

# **VERITAS File System 4.1 Release Notes**

**HP-UX 11i v2**

**Fourth Edition**



**Manufacturing Part Number: 5991-1139**

**May 2005**

Edition 4

Printed in the United States

© Copyright 2004-2005 Hewlett-Packard Development Company L.P.

---

## Legal Notices

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

### Trademark Notices

UNIX® is a registered trademark in the United States and other countries, licensed exclusively through The Open Group.

VERITAS™ is a registered trademark of VERITAS Software Corporation.

---

## Publication History

The manual publication date and part number indicate its current edition. The publication date will change when a new edition is released. The manual part number will change when extensive changes are made.

To ensure that you receive the new editions, you should subscribe to the appropriate product support service. See your HP sales representative for details.

- First Edition: December 2002, 5187-1878, HP-UX 11i Version 1(B.11.11)  
CD-ROM (Software Pack 11i December 2002) and Web (<http://docs.hp.com>)
- Second Edition: August 2003, 5971-4774, HP-UX 11i Version 2 (B.11.23)  
Web (<http://docs.hp.com>)
- Third Edition: December 2003, B3929-90015, HP-UX 11i Version 2 (B.11.23)  
Web (<http://docs.hp.com>)
- Fourth Edition: June 2005, 5991-1139, HP-UX 11i version 2

## Preface

### 1. VERITAS File System 4.1 Release Notes

New Features and Changes.....	7
End of Product Support.....	10
Unsupported Features.....	10
Compatibility With Previous Versions of VxFS.....	11
Patches Required.....	11
Documentation.....	12
Online Manpages.....	12
Command Summary.....	14
Software Issues in VxFS 4.1.....	15
Product Licensing.....	18

---

# Contents

---

## Preface

The document provides information on VERITAS File System 4.1. This version of VxFS operates only on the HP-UX 11i v2 September 2004 (or later) release installed.

## Conventions

The following table describes the typographic conventions used in this guide.

**Table 1**

Typeface	Usage	Examples
monospace	Computer output, file contents, files, directories, software elements such as command options, function names, and parameters	Read tunables from the <code>/etc/vx/tunefstab</code> file. See the <code>ls (1)</code> manual page for more information.
<i>italic</i>	New terms, book titles, emphasis, variables to be replaced by a name or value	See the <i>VERITAS File System 4.1 Release Notes</i> for details.
<b><i>monospace</i></b> <i>(bold and italic)</i>	Variables to be replaced by a name or value in user input	<code># mount -F fstype mount_point</code>

Symbol	Usage	Examples
%	C shell prompt	Not applicable
\$	Bourne/Korn/Bash shell prompt	Not applicable

**Table 1** (Continued)

Typeface	Usage	Examples
#	Superuser prompt (all shells)	Not applicable
\	Continued input on the following line	# <b>mount -F vxfs \</b> <b>/h/filesys</b>
[]	In a command synopsis, brackets indicate an optional argument	ls [ -a ]
	In a command synopsis, a vertical bar separates mutually exclusive arguments	mount [suid   nosuid ]

## Getting Help

For license information (U.S. and Canadian Customers) contact:

- Phone: 650-960-5111
- Email: [hplicense@mayfield.hp.com](mailto:hplicense@mayfield.hp.com)

For license information (Europe) contact:

- Phone :+33.(0)4.76.14.15.29
- Email: [codeword\\_europe@hp-france-gen1.om.hp.com](mailto:codeword_europe@hp-france-gen1.om.hp.com)

For the latest patch information contact:

- <http://itrc.hp.com>

For Technical Support contact:

- <http://welcome.hp.com/country/us/en/support.html>

# 1 VERITAS File System 4.1 Release Notes

This document provides information on VERITAS File System Release 4.1. HP recommends you review the entire document before installing Base-VXFS 4.1 SD-Bundle. The following topics are discussed in this guide:

- “New Features and Changes” on page 7
- “End of Product Support” on page 10
- “Unsupported Features” on page 10
- “Patches Required” on page 11
- “Documentation” on page 12
- “Online Manpages” on page 12
- “Command Summary” on page 14
- “Software Issues in VxFS 4.1” on page 15

---

## New Features and Changes

VxFS 4.1 has the following new features and changes:

- **Applications using `statvfsdev(3C)` interface need to relink with new LIBC**

The applications which use `statvfsdev(3C)` interface on a device name, need to relink with the new LIBC to be able to understand disk layout Version 6, supported by VxFS 4.1. If the applications are not relinked, they will not be able to recognize disk layout version 6 filesystem. However, they would continue to recognize disk layout Version 4 and 5 filesystem.

- **Large File Systems and Large File Support**

This release is the first release to support a new disk layout (Version 6) that enables the creation of files and file systems up to 256 TB in size. File systems created on VxFS 4.1 will by default use the Version 6 disk layout. File systems larger than 1TB can be created

only on 64-bit kernels and must be created on VERITAS Volume Manager volumes. The 64 TB file system requires a 4 K block size, and a 128 TB or larger file system requires an 8 K block size. File systems over 32 terabytes are called extra large file systems and to create them the user must have a license.

You can use the online conversion utility `vxupgrade`, to upgrade existing disk layouts to Version 6 on mounted file systems.

---

**NOTE** In this release of Base-VXFS 4.1 file sizes can grow only up to 2 TB.

---

- **Reverse Path Name Lookup**

The reverse path name lookup feature obtains the full path name of a file or directory by providing an inode number as an argument to the `vxlsino` command. Reverse path name lookup can be useful in a variety of applications, such as:

- Clients of the VxFS 4.1
- File Change Log feature
- backup and restore utilities.
- In replication products

There is a new command associated with this feature:

`vxlsino` VxFS 4.1 reverse path name lookup utility.

The reverse name lookup feature includes the following application programming interface (API):

`vxfs_inotopath` Returns path names for a given inode number.

`vxfs_inotopath_gen` Returns path names for a given inode number with a specific generation count.

- **Intent Log Resize**

The `log` option of the `fsadm` command for Version 6 disk layout, can be used to resize the intent log. Dynamically changing the intent log size can improve system performance. The maximum intent log size is now 2 gigabytes. See the *VERITAS File System 4.1 Administrator's Guide* and the `mkfs_vxfs(1M)` and `fsadm_vxfs(1M)` manpages for more information on intent log size.

- **Named Data Streams**

Traditional UNIX files have an inode identifier and a single stream of file data. Use named data streams to retain and associate inodes with multiple data streams. In Base-VXFS 4.1, the original data stream is accessed in the same way as in previous



releases, but other data streams are referenced through a new directory inode associated with the file. The directory inode points to the new inodes, one per stream. The directory containing the named streams is not directly visible to the user.

VxFS 4.1 named data stream functionality is available through the following application programming interface (API) functions:

<code>vxfs_nattr_link</code>	Links to a named data stream.
<code>vxfs_nattr_open</code>	Open a named data stream.
<code>vxfs_nattr_rename</code>	Renames a named data stream.
<code>vxfs_nattr_unlink</code>	Removes a named data stream.

- **Other VxFS 4.1 Enhancements**

- The `fsadm` command now allows fragmentation reporting and defragmentation operations on an individual file or directory, or on Storage Checkpoints.
- Disk layout versions prior to Version 4 cannot be mounted. You must use the `vxfsconvert` command to convert them to a disk layout version that can be mounted.
- The `vxupgrade` command upgrades only Version 4 and Version 5 disk layouts. Disk layout versions prior to Version 4 cannot be mounted.
- The `largefiles` option is the default file size option for the `mount` command and `mkfs` command.
- The `delaylog` option is the default intent logging mount option. Use the `mount` command `delaylog` option to cache data to improve performance. See the *VERITAS File System 4.1 Administrator's Guide* for more information.
- You can remove the default storage checkpoint creation.
- Operation of the `intent log` replay has been improved to increase the speed of recovery after a file system failure.
- The `histlog` function has been implemented in the `fsdb_vxfs` command. The history log records structural changes to the file system to aid in product support.
- Two new tunable parameters, `inode_aging_count` and `inode_aging_size`, for use with the Storage Checkpoint API, have been added to the `vxtunefs` command. See the *vxtunefs(1M)* manpage for more information.

More VxFS 4.1 functions can be performed from the VERITAS Enterprise Administrator GUI.

The `vxfsu_get_ioffsets` library call has been renamed `vxfs_get_ioffsets`.

- **VxFS Enhancement**

The `vx_ninode` and `vxfs_bc_bufhwm` can be dynamically tuned without rebooting the machine. The changes take effect immediately, although they can be optionally made to take effect only after a reboot. See the *VERITAS File System 4.1 Administrator's Guide* for more information.

---

## End of Product Support

The VxFS 4.1 release does not support the creation or mounting of file systems using the Version 3 disk layouts. VxFS 4.1 still provides the capability to upgrade to Version 3 disk layouts using the `vxfsconvert` utility. VERITAS recommends that you upgrade all older disk layouts.

The next major release will be the last to support the Version 4 and Version 5 disk layouts. HP recommends that you upgrade file systems using these older disk layouts to Version 6. See “Compatibility With Previous Versions of VxFS” on page 11 for more information on upgrading.

The QuickLog feature is not available in the VxFS 4.1 release. The multi-volume support feature replaces most of the functionality provided by QuickLog. See the *VERITAS File System 4.1 Administrator's Guide* for information on migrating QuickLog devices to the multi-volume support function.

This is the last release to support shared (cluster) mounts on disk layout Version 5 cluster file systems. You must upgrade all current cluster file system disk layouts to Version 6 to enable new VERITAS File System features and to support upgrades in future releases.

---

## Unsupported Features

The following features are not supported in this release:

- VERITAS Quick I/O for Databases
- VERITAS QuickLog
- VERITAS FlashSnap
- Storage Checkpoints

- Cluster File Systems
- Moving VERITAS Volume Manager Volumes or VERITAS File Systems from HP-UX 11i Version 1 to HP-UX 11i Version 2 systems has not yet been tested and is not supported.

---

## Compatibility With Previous Versions of VxFS

VxFS 4.1 file systems employ disk layout Version 6. With the VxFS 4.1 release, you can no longer create or mount Version 3 disk layout file systems. Base-VxFS 4.1 is the last major release to support disk layout Version 4 and Version 5. HP recommends upgrading any previously installed VxFS file system to Version 6 disk layout available with VxFS 4.1.

Before or after installing VxFS 4.1, you can upgrade the disk layout on mounted file systems using the `vxupgrade` command (see the `vxupgrade(1M)` manual page for details), or `vxfsconvert` (see the `vxfsconvert(1M)` manual page) to upgrade the disk layout of unmounted file systems post installation.

You must read this guide for more information on upgrading from previous file system disk layout versions.

---

**CAUTION** VxFS file systems must be cleanly unmounted before upgrading to the VERITAS File System 4.1 release from any previous release.

---

---

## Patches Required

For a list of the required HP-UX patches, see the *VERITAS 4.1 Installation Guide*.

---

**NOTE** If HP Logical Volume Manager (LVM) is used with disk layout Version 6 and Base-VXFS 4.1, the LVM patch PHCO\_33308 needs to be installed. This patch provides the correct `LIBC` relink with LVM.

---

If LVM is used the user needs to install the following patches as a safeguard for potential data corruption.

- **LVM**

The default disk layout of VxFS 4.1 is version 6. When reducing a volume, LVM checks on the volume to prevent part of the existing filesystem from getting lost. Likewise, when an LVM physical volume is created, LVM checks to make sure that it is not clobbering an existing filesystem. LVM does not recognize a VxFS 4.1 filesystem with disk layout version 6. To make LVM work with disk layout version 6, PHCO\_33308 needs to be installed.

- **swapon**

PHCO\_33238 must be installed on the system for swapon command to enable VxFS 4.1 filesystem for paging.

---

## Documentation

The *VERITAS File System 4.1 Administrator's Guide* and the *VERITAS File System 4.1 Release Notes* are available for download from <http://www.software.hp.com>. These documents are provided as PDF files. To view or print PDF documents, you need the Adobe Acrobat Reader.

---

## Online Manpages

This release includes the following online manual pages:

**Table 1-1 VERITAS Filesystem 4.1 Manpages**

Section 1M	Description
extendfs_vxfs	Extend the size of VxFS file system.
getext	Gets extent attributes for a VxFS file system.
setext	Sets extent attributes on a file in a VxFS file system.
df_vxfs	Reports the number of free disk blocks and inodes for a VxFS file system.

**Table 1-1 VERITAS Filesystem 4.1 Manpages (Continued)**

ff_vxfs	Lists file names and inode information for a VxFS file system.
fsadm_vxfs	Resize or defragment a VxFS file system.
fscat_vxfs	Cats a VxFS file system.
fsck_vxfs	Checks and repairs a VxFS file system.
fsdb_vxfs	VxFS file system debugger.
fstyp_vxfs	Returns the type of file system on a specified disk partition.
mkfs_vxfs	Constructs a VxFS file system.
mount_vxfs	Mounts a VxFS file system.
ncheck_vxfs	Generates path names from inode numbers for a VxFS file system.
vxdump	Incremental file system dump.
vxfsconvert	Converts an unmounted file system to VxFS or upgrades a VxFS disk layout version.
vxfsstat	Displays file system statistics.
vxrestore	Restores a file system incrementally.
vxtunefs	Tunes a VxFS file system.
vxupgrade	Upgrades the disk layout of a mounted VxFS file system.
<b>Section 4</b>	<b>Description</b>
fs_vxfs	Format of a VxFS file system volume.
inode_vxfs	Format of a VxFS file system inode.
tunefstab	VxFS file system tuning parameters table.
<b>Section 7</b>	<b>Description</b>
vxfsio	VxFS file system control functions.
vxlsino	VxFS reverse path name lookup utility.
vxumount	Unmount a VxFS File System

---

## Command Summary

This release includes the following commands:

**Table 1-2 VERITAS File System 4.1 Commands**

<b>Command</b>	<b>Description</b>
df	Reports the number of free disk blocks and inodes for a VxFS file system.
extendfs	Extend a VxFS file system size.
ff	Lists file names and inode information for a VxFS file system.
fsadm	Resize or defragment a VxFS file system.
fscat	Cats a VxFS file system.
fsck	Checks and repairs a VxFS file system.
fsdb	VxFS file system debugger.
fstyp	Returns the type of file system on a specified disk partition.
getext	Gets extent attributes for a VxFS file system.
mkfs	Constructs a VxFS file system.
mount	Mounts a VxFS file system.
ncheck	Generates path names from inode numbers for a VxFS file system.
newfs	Creates a new VxFS file system.
setext	Sets extent attributes on a file in a VxFS file system.
vxdump	Incremental file system dump.
vxedquota	Edits user quotas for a VxFS file system.
vxfstocvt	Converts an unmounted file system to VxFS or upgrades a VxFS disk layout version.
vxfststat	Displays file system statistics.

**Table 1-2 VERITAS File System 4.1 Commands (Continued)**

Command	Description
vxlsino	VxFS reverse path name lookup utility.
vxrestore	Restores a file system incrementally.
vxtunefs	Tunes a VxFS file system.
vxupgrade	Upgrades the disk layout of a mounted VxFS file system.
vxumount	Unmount a VxFS File System

---

## Software Issues in VxFS 4.1

- **Systems With Less Cache Memory Hang**

VxFS 4.1 allocates a fixed amount of memory. You can change the amount of memory allocated using the tunables `vx_ninode` and `vx_bc_bufhwm`. After long hours of operation, low-memory systems (that is, systems having less physical RAM), may slow down or hang due to memory pressure. To alleviate this problem, lower the values of `vx_ninode` and `vx_bc_bufhwm` to limit VxFS 4.1 memory consumption. The `vx_ninode` and `vxfs_bc_bufhwm` can now be dynamically tuned without rebooting the machine.

---

**NOTE** VxFS 4.1 consumes a fixed percentage of memory for storing the default values of tunables such as `vx_ninode` and `vx_bc_bufhwm`. For the default value value of `vx_ninode` VxFS consumes nearly 10% of total physical memory. If the system has 512 MB of memory, VxFS 4.1 will require up to 52 MB to store the VxFS 4.1 inode cache. Therefore it is necessary to tune the tunables according to the type of load on the system.

---

- **Tuning Down the Inode Cache**

VxFS 4.1 file systems allocate and release inodes based on the file system load. Generally, larger inode caches help the file systems to perform better in case of a file server or web server load. The global dynamic tunable, `vx_ninode`, determines the maximum possible size of the VxFS 4.1 inode cache. If the default value of `vx_ninode` is set to zero, VxFS 4.1 automatically tunes the size of the inode cache at boot time, based on the size of the

physical memory on the system. On systems with a RAM size of less than or equal to 1 GB per CPU, you can manually tune down `vx_ninode` to a value not less than that of `nfile` (`nfile` is an HP-UX tunable that represents the maximum number of file descriptors).

- **Tuning Down the Buffer Cache**

VxFS 4.1 implements a private buffer cache to use exclusively for metadata. The allocations to this buffer cache vary during system usage based on file system load and the maximum cache size specified in the global (static) tunable, `vx_bc_bufhwm`. If the value of `vx_bc_bufhwm` is set to zero, VxFS 4.1 automatically tunes the maximum size of the metadata buffer cache at boot time, based on the system memory size. On systems with a RAM size of 1 GB per CPU or less, you can manually tune down the value of `vx_bc_bufhwm` to a minimum of 6144 (6MB).

- **Full File System Cannot Be Resized**

The `fsadm` command cannot resize a 100% full file system due to lack of space for updating structural information. You must check VxFS file systems on a regular basis and increase their size if they approach 100% capacity. If a file system is busy or too fragmented, the resize operation can also fail.

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

Data integrity problems occur in disk drives and disk arrays after power failure or SCSI bus reset. This problem occurs as the cached data has not been written to non-volatile. Contact your disk drive or disk array manufacturer to determine whether your disk drives use a write-back cache, and if the configuration can be changed to disable write-back caching.

- **The command `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.

- **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 displays information on super-block. For the Version 4 and Version 5 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 fields that are not part of super-block. See the *mkfs\_vxfs(1M)* and *mkfs(1M)* manual pages for more information.

- **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 Base-VXFS 4.1, and that file system is subsequently used on an earlier version of Base-VXFS 4.1, 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 4.1 unless the file system is dirty, in which case a full `fsck` is required.

- **Non-standard Command Behavior When Using Access Control Lists (ACLs)**

The output of the `ls -l` command on VxFS 4.1 file systems shows `mask /CLASS_OBJ` in place of group permissions if ACLs are in use on a file or a directory. You can determine the effective group permissions by using the `getfacl` command.

The `chmod` command changes `mask /CLASS_OBJ` instead of the group permissions if ACLs are in use on a file or a directory. `GROUP_OBJ` is not changed by `chmod`, and because effective group permissions are determined by `GROUP_OBJ` and `CLASS_OBJ`, the default group may not receive the permissions specified by `chmod`. Because `ls -l` shows `mask` only (which is changed by `chmod`), it only appears that the group permissions are changed as specified in `chmod`. Use of `chmod` command is not advisable on files with ACLs. Instead, use the `getfacl` command to manipulate permissions.

In the next release of VxFS, the behavior of both `mask /CLASS_OBJ` and `GROUP_OBJ` will change after executing the `chmod` command. Even in that case, using the `getfacl` command is advisable for manipulating permissions.

See the following manual pages for ACL-related information: *aclcheck(3)*, *aclsort(3)*, *chmod(1)*, *getfacl(1)*, *ls(1)*, *setfacl(1)*, and *uname(1)*.

- **The `newfs -R` Command Allows Reserving Swap Space Larger Than the Device**

The `newfs -R` command will reserve swap space greater than is available on the underlying device. This problem will be addressed in a future patch release from HP.

## Product Licensing

The Base-VXFS 4.1 SD-Bundle includes all features that are part of the full VxFS. Users can enable the related features with existing VxFS 3.5 licenses. However you have reapply the VxFS 3.5 licences as follows:

```
# /sbin/fs/vxfs4.1/vxenablef -a
```