

The Leading Parallel File System

BeeGFS transparently spreads user data across multiple servers. By increasing the number of servers and disks in the system, you can simply scale performance and capacity of the file system to the level that you need, seamlessly from small clusters up to enterprise-class systems with thousands of nodes.

Distributed File Contents and Metadata

One of the most fundamental advantages of BeeGFS is the strict avoidance of architectural bottlenecks or locking situations in the cluster, through the user space architecture.

This concept allows to scale non-disruptive and linear on metadata & the storage level. BeeGFS is striping the file content on multiple storage nodes, plus it is distributing filesystem metadata (e.g., directory information) across multiple metadata servers. BeeGFS is releasing the full bandwidths of the any hardware components & delivers wire speed through any interconnect. Small, large systems (up exabyte) with metadata intensive applications, in general, can greatly profit from the latter feature without changes to the given workflows.

HPC Technologies

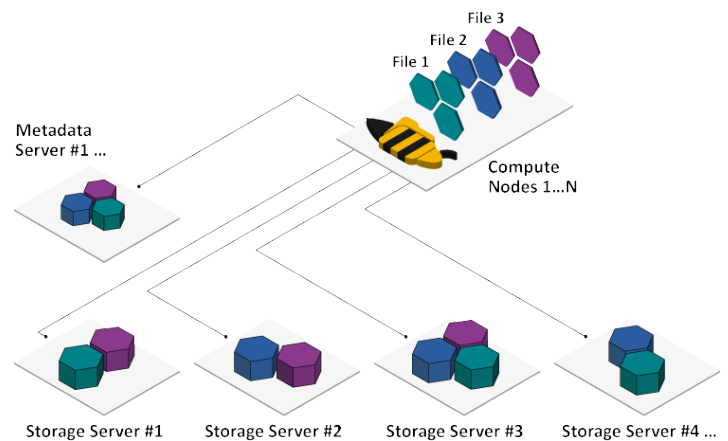
BeeGFS is built on highly efficient and scalable multithreaded core components with native RDMA support. File system nodes can serve RDMA (InfiniBand, (Omni-Path), RoCE and TCP/IP) network connections at the same time and automatically switch to a redundant connection path in case

Easy to Use

BeeGFS requires no kernel patches (the client is a patchless kernel module, the server components are user space daemons), it comes with graphical Grafana dashboards and allows to add more clients and servers to a productive system whenever required.

Optimized for Highly Concurrent Access

Simple remote file systems like NFS do not only have serious performance problems in case of highly concurrent access, they can even corrupt data when multiple clients write to the same shared file, which is a typical use-case for cluster applications. BeeGFS was specifically designed with such use-cases in mind to deliver optimal robustness and performance in situations of any high I/O loads or pattern.



Client and Server on any Machine

No specific enterprise Linux distribution or other special environment is required to run BeeGFS. It uses existing partitions, formatted with any of the standard Linux file systems, e.g., XFS, ext4 or ZFS, which allows different use cases.

1. BeeGFS client and servers can even run on the same machine as converged system to enable performance increases for small clusters or networks.
2. For larger networks (classical HPC ecosystems), it is possible to create several distinct BeeGFS filesystem partitions with different configurations.
3. To decouple, speed-up very IO intensive pattern & workloads (temp jobs, random IO 4k blocks, GPU pattern) from the persistent filesystem infrastructure it is possible to enable a BeeGFS filesystem instances (BeeOND) on the flight & use available SSDs/NVMe components on the client nodes for burst buffering.