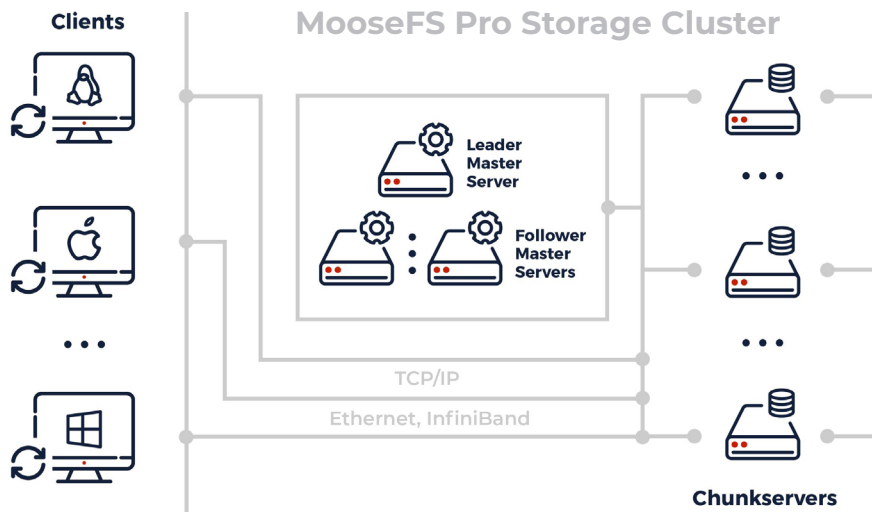


MOOSEFS BY TUXERA ARCHITECTURE & FEATURES

MooseFS by Tuxera is a fault-tolerant distributed parallel file system designed from the ground up for extremely demanding data workloads. It offers exceptional performance and scalability on commodity hardware—with no external dependencies.

How MooseFS by Tuxera works



MooseFS by Tuxera distributes data in parallel over a virtually unlimited number of servers that act as one disk to the client devices accessing it. A MooseFS storage cluster is comprised of Master and Chunkservers. The Master servers hold file metadata, while the Chunkservers redundantly store chunks of data. The Leader Master manages the Chunkservers, ensuring data is always available or recovered should any Chunkservers fail. In the event a Leader server fails, the Chunkservers will “elect” a new Leader from the remaining Follower Master servers. This is how the system self-heals and maintains data integrity.

Features and specifications



High performance

- **Store up to 16 exabytes** and more than 2 billion files on a single cluster
- **High throughput** across the entire storage cluster
- **Parallelism**—performs all I/O operations in parallel threads of executions to deliver high read-write operations performance
- **Compute on nodes** efficiently uses idle CPU, GPU and memory resources for low-latency data processing and storage on same machine
- **Direct drive mounts** means a significant performance boost over traditional RAID arrays—each chunkserver creates a mount-point for each drive, allowing MooseFS to intelligently decide how data is written to individual drives
- **Native clients and MooseFS (MFS) protocol**—dedicated client component for Linux, FreeBSD, Windows, and Mac OS X systems for max performance
- **Pre-fetch and read-ahead algorithms** enhance HDD disk cluster performance—improves stream-like cluster access patterns including storing and serving video files, logging, etc.



Features and specifications



High availability and reliability

- Runs as **single file system volume**, no matter the cluster size
- **SPOF-less configuration** (no single point of failure)
- **Metadata redundancy** with two or more raw copies on physically redundant master servers for rapid recovery
- **Distributed user data**—data divided into “chunks” and redundantly spread across storage servers
- **Redundancy in all system components**—master servers and chunkservers
- **Transparent automatic failover**—the system instantly initiates parallel data replication from redundant copies to other resources
- **Fast disk recovery**—as little as 15 minutes for a 1 TB drive
- **Cyclic redundancy checks**—each chunk of data in a cluster is followed by CRC to assure 100% consistency



Competitive TCO

- **Completely hardware- and OS-agnostic**—mix existing and new hardware for a cost-effective solution
- **Erasur coding** using error correction code algorithms (with up to 9 parity sums) ensures redundancy on less raw space
- **POSIX compliance**—meets IEEE-specified standards that define uniform APIs for Unix-like operating systems, meaning no need to replace existing architecture
- **Tiered storage**—move hot data to fast SSDs, cold data on cheaper hardware
- **Convenient licensing model** that includes software, all features, updates, and support



Easy deployment

- **Software-defined storage solution** that you can set up in less than 30 minutes
- Use **commodity hardware** or purchase **turnkey appliance** from Tuxera
- **On-the-fly upgrades**—one node at a time—means zero downtime
- **Unmatched deployment flexibility**—on-premises, cloud, hybrid, and geographically distributed options



Security and management

- **Granular quota limits and storage policies** enable full control over how data is stored and accessed: number of i-nodes, capacity on directory level
- **Access control** based on a standard Unix access control lists (ACLs) model for enhanced data security
- **Global trash** makes it easy to recover accidentally deleted data—configurable for each individual individual file and directory
- **Rich set of admin tools:** command-line, web-based, and SNMP-based interfaces
- **Advanced storage policies and tiering:**
 - Tier between unlimited number of storage pools like HPC and archive
 - Files automatically moved between storage tiers based on triggers set by admin
 - Move files on-the-fly between different types and levels of replication



Advanced features

- **Atomic snapshots** provide instantaneous, uninterrupted provisioning of the file system at any point in time—ideal for online backup solutions
- Access storage with any major industry-standard protocol for file access including **SMB, NFS, FTP, and HTTP**
- **Ethernet-based network (1, 10, 40, 100, or 200 GB/s)** used for all the communication with support for LACP configurations

Get in touch to start your evaluation of MooseFS by Tuxera: sales@tuxera.com