

VERITAS File System™ 3.4

Release Notes

Solaris

December 2000
32-000010-399


VERITAS

Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

Copyright

Copyright © 2000 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the U.S. and other countries. The VERITAS logo, VERITAS File System, and VxFS are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Printed in the USA, November 2000.

VERITAS Software Corporation
1600 Plymouth St.
Mountain View, CA 94043
Phone 650-527-8000
Fax 650-527-8050
<http://www.veritas.com>



VERITAS File System Release Notes

This guide provides information on VERITAS File System™ (VxFS™) release 3.4, including information on the new Cluster File System (CFS) feature (see “[Cluster File System](#)” on page 3). The 3.4 release of VxFS operates on Solaris 2.6, Solaris 7 (32-bit and 64-bit), and Solaris 8 (32-bit and 64-bit) operating systems. VxFS Cluster File System is supported only on Solaris 8.

Note For VxFS to work reliably on Solaris 8, you must have Sun patch ID 108528-02 installed, or upgrade to Solaris 8 Update 2.

Review this entire document before installing VxFS.

The VERITAS File System package includes VxFS software, documentation, and the optional VERITAS Quick I/O™ for Databases and VERITAS QuickLog™ features. Topics in this guide include:

- ◆ [Getting Help](#)
- ◆ [Licensing and Support From Sun Microsystems](#)
- ◆ [New Features](#)
- ◆ [End of Product Support](#)
- ◆ [Using VERITAS Quick I/O and VERITAS QuickLog](#)
- ◆ [Compatibility With Previous Versions of VxFS](#)
- ◆ [Installing and Upgrading VxFS](#)
- ◆ [Documentation](#)
- ◆ [Cluster File System Documentation](#)
- ◆ [Online Manual Pages](#)
- ◆ [Command Directory Locations](#)
- ◆ [Software Problems and Limitations in CFS and VxFS](#)
- ◆ [Using VxFS in VCS and Other HA Environments](#)



Getting Help

For assistance with any of the VERITAS products, contact VERITAS Technical Support:

- ◆ U.S. and Canadian Customers: 1-800-342-0652
- ◆ International: +1-650-527-8555
- ◆ Email: support@veritas.com

For license information:

- ◆ Phone: 1-925-931-2464
- ◆ Email: license@veritas.com
- ◆ Fax: 1-925-931-2487

For software updates:

- ◆ Email: swupdate@veritas.com

For information on purchasing VERITAS products:

- ◆ Phone: 1-800-258-UNIX (1-800-258-8649) or 1-650-527-8000
- ◆ Email: vx-sales@veritas.com

For additional information about VERITAS and VERITAS products, visit the Web site at:

<http://www.veritas.com>

For software updates and additional technical support information, such as TechNotes, product alerts, and hardware compatibility lists, visit the VERITAS Technical Support Web site at:

<http://support.veritas.com>

Licensing and Support From Sun Microsystems

When you buy the VERITAS File System through Sun Microsystems, you must also purchase a license kit from Sun for each feature. For support and licensing information, refer directly to the license kits, *not* the contact information provided above and in the VERITAS File System documentation.

New Features

VxFS Release 3.4 has the following new features and changes.

▼ Cluster File System

Cluster File System (CFS) is the file system clustering functionality of VxFS. CFS is a separately licensable feature of VxFS, however, installing VxFS and enabling the cluster feature does not create a cluster file system configuration. File system clustering requires several other VERITAS products to enable communication services and provide disk storage resources. The VERITAS Cluster Server™ release 1.3.0 and VERITAS Volume Manager™ release 3.1 are packaged with VxFS as the Cluster File System to provide a complete clustering environment. Cluster File System operates only on Solaris 8 Update 2.

File system clustering uses a *master/slave* model for shared file systems, that is, one node is configured as a primary server for the file system, and the other members are configured as secondaries. The primary node logs all transactions for the shared file system. All servers, both primary and secondary, access the shared disks directly for file data operations. If the primary server fails, one of the secondary servers takes over as the primary for the remaining members of the cluster. There can be up to four nodes per cluster.

There is a new guide, the *VERITAS Cluster File System Installation and Configuration Guide*, that provides details on installing and using VxFS cluster functionality (see “[Cluster File System Documentation](#)” on page 10).

▼ Storage Checkpoints

A Storage Checkpoint is a *frozen image* of a mounted file system. The frozen image, or *checkpoint*, initially consists only of pointers to the file system’s data, so Storage Checkpoints require minimal space. As data blocks subsequently change in the file system, the Storage Checkpoint keeps track of the changes. A Storage Checkpoint therefore provides a consistent representation of a file system at a specific point-in-time by identifying modified data blocks and incorporating the original data into its own directory structure.

Storage Checkpoints serve as the enabling technology for two other VERITAS products: *Block-Level Incremental Backups* and *Storage Rollback*, which are used extensively for backing up databases. (For information on how to obtain these products, contact your VERITAS sales representative). Until the VxFS 3.4 release, Storage Checkpoint technology could be used only through these other products. With this release, VxFS introduces a new administrative model that allows Storage Checkpoints to be treated like regular file systems, so that all VxFS users can take direct advantage of this technology. New features and other general improvements include the following:

- ◆ Storage Checkpoints are now writable, and can be created, mounted, and removed with the new `fsckptadm` utility and an added option to the `mount` command.



- ◆ Performance enhancements in maintaining *data Storage Checkpoints* (Storage Checkpoints that are complete images of the file system) makes using the *Storage Rollback* feature easier and more efficient, therefore more viable for backing up large databases.
- ◆ Multi-file system Storage Checkpoint creation allows database backups without having to shut down the database.

Storage Checkpoint functionality is licensable only with the VERITAS Database Edition.

▼ Support for Oracle Disk Manager (ODM)

VxFS 3.4 supports the Oracle Disk Manager driver. ODM is a custom storage interface for files and raw devices that achieves a higher degree of file administration and increased database performance. ODM features include atomic file naming, creation and deletion, asynchronous file I/O, reduction of system overhead by multiplexing requests and completions in one I/O system call, and the ability to determine file I/O attributes.

The ODM Application Programming Interface (API) is targeted for the Oracle9i release, and will be available only with the VERITAS Database Edition.

▼ Forced Unmounts

The `umount_vxfs` command was added in this release to perform forced unmounts (`umount_vxfs -o force`) of VxFS file systems. This is useful in a variety of situations such as High Availability environments where a mounted file system could prevent timely failover. Any active process with I/O operations pending on an unmounted file system receive an I/O error (EIO). This command can cause data loss and must be used carefully.

▼ File Access Time Disabling

The `noatime` option was added to the `mount` command to disable access time updates. This improves performance in read-only environments by eliminating unneeded write operations.

▼ Parallel Log Replay

The `-o p` option was added to the `fsck` command to allow a log replay on multiple file systems in parallel. When incorporated into run command scripts, the new parallel `fsck` functionality improves system start-up times. On Solaris 2.6 and Solaris 7, VxFS uses the `rc` script `/etc/rc2.d/S01CHECKVXFSYS` to manage the log replays. On Solaris 8 Update 2, VxFS uses the script `/usr/lib/fs/vxfs/fsckall` to manage the log replays. The parallel `fsck` functionality does not work on earlier versions of Solaris 8.

▼ VERITAS QuickLog

The interface to the VERITAS QuickLog (formerly VERITAS Accelerator for NFS) feature was completely redesigned, eliminating or combining several commands, improving performance, and making QuickLog easier to use. The command name changes are described in the QuickLog chapter of the *VERITAS File System Administrator's Guide*.

Other Changes

▼ VxFS Packaging

The optionally licensable features QuickLog and Quick I/O for Databases are no longer provided in separate packages and have been incorporated into the VERITAS File System package `VRTSvxfs`. They still require their own license keys to function.

To conform to Sun Microsystems Architectural Review Committee (ARC) standards, VERITAS-specific commands are now installed in the `/opt/VRTSvxfs/sbin` directory. Other commands remain in the `/usr/lib/fs/vxfs` directory and `/etc/fs/vxfs`, so all three must be specified in the `PATH` environment variable to be accessible (see the table under “[Command Directory Locations](#)” on page 13 for a list of VxFS commands and their directory locations). The online manual pages are now installed in the `/opt/VRTS/man` directory. This directory must be added to the `MANPATH` environment variable.

▼ VERITAS License Facility Now a Separate Package

The VERITAS license facility is no longer part of the `VRTSvxfs` package. The `vxlicense` command and `vxlicense(1M)` manual page are now shipped in a separate package named `VRTSlic`. You must install this package on your system to properly license VxFS features.

▼ API for Manipulating Disk Quotas

VxFS now implements the quota API documented in the Solaris `quotactl(7I)` manual page. Users who have written their own quota tools based on the `Q_QUOTACTL` ioctl can now use those tools on VxFS file systems. However, you cannot administer VxFS file system quotas using the `Q_QUOTACTL` ioctl from a client which mounts VxFS over NFS. This capability will not be available until a modification to the RPC quota daemon (enabling quotas on file systems other than UFS) is implemented on the Solaris operating system.

▼ VERITAS File System Quick Start Guide No Longer a Separate Manual

For convenience and ease of reference, the *VERITAS File System Quick Start Guide* was incorporated into the *VERITAS File System Administrator's Guide* as an appendix.



End of Product Support

In VxFS 3.4, the directories where command executables reside were changed. There are now symbolic links from the old directories to the new directories, but the links will be removed in the VxFS 3.5 release. See “[Command Directory Locations](#)” on page 13 for the new directories to add to your PATH environment variable.

VxFS 3.4 does not operate on Solaris 2.5.1.

With release 3.4, the following VxFS functionality is no longer supported:

- ◆ The `nolog` option of the VxFS mount command.
- ◆ The Quick I/O `vxmkcdev` and `vxfsddstat` commands. Use the `qiomkfile` and `qiostat` commands instead.
- ◆ All the interface commands to the VERITAS Accelerator *for NFS* (now VERITAS QuickLog) beginning with the letters “vxld.” See the QuickLog chapter of the *VERITAS File System Administrator’s Guide* for information on the revamped administrative utilities.

The next VxFS feature release will be the last to support the VxFS Version 1 and Version 2 disk layouts. VERITAS recommends that you begin upgrading file systems using these older disk layouts to Version 4. The following are issues to consider when planning disk layout upgrades:

- ◆ Version 1 disk layout file systems can support more than 8 million inodes, while Version 2 disk layout file systems have an 8 million inode limit.
- ◆ The Version 1 disk layout provides finer control of disk geometry than subsequent disk layouts. This finer control is not relevant on disks employing newer technologies, but can still be applicable on older hardware. If you are using Version 1 disk layout file systems on older hardware that needs fine control of disk geometry, a disk layout upgrade may be problematic.
- ◆ Images of Version 1 or Version 2 disk layout file systems created by copy utilities, such as `dd` or `volcopy`, will become unusable after a disk layout upgrade. Offline conversions tools will be provided in the next VxFS feature release to aid in migrating volume-image backup copies of Version 1 and Version 2 disk layout file systems to a Version 4 disk layout.



Using VERITAS Quick I/O and VERITAS QuickLog

The VERITAS File System package, `VRTSvxfs`, now includes the two optionally licensable add-on products, VERITAS Quick I/O for Databases and the VERITAS QuickLog. These features are described in the *VERITAS File System Administrator's Guide*. Quick I/O and QuickLog are available only with VERITAS Editions products.

See “[Getting Help](#)” on page 2 for contact information on these products.

Compatibility With Previous Versions of VxFS

Note VERITAS recommends upgrading any previously installed VxFS file system to VxFS 3.4.

VERITAS 3.x file systems employ disk layout Version 4. To ensure the best performance, upgrade any Version 1 and Version 2 disk layouts to Version 4. You can do the upgrade online using the `vxupgrade` command (see the `vxupgrade(1M)` manual page for details).

VERITAS 3.4 file systems support all previous VxFS disk layouts, but the contents of intent logs created on previous layout versions cannot be used by VxFS 3.4. So the *first* time you mount an older file system on VxFS 3.4 *and* a file system check is required, you must run `fsck -o full` to repair it (see the `fsck_vxfs(1M)` manual page for details).

Cluster file systems and Storage Checkpoints require the Version 4 disk layout.



Installing and Upgrading VxFS

See the *VERITAS File System Installation Guide* for complete instructions on how to install VxFS, how to upgrade, and how to install and configure your system to support a cluster environment.

The VERITAS CD-ROM purchased from VERITAS contains the following file system packages. See the *VERITAS File System Installation Guide* for the list of VERITAS packages required to support a cluster environment.

- ◆ `VRTSvxfs`—VxFS software and online manual pages
- ◆ `VRTSfsdoc`—VxFS Documentation
- ◆ `VRTSlic`—VERITAS products licensing facility

Note VxFS is a licensed product; you must obtain a license key before installing it. License keys valid for VxFS 2.3.x and other 3.x File Systems are also valid for VERITAS 3.4 File Systems. For information on obtaining a license key, see the *VERITAS File System Installation Guide*.

Documentation

The following documents accompany this VxFS release as PostScript and PDF files:

- ◆ *VERITAS File System Installation Guide*
- ◆ *VERITAS File System Administrator's Guide*

The `VRTSvxfs` package contains manual pages for VxFS commands and utilities.

Displaying Documentation Online

The VERITAS File System guides are provided on the CD-ROM under the `pkgs/VRTSfdoc/root/opt/VRTSfdoc` directory. See the *VERITAS File System Installation Guide* for VRTSfdoc package installation information.

PostScript Format

You can use the Solaris Image Tool (`/usr/openwin/bin/imagetool`) or another PostScript viewer to display the following VxFS guides in their PostScript format:

- ◆ *VERITAS File System Installation Guide*

After installing the VRTSfdoc package, you can access this guide in the directory `/opt/VRTSfdoc/install/fsinstall.ps`.

- ◆ *VERITAS File System Administrator's Guide*

After installing the VRTSfdoc package, you can access this guide in the directory `/opt/VRTSfdoc/sys_admin/fssag.ps`.

PDF Format

Adobe Portable Document Format (PDF) versions of the online manuals mentioned above are installed in the same directory locations under `/opt/VRTSfdoc`. To view or print PDF documents, you need the Adobe Acrobat Reader. You can use Acrobat Reader as a stand-alone application, or as a plug-in to your web browser.



Printing PostScript Documentation

To print the PostScript versions, you need access to a PostScript printer or print facilities that print PostScript documents. You can print the PostScript documentation in two ways:

- ◆ Use the print option in your PostScript viewer to print one or more pages.
- ◆ Print the entire document using the `lp` command.

For example, you can print the System Administrator's guide by going to the directory `/opt/VRTSfdoc/sys_admin` and entering:

```
$ lp -d printer_name fssag.ps
```

Documentation Notes

The *VERITAS File System Installation Guide*, *VERITAS File System Administrator's Guide*, and the online manual pages were updated for the 3.4 release.

Cluster File System Documentation

There is a separate documentation package named `VRTScfsdc` for the Cluster File System. The `VRTScfsdc` documentation package contains the same *VERITAS File System Administrator's Guide* as `VRTSfdoc`, but has a different installation guide, the *VERITAS Cluster File System Installation and Configuration Guide*.

The online manual pages available with the VERITAS File System (listed under “[Online Manual Pages](#)” on page 11) are also available with the Cluster File System. There is a section in each VxFS manual page detailing any cluster functionality issues associated with that command.

Online Manual Pages

This release includes the following online manual pages as part of the `VRTSvxfs` package. The `pkgadd` command installs these in the appropriate directories under `/opt/VRTS/man` (add this to your `MANPATH` environment variable), but does not update the `windex` database. To ensure that new VxFS manual pages display correctly, update the `windex` database after installing `VRTSvxfs`. See the `catman(1M)` manual page for more information

Section 1	Description
<code>cp_vxfs</code>	VxFS-specific copy command.
<code>cpio_vxfs</code>	VxFS-specific cpio command.
<code>getext</code>	Gets extent attributes for a VxFS file system.
<code>ls_vxfs</code>	VxFS-specific list command.
<code>mv_vxfs</code>	VxFS-specific move command.
<code>qioadmin¹</code>	VxFS Quick I/O for Databases cache administration utility.
<code>qiomkfile¹</code>	Creates a VxFS Quick I/O device file.
<code>qiostat¹</code>	VxFS Quick I/O for Databases statistics utility.
<code>setext</code>	Sets extent attributes on a file in a VxFS file system.
Section 1M	Description
<code>df_vxfs</code>	Reports the number of free disk blocks and inodes for a VxFS file system.
<code>ff_vxfs</code>	Lists file names and inode information for a VxFS file system.
<code>fsadm_vxfs</code>	Resizes or reorganizes a VxFS file system.
<code>fscat_vxfs</code>	Cats a VxFS file system.
<code>fsck_vxfs</code>	Checks and repairs a VxFS file system.
<code>fsckptadm⁵</code>	VxFS Storage Checkpoint administration utility.
<code>fsclustadm^{3,5}</code>	Manages cluster-mounted VxFS file systems.
<code>fsdb_vxfs</code>	VxFS file system debugger.
<code>fstyp_vxfs</code>	Returns the type of file system on a specified disk partition.
<code>glmconfig³</code>	Group Lock Manager (GLM) configuration utility.
<code>mkfs_vxfs</code>	Constructs a VxFS file system.
<code>mount_vxfs</code>	Mounts a VxFS file system.
<code>ncheck_vxfs</code>	Generates path names from inode numbers for a VxFS file system.
<code>qlogadm²</code>	Low level ioctl utility for the QuickLog driver.



qlogattach ²	Attaches a previously formatted QuickLog volume to a QuickLog device.
qlogck ²	Recovers QuickLog devices during the boot process.
qlogdetach ²	Detaches a QuickLog volume from a QuickLog device.
qlogdisable ²	Remounts a VxFS file system with QuickLog logging disabled.
qlogenable ²	Remounts a VxFS file system with QuickLog logging enabled.
qlogmk ²	Creates and attaches a QuickLog volume to a QuickLog device.
qlogprint ²	Displays records from the QuickLog configuration.
qlogrec ²	Recovers the QuickLog configuration file during a system failover.
qlogrm ²	Removes a QuickLog volume from the configuration file.
qlogstat ²	Prints statistics for running QuickLog devices, QuickLog volumes, and VxFS file systems.
qlogtrace ²	Prints QuickLog tracing.
umount_vxfs ⁵	Unmounts a VxFS file system.
vxdump	Incremental file system dump.
vxedquota	Edits user quotas for a VxFS file system.
vxlicense ⁴	VERITAS licensing key utility.
vxquot	Displays file system ownership summaries for a VxFS file system.
vxquota	Displays user disk quotas and usage on a VxFS file system.
vxquotaoff vxquotaon	Turns quotas on and off for a VxFS file system.
vxrepquota	Summarizes quotas for a VxFS file system.
vxrestore	Restores a file system incrementally.
vxtunefs	Tunes a VxFS file system.
vxupgrade	Upgrades the disk layout of a VxFS file system.

Section 4	Description
fs_vxfs	Format of a VxFS file system volume.
inode_vxfs	Format of a VxFS file system inode.
qlog_config ²	QuickLog configuration file.
tunefstab	VxFS file system tuning parameters table.
Section 7	Description
qlog ²	VERITAS QuickLog device driver.
vxfsio	VxFS file system control functions.
¹ Functionality available only with VERITAS Quick I/O for Databases feature ² Functionality available only with VERITAS QuickLog feature ³ Functionality available only with VERITAS Cluster file system feature ⁴ Functionality available only with installation of VRTSlic package ⁵ New in VxFS 3.4	

Command Directory Locations

With the 3.4 release, VxFS commands are installed in the directories shown in the table. Put these directories in your PATH to access the commands:

- ◆ /opt/VRTSvxfs/sbin
- ◆ /usr/lib/fs/vxfs
- ◆ /etc/fs/vxfs

Command	Executable Directory	Description
cp	/opt/VRTSvxfs/sbin	VxFS-specific copy command.
cpio	/opt/VRTSvxfs/sbin	VxFS-specific cpio command.
df	/usr/lib/fs/vxfs	Reports the number of free disk blocks and inodes for a VxFS file system.
ff	/usr/lib/fs/vxfs	Lists file names and inode information for a VxFS file system.
fsadm	/opt/VRTSvxfs/sbin	Resizes or reorganizes a VxFS file system.
fscat	/opt/VRTSvxfs/sbin	Cats a VxFS file system.
fsck	/usr/lib/fs/vxfs	Checks and repairs a VxFS file system.
fsckptadm ³	/opt/VRTSvxfs/sbin	VxFS Storage Checkpoint administration utility.
fsclustadm ^{3,5}	/opt/VRTSvxfs/sbin	Manages cluster-mounted VxFS file systems.



Command	Executable Directory	Description
fsdb	/usr/lib/fs/vxfs	VxFS file system debugger.
fstyp	/usr/lib/fs/vxfs	Returns the type of file system on a specified disk partition.
getext	/opt/VRTSvxfs/sbin	Gets extent attributes for a VxFS file system.
glmconfig ³	/sbin	Group Lock Manager (GLM) configuration utility.
ls	/opt/VRTSvxfs/sbin	VxFS-specific list command.
mkfs	/usr/lib/fs/vxfs	Constructs a VxFS file system.
mount	/etc/fs/vxfs	Mounts a VxFS file system.
mv	/opt/VRTSvxfs/sbin	VxFS-specific move command.
ncheck	/usr/lib/fs/vxfs	Generates path names from inode numbers for a VxFS file system.
qioadmin ¹	/opt/VRTSvxfs/sbin	VxFS Quick I/O for Databases cache administration utility.
qiomkfile ¹	/opt/VRTSvxfs/sbin	Creates a VxFS Quick I/O device file.
qiostat ¹	/opt/VRTSvxfs/sbin	VxFS Quick I/O for Databases statistics utility.
qlogadm ²	/opt/VRTSvxfs/sbin	Low level IOCTL utility for the QuickLog driver.
qlogattach ²	/etc/fs/vxfs	Attaches a previously formatted QuickLog volume to a QuickLog device.
qlogck ²	/etc/fs/vxfs	Recovers QuickLog devices during the boot process.
qlogdb ²	/opt/VRTSvxfs/sbin	QuickLog debugging tool.
qlogdetach ²	/opt/VRTSvxfs/sbin	Detaches a QuickLog volume from a QuickLog device.
qlogdisable ²	/opt/VRTSvxfs/sbin	Remounts a VxFS file system with QuickLog logging disabled.
qlogenable ²	/opt/VRTSvxfs/sbin	Remounts a VxFS file system with QuickLog logging enabled.
qlogmk ²	/opt/VRTSvxfs/sbin	Creates and attaches a QuickLog volume to a QuickLog device.
qlogprint ²	/opt/VRTSvxfs/sbin	Displays records from the QuickLog configuration.
qlogrec ²	/etc/fs/vxfs	Recovers the QuickLog configuration file during a system failover.
qlogrm ²	/opt/VRTSvxfs/sbin	Removes a QuickLog volume from the configuration file.
qlogstat ²	/opt/VRTSvxfs/sbin	Prints statistics for running QuickLog devices, QuickLog volumes, and VxFS file systems.



Command	Executable Directory	Description
qlogtrace ²	/opt/VRTSvxfs/sbin	Prints QuickLog tracing.
setext	/opt/VRTSvxfs/sbin	Sets extent attributes on a file in a VxFS file system.
umount_vxfs ⁵	/usr/lib/fs/vxfs	Unmounts a VxFS file system.
vxdump	/opt/VRTSvxfs/sbin	Incremental file system dump.
vxlicense ⁴	/opt/VRTSlic/bin	VERITAS licensing key utility.
vxquot	/opt/VRTSvxfs/sbin	Displays file system ownership summaries for a VxFS file system.
vxquota	/opt/VRTSvxfs/sbin	Displays user disk quotas and usage on a VxFS file system.
vxquotaoff vxquotaon	/opt/VRTSvxfs/sbin	Turns quotas on and off for a VxFS file system.
vxrepquota	/opt/VRTSvxfs/sbin	Summarizes quotas for a VxFS file system.
vxrestore	/opt/VRTSvxfs/sbin	Restores a file system incrementally.
vxtunefs	/opt/VRTSvxfs/sbin	Tunes a VxFS file system.
vxupgrade	/opt/VRTSvxfs/sbin	Upgrades the disk layout of a VxFS file system.
¹ Functionality available only with VERITAS Quick I/O for Databases feature ² Functionality available only with VERITAS QuickLog feature ³ Functionality available only with VERITAS Cluster file system feature ⁴ Functionality available only with installation of VRTSlic package ⁵ New in VxFS 3.4		



Software Problems and Limitations in CFS and VxFS

▼ Cluster File System: Panics Due to a Network Failure

If, due to a network failure, a node has only one LLT private network connection, the node can panic. Use at least three LLT private networks between the cluster nodes to reduce the probability of encountering this problem.

▼ Cluster File System: Incomplete Removal of VERITAS Software Packages

In some circumstances, the `cfsdeinstall` script fails to completely remove the VERITAS Volume Manager (`VRTSvxvm`) and VERITAS File System (`VRTSvxfs`) packages, and partially installed packages cannot be overwritten by a subsequent running of the `cfsinstall` script. For the Volume Manager, incomplete removal occurs if the root disk is encapsulated. And `VRTSvxfs` cannot be removed non-interactively if the `vxfs` module is not unloaded.

To avoid this problem, after running the `cfsdeinstall` script, check each node in the cluster to determine if the `VRTSvxfs` and `VRTSvxvm` packages are still installed:

```
# pkginfo -x | grep VRTS
```

If so, remove them manually. Follow the instructions described in the “Deinstalling CFS Packages” section of the *VERITAS Cluster File System Installation and Configuration Guide*.

▼ Cluster File System: VMSA Does Not Support CFS

The 3.1 release of the VERITAS Volume Manager Storage Administrator (VMSA) GUI, currently included in the Cluster File System product, does not support CFS administration. You can create volumes and shared volumes, but you cannot mount or unmount shared file systems. CFS functionality will be supported in a future VMSA release.

▼ Cluster File System: VCS Resources for CFS Agents Become Unavailable

VCS resources for CFS and CVM can get stuck in the `ONLINE | UNABLE TO OFFLINE` state. This can occur if you try to unmount (offline) the primary node when there are mounted (online) secondary nodes for the `CFSMount` resource, or if you try to take a `CVMVolume/CVMDiskGroup` resource on a CVM slave node offline.

To prevent this condition, you must stop and start VCS without saving the configuration:

```
# hastop -sys sysname -force  
# hastart
```

▼ Cluster File System: The `cvminstall` Script was Renamed `vxcvmconfig`

The `vxclustadm(1M)` online manual page, available in the `VRTSvmman` package, refers to the `cvminstall` command. This command was renamed `vxcvmconfig`. (The `vxcvmconfig` command adds cluster configuration information to the VERITAS Cluster Server configuration file `/etc/VRTSvcs/conf/config/main.cf`, which is used for cluster file systems).

▼ Panics on Solaris 8 Operating Systems

There are possible stack overflow problems on VxFS file systems that are NFS exported. The problem occurs only on Solaris 8, 32-bit systems. If you are using Solaris 8 Update 2 or earlier, follow these steps to increase the NFS thread stack size:

1. Apply Sun Microsystems patch number 108901-03 or later. See the `patchadd(1M)` manual page for information on installing patches.
2. Add the following line to the file `/etc/system`:

```
set rpcmod:svc_default_stksize=0x4000
```
3. Reboot the system.

Note For Solaris 8 Update 2 or earlier, you *must* apply this patch on *both* 32-bit and 64-bit operating systems before altering the default stack size parameter, otherwise the system will panic on reboot.

If you are using Solaris 8 Update 3, follow these steps to increase the stack size:

1. Add the following line to the file `/etc/system`:

```
set rpcmod:svc_default_stksize=0x4000
```
2. Reboot the system.

▼ Quick I/O Files Cannot Be Sparse Files

If you try to convert a sparse file to a Quick I/O file, the Oracle instance can fail if Oracle tries to write into an unallocated block. Specifically, datafiles used by the Oracle8i and Oracle9i temporary tablespace may be sparse files, so do not convert these to Quick I/O files. See the *VERITAS Database Edition for Oracle Database Administrator's Guide* for more information.



▼ **DMAPI Not Supported on Version 1 Disk Layouts**

Use DMAPI only on VxFS Version 2 or higher disk layouts.

▼ **Data Integrity Issues with Disks and Disk Arrays with Write-back Caches**

Disk drives configured to use a write-back cache, or disk arrays configured with a volatile write-back cache, exhibit data integrity problems. The problems occur after a power failure, SCSI bus reset, or other event in which the disk has cached data, but has not yet written it to non-volatile storage. Contact your disk drive or disk array manufacturer to determine whether your system disk drives use a write-back cache, and if the configuration can be changed to disable write-back caching.

▼ **Increased Kernel Stack Size Required**

VxFS often requires more than the default 8K kernel stack size, so during the `VRTSvxfs` installation, entries are added to `/etc/system` to increase the kernel thread stack size to 16K.

▼ **The `vxupgrade` Command Cannot Upgrade Some Older File Systems Directly to Version 4**

The `vxupgrade` command cannot upgrade a Version 1 file system disk layout directly to Version 4. You must first upgrade to Version 2, then to Version 4.

▼ **100% Full File System Cannot Be Resized**

In some circumstances, the `fsadm` command cannot resize a 100% full file system due to lack of space for updating structural information. Check VxFS file systems on a regular basis; increase their size if they approach 100% capacity.

▼ **Under Some Conditions, `fsadm` Cannot Truncate a Directory**

The `fsadm` command cannot truncate a directory if it has only one extent that is more than two blocks in length, even if all the directory entries are deleted.

▼ **Must Reboot After Running the `pkgadd` Command**

When you upgrade to a new `VRTSvxfs` package, reboot the system. New kernel modules are not loaded by the `pkgadd` command, so a reboot is required.

▼ A Change in the Method of Computing CUT Values May Cause Misleading Error Messages to Display.

In this release, the method for computing the Current Usage Table (CUT) values for a Version 2 file system has changed.

If a Version 2 file system is mounted on a system running VxFS 3.4, and that file system is subsequently used on an earlier version of VxFS, the following messages may display when performing a full `fsck`:

```
vxfs fsck: incorrect CUT entry for filest 1, fix? (ynq)
vxfs fsck: incorrect CUT entry for filest 999, fix? (ynq)
```

This is expected and does not indicate file system corruption. Answer **y** to both questions. There is no need to perform a full `fsck` when moving such a file system to and from different versions of VxFS unless the file system is dirty, in which case a full `fsck` is required.

▼ Inode Limitation on File Systems Without Large File Support

For a file system to have more than 8 million inodes, you must create it using the `largefiles` option of `mkfs` (the `fsadm` utility can also be used to set the `largefiles` flag on the file system). See the `mkfs_vxfs(1M)` and `fsadm_vxfs(1M)` manual pages for details.

▼ Some Fields Not Displayed by the `fstyp` Command

The `fstyp -v` option shows the super-block. For the Version 4 disk layout, some information is no longer in the super-block, so fields such as `nau`, `logstart`, or `logend` display zeros. `nau` can be computed using the following formula:

$$\text{nau} = (\text{size} + \text{aulen} - 1) / \text{aulen}$$

`fstyp -v` displays the `size` and `aulen` fields. You can use `mkfs -F vxfs -m raw_device_file` to display many fields that are not part of super-block. See the `mkfs_vxfs(1M)` and `mkfs(1M)` manual pages for more information.



Using VxFS in VCS and Other HA Environments

The VERITAS File System can be used in VERITAS Cluster Server™ and other High Availability environments. Because the VxFS driver is loadable, it is not guaranteed to occupy the same position in each system's virtual file system switch (vfssw) table. To ensure reliable failover of a VxFS file system between hosts, add the following line in the same position to each host's `/etc/system` file:

```
forceload: fs/vxfs
```

When using VxFS in a High Availability environment, make sure that all systems in the cluster are running the same version of VxFS. Systems running different versions of VxFS cannot failover.

Note The VxFS cluster feature works only with VERITAS Cluster Server.
