LUSTRE™
BECAUSE INNOVATION MATTERS.

A High Performance Cluster File System

Lustre is an open-source, high performance storage architecture and scalable shared file system for clusters. In addition to small clusters, the Lustre file system can be used on large clusters containing thousands of nodes where high-speed I/O throughput is required. The Lustre file system is designed, developed, and maintained by Sun Microsystems and you can download the latest version at: sun.com/software/products/lustre.

The Lustre file system is currently available for Linux and provides a POSIX-compliant UNIX® platform interface.

High Performance Storage for High Performance Computing

Traditionally, much of the focus on high performance computing (HPC) has centered on CPU performance. However, as computing requirements have grown, HPC clusters demand increasingly higher rates of aggregate data throughput. Today's clusters feature larger numbers of nodes with increased compute speeds. Higher clock rates and operations per clock cycle create an increased demand for local data on each node. In addition, InfiniBand and other high-speed, low-latency interconnects increase the data throughput available to each node.

Traditional shared file systems, such as NFS, have not been able to scale to meet this growing demand for data throughput on large HPC clusters. Scalable cluster file systems that offer parallel data access to hundreds of nodes and petabytes of storage are needed to provide the high-data throughput required by large HPC applications in areas such as manufacturing, electronic design, and research.

The Lustre file system is used with 7 of the 10 largest HPC clusters in the world, which contain tens of thousands of client systems, petabytes of storage, and hundreds of gigabytes per second of I/O throughput. Many HPC sites use Lustre as a site-wide global file system, servicing dozens of clusters on an unprecedented scale.

Lustre Clusters

Lustre clusters contain three kinds of systems:

- File system clients, which can be used to access the file system
- Object storage servers (OSS), which provide file I/O service
- Metadata servers (MDS), which manage the names and directories in the file system

The storage attached to the servers is partitioned, optionally organized with logical volume management, and formatted as file systems. The OSS and MDS read, write, and modify data in the format imposed by these file systems. Each OSS can be responsible for multiple object storage targets, one for each volume, and I/O traffic is load-balanced against servers and targets. Depending on the server's hardware, an OSS typically serves between 2 and 25 targets, each of which can store up to 8 terabytes. The capacity of a Lustre file system is the sum of the capacities provided by the targets.

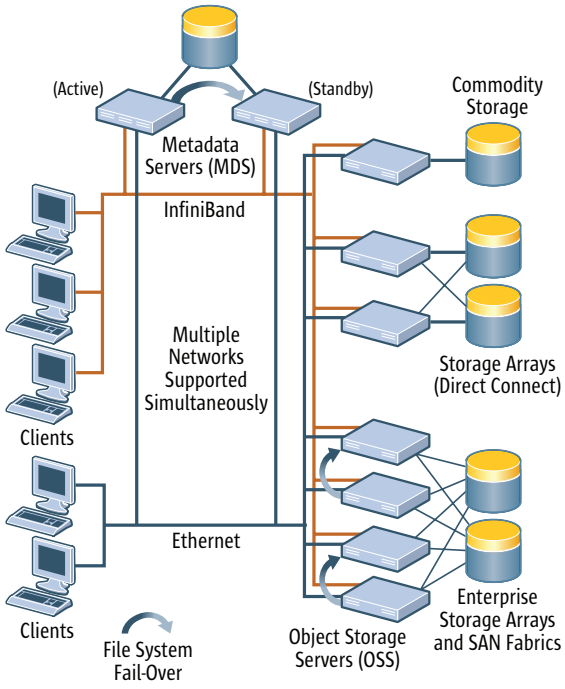

Featured Partners

Because Lustre is open-source software, it has been adopted by a number of other computing companies and integrated with their offerings. For example, RedHat and Novell offer kernels with Lustre patches for easy deployment.

Other Sun partners that use the Lustre file system as a component in their cluster offerings include Cray, Dell, EMC, HP, Hitachi Data Systems, LSI, and SGI.

Benefits to Students

Familiarity with Lustre technology provides students with a working knowledge of a file system used by a number of different companies that work with Linux and HPC. As the size and complexity of clusters increase, the need for people familiar with the Lustre file system will increase as well.

Learn More

For more information about the Lustre file system, visit sun.com/lustre, docs.sun.com, and lustre.org. Lustre documentation is available at manual.lustre.org and docs.sun.com/app/docs/doc/820-3681.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com