadm command on the SFCFS cluster gives connection error	<p>If you receive a connection error message after issuing the <code>cpsadm</code> command on the SFCFS cluster, perform the following actions:</p> <ul style="list-style-type: none"> ■ Ensure that the CP server is reachable from all the SFCFS cluster nodes. ■ Check that the SFCFS cluster nodes use the correct CP server virtual IP or virtual hostname and the correct port number. Check the <code>/etc/vxfenmode</code> file. ■ Ensure that the running CP server is using the same virtual IP/virtual hostname and port number.
Authorization failure	<p>Authorization failure occurs when the CP server's nodes or users are not added in the CP server configuration. Therefore, fencing on the SFCFS cluster (client cluster) node is not allowed to access the CP server and register itself on the CP server. Fencing fails to come up if it fails to register with a majority of the coordination points.</p> <p>To resolve this issue, add the CP server node and user in the CP server configuration and restart fencing.</p>

Table 7-1 Fencing startup issues on SFCFS cluster (client cluster) nodes
(continued)

Issue	Description and resolution
Authentication failure	<p>If you had configured secure communication between the CP server and the SFCFS cluster (client cluster) nodes, authentication failure can occur due to the following causes:</p> <ul style="list-style-type: none"> ■ Symantec Product Authentication Services (AT) is not properly configured on the CP server and/or the SFCFS cluster. ■ The CP server and the SFCFS cluster nodes use the same root broker but the certificate hash of the root broker is not same on the SFCFS cluster and the CP server. Run the following command on both the CP server and the SFCFS cluster to see the certificate hash: <pre># cpsat showalltrustedcreds</pre> ■ The CP server and the SFCFS cluster nodes use different root brokers, and trust is not established between the authentication brokers: ■ The hostname of the SFCFS cluster nodes is not the same hostname used when configuring AT. The hostname of the SFCFS cluster nodes must be set to the hostname used when configuring AT. You can view the fully qualified hostname registered with AT using the <code>cpsat showcred</code> command. After entering this command, the hostname appears in the User Name field. ■ The CP server and SFCFS cluster do not have the same security setting. In order to configure secure communication, both the CP server and the SFCFS cluster must have same security setting. In order to have the same security setting, the security parameter must have same value in the <code>/etc/vxcps.conf</code> file on CP server and in the <code>/etc/vxfenmode</code> file on the SFCFS cluster (client cluster) nodes.

Table 7-1 Fencing startup issues on SFCFS cluster (client cluster) nodes
(continued)

Issue	Description and resolution
Preexisting split-brain	<p>Assume the following situations to understand preexisting split-brain in server-based fencing:</p> <ul style="list-style-type: none"> ■ There are three CP servers acting as coordination points. One of the three CP servers then becomes inaccessible. While in this state, also one client node leaves the cluster. When the inaccessible CP server restarts, it has a stale registration from the node which left the SFCFS cluster. In this case, no new nodes can join the cluster. Each node that attempts to join the cluster gets a list of registrations from the CP server. One CP server includes an extra registration (of the node which left earlier). This makes the joiner node conclude that there exists a preexisting split-brain between the joiner node and the node which is represented by the stale registration. ■ All the client nodes have crashed simultaneously, due to which fencing keys are not cleared from the CP servers. Consequently, when the nodes restart, the vxfen configuration fails reporting preexisting split brain. <p>These situations are similar to that of preexisting split-brain with coordinator disks, where the problem is solved by the administrator running the <code>vxfenclearpre</code> command. A similar solution is required in server-based fencing using the <code>cpsadm</code> command.</p> <p>Run the <code>cpsadm</code> command to clear a registration on a CP server:</p> <pre># cpsadm -s cp_server -a unreg_node -c cluster_name -n nodeid</pre> <p>where <code>cp_server</code> is the virtual IP address or virtual hostname on which the CP server is listening, <code>cluster_name</code> is the VCS name for the SFCFS cluster, and <code>nodeid</code> specifies the node id of SFCFS cluster node. Ensure that fencing is not already running on a node before clearing its registration on the CP server.</p> <p>After removing all stale registrations, the joiner node will be able to join the cluster.</p>

Checking the connectivity of CP server

You can test the connectivity of CP server using the `cpsadm` command.

You must have set the environment variables `CPS_USERNAME` and `CPS_DOMAINTYPE` to run the `cpsadm` command on the SFCFS cluster (client cluster) nodes.

To check the connectivity of CP server

- ◆ Run the following command to check whether a CP server is up and running at a process level:

```
# cpsadm -s cp_server -a ping_cps
```

where *cp_server* is the virtual IP address or virtual hostname on which the CP server is listening.

Issues during online migration of coordination points

During online migration of coordination points using the `vxfenswap` utility, the operation is automatically rolled back if a failure is encountered during validation of coordination points from all the cluster nodes.

Validation failure of the new set of coordination points can occur in the following circumstances:

- The `/etc/vxfenmode` file is not updated on all the SFCFS cluster nodes, because new coordination points on the node were being picked up from an old `/etc/vxfenmode` file.
- The coordination points listed in the `/etc/vxfenmode` file on the different SFCFS cluster nodes are not the same. If different coordination points are listed in the `/etc/vxfenmode` file on the cluster nodes, then the operation fails due to failure during the coordination point snapshot check.
- There is no network connectivity from one or more SFCFS cluster nodes to the CP server(s).
- Cluster, nodes, or users for the SFCFS cluster nodes have not been added on the new CP servers, thereby causing authorization failure.

Vxfen service group activity after issuing the vxfenswap command

After issuing the `vxfenswap` command, the Coordination Point agent reads the details of coordination points from the `vxfenconfig -l` output and starts monitoring the registrations on them.

During `vxfenswap`, when the `vxfenmode` file is being changed by the user, the Coordination Point agent does not move to FAULTED state but continues monitoring the old set of coordination points.

As long as the changes to `vxfenmode` file are not committed or the new set of coordination points are not re-elected in `vxfenconfig -l` output, the Coordination

Point agent continues monitoring the old set of coordination points it read from `vxfenconfig -l` output in every monitor cycle.

The status of the Coordination Point agent (either ONLINE or FAULTED) depends upon the accessibility of the coordination points, the registrations on these coordination points, and the fault tolerance value.

When the changes to `vxfenmode` file are committed and reflected in the `vxfenconfig -l` output, then the Coordination Point agent reads the new set of coordination points and proceeds to monitor them in its new monitor cycle.

Troubleshooting server-based I/O fencing in mixed mode

Use the following procedure to troubleshoot a mixed I/O fencing configuration (configuration which uses both coordinator disks and CP server for I/O fencing).

This procedure uses the following commands to obtain I/O fencing information:

- To obtain I/O fencing cluster information on the coordinator disks, run the following command on one of the cluster nodes:

```
# vxfenadm -s diskname
```

Any keys other than the valid keys used by the cluster nodes that appear in the command output are spurious keys.

- To obtain I/O fencing cluster information on the CP server, run the following command on one of the cluster nodes:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster name* is the VCS name for the SFCFS cluster. Nodes which are not in GAB membership, but registered with CP server indicate a pre-existing network partition.

Note that when running this command on the SFCFS cluster nodes, you need to first export the `CPS_USERNAME` and `CPS_DOMAINTYPE` variables.

The `CPS_USERNAME` value is the user name which is added for this node on the CP server.

- To obtain the user name, run the following command on the CP server:

```
# cpsadm -s cp_server -a list_users
```

where *cp server* is the virtual IP address or virtual hostname on which the CP server is listening.

The `CPS_DOMAINTYPE` value is `vx`.

The following are export variable command examples:

```
# export CPS_USERNAME=_HA_VCS_test-system@HA_SERVICES@test-system.symantec.com

# export CPS_DOMAINTYPE=vx
```

Once a pre-existing network partition is detected using the above commands, all spurious keys on the coordinator disks or CP server must be removed by the administrator.

To troubleshoot server-based I/O fencing configuration in mixed mode

- 1 Review the current I/O fencing configuration by accessing and viewing the information in the vxfenmode file.

Enter the following command on one of the SFCFS cluster nodes:

```
# cat /etc/vxfenmode

vxfen_mode=customized
vxfen_mechanism=cps
scsi3_disk_policy=dmp
security=0
cps1=[10.140.94.101]:14250
vxfendg=vxfencoorddg
```

- 2 Review the I/O fencing cluster information.

Enter the vxfenadm -d command on one of the cluster nodes:

```
# vxfenadm -d

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

Fencing Protocol Version: 201
Fencing Mode: Customized
Fencing Mechanism: cps
Cluster Members:

    * 0 (galaxy)
      1 (nebula)

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

- 3 Review the SCSI registration keys for the coordinator disks used in the I/O fencing configuration.

The variables *disk_7* and *disk_8* in the following commands represent the disk names in your setup.

Enter the `vxfenadm -s` command on each of the SFCFS cluster nodes.

```
# vxfenadm -s /dev/vx/rdmp/disk_7
```

```
Device Name: /dev/vx/rdmp/disk_7
Total Number Of Keys: 2
key[0]:
    [Numeric Format]: 86,70,66,69,65,68,48,48
    [Character Format]: VFBEAD00
    [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
    [Numeric Format]: 86,70,66,69,65,68,48,49
    [Character Format]: VFBEAD01
*    [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

Run the command on the other node:

```
# vxfenadm -s /dev/vx/rdmp/disk_8
```

```
Device Name: /dev/vx/rdmp/disk_8
Total Number Of Keys: 2
key[0]:
    [Numeric Format]: 86,70,66,69,65,68,48,48
    [Character Format]: VFBEAD00
    [Node Format]: Cluster ID: 57069 Node ID: 0 Node Name: galaxy
key[1]:
    [Numeric Format]: 86,70,66,69,65,68,48,49
    [Character Format]: VFBEAD01
*    [Node Format]: Cluster ID: 57069 Node ID: 1 Node Name: nebula
```

- 4 Review the CP server information about the cluster nodes. On the CP server, run the `cpsadm list nodes` command to review a list of nodes in the cluster.

```
# cpsadm -s cp_server -a list_nodes
```

where *cp_server* is the virtual IP address or virtual hostname on which the CP server is listening.

- 5 Review the CP server list membership. On the CP server, run the following command to review the list membership.

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp_server* is the virtual IP address or virtual hostname on which the CP server is listening, and *cluster_name* is the VCS name for the SFCFS cluster.

For example:

```
# cpsadm -s 10.140.94.101 -a list_membership -c gl-ss2
```

```
List of registered nodes: 0 1
```

Checking keys on coordination points when `vxfen_mechanism` value is set to `cps`

When I/O fencing is configured in customized mode and the `vxfen_mechanism` value is set to `cps`, the recommended way of reading keys from the coordination points (coordinator disks and CP servers) is as follows:

- For coordinator disks, the disks can be put in a file and then information about them supplied to the `vxfenadm` command.

For example:

```
# vxfenadm -s all -f file_name
```

- For CP servers, the `cpsadm` command can be used to obtain the membership of the SFCFS cluster.

For example:

```
# cpsadm -s cp_server -a list_membership -c cluster_name
```

where *cp_server* is the virtual IP address or virtual hostname on which CP server is configured, and *cluster_name* is the VCS name for the SFCFS cluster.

Troubleshooting CVM

This section discusses troubleshooting CVM problems.

CVM group is not online after adding a node to the cluster

The possible causes for the CVM group being offline after adding a node to the cluster are as follows:

- The `cssd` resource is configured as a critical resource in the `cvm` group.
- Other resources configured in the `cvm` group as critical resources are not online.

To resolve the issue if `cssd` is configured as a critical resource

- 1 Log onto one of the nodes in the existing cluster as the root user.
- 2 Configure the `cssd` resource as a non-critical resource in the `cvm` group:

```
# haconf -makerw
# hares -modify cssd Critical 0
# haconf -dump -makero
```

To resolve the issue if other resources in the group are not online

- 1 Log onto one of the nodes in the existing cluster as the root user.
- 2 Bring the resource online:

```
# hares -online resource_name -sys system_name
```

- 3 Verify the status of the resource:

```
# hastatus -resource resource_name
```

- 4 If the resource is not online, configure the resource as a non-critical resource:

```
# haconf -makerw
# hares -modify resource_name Critical 0
# haconf -dump -makero
```

Shared disk group cannot be imported

If you see a message resembling:

```
vxvm:vxconfigd:ERROR:vold_pgr_register(/dev/vx/rdmp/disk_name):
local_node_id<0
```

Please make sure that CVM and vxfen are configured and operating correctly

First, make sure that CVM is running. You can see the CVM nodes in the cluster by running the `vxclustadm nidmap` command.

```
# vxclustadm nidmap
Name          CVM Nid    CM Nid     State
galaxy        1          0          Joined: Master
nebula        0          1          Joined: Slave
```

This above output shows that CVM is healthy, with system galaxy as the CVM master. If CVM is functioning correctly, then the output above is displayed when CVM cannot retrieve the node ID of the local system from the `vxfen` driver. This usually happens when port `b` is not configured.

To verify vxfen driver is configured

- ◆ Check the GAB ports with the command:

```
# gabconfig -a
```

Port `b` must exist on the local system.

Error importing shared disk groups

The following message may appear when importing shared disk group:

```
VxVM vxdg ERROR V-5-1-587 Disk group disk group name: import
failed: No valid disk found containing disk group
```

You may need to remove keys written to the disk.

For information about removing keys written to the disk:

See [“Removing preexisting keys”](#) on page 85.

Unable to start CVM

If you cannot start CVM, check the consistency between the `/etc/llthosts` and `main.cf` files for node IDs.

You may need to remove keys written to the disk.

For information about removing keys written to the disk:

See [“Removing preexisting keys”](#) on page 85.

CVMVolDg not online even though CVMCluster is online

When the CVMCluster resource goes online, then all shared disk groups that have the auto-import flag set are automatically imported. If the disk group import fails for some reason, the CVMVolDg resources fault. Clearing and taking the CVMVolDg type resources offline does not resolve the problem.

To resolve the resource issue

- 1 Fix the problem causing the import of the shared disk group to fail.
- 2 Offline the cvm group containing the resource of type CVMVolDg as well as the service group containing the CVMCluster resource type.
- 3 Bring the cvm group containing the CVMCluster resource online.
- 4 Bring the cvm group containing the CVMVolDg resource online.

Creating a starter database

This appendix includes the following topics:

- [Creating a database for Oracle 10g or 11g](#)

Creating a database for Oracle 10g or 11g

Create a database tablespace for Oracle 10g or 11g using one of the two options:

- Option 1: on shared raw VxVM volumes
- Option 2: on cluster file system (CFS)

Before you begin, take note of the following prerequisites:

- CRS daemons must be running. To verify the status of CRS, enter:

```
# $CRS_HOME/bin/crs_stat
```

- Use the `ping` command to verify that all private IP addresses on each node are up.

Creating database tablespace on shared raw VxVM volumes (option 1)

This section describes how to create database tablespace on shared raw VxVM volumes (option 1).

To create database tablespace on shared raw VxVM volumes (option 1)

- 1 On any cluster node, log in as `root`. Find out the spare disks that can be used for creating shared disk group for Oracle database tablespaces, enter:

```
# vxdisk -o alldgs list
DEVICE TYPE          DISK  GROUP    STATUS
sda   auto:none         -    -        online invalid
sdb   auto:none         -    -        online invalid
sdc   auto:cdsdisk     -    tempdg   online shared
sdd   auto:cfsdisk     -    ocrvotedg online shared
sde   auto:cdsdisk     -    -        online shared
sdf   auto:cdsdisk     -    -        online shared
```

The above sample output indicates that shared disks `sde` and `sdf` are free and can be used for Oracle database tablespaces.

- 2 Create a shared disk group:

```
# vxdg -s init oradatadg sde sdf
```

- 3 Create a volume in the shared group for each of the required tablespaces. See the *Oracle* documentation specific to the Oracle database release to determine the tablespace requirements.

For example, enter:

```
# vxassist -g oradatadg make VRT_system01 1000M
# vxassist -g oradatadg make VRT_system02 10M
.
.
.
```

- 4 Define the access mode and permissions for the volumes storing the Oracle data. For each volume listed in `$ORACLE_HOME/raw_config`, use the `vxedit` command:

```
# vxedit -g disk_group set group=group user=user mode=660
    volume
```

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

For example, enter:

```
# vxedit -g oradatadg set group=oinstall user=oracle mode=660 \
VRT_system01
```

In this example, `VRT_system01` is the name of one of the volumes. Repeat the command to define access mode and permissions for each volume in the `oradatadg`.

- 5 Create the database.

See the *Oracle* documentation.

Creating database tablespace on CFS (option 2)

This section describes how to create database tablespace on CFS (option 2). If you plan to use a cluster file system to store the Oracle database, use the following procedure to create the file system.

To creating database tablespace on CFS (option 2)

- 1 On any cluster node, log in as `root`. Find out the spare disks that can be used for creating shared disk group for Oracle database tablespaces, enter:

```
# vxdisk -o allldgs list
DEVICE TYPE          DISK  GROUP      STATUS
sda    auto:none         -    -          online invalid
sdb    auto:none         -    -          online invalid
sdc    auto:cdsdisk      -    tempdg    online shared
sdd    auto:cfsdisk      -    ocrvotedg online shared
sde    auto:cdsdisk      -    -          online shared
sdf    auto:cdsdisk      -    -          online shared
```

The above sample output indicates that shared disks `sde` and `sdf` are free and can be used for Oracle database tablespaces.

- 2 Create a shared disk group. For example, enter:

```
# vxdg -s init oradatadg sde
```

- 3 Create a single shared volume that is large enough to contain a file system for all tablespaces.

See the *Oracle* documentation specific to the Oracle database release for tablespace sizes.

Assuming 6.8 GB are required for the tablespaces, enter:

```
# vxassist -g oradatadg make oradatavol 6800M
```

- 4 Create a VxFS file system in this volume, enter:

```
# mkfs -F vxfs /dev/vx/rdisk/oradatadg/oradatavol
```

- 5 Create a mount point for the shared file system, enter:

```
# mkdir /oradata
```

- 6 From the same node, mount the file system, enter:

```
# mount -F vxfs -o cluster /dev/vx/dsk/oradatadg/oradatavol \  
/oradata
```

- 7 Set `oracle` as the owner of the file system, and set `775` as the permissions:

```
# chown oracle:oinstall /oradata  
# chmod 775 /oradata
```

- 8 On the other node(s), complete steps 5 through 6.

- 9 Create the Oracle database.

See the *Oracle* documentation.

Glossary

ACL (access control list)	The information that identifies specific users or groups and their access privileges for a particular file or directory.
agent	A process that manages predefined Veritas Cluster Server (VCS) resource types. Agents bring resources online, take resources offline, and monitor resources to report any state changes to VCS. When an agent is started, it obtains configuration information from VCS and periodically monitors the resources and updates VCS with the resource status.
allocation unit	A group of consecutive blocks on a file system that contain resource summaries, free resource maps, and data blocks. Allocation units also contain copies of the super-block.
API	Application Programming Interface.
asynchronous writes	A delayed write in which the data is written to a page in the system's page cache, but is not written to disk before the write returns to the caller. This improves performance, but carries the risk of data loss if the system crashes before the data is flushed to disk.
atomic operation	An operation that either succeeds completely or fails and leaves everything as it was before the operation was started. If the operation succeeds, all aspects of the operation take effect at once and the intermediate states of change are invisible. If any aspect of the operation fails, then the operation aborts without leaving partial changes.
BLI (block-level incremental) backup	A Veritas backup capability that does not store and retrieve entire files. Instead, only the data blocks that have changed since the previous backup are backed up.
boot disk	A disk that is used for the purpose of booting a system.
boot disk group	A private disk group that contains the disks from which the system may be booted.
buffered I/O	A mode of I/O operation (where I/O is any operation, program, or device that transfers data to or from a computer) that first transfers data into the Operating System buffer cache.
bootdg	A reserved disk group name that is an alias for the name of the boot disk group.
cluster mounted file system	A shared file system that enables multiple hosts to mount and perform file operations on the same file. A cluster mount requires a shared storage device that can be accessed by other cluster mounts of the same file system. Writes to the

shared device can be done concurrently from any host on which the cluster file system is mounted. To be a cluster mount, a file system must be mounted using the `mount -o cluster` option.

Cluster Services	The group atomic broadcast (GAB) module in the SFCFS stack provides cluster membership services to the file system. LLT provides kernel-to-kernel communications and monitors network communications.
contiguous file	A file in which data blocks are physically adjacent on the underlying media.
CVM (cluster volume manager)	The cluster functionality of Veritas Volume Manager.
CVM Master	The cluster volume manager has a master node that records changes to the volume configuration.
data block	A block that contains the actual data belonging to files and directories.
data synchronous writes	A form of synchronous I/O that writes the file data to disk before the write returns, but only marks the inode for later update. If the file size changes, the inode will be written before the write returns. In this mode, the file data is guaranteed to be on the disk before the write returns, but the inode modification times may be lost if the system crashes.
defragmentation	The process of reorganizing data on disk by making file data blocks physically adjacent to reduce access times.
direct extent	An extent that is referenced directly by an inode.
direct I/O	An unbuffered form of I/O that bypasses the kernel's buffering of data. With direct I/O, the file system transfers data directly between the disk and the user-supplied buffer.
discovered direct I/O	Discovered Direct I/O behavior is similar to direct I/O and has the same alignment constraints, except writes that allocate storage or extend the file size do not require writing the inode changes before returning to the application.
encapsulation	A process that converts existing partitions on a specified disk to volumes. If any partitions contain file systems, <code>/etc/fstab</code> entries are modified so that the file systems are mounted on volumes instead. Encapsulation is not applicable on some systems.
extent	A group of contiguous file system data blocks treated as a single unit. An extent is defined by the address of the starting block and a length.
extent attribute	A policy that determines how a file allocates extents.
external quotas file	A quotas file (named <code>quotas</code>) must exist in the root directory of a file system for quota-related commands to work.

file system block	The fundamental minimum size of allocation in a file system. This is equivalent to the fragment size on some UNIX file systems.
fileset	A collection of files within a file system.
fixed extent size	An extent attribute used to override the default allocation policy of the file system and set all allocations for a file to a specific fixed size.
fragmentation	The on-going process on an active file system in which the file system is spread further and further along the disk, leaving unused gaps or fragments between areas that are in use. This leads to degraded performance because the file system has fewer options when assigning a file to an extent.
GB (gigabyte)	2^{30} bytes or 1024 megabytes.
hard limit	The hard limit is an absolute limit on system resources for individual users for file and data block usage on a file system.
heartbeat	Heartbeat messages are sent over the private link to obtain information on cluster membership changes. If a node does not send a heartbeat for 16 seconds, it is removed from the membership. The command <code>lltconfig</code> is used for information on the various heartbeat parameters. The low latency transport (LLT) module provides communication services across the cluster.
indirect address extent	An extent that contains references to other extents, as opposed to file data itself. A single indirect address extent references indirect data extents. A double indirect address extent references single indirect address extents.
indirect data extent	An extent that contains file data and is referenced via an indirect address extent.
inode	A unique identifier for each file within a file system that contains the data and metadata associated with that file.
inode allocation unit	A group of consecutive blocks containing inode allocation information for a given fileset. This information is in the form of a resource summary and a free inode map.
intent logging	A method of recording pending changes to the file system structure. These changes are recorded in a circular intent log file.
internal quotas file	VxFS maintains an internal quotas file for its internal usage. The internal quotas file maintains counts of blocks and indices used by each user.
KB (kilobyte)	2^{10} bytes or 1024 bytes.
large file	A file larger than two terabytes. VxFS supports files up to 8 exabytes in size.
large file system	A file system larger than two terabytes. VxFS supports file systems up to 8 exabytes in size.
latency	For file systems, this typically refers to the amount of time it takes a given file system operation to return to the user.

local mounted file system	A file system mounted on a single host. The single host mediates all file system writes to storage from other clients. To be a local mount, a file system cannot be mounted using the <code>mount -o cluster</code> option.
metadata	Structural data describing the attributes of files on a disk.
MB (megabyte)	2 ²⁰ bytes or 1024 kilobytes.
mirror	A duplicate copy of a volume and the data therein (in the form of an ordered collection of subdisks). Each mirror is one copy of the volume with which the mirror is associated.
multi-volume file system	A single file system that has been created over multiple volumes, with each volume having its own properties.
MVS (multivolume support)	
node	One of the hosts in a cluster.
node abort	A situation where a node leaves a cluster (on an emergency basis) without attempting to stop ongoing operations.
node join	The process through which a node joins a cluster and gains access to shared disks.
OLT (object location table)	The information needed to locate important file system structural elements. The OLT is written to a fixed location on the underlying media (or disk).
page file	A fixed-size block of virtual address space that can be mapped onto any of the physical addresses available on a system.
preallocation	A method of allowing an application to guarantee that a specified amount of space is available for a file, even if the file system is otherwise out of space.
primary fileset	The files that are visible and accessible to the user.
quotas	Quota limits on system resources for individual users for file and data block usage on a file system.
quotas file	The <code>quotas</code> commands read and write the external <code>quotas</code> file to get or change usage limits. When quotas are turned on, the quota limits are copied from the external <code>quotas</code> file to the internal <code>quotas</code> file.
reservation	An extent attribute used to preallocate space for a file.
SFCFS (Storage Foundation Cluster File System)	
SFCFS Primary	There is a primary node for each file system in the cluster responsible for updating metadata in the file system.

shared disk group	A disk group in which the disks are shared by multiple hosts (also referred to as a cluster-shareable disk group).
shared volume	A volume that belongs to a shared disk group and is open on more than one node at the same time.
snapshot file system	An exact copy of a mounted file system at a specific point in time. Used to do online backups.
snapped file system	A file system whose exact image has been used to create a snapshot file system.
soft limit	The soft limit is lower than a hard limit. The soft limit can be exceeded for a limited time. There are separate time limits for files and blocks.
Storage Checkpoint	A facility that provides a consistent and stable view of a file system or database image and keeps track of modified data blocks since the last Storage Checkpoint.
structural fileset	The files that define the structure of the file system. These files are not visible or accessible to the user.
super-block	A block containing critical information about the file system such as the file system type, layout, and size. The VxFS super-block is always located 8192 bytes from the beginning of the file system and is 8192 bytes long.
synchronous writes	A form of synchronous I/O that writes the file data to disk, updates the inode times, and writes the updated inode to disk. When the write returns to the caller, both the data and the inode have been written to disk.
TB (terabyte)	2 ⁴⁰ bytes or 1024 gigabytes.
transaction	Updates to the file system structure that are grouped together to ensure they are all completed.
throughput	For file systems, this typically refers to the number of I/O operations in a given unit of time.
unbuffered I/O	I/O that bypasses the kernel cache to increase I/O performance. This is similar to direct I/O, except when a file is extended; for direct I/O, the inode is written to disk synchronously, for unbuffered I/O, the inode update is delayed.
VCS (Veritas Cluster Server)	
volume	A virtual disk which represents an addressable range of disk blocks used by applications such as file systems or databases.
volume set	A container for multiple different volumes. Each volume can have its own geometry.
vxfs	The Veritas file system type. Used as a parameter in some commands.
VxFS	The Veritas File System.

VxVM

The Veritas Volume Manager.

Index

Symbols

- /etc/default/vxdg file 51
- /etc/fstab file 61
- /etc/gabtab
 - VCS 22
- /etc/llthosts
 - VCS 22
- /etc/llttab
 - VCS 22

A

- Actions 146
- Administering
 - Clustered NFS 148
- Administration
 - I/O fencing 72
 - SFCFS 55
- Agents
 - CVM 53
- Applications
 - SFCFS 17
- Architecture
 - SFCFS 13
 - VCS 22
- Asymmetric mounts 24
 - mount_vxfs(1M) 24

B

- Backup strategies
 - SFCFS 27
- Basic design
 - Clustered NFS 144
- Benefits
 - SFCFS 17

C

- CFS file system
 - Sharing 170
- CFS primaryship
 - determining 25

- CFS primaryship (*continued*)
 - moving 25
- cfscluster command 58
- cfsdadm command 58
- cfsmntadm command 58
- cfsmount command 58
- cfsnfssg_dummy
 - service group 147
- cfsshare
 - manual page 146
- cfsshare (1M) 162
- cfsshare command 59
- cfsuamount command 58
- Changing the CVM master 65
- Checkpoints 61
- CIFS
 - Common Internet File System 161
 - Debugging 175
- Cluster file systems
 - VxFS
 - unsupported features 16
- cluster file systems
 - support features
 - VxFS 14
- Cluster master node
 - changing 65
- Cluster snapshot
 - characteristics 62
- Clustered NFS 143
 - Administering 148
 - Basic design 144
 - Configure 146
 - Debugging 160
 - Functionality 144
 - Unconfigure 148
 - Understanding 144
- clusters
 - private networks 48
- CNFS 143
- commands
 - cfscluster 58
 - cfsdadm 58

commands (*continued*)

- cfsmntadm 58
- cfsmount 58
- cfsshare 59
- cfsu mount 58
- Common Internet File System
 - CIFS 161
- Configure
 - Clustered NFS 146
- Configuring
 - low priority link 30
- Configuring a CNFS
 - samples 152, 154
- Connectivity policy
 - shared disk groups 52
- coordinator disks
 - DMP devices 35
 - for I/O fencing 35
- copy-on-write technique 26
- CP server
 - deployment scenarios 106
 - migration scenarios 106
- CP server database 46
- CP server user privileges 46
- Creating
 - database for Oracle 10g 199
 - database for Oracle 11g 199
 - snapshot
 - SFCFS 63
- CVM 47
 - agents 53
 - functionality 54
- CVM master
 - changing 65
- CVM master node 48

D

- data corruption
 - preventing 33
- data disks
 - for I/O fencing 35
- Debugging CIFS 175
- Debugging Clustered NFS 160
- Determining
 - CFS primaryship 25
- disk groups
 - private 49
 - shared 49
- disk groups types 49

Disk layout version 16

E

Environment

- public network 30
- single private link 30

F

Failover

- primary 14
- secondary 14

Fenced configurations

- troubleshooting 177

file system

- mount 160

File system tunables

- tunefs(1M) 26
- tunefstab(4) 26
- vxtunefs(1M) 26

fire drills

- about 122
- disaster recovery 122
- for global clusters 122

Freeze 15

fsadm_vxfs(1M)

- manual page 59

fsclustadm 144

fsclustadm(1M)

- manual page 59

Functionality

- Clustered NFS 144

G

GAB

- VCS 22

GLM 14

- SFCFS 23

Growing

- file system
 - SFCFS 60

I

I/O error handling 29

I/O failures

- Recovering 30

I/O fencing

- administration 72
- operations 34

I/O fencing (*continued*)

- preventing data corruption 33
- testing and scenarios 40

J

- Jeopardy 31
 - handling 32
 - state 31

L

- Limitations
 - shared disk groups 53
- LLT
 - VCS 22
- Locking 16
- log files 187
- Low priority link
 - configuring 30

M

- main.cf file 156
- Manual page
 - fsadm_vxfs(1M) 59
 - fsclustadm(1M) 59
 - mount(1M) 59
- manual page
 - cfsshare 146
- Master node
 - changing 65
- master node 48
- Memory mapping 16
- mount
 - nfs-exported file system 160
- mount(1M)
 - manual page 59
- mount_vxfs(1M)
 - asymmetric mounts 24
- Moving
 - CFS primaryship 25

N

- name space
 - preserved by Storage Checkpoints 27
- Nested mounts 15
- network partition 48
- NFS mounts 15
- nfs-exported file system
 - mount 160

NTP

- network time protocol daemon 26, 60

O

- OMF 131
 - working with Oracle Disk Manager 131
- Oracle Disk Manager 127
 - benefits 128
 - converting Quick I/O files 134
 - disabling 138
 - migrating files to 134
 - preparing existing databases for use with 134
 - setting up 134
- Oracle Managed Files 131
 - working with Oracle Disk Manager 131

P

- Parallel fsck threads 26
- Parallel I/O 29
- Performance
 - SFCFS 62
- Port membership
 - SFCFS 23
- postoffline
 - trigger 145
- postonline
 - trigger 145
- preonline
 - trigger 144
- Primary
 - failover 14
- primary and secondary
 - terminology 25
- primary fails
 - SFCFS 61
- primaryship
 - setting with fsclustadm 61
- private networks in clusters 48

Q

- Quick I/O
 - converting files to Oracle Disk Manager 134
- Quick I/O for databases 15
- Quotas 15

R

- Recovering
 - I/O failures 30

- reservations
 - description 33

S

- SCSI-3 PR 33
- Secondary
 - failover 14
- server-based fencing
 - replacing coordination points
 - online cluster 101
- service group
 - cfnsfsg_dummy 147
- Setting
 - parallel fsck threads 26
 - primaryship
 - fsclustadm 61
- SFCFS
 - administration 55
 - applications 17
 - architecture 13
 - backup strategies 27
 - benefits 17
 - environments 30
 - features 18
 - GLM 23
 - growing
 - file system 60
 - load distribution 61
 - performance 62
 - port membership 23
 - primary fails 61
 - snapshots 28, 61
 - synchronize time 25
 - usage 18
- Shared CFS file system
 - unsharing 171
- Shared disk groups 48
 - allowed
 - conflicting 51
 - connectivity policy 52
 - limitations 53
 - off
 - default 50
- shared disk groups
 - activation modes 50
- Sharing
 - CFS file system 170
- slave nodes 48

- Snapshot
 - characteristics
 - cluster 62
 - creating
 - SFCFS 63
- Snapshots 15
 - SFCFS 28, 61
- Split-brain 31
- Storage Checkpoints 15, 61
 - definition of 27
- Switching the CVM master 65
- Symmetric architecture 13
- Synchronize time
 - SFCFS 25

T

- terminology
 - primary and secondary 25
- Thaw 15
- Time synchronization
 - cluster file systems 60
- trigger
 - postoffline 145
 - postonline 145
 - preonline 144
- Troubleshooting
 - fenced configurations 177
- troubleshooting
 - CVMVolDg 198
 - overview of topics 196
 - SCSI reservation errors during bootup 180
 - shared disk group cannot be imported 196

U

- Unconfigure
 - Clustered NFS 148
- Understanding
 - Clustered NFS 144
- Unsharing
 - shared CFS file system 171
- upgrade
 - from raw devices 135
- Usage
 - SFCFS 18

V

- VCS
 - /etc/gabtab 22

VCS (*continued*)

- /etc/llthosts 22

- /etc/llttab 22

- architecture 22

- GAB 22

- LLT 22

VxFS

- supported features

- cluster file systems 14

- unsupported features

- cluster file systems 16