

VERITAS File System™ 3.4 Patch 02

Release Notes

Solaris

August 2001
30-000384-011


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 © 2001 VERITAS Software Corporation. All Rights Reserved. VERITAS, VERITAS SOFTWARE, the VERITAS logo, *Business Without Interruption*, VERITAS The Data Availability Company, and VERITAS File System are trademarks or registered trademarks of VERITAS Software Corporation in the U.S. and/or other countries. Other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

Printed in the USA, August 2001.

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 the VERITAS File System™ (VxFS™) release 3.4 Patch 02. The 3.4 Patch 02 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. References in this document to VxFS 3.4 and VxFS Patch 01 regarding new features, end of product support, compatibility, upgrading, and software limitations apply to VxFS 3.4 Patch 02. Review this entire document before installing VxFS.

Note For VxFS to work reliably on Solaris 8, you must have Sun patch ID 108528-02 and 108901-03 installed.

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](#)
- ◆ [Changes in VxFS Release 3.4 Patch 02](#)
- ◆ [Changes in VxFS Release 3.4 Patch 01](#)
- ◆ [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](#)
- ◆ [Online Manual Pages](#)
- ◆ [Command Directory Locations](#)
- ◆ [Software Problems and Limitations in VxFS](#)
- ◆ [Software Fixes and Enhancements in This Release](#)
- ◆ [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 website 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 website 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.

Changes in VxFS Release 3.4 Patch 02

VxFS Release 3.4 Patch 02 has the following new features and changes.

▼ 16-node Cluster File Systems

This release of VxFS supports cluster file systems on up to 16 nodes.

▼ New I/O Error Handling Policy

The `ioerror` option was added to the `mount` command to provide four different ways to handle system I/O errors. These policies were implemented on VxFS in response to evolving storage technology functionality for which a single method of I/O error processing was no longer sufficient. One major behavior change is to disable a file system instead of marking inodes bad on disk—this will prevent inadvertent data loss. See the `mount(1M)` manual page for more information.

▼ New Default Intent Log Mode

To increase performance, the `mount` command `delaylog` option has replaced the `log` option as the default for VxFS file systems. See the `mount(1M)` manual page for more information.

▼ New Default File System Block Size

The default block size is now 1024 bytes for all VxFS file systems, no matter how large the file system you create. Previously, the default block size increased with larger file systems. You can still select a larger block size using the `mkfs bsize` option. See the `mkfs(1M)` manual page for more information.

▼ Quick I/O Not Automatically Enabled on Cluster File Systems

If you have the Quick I/O for Databases feature licensed on your system, it is automatically enabled when mounting file systems in local mode. On cluster file systems, you must now explicitly select it using the `mount qio` option.

▼ Added Storage Checkpoint Functionality

More VxFS commands support operations on Storage Checkpoints. You can now specify a Storage Checkpoint name with the `df` and `ncheck` commands, and specify a Storage Checkpoint as the mount point for `fsadm`. See the `df_vxfs(1M)`, `fsadm_vxfs(1M)`, and `ncheck_vxfs(1M)` manual page for more information.



▼ Improved QuickLog Functionality

The VERITAS QuickLog feature was modified to check that enough space is available on a QuickLog device before allowing a VxFS file system to begin logging to that device. The automatic checking will prevent system hangs due to insufficient space, and alleviates the need for system administrators to monitor log space. The minimum QuickLog device size was also increased from 32 MB to 41 MB.

▼ API for Obtaining VxFS Inode Field Offsets

For open source utilities, such as LSOE, to support VxFS 3.4 file systems, the library `vxfsutil` was introduced. The library provides an API for accessing VxFS header file information because the header files are no longer public. See the `vxfsu_get_ioffsets(3)` manual page for details.

Changes in VxFS Release 3.4 Patch 01

VxFS Release 3.4 Patch 01 has the following new features and changes.

▼ 8-node Cluster File Systems

This release of VxFS supports cluster file systems on up to eight nodes.

▼ Cluster File System Functionality on VMSA

The 3.1.1 version of the VERITAS Volume Manager Storage Administrator™ (VMSA), included in the current release of CFS packages, supports operations on cluster file systems. The new operations include mounting/unmounting, resizing, setting a node to primary, and removing cluster file system configuration information. See the *VERITAS Volume Manager Storage Administrator Administrator's Guide* for instructions on where to access the new functions on the VMSA GUI menus.

▼ Improvements to the CFS Global Lock Manager

The Global Lock Manager (`VRTSg1m`) included in the current release of CFS packages was enhanced to allow for better scalability and performance.

▼ Converting UFS File Systems to VxFS File Systems

The `vxfsconvert` command is supported in the 3.4 Patch 01 release. Using `vxfsconvert`, you can convert your existing UFS file systems on Solaris to VxFS file systems. See the `vxfsconvert(1M)` online manual page for information on usage, options, and space requirements.

▼ Internationalized Commands

Starting with this release, VxFS is delivered with internationalization (I18N) support enabled for messages displayed by VERITAS File System commands. Commands such as `fsck_vxfs` that prompt users for input are not internationalized.

▼ VERITAS FastResync Option

The VERITAS FastResync Option is a separately licensable feature available in this release. FastResync implements off-host processing solutions for offline and online backup of databases and cluster file systems, decision support, report generation, and database error recovery on enterprise clustered systems. For information on how to use this feature, visit the VERITAS website:

<http://seer.support.veritas.com/docs/235066.htm>

For information on purchasing the product, contact your sales representative.

▼ VxFS System Activity Reporter

The `vxfsstat` command displays VxFS file system statistics which can be used to analyze performance and aid in tuning. Similar to the `sar(1)` command, `vxfsstat` gathers statistics on buffer cache, inode cache, and per-CPU usage. See the `vxfsstat(1M)` online manual page for detailed information.

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. 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 16 nodes per cluster.



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 storage resources. The VERITAS Cluster Server™ release 2.0 and VERITAS Volume Manager™ release 3.2 are packaged with VxFS as the Cluster File System to provide a complete clustering environment. The Cluster File System is marketed by VERITAS under the product name SANPoint Foundation Suite-HA (SPFS HA). The SPFS HA product includes the *VERITAS Cluster File System Installation and Configuration Guide*, which provides details on installing and using VxFS cluster functionality. The documentation on all the constituent products is also included in SPFS HA.

▼ 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 features: *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

A VxFS-specific `umount` command was added in this release to perform forced unmounts (`umount -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 receives an I/O error (EIO). This command can cause data loss and must be used carefully.

▼ File Access Time Disabling

The `-o 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 16 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, including VxFS mounts over NFS.

▼ 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

This is the last release to include PostScript versions of the VxFS guides and Release Notes. Subsequent releases will supply only PDF files on the CD and in the documentation packages.

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 16 for the new directories to add to your PATH environment variable.

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.

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 `vxfddstat` 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.



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 with VERITAS Editions products, QuickLog is also available in the VERITAS Foundations Suite.

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
- ◆ `VRTSfdoc`—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/VRTSfsdoc/reloc/VRTSfsdoc` directory. See the *VERITAS File System Installation Guide* for `VRTSfsdoc` 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 `VRTSfsdoc` package, you can access this guide in the directory `/opt/VRTSfsdoc/install/fsinstall.ps`.
- ◆ *VERITAS File System Administrator's Guide*
After installing the `VRTSfsdoc` package, you can access this guide in the directory `/opt/VRTSfsdoc/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/VRTSfsdoc`. 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/VRTSfsdoc/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.

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.
vxfsconvert	Convert a UFS file system to VxFS.
vxfsstat	Display VxFS file system statistics.
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 3	Description
vxfsu_get_ioffsets ⁶	Obtains VxFS inode field offsets.

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 <code>VRTSlic</code> package ⁵ New in VxFS 3.4 ⁶ New in VxFS 3.4 Patch 02	



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
- ◆ /opt/VRTSlic/sbin (for vxlicense command)

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 ⁹	/opt/VRTSvxfs/sbin	VxFS Storage Checkpoint administration utility.
fsclustadm ^{3,5}	/opt/VRTSvxfs/sbin	Manages cluster-mounted VxFS file systems.
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.



Command	Executable Directory	Description
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.
qlogtrace ²	/opt/VRTSvxfs/sbin	Prints QuickLog tracing.
setext	/opt/VRTSvxfs/sbin	Sets extent attributes on a file in a VxFS file system.
umount ⁵	/usr/lib/fs/vxfs	Unmounts a VxFS file system.
vxdump	/opt/VRTSvxfs/sbin	Incremental file system dump.
vxedquota	/opt/VRTSvxfs/sbin	Incremental file system dump.
vxfsconvert	/opt/VRTSvxfs/sbin	Convert a UFS file system to VxFS.
vxfsstat	/opt/VRTSvxfs/sbin	Display VxFS file system statistics.
vxlicense ⁴	/opt/VRTSlic/bin	VERITAS licensing key utility.



Command	Executable Directory	Description
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 VxFS

▼ Forced Unmount Does Not Work on File Systems with Active Quick I/O Files

A system panic can occur if a forced unmount, the `umount -o force` command, is performed on a VxFS file system that has active Quick I/O files. Stop all database applications using the Quick I/O interface before using the forced unmount option.

▼ ACL Behavior is Different on VxFS and UFS

VxFS differs from UFS in the value returned for the group owner permission bits on certain ACLs. This does not affect the behavior of permission checking against the ACL, which is identical on VxFS and UFS, it affects only the values returned by the `stat` system call, which is used by the `ls` command (see `stat(2)` and `ls(1)` manual pages).

VxFS currently returns the ACL mask (the `CLASS_OBJ` object), while UFS returns the intersection (bitwise AND) of the ACL group owner permissions and ACL mask entries (see the `setfacl(1)` and `aclcheck(3)` manual pages).

When setting ACLs that contain no additional ACL entries besides the standard `unix|group|other` permissions and an ACL mask value, both VxFS and UFS discard the mask and group permission values and store only the intersection (bitwise AND) of the two values.

The behavior of VxFS may change in a future release to retain both values. This will not change the behavior of permission checking of VxFS, only what is stored and returned.

▼ 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 on 32-bit Kernels

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. Sun patch ID 108901-03 resolves this problem.

▼ 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.



Software Fixes and Enhancements in This Release

This section lists VxFS and CFS enhancements and problems fixed since VERITAS File System release 3.4 Patch 01.

Incident	Description
18142	The VxFS error handling policy was changed to not mark inodes bad on disk because there are some situations, such as intermittent hardware problems, where correct metadata can be considered corrupted.
38669	The <code>vxrestore(1M)</code> online manual page was updated to better explain error conditions.
52408	On Solaris 8, the <code>utmp</code> and <code>wtmp</code> database files were superseded by the <code>utmpx</code> and <code>wtmpx</code> database files. For the <code>vxdump</code> command to work correctly, it was modified to use the <code>getutxent(3C)</code> function to access <code>utmpx</code> .
53075	Because the file size was not a multiple of a 512-byte sector, a Quick I/O file could not be resized to one byte less than 2 GB.
57603	System panicked during an extent reorganization because of a race condition between the <code>vx_ialloc</code> and <code>vx_rename1</code> functions.
57694	When the <code>setfacl</code> or <code>acl</code> system call was used to set file permissions and a mask (<code>CLASS_OBJ</code>) value was specified, if the mask value was different from the group permissions, VxFS would set the wrong value for the group permission bits.
59984	System hang occurred due to a deadlock condition between an inode blockmap lock and a disk block buffer.
62971	VxFS file system buffer pages were being allocated as kernel pages instead of user pages.
62978	On Solaris 8 systems, the remote host's name was not correct in the output displayed from the <code>vxdump</code> command.
63546	Even though the <code>mount</code> command <code>datainlog</code> option was selected, it was not being reported correctly in the <code>/etc/mnttab</code> file.
64192	On systems with VxFS 3.4 installed, panics occurred when trying to <code>modunload</code> the Global Atomic Broadcast (GAB) software.



Incident	Description
68523	The warning message “no entries found in /etc/mnttab” incorrectly displayed when trying to umount VxFS file systems using a script that ran <code>umount</code> commands in the background.
69366	A system panic occurred during failover due to a NULL pointer de-reference in QuickLog. This happened because the <code>qllogadm</code> command was invoked when there were no VxFS file systems enabled on the QuickLog device. QuickLog now checks that a file system volume exists before trying to remove it.

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.

The file system switch table lists all the file system types that are running on a system. Each file system type creates an entry in this table the first time it is loaded, typically when a `mount` command is issued after a system reboot. The VxFS kernel module loads automatically when the first VxFS file system is mounted, which puts the module in the next available slot in the `vfssw` table.

In HA environments, where VxFS resources are under cluster control, the VxFS module must always occupy the same `vfssw` position to ensure reliable failover. So to guarantee the correct failover of a VxFS file system between hosts, put 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 VERITAS Cluster File System works only with VERITAS Cluster Server.



