XFS™ for IRIX™

sgi™

The World's Most Scalable Journaling-Filesystem Technology, Ready Today for the 21st Century

XFS: Scalable, Reliable, Fast

XFS combines advanced journaling technology with full 64-bit addressing, scalable structures, and scalable algorithms. This combination delivers the most scalable, reliable, and high-performance filesystem in the world—ready now for the 21st century. Its combination of scalability, performance, quick recovery, and guaranteed rate I/O (GRIO) makes XFS well suited for servers. This includes web servers, video servers, graphics servers, large and fast data servers, high-availability (HA) servers, high-performance computing (HPC), and other demanding filesystem applications.

21st Century Scalability

XFS provides a graceful path for managing disk growth for decades to come. Multiple terabyte filesystems are common today. Should disk densities continue increasing at 70% per year, in 10 years, one-thousand terabyte disk systems would be routine. In 20 years one-million terabyte disk systems would be common. XFS is a full 64-bit filesystem, ready today to handle these projected **million-terabyte filesystems**.

2^{63}-1 ≈ 9 x 10^{18} = 9 million-terabytes = 9 exabytes

As the disk sizes grow, the structures and algorithms also need to scale. XFS is ready today with the designed-in technologies needed for this scalability.

Journaling: Quick Recovery & High Reliability

The XFS journaling technology guarantees filesystem consistency, **while restarting in less than a second** (after an unexpected interruption), regardless of the number of files it is managing. Traditional filesystems use slow filesystem checks after an interruption, which can take many hours to complete. XFS journaling avoids these lengthy filesystem checks.

XFS technology is field-proven—running reliably on hundreds of thousands of IRIX systems for many years.

Fast Transactions and Searches

XFS optimizes the speed of read and write transactions. Its journaling and other advanced structures and algorithms are tuned to log the transactions more rapidly than with general UNIX® filesystem technology.

XFS uses efficient table structures (b-trees) for fast searches and rapid space allocation. XFS continues to deliver rapid response times, even for directories with tens of thousands of entries. The performance of most UNIX filesystems significantly degrades as the number of entries per directory grows—not so with XFS.

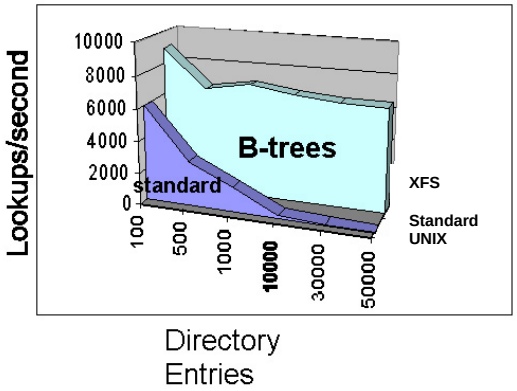

XFS Bandwidth

XFS delivers near-raw I/O performance. This was demonstrated on the largest disk configuration SGI was able to assemble: **over 4 gigabytes per second** (read and write) on 704 disks to one multithreaded process on an Origin™ 2000 IRIX® system. XFS delivers excellent bandwidth on production filesystems by routinely and automatically allocating **large contiguous extents**, avoiding fragmentation.

Guaranteed Rate I/O (GRIO)

XFS supports guaranteed bandwidths for applications, such as video servers, which need guaranteed rate I/O (GRIO).

Coming Soon to XFS:

Open Source XFS on Linux™: SGI is contributing XFS (one of its core competencies) to boost Linux to an enterprise-preferred solution.

Clustered XFS (CXFS): Allows multiple IRIX, Linux, and NT hosts to share a common set of disks, with hundreds or thousands of MB/s shared-clustered bandwidth.

XFS, the most scalable and reliable file system, is ready today to manage the rapid data growth expected over the next few decades.

Datasheet

XFS features

- Scalability which will support disk growth for decades to come:
  - Scalable file sizes (9 million terabytes)
  - Scalable filesystems (18 million terabytes)
  - Scalable algorithms for high performance, even with huge file systems
    - Large numbers of files
    - Large files (including sparse files)
    - Large directories
    - Advanced algorithms for fast performance on huge filesystems

- Journaling and Reliability
  - Guarantees filesystem consistency for high reliability
  - Rapid restarts (less than a second) after unexpected interrupts
  - Proven technology: Hundreds of thousands of systems in the field have use XFS reliably for years

- High performance
  - Extremely fast transaction rates
  - Extremely high bandwidths
  - Extremely fast directory searches
  - Extremely fast space allocation

- Advanced features
  - Video streaming support
    - Guaranteed Rate IO (GRIO)
  - Hierarchical storage (HSMs, DMF)
    - Off-line storage appears as if it is on-line

- Compatibility
  - Backup with popular commercial packages such as Legato NetWorker® for IRIX and Veritas NetBackup™ or with dump/restore, bru, cpio, or tar.
  - Support for multiple HSMs, including SGI DMF™ and Veritas™ HSM through the DMIG-DMAPI interface
  - NFS™ Compatibility
    With NFS version 3, 64-bit filesystems can be exported to other systems that support the NFS V3 protocol. Systems that use NFS V2 protocol may access XFS filesystems within the 32-bit limit imposed by the protocol.
  - Windows NT® Compatibility
    SGI uses the Open Source Samba server to export XFS filesystems to Windows® and Windows NT systems. Samba speaks the SMB (Server Message Block) and CIFS (Common Internet File System) protocols.

Technical Specifications

Technology
Journaled 64-bit filesystem for IRIX (and soon Linux) with guaranteed filesystem consistency.

Product Span
IRIX 5.3 or later (32 MB minimum memory).

Filesystem Block Size
512 bytes to 64 KB for normal data and up to 1 MB for real-time data.
Filesystem extents (contiguous data) are configurable up to 4 GB in size.
512 byte physical disk sectors supported.

Swap to Files
Swap to files is supported.

Guaranteed Rate I/O (GRIO)
Hard and soft guarantees are supported. Hard guarantees require turning off disk-drive self diagnostics. Guarantees are expressed as a file descriptor, data rate, duration, and start time. GRIO for more than four streams is optional.

Performance
Throughput in excess of 7 GB per second has been demonstrated on a single filesystem using a 32-processor Origin2000 server.
Single file reads and writes exceeded 4 GB per second.

Dump Interchangeability
Restore can restore EFS dumps to either EFS or XFS filesystems.
xfs_restore can restore XFS dumps to either XFS or EFS filesystems.
Dumps of active XFS filesystems are supported.

SGI corporate office: www.sgi.com (650) 960-1980
1600 Amphitheater Parkway Mountain View, CA 94403

KentK (7/99)