

# VERITAS File System™ 3.4 for Solaris 9

---

## Release Notes

Solaris 9

May 2002  
30-000727-011

  
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 © 2002 VERITAS Software Corporation. All rights reserved. VERITAS, VERITAS Software, the VERITAS logo, VERITAS File System, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation in the USA and/or other countries. Other product names mentioned herein may be trademarks or registered trademarks of their respective companies.

May 2002.

VERITAS Software Corporation  
350 Ellis Street  
Mountain View, CA 94043  
USA  
Phone 650-527-8000  
Fax 650-527-2908  
<http://www.veritas.com>



# VERITAS File System Release Notes

---

This guide provides information on the VERITAS File System™ (VxFS™) Release 3.4 for Solaris 9.

The VxFS 3.4 Release for Solaris 9 operates only on Solaris 9 (32-bit and 64-bit) operating systems. This release contains all bug fixes and functionality changes made in the VxFS 3.4 Patch 01, VxFS 3.4 Patch 02, and VxFS 3.4 Patch 03 maintenance releases.

References in this document to VxFS 3.4, VxFS 3.4 Patch 01, VxFS 3.4 Patch 02, and VxFS Patch 03 regarding new features, end of product support, compatibility, upgrading, and software limitations apply to VxFS 3.4 for Solaris 9. Review this entire document before installing VxFS.

---

**Caution** You cannot install this release of VxFS on any Solaris operating system other than Solaris 9. See “[Software Problems and Limitations in VxFS](#)” on page 19 for more information on installing VxFS.

---



---

Topics in this guide include:

- ◆ [Getting Help](#)
- ◆ [Licensing and Support From Sun Microsystems](#)
- ◆ [Changes in VxFS Release 3.4 for Solaris 9](#)
- ◆ [Changes in VxFS Release 3.4 Patch 03](#)
- ◆ [Changes in VxFS Release 3.4 Patch 02](#)
- ◆ [Changes in VxFS Release 3.4 Patch 01](#)
- ◆ [New Features](#)
- ◆ [End of Product Support](#)
- ◆ [Compatibility With Previous Versions of VxFS](#)
- ◆ [Installing and Upgrading VxFS](#)
- ◆ [Documentation](#)
- ◆ [Online Manual Pages](#)
- ◆ [Command Directory Locations](#)
- ◆ [Software Problems and Limitations in 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](mailto:support@veritas.com)

For license information:

- ◆ Phone: 1-925-931-2464
- ◆ Email: [license@veritas.com](mailto:license@veritas.com)
- ◆ Fax: 1-925-931-2487

For software updates:

- ◆ Email: [swupdate@veritas.com](mailto: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](mailto: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 for Solaris 9

VxFS Release 3.4 for Solaris 9 has the following new features and changes.

### ▼ Solaris 9 Support

VERITAS File System now operates on Solaris 9 in 32-bit and 64-bit mode.

### ▼ VxFS Features Not Supported on Solaris 9

The following VxFS features are not supported on this release:

- ◆ VERITAS Quick I/O for Databases
- ◆ VERITAS QuickLog
- ◆ VERITAS Cluster File System

## Changes in VxFS Release 3.4 Patch 03

VxFS Release 3.4 Patch 03 has the following new features and changes.

---

**Note** The changes for the various VxFS patch releases described in the following sections were incorporated into the VxFS 3.4 Release for Solaris 9. You do not have to install any patches for this release.

---

### ▼ Storage Checkpoint Performance Enhancements

In VxFS 3.4 Patch 02 and VxFS Patch 03, Storage Checkpoints can be taken in environments with a large number of files (for example, file servers with millions of files) with little adverse impact on performance. Unlike previous releases, the file system no longer remains frozen during Storage Checkpoint creation, which allows applications to access the file system even while the Storage Checkpoint is taken. Storage Checkpoint creation, however, may take several minutes to complete depending on the number of files in the file system.

In previous releases of VxFS, Storage Checkpoints were only licensed through the VERITAS Editions product and were used only in database environments, typically with only a few large files. Storage Checkpoints are now also available with the VERITAS FlashSnap™ product.

## Changes in VxFS Release 3.4 Patch 02

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

### ▼ 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 was 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.

### ▼ 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.

### ▼ 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_iooffsets(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.

### ▼ 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, 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.

### ▼ 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.

### ▼ 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.



### ▼ 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.

## Other Changes

### ▼ VxFS Packaging

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.
- ◆ 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.

## 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).

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. See “[Documentation Issues](#)” on page 19 for further information on installing VxFS 3.4 for Solaris 9. This release contains the following file system packages:

- ◆ `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.

## Displaying Documentation Online

The VERITAS File System guides are provided 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 <code>cpio</code> 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>qioadmin</code> <sup>1</sup>	VxFS Quick I/O for Databases cache administration utility.
<code>qiomkfile</code> <sup>1</sup>	Creates a VxFS Quick I/O device file.
<code>qiostat</code> <sup>1</sup>	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>fsckptadm</code> <sup>5</sup>	VxFS Storage Checkpoint administration utility.
<code>fsclustadm</code> <sup>3, 5</sup>	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>glmconfig</code> <sup>3</sup>	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>qlogadm</code> <sup>2</sup>	Low level <code>ioctl</code> utility for the QuickLog driver.



qlogattach <sup>2</sup>	Attaches a previously formatted QuickLog volume to a QuickLog device.
qlogck <sup>2</sup>	Recovers QuickLog devices during the boot process.
qlogdetach <sup>2</sup>	Detaches a QuickLog volume from a QuickLog device.
qlogdisable <sup>2</sup>	Remounts a VxFS file system with QuickLog logging disabled.
qlogenable <sup>2</sup>	Remounts a VxFS file system with QuickLog logging enabled.
qlogmk <sup>2</sup>	Creates and attaches a QuickLog volume to a QuickLog device.
qlogprint <sup>2</sup>	Displays records from the QuickLog configuration.
qlogrec <sup>2</sup>	Recovers the QuickLog configuration file during a system failover.
qlogrm <sup>2</sup>	Removes a QuickLog volume from the configuration file.
qlogstat <sup>2</sup>	Prints statistics for running QuickLog devices, QuickLog volumes, and VxFS file systems.
qlogtrace <sup>2</sup>	Prints QuickLog tracing.
umount_vxfs <sup>5</sup>	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.
<b>Section 3</b>	<b>Description</b>
vxfsu_get_ioffsets <sup>6</sup>	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 <sup>2</sup>	QuickLog configuration file.
tunefstab	VxFS file system tuning parameters table.
Section 7	Description
qlog <sup>2</sup>	VERITAS QuickLog device driver.
vxfsio	VxFS file system control functions.
<sup>1</sup> Functionality available only with VERITAS Quick I/O for Databases feature	
<sup>2</sup> Functionality available only with VERITAS QuickLog feature	
<sup>3</sup> Functionality available only with VERITAS Cluster File System feature	
<sup>4</sup> Functionality available only with installation of VRTSlic package	
<sup>5</sup> New in VxFS 3.4	
<sup>6</sup> 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 <sup>5</sup>	/opt/VRTSvxfs/sbin	VxFS Storage Checkpoint administration utility.
fsclustadm <sup>3,5</sup>	/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 <sup>3</sup>	/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.
qiadmin <sup>1</sup>	/opt/VRTSvxfs/sbin	VxFS Quick I/O for Databases cache administration utility.

Command	Executable Directory	Description
qiomkfile <sup>1</sup>	/opt/VRTSvxfs/sbin	Creates a VxFS Quick I/O device file.
qiostat <sup>1</sup>	/opt/VRTSvxfs/sbin	VxFS Quick I/O for Databases statistics utility.
qlogadm <sup>2</sup>	/opt/VRTSvxfs/sbin	Low level ioctl utility for the QuickLog driver.
qlogattach <sup>2</sup>	/etc/fs/vxfs	Attaches a previously formatted QuickLog volume to a QuickLog device.
qlogck <sup>2</sup>	/etc/fs/vxfs	Recovers QuickLog devices during the boot process.
qlogdb <sup>2</sup>	/opt/VRTSvxfs/sbin	QuickLog debugging tool.
qlogdetach <sup>2</sup>	/opt/VRTSvxfs/sbin	Detaches a QuickLog volume from a QuickLog device.
qlogdisable <sup>2</sup>	/opt/VRTSvxfs/sbin	Remounts a VxFS file system with QuickLog logging disabled.
qlogenable <sup>2</sup>	/opt/VRTSvxfs/sbin	Remounts a VxFS file system with QuickLog logging enabled.
qlogmk <sup>2</sup>	/opt/VRTSvxfs/sbin	Creates and attaches a QuickLog volume to a QuickLog device.
qlogprint <sup>2</sup>	/opt/VRTSvxfs/sbin	Displays records from the QuickLog configuration.
qlogrec <sup>2</sup>	/etc/fs/vxfs	Recovers the QuickLog configuration file during a system failover.
qlogrm <sup>2</sup>	/opt/VRTSvxfs/sbin	Removes a QuickLog volume from the configuration file.
qlogstat <sup>2</sup>	/opt/VRTSvxfs/sbin	Prints statistics for running QuickLog devices, QuickLog volumes, and VxFS file systems.
qlogtrace <sup>2</sup>	/opt/VRTSvxfs/sbin	Prints QuickLog tracing.
setext	/opt/VRTSvxfs/sbin	Sets extent attributes on a file in a VxFS file system.
umount <sup>5</sup>	/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 <sup>4</sup>	/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.
<sup>1</sup> Functionality available only with VERITAS Quick I/O for Databases feature <sup>2</sup> Functionality available only with VERITAS QuickLog feature <sup>3</sup> Functionality available only with VERITAS Cluster file system feature <sup>4</sup> Functionality available only with installation of VRTSlic package <sup>5</sup> New in VxFS 3.4		



## Software Problems and Limitations in VxFS

### ▼ Documentation Issues

The *VERITAS File System Installation Guide* included in the `VRTSfsdoc` package for this release does not contain information about Solaris 9. The information in the guide is appropriate for installing VxFS on Solaris 9 except that the upgrade procedure is not valid.

The “Storage Checkpoints” chapter of the *VERITAS File System Administrator’s Guide* shipped with the Patch 03 release states that it is not advisable to use Storage Checkpoints in environments that use a large number of files and directories. This is no longer true. The guide also states that Storage Checkpoints are best suited for file systems containing database files. With the recent performance enhancements, Storage Checkpoints are now useful in other environments. Specifically those with high availability needs, but where only a subset of blocks is likely to change. For example, CAD/CAM, IC design tools, and software development. See “[Storage Checkpoint Performance Enhancements](#)” on page 4 for more information.

### ▼ 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.

### ▼ 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.



## 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.