

Veritas File System 5.0.1 Release Notes

HP-UX 11i v3

HP Part Number: 5900-0032
Published: November 2009
Edition: 1.0



© Copyright 2009 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.

UNIX is a registered trademark of The Open Group.

Veritas is a registered trademark of Symantec Corporation.

Copyright © 2009 Symantec Corporation. All rights reserved. Symantec, the Symantec Logo, Veritas, and Veritas Storage Management are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

Table of Contents

1 Veritas File System 5.0.1 Release Notes.....	5
Overview.....	5
New Features.....	5
System Requirements.....	7
Compatibility With Previous Versions of VxFS.....	7
Tuning VxFS for Performance.....	7
Product Licensing.....	8
End of Support.....	11
Features no Longer Supported.....	11
Known Limitations.....	12
Known Problems and Workarounds.....	12

List of Tables

1-1	Features Enabled by HP OnlineJFS.....	8
1-2	Features Enabled by HP Serviceguard Storage Management Licenses.....	9

1 Veritas File System 5.0.1 Release Notes

This chapter describes the new features, licensing, system requirements, compatibility with previous releases and known problems with the Veritas File System 5.0.1 on HP-UX 11i v3.

This chapter addresses the following topics:

- “Overview” (page 5)
- “New Features” (page 5)
- “System Requirements” (page 7)
- “Compatibility With Previous Versions of VxFS” (page 7)
- “Tuning VxFS for Performance” (page 7)
- “Product Licensing” (page 8)
- “End of Support” (page 11)
- “Features no Longer Supported” (page 11)
- “Known Limitations” (page 12)
- “Known Problems and Workarounds” (page 12)

Overview

The Veritas File System (VxFS) is an extent-based, intent logging file system. VxFS is designed for use in UNIX environments that require high performance and availability, and deal with large volumes of data.

Currently, VxFS 5.0.1 is available on independent media only but may later be integrated into the HP-UX 11i v3 Operating Environment, replacing the currently available VxFS 5.0. version.

New Features

The following new features are supported with VxFS 5.0.1 on the HP-UX 11i v3 operating system:

- **The Direct I/O Feature**
This feature now does not require the HP OnlineJFS license and is available in all HP-UX Operating environments with the Base-VxFS bundle.
- **The Concurrent I/O Feature**
The concurrent I/O database acceleration feature was previously available with the SG-SMS license. Starting with the VxFS 5.0.1 release, this feature is also available with the HP OnlineJFS license and will continue to be available with the SG-SMS license.
- **Setting Extent Attributes**
This disk space management feature now does not require the HP OnlineJFS license and is available in all HP-UX Operating environments with the Base-VxFS bundle
- **Increased File System Size**
The HP OnlineJFS license previously supported file systems of size up to 32 TB. Starting with this release of VxFS 5.0.1, the HP OnlineJFS license supports file systems of size up to 256 TB.
- **The SmartMove™ Feature**
SmartMove reduces the time and I/O required to attach or reattach a plex to an existing VxVM volume, in the specific case where a VxVM volume has a VxFS file system mounted on it. The SmartMove feature uses VxFS information to detect free extents and avoids copying them.

SmartMove provides the following benefits:

- Less I/O is sent through the host, through the storage network, and to the disks or LUNs.
- Faster plex creation results in faster array migrations.
- Ability to migrate from a traditional LUN to a thinly provisioned LUN. This removes unused space in the process.

In this release, the SmartMove feature is turned off by default. You can enable this feature by specifying the `usefssmartmove=on` option in the `/etc/default/vxsf` file. To use the SmartMove feature, VxVM and VxFS version 5.0.1 or later is required.



NOTE: The Smartmove feature is available only with the SG-SMS license.

- **Dynamic Storage Tiering Enhancements**

The Dynamic Storage Tiering (DST) feature provides the following enhancements:

- Enhanced DST APIs to provide a new interface for managing allocation policies of storage checkpoints during creation and later, and for managing named data stream allocation policies.
- `fspadm` support for user ID (UID), group ID (GID), and tagging (TAG) elements in the placement policy XML file.
- Improved scan performance in the `fspadm` command
- Suppressed processing of the chosen RULE.
- Parser support for UID, GID, and TAG elements in a DST policy
- Storage Checkpoint data placement support in a DST policy
- Shared DB thread handle support
- CPU and I/O throttling support for DST scans
- New command, `fstag`, for file tagging
- New command, `fspm`, for creating XML policies



NOTE: To ensure proper functioning of DST enhancements with databases, upgrade the SQLite version to 3.3.9 or later.

The DST feature is available only with the SG-SMS license.

- **Availability of the `mntlock` and `mntunlock` Mount Options**

You can specify the `mntlock` option with the `mount` command to prevent a file system from being unmounted by an application. This option is useful for an application to prevent file systems (that it monitors) from being improperly unmounted by other applications or administrators. The `mntlock` option is useful particularly for clustering applications where multiple file systems are monitored by different applications or administrators.

The `mntunlock` option of the `vxumount` command disables the `mntlock` operation if you previously locked the file system.



NOTE: These features are available only with the HP OnlineJFS license.

System Requirements

The operating system version supported and the patches required for VxFS 5.0.1 on HP-UX 11i v3 are as follows:

- **Operating System Version**

HP-UX 11i v3 0903 OEUR (or later)

- **Required Patches**

For information on the required and recommended patches, see *Veritas 5.0.1 Installation Guide HP-UX 11i v3*.

Compatibility With Previous Versions of VxFS

The following section discusses the compatibility of VxFS 5.0.1 with the previous VxFS releases:

- The VxFS 5.0.1 release does not support creating or mounting disk layout version 1, 2, or 3 file systems. HP recommends that you upgrade any previously installed VxFS file system to Disk Layout Version (DLV) 7 available with VxFS 5.0.1.

An online conversion utility `vxupgrade` is provided to upgrade existing disk layouts to Disk Layout Version 7 on mounted file systems.



CAUTION: Do not upgrade `/stand` and `/` file systems to Disk Layout Version 6 or Disk Layout 7. The HP-UX boot loader does not understand this layout. The HP-UX boot loader does not recognize DLV greater than 5.

- The disk layout of a VxFS file system can be upgraded after installing VxFS 5.0.1. Use the `vxupgrade` (1M) command to upgrade the disk layout from Version 4 or 5 to disk layout Version 6, and then upgrade to Version 7 on a mounted file system. Use the `vxfsconvert` (1M) command to convert an unmounted file system to VxFS Disk Layout Version 7.



NOTE: A Disk Layout Version 7 or Disk Layout Version 6 file system created with VxFS 5.0.1 software is not accessible if the VxFS 5.0.1 file system software is removed and the system is reverted to VxFS 4.1, which is the default supported with HP-UX 11i v3 release.

Tuning VxFS for Performance

- **Tuning down the inode cache**

VxFS allocates and releases inodes based on the file system load. Generally, larger inode caches help the file system to perform better in the case of a file server or web server load. The global tunable, `vx_ninode`, determines the maximum possible size of the VxFS inode cache. If the default value of `vx_ninode` is set to zero, VxFS 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, you can manually tune down `vx_ninode` to a value not less than that of `ninode`. If the user does not manually tune the value of `vx_ninode` then VxFS calculates `vx_ninode` value based on system memory. VxFS ensures that the value of `vx_ninode` is not less than `ninode`. For more information on these commands, see `vx_ninode(5)`.

- **Tuning down the buffer cache**

VxFS 5.0.1 implements a private buffer cache to use exclusively for metadata. The allocations to this buffer cache vary during system usage based on the file system load and the maximum

cache size specified in the global (static) tunable, `vxfs_bc_bufhwm`. If the value of `vxfs_bc_bufhwm` is set to zero, VxFS 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 or less, you can manually tune down the value of `vxfs_bc_bufhwm` to a minimum of 6144 (6 MB). If the user does not manually tune the value of `vxfs_bc_bufhwm` then VxFS calculates `vxfs_bc_bufhwm` value based on system memory. VxFS ensures that the value of `vxfs_bc_bufhwm` is not less than 800 KB.

For more information on tuning VxFS, see the following documents available at <http://www.docs.hp.com>:

- *Common Misconfigured HP-UX Resources*
- *Veritas File System 5.0.1 Administrator's Guide*



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

Product Licensing

Following are the features enabled by the HP OnlineJFS and HP Serviceguard Storage Management licenses:

- **HP OnlineJFS License**

Table 1-1 lists the features enabled by the HP OnlineJFS license.

Table 1-1 Features Enabled by HP OnlineJFS

Feature	License
Veritas File System	HP OnlineJFS
File Change Log – Logs file system activity for fast incremental backup and auditing.	HP OnlineJFS
Online FileSystem Operations (<code>fsadm</code>) <ul style="list-style-type: none"> – File system resize – File system defragmentation – Single file or directory defragmentation – File System intent log resize 	HP OnlineJFS
Mounting <ul style="list-style-type: none"> – Forced unmount of a file system (<code>vxumount -oforce</code>) – Snapshot mounts (<code>mount -osnapof</code>) – Mount with special mount options such as <code>convosync</code>, <code>unbuffered</code>, <code>direct</code> and so on. – Nested Mounts 	HP OnlineJFS
DMAPI – A VxFS interface that enables Hierarchical Storage Management (HSM) also known as XDSM.	HP OnlineJFS
File systems up to 256 TB	HP OnlineJFS and SG-SMS
Concurrent I/O database accelerator	HP OnlineJFS
Named Streams	HP OnlineJFS

For information on supported file systems and large file support, see the *Supported File and File System Sizes for HFS and JFS* whitepaper available at:

<http://www.docs.hp.com>

- **HP Serviceguard Storage Management Licenses**

Table 1-2 lists the features enabled by the HP Serviceguard Storage Management Licenses.

Table 1-2 Features Enabled by HP Serviceguard Storage Management Licenses

Feature	T2771DB	T2773DB	T2774DB	T2775DB	T2777DB
Veritas File System	Yes	Yes	Yes	Yes	Yes
Multi-Volume File System - Allows more than 1 volume to comprise a file system. The VxFS intent log and other metadata can be placed on a separate volume(s), and files can be dynamically striped across multiple volumes, avoiding volume relayouts when adding storage. This does not license Dynamic Storage Tiering lthough both share many underlying technical components.	Yes	Yes	Yes	Yes	Yes
Veritas Volume Manager (Full)	Yes	Yes	Yes	Yes	Yes
Smartmove (volume resynchronization)	Yes	Yes	Yes	Yes	Yes
Thin Provisioning	Yes	Yes	Yes	Yes	Yes
Dynamic Multipathing (DMP) - Balance I/O across multiple paths between the server and the storage array to improve performance and availability, active/passive (A/P) failover for root disk	Yes	Yes	Yes	Yes	Yes
Improved usability for I/O statistics – More visibility for I/O statics per LUN, enclosure, and array.	Yes	Yes	Yes	Yes	Yes
ALUA support	Yes	Yes	Yes	Yes	Yes
Enclosure-Based Naming – Name LUNs (array volume ID)	Yes	Yes	Yes	Yes	Yes
Hot relocation – Automatically migrates data from failing disks to healthy disks, online.	Yes	Yes	Yes	Yes	Yes
Online Administration – Limits the amount of time storage required to be offline for maintenance by performing volume resizing (including shrinking), domain reconfiguration, backup, and off-host processing while the data remains online and available (online migraiton, volume resizing, online relayout). Provides task monitoring for VxVM tasks.	Yes	Yes	Yes	Yes	Yes
Oracle SmartSync Support – Improves Oracle runtime and recovery performance with mirrored volumes. Raw volume configuration.	Yes	Yes	Yes	Yes	Yes
Oracle SmartSync Support - improves Oracle runtime and recovery performance with mirrored volumes. File system through Oracle Disk Manager (ODM) configuration.	No	Yes	Yes	Yes	Yes

Table 1-2 Features Enabled by HP Serviceguard Storage Management Licenses *(continued)*

Feature	T2771DB	T2773DB	T2774DB	T2775DB	T2777DB
Portable Data Containers (cross-platform data sharing) – Easily and quickly converts data for use on different operating systems. Makes it easy to migrate to new operating systems.	Yes	Yes	Yes	Yes	Yes
Site awareness / remote mirrors for campus clusters	No	No	Yes	Yes	Yes
Storage Foundation Management Server – Centralizes management of multiple servers to provide complete visibility to application, server, and storage resources. ¹	Yes	Yes	Yes	Yes	Yes
Veritas Enterprise Administrator (VEA)	Yes	Yes	Yes	Yes	Yes
Storage Expert – Script-based tool to check for unusual or non-recommended configurations	Yes	Yes	Yes	Yes	Yes
Dynamic LUN Expansion – Online LUN resize and automatic volume growth	Yes	Yes	Yes	Yes	Yes
Import Cloned LUN on the same host as the original LUN used with ShadowCopies, BCVs, and so on.	No	No	Yes	Yes	Yes
Dynamic Storage Tiering - Allows the administrator to identify and move infrequently used files online to less expensive storage, transparently to users and applications. ²	No	No	Yes	Yes	Yes
Storage Checkpoints and Storage Rollback - Instantly create disk-based backups or file systems without adding storage. The backup can be easily restored by the user. ²	Yes	Yes	Yes	Yes	Yes
Database Dynamic Storage Tiering – SF Enterprise 5.0	No	No	Yes	No	Yes
Flash Snap – Takes instant, full volume, or space-optimized snapshots of data for off-host processing, disk-based recovery, and backup. Only resynchronizes changed blocks for fast resynchronization. Online split/join of disk groups.	No	No	Yes	Yes	Yes
Database FlashSnap ²	No	No	Yes	No	Yes
Storage Mapping	Yes	Yes	Yes	Yes	Yes
Extent Balanced File System	No	No	Yes	No	Yes
Cluster Volume Manager	No	No	No	Yes	Yes
Oracle RAC Extensions	No	No	No	No	Yes
I/O fencing	No	No	No	Yes	Yes
Mounting	Yes	Yes	Yes	Yes	Yes
– Forced unmount of a file system (vxumount -oforce)					
– Snapshot mounts (mount -osnapof)					
– Mount with special mount options, such as convosync, unbuffered, direct and so on.					

Table 1-2 Features Enabled by HP Serviceguard Storage Management Licenses *(continued)*

Feature	T2771DB	T2773DB	T2774DB	T2775DB	T2777DB
File Change Log — Logs file system activity for fast incremental backup and auditing	Yes	Yes	Yes	Yes	Yes
Online FileSystem Operations (<code>fsadm</code>) — File system resize — File system defragmentation — Single file or directory defragmentation — File System intent log resize	Yes	Yes	Yes	Yes	Yes
DMAPI - A VxFS interface that enables Hierarchical Storage Management (HSM). Also known as XDSM.	Yes	Yes	Yes	Yes	Yes
Big File Systems - A single file system of size greater than 2 TB and lesser than or equal to 32 TB.	Yes	Yes	Yes	Yes	Yes
File systems up to 256 TB	Yes	Yes	Yes	Yes	Yes
ODM, Quick I/O (QIO); Concurrent I/O (CIO)	Yes	Yes	Yes	Yes	Yes
Cluster File System	No	No	No	Yes	Yes

1 HP-UX client support.

2 Defaults to standalone support of Oracle database features.

End of Support

The following features are no longer supported:

- VxFS 5.0.1 does not support creating or mounting file systems using Version 1, 2, or 3 disk layout. HP recommends that you upgrade all older disk layouts to Version 7.
- Future major releases of VxFS may discontinue support for DLV 4 and DLV 5.
- The `nau`, `ausize`, `aufirst`, `aupad` and `ninode` options to the `mkfs` command are not supported. The `mkfs` command accepts these options but ignores them without a warning. These options can be omitted from the `mkfs` command in future releases.
- The default version for the File Change Log (FCL) in VxFS 5.0.1 is version 4. VxFS 5.0.1 is the last major release to support FCL version 3. Upgrade FCL to version 4 to make use of the advanced logging facilities provided by Version 4 and for facilitating future upgrades.
- The `fscdstask` command, used for performing online Cross-Platform Data Sharing (CDS) operations, has been replaced by the `fscdsadm` command. Support for the `fscdstask` command might be dropped in future VxFS releases.
- The functionality provided by the Veritas multi-volume support feature replaces most of the functionality provided by the Quicklog feature.
- The snapshot feature will not be available in future Veritas releases.

Features no Longer Supported

- **Version 4 and Version 5 file system disk layouts for shared (cluster) mounts**

VxFS 5.0.1 supports shared (cluster) mount only for file systems with disk layout 6 or 7. If you have a shared (cluster) mounted file system with disk layout versions 4 or 5, you must

mount this file system locally (without the `-o cluster` mount option) and upgrade the disk layout to version 6 or 7 before mounting it with the `-o cluster` mount option.

- **I/O error mount options**

In this and future releases, only the `ioerror=disable` and `ioerror=mdisable` mount options are supported for clustered file systems. Other `ioerror` mount options are not supported. For more information, see *mount_vxfs(1M)*.

Known Limitations

The following are the software limitations in Veritas File System 5.0.1:

- **Quick I/O, ODM, `mount -o cio`, and the `VX_CONCURRENT` advisory are mutually exclusive**

The `VX_CONCURRENT` advisory cannot be set on a file that is actively open by Quick I/O or ODM. A file that has the `VX_CONCURRENT` advisory set may not be concurrently opened by Quick I/O or ODM. Quick I/O and ODM access are not allowed for any files on a file system that is mounted with the `mount -o cio` option.

Known Problems and Workarounds

Following known problems exist for VxFS 5.0.1:



NOTE: For information on the Known Problems and Workarounds in VxFS 5.0 in the HP-UX 11i v3 operating system, see *Veritas File System 5.0 Release Notes, HP-UX 11i v3, First Edition, May 2008* available at <http://www.docs.hp.com>.

- **Problem**

File Change Log tunable setting for proper functioning of Dynamic Storage Tiering applications

Workaround

If the active placement policy of a given file system uses I/O or access temperatures, after the policy becomes active by being assigned, you must tune the `fcl_malloc` tunable of the file system with the following command:

```
# vxtunefs -o fcl_maxalloc=0 mount_point
```

However, if any application other than DST uses FCL, this setting may conflict with those applications.

- **Problem**

If an ANSI C compiler is installed on a system before installing the 5.0.1 bundles, the following `swverify` errors may be displayed after installing the 5.0.1 bundles.

```
=====  
ERROR: Symlink "/usr/bin/cc" should have value "/usr/ccs/bin/cc" but  
       the actual value is "/opt/ansic/bin/cc".  
ERROR: Fileset "OS-Core.C-KRN,l=/,r=B.11.31" had file errors.  
=====
```

Workaround

You can ignore these error messages because these messages are harmless messages and do not impact any functionality on the system.

- **Problem**

When you upgrade from any old HP-UX 11i version to the HP-UX 11i v3 March 2009 OEUR release together with the Veritas 5.0.1 bundles, the upgrade fails with the following error message:

```
-----  
ERROR "You must select both Base-VxFS and Base-VxVM bundles for the upgrade to proceed."  
-----
```

The upgrade halts even if the correct Veritas 5.0.1 bundles are selected for upgrade.

Workaround

To upgrade successfully, complete the following steps:

1. Cold install or update the system to the HP-UX 11i v3 March 2009 OEUR release.
2. Install the Veritas 5.0.1 bundles using the installation procedure discussed in the *Veritas 5.0.1 Installation Guide*.



NOTE: This limitation is resolved in the HP-UX 11i v3 September 2009 OEUR or later releases.
