

AN EMERGING ORANGEFS

THE NEXT EVOLUTION IN PARALLEL FILE SYSTEMS

OrangeFS is the scalable software storage solution for High Performance Computing and Big Data.

Data in a parallel file system is distributed, residing across multiple servers instead of just one. This distribution is invisible to the user (client), who only sees one coherent and consistent file system. Distributed file data ensures better scalability in performance and capacity.

GROWTH IS ONLY LIMITED BY SCALABILITY

The unique class of distributed, parallel file systems ensures scalable, virtually unlimited growth. Simply take the highest performing single storage system at your site and add more. Clients now seamlessly write data across all of them.

WHERE ORANGEFS STANDS OUT

In the case of OrangeFS, thanks to its unique architecture, the performance of additional systems translates into proportional increases in capacity and throughput. OrangeFS clients also leverage the existence of multiple storage servers in the file system. In fact, the order and number of storage servers to use when writing a file can be configured per directory or at file create time. This allows the client to stripe a single file across multiple storage servers. The metadata can also be distributed across multiple servers.

LARGE FILE SYSTEM ALTERNATIVES

Alternative mass storage file systems require a software or hardware solution to aggregate and present multiple disks as a single device to the operating system. But as storage demands grow, gains in performance become less attainable; data transfer are only as fast as the bus or back-plane of a single system, and the chassis of the server can only hold so many disks. I/O growth inevitably tops out, reaching the maximum performance of a single system.

WHAT CAN YOU USE ORANGEFS FOR?

HPC Parallel Jobs OrangeFS interfaces with MPI, allowing HPC parallel jobs to use the OrangeFS system seamlessly.

Big-Data Aggregate Storage For years, PVFS has excelled in handling extremely large data sets. OrangeFS is ideal for the growing demands of commercial big data, where the ability to capture, manage and process huge data sets within tolerable times has spawned an entirely new field in information management.

Guest Virtualization Guest VMs can use this shared, high-performance file system to achieve I/O similar to a dedicated system. Traditional NFS or CIFS file systems with this ability fall short of OrangeFS flexibility and performance.

Capacity and Throughput OrangeFS handles heavy medium-to-large-file I/O capacity without the costs and difficulty of local, directly attached, or other network file systems. Add flexibility when you need to use storage quickly or shuffle it among hosts or file systems. Capacity is simply a matter of scaling out—adding systems as growth dictates. More servers equal more throughput.

Supports Many Client Interfaces

- Kernel module (VFS); no kernel patches necessary
- Windows Client support
- MPI-IO interface
- Shared standard I/O library
- WebDAV and S3 support (Web Pack)
- FUSE support for Mac
- Direct Interface (bypasses Linux Kernel)
- Hadoop (via VFS)

OTHER THINGS ORANGEFS DOES WELL

- High performance video storage.
- Enables multiple clients shared access to a unified name space.
- Provides unique object-based file data transfer, allowing clients to work on objects and not handle underlying storage details, such as data blocks.
- Enables distribution of both file data and file system metadata to storage servers.
- Utilizes commodity hardware running Linux for the server, with access by various clients, including Windows.
- Empowers heterogeneous systems to achieve and exceed performance of traditional single host storage systems.
- Allows use of any Linux file system as underlying local storage on each connected server.

ANY VENTURE CAN PROFIT FROM ORANGEFS

Imagine the ability to conjure up as much disk space as you need to support any high performance application. With OrangeFS, you can leverage two Linux computers or a room of storage servers to build one contiguous single-name-space file system. *Today.*