

## **VERITAS FILE SYSTEM 4.0 VERSUS EXT3**

**A Competitive File System Comparison**

# TABLE OF CONTENTS

Introduction .....	3
Features.....	3
Detailed Feature Comparison.....	5
Platform support .....	5
Maximum File system size .....	5
Resizing the file system.....	6
File system snapshots.....	6
Journaling technology and Maximum LOG size.....	6
Defragmentation .....	7
Supported for root.....	7
Management tools .....	7
Quality of Storage Services .....	7
Portable Data Containers .....	8
Advance online tuning options .....	8
Non-VERITAS File System Conversion Utility .....	8
Upgrade from previous version .....	8
Conclusion .....	9

## INTRODUCTION

Ext3 is a journaling file system developed for Linux operating environments and, like other Linux components, is freely available. Due to its mature maintenance utilities and easy upgrade path from Ext2, the default Linux file system, Ext3 has proven to be a popular file system for Linux configurations.

VERITAS File System (VxFS) 4.0 is a component of VERITAS Foundation Suite 4.0 and is supported on the Red Hat Enterprise Linux 3 distribution of Linux. VxFS 4.0 provides the sophisticated feature set that is to be expected of an enterprise-class commercial product. Support for online growing and shrinking of the file system, non-disruptive defragmentation, and a user-friendly graphical user interface (GUI) significantly ease file system administration and increase administrator productivity.

Performing a side-by-side comparison of the technical features of Ext3 and VxFS 4.0 gives administrators the information they need to decide which product is more appropriate for their environment. The analysis, detailed below, shows VERITAS VxFS 4.0 to be the superior file system when enterprise-class features and performance are demanded.

## FEATURES

Product feature	VxFS 4.0	Ext3
Brief description	VERITAS File System 4.0	Ext2 with added journaling support
Linux kernel support	Red Hat EL 3	Most Linux kernels
Included in distribution		Red Hat, SUSE
Platform support (other than Linux)*	Solaris, HP-UX, AIX, Linux	Linux only
Max file system size*	2 TB (32-bit)	2 TB (32-bit)
Max file size	2 TB (normal file), 16 TB (sparse file)	2 TB (normal file), 16 TB (sparse file)
Supported block sizes	1024, 2048, 4096, 8192	1024, 2048, 4096
Address range (code support)	32 and 64 bit	32 and 64 bit
Online resize*	Growing and shrinking. Integrated with VERITAS Volume Manager	Growing only.
Snapshots (non-persistent)*	Yes	No

Product feature	VxFS 4.0	Ext3
Snapshots, persistent(max number of snapshots)*	Yes (16 million). Can be mounted read-only or read-write	No
Journaling technology*	Transaction based. Metadata and/or data operations are logged. Log can be kept on a different device	Block type, stores complete data and meta data blocks in memory and on disk journal. Uses JBD (generic journaling technology)
Max log (journal) size*	256 MB (internal or external journal)	102,400 blocks. Journal can be put on a separate device
File System error handling mount options	disable, nodisable, wdisable, mwdisable	continue, remount-ro, panic
File allocation technology	Extent based, dynamic inode allocation	Inode based (all inodes are created at file system creation time)
Defragmentation*	Yes, online	No
User quota support	Yes	Yes
Group quota support	Yes	No
Extended attributes	Yes	No
Unix ACL support	No	No
Supported for root ( / )*	No	Yes
DMAPI Support	No	No
Named Data Streams	Only supported with VxFS API's	No
File System Event History Log	Yes	No
Management tools*	GUI/CLI	CLI
Posix compliant	Yes	Unknown
Quality of Storage Services*	Yes	No

Product feature	VxFS 4.0	Ext3
Portable Data Containers*	Yes	No
File system statistics (buffer cache, counters ...)	Yes	No
Direct IO support	Yes (programmatic and mount option)	Yes, programmatic interface only
Oracle Disk Manager (ODM)	Yes (Storage Foundation for Oracle RAC)	No
Advance online tuning options (read-ahead, preferred read/write size, number of read/write threads)*	Yes. Will automatically interact with VxVM to default to the most optimal settings	No tuning but change of options. Some of them online
Conversion utility from other File System	Yes. Convert from ext2 or ext3	No
File System layout upgrade from previous version*	Yes	No upgrade required

Features with an \* are described in greater detail below.

## DETAILED FEATURE COMPARISON

### Platform support

VERITAS VxFS 4.0 supports multiple platforms, including Sun (Solaris), HP (HP-UX), and IBM (AIX), as well as Linux. Currently, Ext3 is only available for Linux environments.

VxFS 4.0 offers a consistent management interface across all platforms, easing the System Administrator's learning curve. The command set for VERITAS products is also consistent across platforms, allowing scripts to be easily ported between UNIX and Linux environments.

### Maximum File system size

Linux is a 32-bit operating system with a practical file system size limit of two TB. The Linux VFS (Virtual File System) layer has a compliance-defined standard interface and a 16 TB address range limitation. In addition, the Linux block device layer has a two TB address limitation, restricting file systems to a maximum size of two TB.

Red Hat Enterprise Linux 3 implements a 32-bit Ext3 file system limited to two TB maximum file system size.

VxFS 4.0 supports the maximum allowable file system size of two TB (minus one KB). As mentioned previously, this is restricted by the block device layer's inability to address data

beyond two TB. When this limit is lifted in future versions of Linux VERITAS will support a higher value (usually limited by the operating system).

### **Resizing the file system**

Ext3 supports online growth of the file system, but shrinking is not supported. Ext3 does not interact with the underlying volume manager, and growing the file system requires the administrator to expand the volume before resizing the file system. This manual process makes it easy to make a mistake and corrupt the file system.

VERITAS VxFS 4.0 supports online growing and shrinking of the file system. Interacting with VERITAS Volume Manager (VxVM), VxFS automatically manages resizing of both the file system and the logical volume to ensure success.

### **File system snapshots**

The Ext3 file system does not have built-in snapshot functionality, however, it is able to utilize the Logical Volume Manager (LVM) to create a block-level snapshot. The block-level snapshot requires a reserve of storage in addition to the space used by the file system. The amount of storage needed depends on how active the file system is — 10 percent to 20% of the file system size is recommended — and the storage must be hosted on a volume separate from the file system. To ensure consistency, the file system must be quiesced before running the snapshot.

The Ext3 snapshot utility is designed to support online backups and must be mounted read-only. The snapshot volume is dependent on the original disks and cannot be exported to another host. However, the snapshot is reboot-persistent.

VERITAS VxFS 4.0 uses two types of snapshot technology: snapshots, and storage checkpoints. Snapshots are comparable to Ext3 snapshots and are non-persistent, read-only copies of the original file system. Designed for use with backups, these snapshots require their own storage device and do not take up free-space in the file system. If the original file system is unmounted the snapshots are removed and no longer available.

Storage checkpoints, based on copy-on-write technology, provide administrators with significantly greater flexibility when compared to snapshots. Each checkpoint is a reboot-persistent copy of the file system, available to be mounted read-write or read-only. Checkpoints are ideal for fast recovery of files, without going through the cumbersome process of restoring data from tape.

Storage checkpoints make use of the free space within the file system and store only modified data blocks. The properties of the checkpoint can be set so that the checkpoint automatically expires if the file system free space gets too low. Checkpoints can also be created from other checkpoints giving the system administrator even greater flexibility.

### **Journaling technology and Maximum LOG size**

The Ext3 file system was created by adding Linux-specific journaling technology — the journaling block device (JBD) — to the standard Ext2 file system. The JBD was formerly known as journaling file system. The system administrator has the option of logging updates to data or both data and metadata, and Ext3 guarantees atomicity of operations.

Ext3 supports a maximum journal size of 102,400 blocks. The JBD log can be put on a separate device, and several file systems can share the same journaling log.

VERITAS VxFS 4.0 has two journaling options. Both allow for a maximum log device size of 256 MB. If simplified management is desired the journal can be kept within the file system.

The latest release of VxFS includes multi-device support. This feature allows the journal to be located on a volume separate from the file system: journal data is usually accessed sequentially and can conflict with the random access typical of a file system. Moving the journal to its own storage volume can provide a substantial performance boost to the file system.

## **Defragmentation**

Fragmentation is a fact of life in all file systems. As users and applications access and modify data, fragments of files and free space become scattered around the disk significantly reducing performance of the file system.

Ext3 does not support defragmentation. Defragmenting an Ext3 file system requires volumes hosting the file system to be backed up to tape and restored. This is a lengthy, inefficient, and disruptive process.

VERITAS VxFS 4.0 supports online defragmentation. This allows administrators to carry out the performance-optimizing task of defragmenting the file system without interrupting end-user access to file system data.

## **Supported for root**

Ext3 can be established as the root file system for the Linux environment when the operating system is installed. The Ext3 binaries are included with many Linux distributions, including Red Hat 7.2, and later.

VERITAS VxFS 4.0 is a commercially available product, and must be purchased and installed separately from the Linux operating system software. VxFS cannot be installed at the same time as the operating system and is not supported as a root file system.

## **Management tools**

The Ext3 file system has no special management infrastructure and is administered using the UNIX command line interface (CLI) tools supplied with the Linux distribution.

VERITAS VxFS 4.0 and VERITAS VxVM offer a fully integrated graphical user interfaces (GUI), in addition to a CLI interface. The VERITAS Enterprise Administrator VEA GUI is java-based interface that can be run from any platform. Commands being executed against the file system are displayed and logged, by the GUI, allowing administrators to incorporate these same command sequences in CLI-based scripts for customized management of the file system. The VEA GUI provides a common interface to all VERITAS products, significantly easing the administrator's learning curve.

## **Quality of Storage Services**

Most commercial file systems today are designed to occupy the storage capacity of a single disk — whether physical or virtual. And storage systems, in turn, are designed so that each disk has a uniform configuration — mirrored, striped, etc. These restrictions limit the flexibility of storage and system administrators because all the files in any given file system occupy one virtual disk, with one virtualization configuration. The VERITAS Storage Foundation Quality of Storage Services (QoS) feature solves this problem by enabling multi-volume support in the VERITAS file system.

Closely integrated with other components of VERITAS Storage Foundation, the QoS feature allows a file system to form a pool of storage from which data is allocated. A multi-volume file system can occupy up to 255 virtual volumes, and this configuration is dynamic, allowing volumes to be added or removed as needed. Each volume in the file system can have a unique combination of capacity, concatenation, striping, mirroring, and RAID configuration to satisfy the requirements of different types of hosted data.

### **Portable Data Containers**

VERITAS Portable Data Containers (PDC) allow fast migration of user data between different computing platforms with minimal downtime for end-users. PDCs unlock user data from the confines of the server giving administrators the flexibility to locate processing resources on the optimum platform.

A VERITAS File System defined for PDC use implements a common disk layout that can be used by any supported operating system platform. Moving data from one platform to another — for example, migrating an Oracle database from a Solaris box to Linux — or sharing data between two platforms — or example, performing off-host backup processing — does not require any movement of the data. File System volumes are simply unmounted from one server and mounted to another. And because no data is physically moved the size of the data on the file system does not matter — a one KB file is moved as quickly as a one GB file.

### **Advance online tuning options**

The VERITAS File System is able to change a variety of tuning options online to improve performance. These options include the preferred IO read and write size and the number of threads reading or writing per file. Environments with large files, and many clients accessing those files, will find that these options make a significant difference in performance.

VxFS can automatically bypass page cache if it determines that an I/O is larger than a value set by the administrator — `discovered_direct_iosz` tunable. This will provide a substantial performance boost in environments with a wide variation in IO sizes — for example, when sharing a file system with a database, setting this tunable parameter will enable VxFS to determine whether to use directio or not.

Most of online tuning parameters are automatically set. However, environments with set up for hardware RAID, where VxVM and VxFS are unable to probe the internal details of the LUN configuration, the parameters must be changed manually.

### **Non-VERITAS File System Conversion Utility**

To minimize the overhead of migrating from ext2 or ext3 to VxFS, VERITAS provides a file system conversion utility. This program saves a substantial amount of time when large amounts of file system data must be migrated to the new VxFS installation.

### **Upgrade from previous version**

The simple upgrade path from Ext2 is the most significant benefit to deploying the Ext3 file system. With only minor differences between the two products, and no on-disk format changes, administrators are able to perform a minimally invasive online upgrade.

VERITAS VxFS 3.4 is online upgradeable to VxFS 4.0, minimizing interruptions for the user. A reboot is required to install the new software and reload kernel drivers. Because VxFS is a



module, and not compiled into the kernel like some version of Ext3, it is an easier and safer upgrade when compared to Ext3.

## **CONCLUSION**

VERITAS VxFS 4.0 offers performance that scales with the load on the system. As Linux servers find their way into more and more mission-critical computing environments, predictable performance, regardless of system load, is becoming an essential requirement. VxFS offers the bullet-proof, enterprise-class features and performance demanded by today's Linux system administrators and users.