



VERITAS® File System
Release Notes
Release 3.3.2 - Binary

READ THIS BEFORE INSTALLING THE SOFTWARE!

Solaris
May 1999
P/N 100-001042

© 1999 VERITAS® Software Corporation. All rights reserved.

TRADEMARKS

VERITAS, VxVM, VxVA, VxFS, and FirstWatch are registered trademarks of VERITAS Software Corporation in the United States and other countries.

VERITAS Volume Manager, VERITAS File System, VERITAS NetBackup, VERITAS HSM, CVM, VERITAS Quick I/O, and, and the VERITAS logo are trademarks of VERITAS Software Corporation.

Other products mentioned in this document are trademarks or registered trademarks of their respective holders.

Introduction

This guide provides information on VERITAS File System™ (VxFS®) Release 3.3.2 for Solaris 2.5.1, 2.6, and 7 (32-bit and 64-bit) operating systems. References in this document to VxFS 3.3 regarding new features, end of product support, compatibility, and software limitations apply to VxFS 3.3.2.

Review this entire document before installing the VERITAS File System.

VERITAS File System packages include VxFS software, documentation, and the optional VERITAS Quick I/O™ for Databases and VERITAS QuickLog™.

Topics in this guide include:

- Getting Help
- New Features
- Changes in VxFS Release 3.3.2
- Using VERITAS Quick I/O or VERITAS QuickLog
- End of Product Support
- Compatibility With Previous Versions of VxFS
- Installing VxFS
- Upgrading VxFS
- Documentation



-
- Software Limitations and Problems in VxFS Release 3.3
 - Software Problems Fixed in This Release
 - Software Problems Fixed in VxFS 3.3.1
 - Using VxFS in FirstWatch and Other HA Environments
 - Using QuickLog With FirstWatch

Getting Help

For license information or information about VERITAS service packages, contact VERITAS Customer Support.

U.S. and Canadian Customers: 1-800-342-0652

International Customers: +1 (650) 335-8555

Fax: (650) 335-8428

Electronic mail: support@veritas.com

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

www.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 package. 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 3.3 has the following new features:

▼ VERITAS File System Storage Checkpoint

VERITAS File System Storage Checkpoint is a snapshot technology that creates a *clone* of a currently mounted VxFS file system. A Storage Checkpoint presents a consistent, point-in-time view of the file system by identifying and maintaining modified file system blocks. Storage Checkpoint serves as an enabling technology for two other VERITAS File System features: *Block-Level Incremental Backups* and *Storage Rollback*.

Block-Level Incremental Backup (BLIB) is a backup method that stores and retrieves only the data blocks changed since the previous backup, not entire files. This saves times, storage space, and computing resources required to backup large databases.

Storage Rollback is an on-disk restore capability for faster recovery from software problems, such as an accidentally deleted file. Because each Storage Checkpoint is a point-in-time image of a file system, Storage Rollback simply restores, or rolls back, a file or file system to a designated Storage Checkpoint.

This file system technology is implemented in other VERITAS products. For information on how to use these features, see the *VERITAS Database Edition for Oracle Database Administrator's Guide*, specifically the chapter on Database Backup and Restore, and the *VERITAS NetBackup Block-Level Incremental Backup Extension for Oracle System Administrator's Guide*.

▼ Cached Quick I/O

The Quick I/O™ for Databases feature provides a raw-device interface to regular files, bypassing the normal file system locking and buffering of file data. The Cached Quick I/O feature was added to 3.3 to make Quick I/O more efficient on machines with large amounts of memory by performing caching in the Quick I/O read path. When the database reads data through the Quick I/O interface, the data is cached in the system page cache. Subsequent reads of the same data can access this cached copy and avoid doing disk I/O. This feature is beneficial only on systems with large amounts of RAM, and must be explicitly enabled by setting the `qio_cache_enable` parameter with the `vxtunefs` utility. There is more information available on this feature in the `vxtunefs(1M)` and `qioadmin(1)` online manual pages.



▼ Year 2000 Compliance

VERITAS File System release 3.3 was thoroughly tested to ensure year 2000 compliance. Visit the VERITAS web site at www.veritas.com for the complete Year 2000 Certification Statement.

▼ Internationalized Commands

VxFS Release 3.3 has internationalized commands for easy portability to languages other than English.

▼ Online Manual Pages

There is a new, online manual page for the Cached Quick I/O feature:

- `qioadmin(1)`

Changes in VxFS Release 3.3.2

VERITAS File System release 3.3.2 operates on Solaris 7 in 32-bit and 64-bit mode.

The name of the optionally licensable VERITAS File System product Accelerator *for NFS*[™] was changed to VERITAS QuickLog[™]. All references to the Accelerator *for NFS* or Accelerator were replaced in the 3.3.2 File System documentation set, and the Accelerator package name, `VRTSvxld`, was changed to `VRTSqlog`. However, QuickLog commands and man pages retain their original names; reference them using the same naming convention as with the Accelerator product (for example `vxld_mklog` and `vxld_mklog(1M)`). Problems or incidents in this document regarding older versions of the product, such as Accelerator *for NFS* Version 1.0.4, still refer to the original name.

The VERITAS license key utility `vxserial` was replaced by a new utility called `vxlicense`. Use the `vxlicense` command to enter a license key for a VERITAS product after installing a package on your system.

VERITAS File System release 3.3.2 is certified to operate with VERITAS HSM (Hierarchical Storage Manager) release 3.2 and VERITAS NetBackup release 3.2. This is the recommended configuration.



Using VERITAS Quick I/O or VERITAS QuickLog

The VERITAS File System has two optionally licensable add-on products: VERITAS Quick I/O™ for Databases and the VERITAS QuickLog (formerly the Accelerator for NFS). These products are described in the *VERITAS File System System Administrator's Guide*. You can purchase these products from VERITAS and VERITAS sales channels, or through Sun Microsystems in VERITAS Editions products.

For information on these products, contact VERITAS Software by phone at 1-800-258-UNIX or 1-650-335-8000, or send email to vx-sales@veritas.com.

End of Product Support

This release of the VERITAS File System no longer supports the `volcopy` or `labelit` utilities. Scripts and applications using either of these utilities will not operate; substitute alternative commands described in the *VERITAS File System System Administrator's Guide*.

The `nolog` option of the `mount` command is now the functional equivalent of the `tmplog` option, and no error is returned if `nolog` is specified. This is the last major release of VxFS that supports this option. In the next release, using the `nolog` option will return an error. See the `mount(1M)` manual page or the *VERITAS File System System Administrator's Guide* for more information on `mount` command options).

The VERITAS Quick I/O for Databases package name was changed from `VRTSfdd` to `VRTSqio`, and the Quick I/O feature has two renamed commands:

- `vxmkcdev` was changed to `qiomkfile`
- `vxfddstat` was changed to `qiostat`

VxFS 3.3 is the last release to support the original command names.



Compatibility With Previous Versions of VxFS

Note: VERITAS recommends upgrading any previously installed VxFS File System to VxFS 3.3.2.

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 3.3.2 file systems support all previous VxFS disk layouts, but the contents of intent logs created on previous layout versions cannot be used by VxFS 3.3.2. So the *first* time you mount an older file system on VxFS 3.3.2, *and* a file system check is required, you must run an `fsck -o full` to repair it (see the `fsck_vxfs(1M)` manual page for details).

Installing VxFS

See the *VERITAS File System Installation Guide* for complete instructions on how to install VxFS using the `pkgadd` command.

The VERITAS CD-ROM purchased from VERITAS contains the following file system packages:

- `VRTSvxfs`—VxFS software and online manual pages
- `VRTSfsdoc`—VxFS Documentation
- `VRTSqio`—VERITAS Quick I/O for Databases
- `VRTSqllog`—VERITAS QuickLog

Note: VxFS is a licensed product; you must obtain a license key before installing VxFS. For information on obtaining a license key, see the *VERITAS File System Installation Guide*.



Upgrading VxFS

See the *VERITAS File System Installation Guide* for instructions on how to upgrade VxFS from a previous release.

Note: License keys valid for VxFS 2.3.x are also valid for VERITAS 3.3.2 File Systems.

Documentation

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

- *VERITAS File System Installation Guide*
- *VERITAS File System Quick Start Guide*
- *VERITAS File System 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/root/opt/VRTSfsdoc` directory. When you do a `pkgadd` for `VRTSfsdoc`, the program gives you the choice of installing the PostScript, PDF, or both formats. 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 Quick Start Guide*

After installing the `VRTSfsdoc` package, you can access this guide in the directory `/opt/VRTSfsdoc/quick_start/qsg.ps`.

- *VERITAS File System 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. To view or print PDF documents, you must use 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 and your PostScript printer.

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 System Administrator's Guide*, the *VERITAS File System Quick Start Guide*, and the online manual pages were updated for the 3.3.2 release.

Manual Pages

This release includes the following online man pages as part of the `VRTSvxfs` package. The `pkgadd` command installs these man pages in the appropriate directories in `/usr/share/man`, but does not update the `windex` database. To ensure that new VxFS commands display correctly, update the `windex` database after installing `VRTSvxfs`. See the `catman(1M)` man page for more information.

Section 1

- `cp_vxfs`
- `cpio_vxfs`
- `getext`
- `ls_vxfs`
- `mv_vxfs`
- `setext`

Section 1M

- `df_vxfs`
- `ff_vxfs`
- `fsadm_vxfs`
- `fscat_vxfs`
- `fsck_vxfs`
- `fsdb_vxfs`
- `fstyp_vxfs`
- `mkfs_vxfs`



-
- mount_vxfs
 - ncheck_vxfs
 - vxdump
 - vxedquota
 - vxlicense
 - vxquot
 - vxquota
 - vxquotaoff
 - vxquotaon
 - vxrepquota
 - vxrestore
 - vxtunefs
 - vxupgrade

Section 4

- fs_vxfs
- inode_vxfs
- tuneftab

Section 7

- vxfsio



Software Limitations and Problems in VxFS Release 3.3

▼ **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**

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 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 8 KB kernel stack size, so during the `VRTSvxfs` installation, entries are added to `/etc/system` to increase the kernel thread stack size to 16 KB.

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

The `vxupgrade` command cannot upgrade a Version 1 file system directly to Version 4. You must first upgrade to Version 2, then to Version 4. Also, a Version 2 file system without quotas cannot be upgraded to Version 4 with quotas.

▼ **Doing a `pkgadd` After a `pkgrm` May Cause a System Crash**

On Solaris operating systems, loading kernel modules with similar symbol tables can cause a system crash. When you upgrade the `VRTSvxfs` or `VRTSqio` package, reboot the system after doing a `pkgrm` and before doing a `pkgadd`. This problem does not exist on VxFS 3.3.2 because the kernel module is not loaded by the `pkgadd`.



▼ **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 percent 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.

▼ **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 changed.

If a Version 2 file system is mounted on a system running VxFS 3.3, and that file system is subsequently used on an earlier version of VxFS, then 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)` man pages for details.



▼ Some Fields Not Displayed by `fstyp` Command

The `fstyp -v` option shows the super-block. For the Version 4 file system, 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:

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

The above fields are displayed by `fstyp -v`. You can use `mkfs -m raw_device_file` to display many fields that are not part of super-block. See the `mkfs_vxfs(1M)` and `mkfs(1M)` man pages for more information.

Software Problems Fixed in This Release

This section lists problems fixed since VxFS Release 3.3.1. The VERITAS incident numbers are in brackets.

- [30117] When using Quick I/O with clones on large multiprocessor machines, a race condition sometimes occurred that resulted in a system panic.
- [30098] Panics occurred when upgrading a Version 2 disk layout whose size was not aligned on a 65536 sector boundary.
- [29888] System would hang on a call to `vx_iget()` during a short period of time when a race condition occurred.
- [29744] The Version 2 disk layout allows an unlimited number of inodes. When upgrading a file system from a Version 1 disk layout to Version 2, the ninode limit (specified by `mkfs`) of the Version 1 disk layout file system was not correctly inherited by the destination Version 2 layout file system.
- [29858] Quick I/O file `mtime` was not being updated, preventing correct database hot backups.
- [29696] File system was corrupted when accessing large files—caused by a problem in Solaris. Implemented code checks for reads and writes larger than two gigabytes to avoid the problem (fixed by Solaris patch 103640-25).



-
- [29515] Performance on VxFS was adversely affected by problems in the automatic read-ahead code.
- [29273] A message printed by the `fsadm -o largefiles` command, “file system dirty, run fsck first,” was misleading in that the file system was not actually marked “dirty” and only needed to be unmounted and mounted. The message was modified and information was added to the `fsadm(1M)` man page.
- [29217] Panics occurred when `fsadm` was doing a structural reorganization on the VxFS extent map file.
- [29072] The VERITAS-specific `cp` and `cpio` commands received errant `permission denied` messages when copying read-only files that had extent attributes or were larger than eight kilobytes.
- [28724] Performance was adversely affected with Oracle installed on VxFS 3.x. when doing calls to the Solaris `vx_page_create()` routine because the routine was deprecated on Solaris 2.6.
- [28865] Added the ETIMEDOUT diagnostic to the `vxfsio(7)` man page.
- [28739] Some fields printed by the `vxrepquota` command were not aligned or spaced correctly.
- [28495] On VxFS file systems that were upgraded from the version 1 or 2 disk layout to the version 4 disk layout, the `ilist` inode extent map was corrupted when the file system size increased to more than two gigabytes. See the `fsadm_vxfs(1M)` man page for more information.
- [28487] Doing a `setfacl -md` command from NFS to VxFS would return incorrect ACL information.
- [28313] Trying to create a VxFS file with a negative user ID through a PC/NFS client caused an error.
- [28214] The `fsadm` command would fail with an EBUSY error when a file in the file system was transferred using FTP.
- [28165] VxFS was getting unexpected return values from some ACL function calls. This prevented some system backup applications from working correctly.



-
- [28130] The description of kernel error message 069 in the *VERITAS File System System Administrator's Guide* incorrectly stated that the message would display only if system tuneables were altered manually.
- [27907] Panics occurred in `vx_attr_findbuf()` when a file system examined an inode's ACL attribute.
- [27699] Panics occurred because VxFS did not handle a control-z signal properly.
- [27579] The snapshot would hang when NetBackup performed direct I/O to the snapshot while the database was active.
- [27294] VxFS was slow doing sequential writes using NFS 3.
- [27251] The `fsadm` command was not shrinking file systems correctly for some combinations of current file system size and final file system size.
- [27120] Writes by HSM to a file system, which should be transparent to applications, were updating files' ctime attribute.
- [26892] System would hang on multiprocessor machines with large physical memory (greater than one gigabyte) when kernel virtual space became fragmented. The problem was addressed by making changes to the VxFS buffer cache code and installing a Solaris kernel patch. The patch minimizes the impact of kernel memory fragmentation when VxFS and the VERITAS Volume Manager are configured together.
- [26755] (Accelerator *for* NFS Version 1.0.4) The `vxldlog` command would do a coredump when attaching a second, non-Volume Manger Accelerator volume to the log device.
- [26436] The `vxrepquota` command printed user IDs instead of account names.
- [26435] The `vxquotaon` command would fail without returning an error message.
- [26280] On Solaris 2.6 systems, panics occurred in `vx_active_common_flush()` code when unmounting a VxFS file system while an `fsck` was in progress.



- [26280] Panics occurred when trying to unmount VxFS file systems.
- [23116] Could not mount a file system that was previously mounted using the `mount -o nolog` option. The `nolog` option is no longer supported; using `nolog` is now equivalent to using the `tmplog` option. See “End of Product Support” on page 5 for more information.
- [20939] Synchronous writes adversely affected VxFS performance.
- [14339] Could not mount a snapshot of a VxFS file system using the block device as the argument to the `-o snapof=` option.

Software Problems Fixed in VxFS 3.3.1

- [27162] Running a full structural `fsck` on a VxFS file system with a Version 2 disk layout having UIDs or GIDs greater than 65,534 resulted in the removal of those files. This occurred only on 3.x releases.
- [26728] System panics occurred when mounting a file system read-only while its NFS exported read-write and being write accessed by NFS clients.
- [26699] A full `fsck` did not rebuild the damaged current usage table (CUT) correctly, which could lead to usable inodes being marked bad.
- [25972] The `vxdump` utility was not reading date information from `/etc/dumpdates` correctly when doing an incremental dump on a snapshot file system.
- [25911] The `vxdump` and `vxrestore` utilities failed to connect to a remote non-UNIX host.
- [25782] There was a memory allocation problem when calls were made using the `VX_MEMNOWAIT` flag.
- [25671] Solaris generated error messages incorrectly identifying VxFS as the cause of system panics. Added code to ignore requests to sync during a system panic.
- [21707] Sequential 8 K reads affected performance on VxFS.



Using VxFS in FirstWatch and Other HA Environments

The VERITAS file system can be used in FirstWatch® and other High Availability environments, much as a UFS file system can. One difference is that since the `vxfs` driver is a loadable driver and the UFS driver is not, the `vxfs` driver 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, the following line should be added in the same position to each host's `/etc/system` file:

```
forceload: fs/vxfs
```

Note: 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.

Using QuickLog With FirstWatch

See the *VERITAS File System Installation Guide* for information on installing and using VERITAS QuickLog with FirstWatch.

