EASY DATA MANAGEMENT FOR EVERYONE

HPC Solution Brief

BeeGFS is the award parallel file system delivering innovative technology with a complete ecosystem of tools for easy data management. Designed to be hardware agnostic and vendor neutral, BeeGFS offers a low entry barrier for an enterprise grade parallel file system empowering research institutions, supercomputing centers, and industry innovators to accelerate scientific discovery, simulation, and analytics by unlocking the full potential of their compute infrastructure.

BeeGFS®

Key Benefits

For HPC environments that demand scalable, low-latency, and high-throughput I/O, BeeGFS is an ideal solution. Its modular architecture supports independent scaling of metadata and storage servers, allowing seamless expansion from small clusters to multi-petabyte systems.

BeeGFS delivers outstanding performance for both large and small files by striping data across multiple storage nodes, making it especially suited for data-intensive workloads such EDA, weather forecasting, CFD, and life sciences.

Unlike traditional file systems, BeeGFS is easy to deploy and manage, enabling your team to unlock the full potential of your infrastructure, allowing them to focus on projects and not on your data management.

With state of the art functionalities including high availability, data protection, real-time performance, and monitoring, our innovative ecosystem of features will provide an easy data management workflow, for you.

High Performance

- Parallel I/O for massive throughput
- Low latency access to small and large files
- Optimized for HPC workloads

Robust

- High availability with BeeGFS on HA
- Fault-tolerant services using Pacemaker/Corosync
- Resilient under heavy metadata and I/O load

Scalable Architecture

- Scale metadata and storage nodes independently
- Seamless scale-out as data grows
- No single point of bottle neck

Ease of Use

- Fast deployment & management tools
- Monitoring via Grafana/ Prometheus
- Integration via Slurm, Kubernetes, Singularity

Data Management

- BeeOND for burst buffer on compute nodes
- Temporary, job-specific scratch space
- Fine-grained control with QoS and striping

Architecture Agnostic

- All kinds of Linux distributions
- Multiple network technologies
- Multitude of CPU/GPU architectures
- All types of storage technologies

Eliminating I/O Bottlenecks in HPC/AI Workloads

I/O Throughput & Bandwidth	Stripes data across multiple storage servers to boost throughput. Enables low-latency, high-speed transfers using RDMA and user-space I/O, minimizing CPU load.
Metadata Performance	Eliminates metadata bottlenecks by distributing operations across nodes—ideal for small file workloads.
Scalability	Add servers or disks to scale linearly. No downtime, No disruption.
Concurrency & Contention	Handles many clients accessing the same data simultaneously, with no single point of contention.
Latency	Bypasses kernel and enables direct memory transfers, reducing latency and CPU overhead.
Cost-Effectiveness & Management	Runs on standard Linux servers with HDDs, SSDs, or NVMe, flexible and cost-effective. BeeOND: Creates fast, temporary file systems on local storage for burst-buffering HPC jobs.

Choose BeeGFS if you value:

- Fast, scalable storage that is simple to deploy and manage
- A trusted technology used in leading HPC sites worldwide
- High-performance for both large and small I/O workloads
- The ability to scale flexibly without vendor lock-in

The BeeGFS Architecture

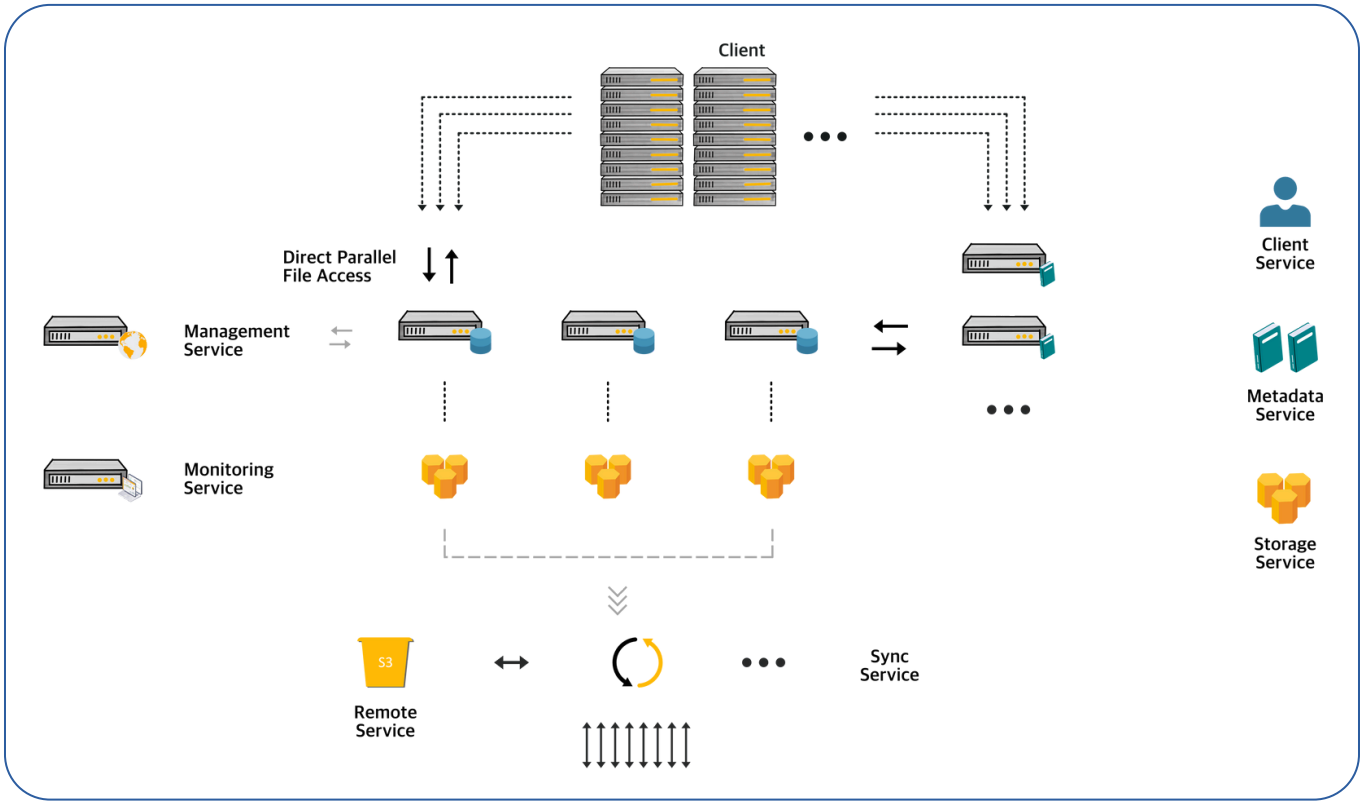

Explore the BeeGFS Hive of Possibilities

ThinkParQ offers BeeGFS as a self-supported Community Edition and a fully supported Enterprise Edition with additional features and functionalities.

Features	BeeGFS Community	BeeGFS Hive Enterprise
Scratch File System	✓	✓
Access Control Lists	✓	✓
Monitoring Tool	✓	✓
Community Mailing List	✓	✓
Commercial License Key		✓
High Availability		✓
Mirroring		✓
Quota Enforcement		✓
Storage Pools		✓
Remote Storage Targets		✓
Copy		✓
Watch		✓
Index		✓

www.beegfs.io