

# 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.



### 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 as 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, real-time performance, and monitoring, our innovative ecosystem of features will provide an easy data management workflow, for you.

|                                                                                                                                                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  <b>High Performance</b> <ul style="list-style-type: none"> <li>Parallel I/O for massive throughput</li> <li>Low latency access to small and large files</li> <li>Optimized for HPC workloads</li> </ul>        |  <b>Robust</b> <ul style="list-style-type: none"> <li>High availability with BeeGFS on HA</li> <li>Fault-tolerant services using Pacemaker/Corosync</li> <li>Resilient under heavy metadata and I/O load</li> </ul>   |  <b>Scalable Architecture</b> <ul style="list-style-type: none"> <li>Scale metadata and storage nodes independently</li> <li>Seamless scale-out as data grows</li> <li>No single point of bottle neck</li> </ul>                               |
|  <b>Ease of Use</b> <ul style="list-style-type: none"> <li>Fast deployment &amp; management tools</li> <li>Monitoring via Grafana/Prometheus</li> <li>Integration via Slurm, Kubernetes, Singularity</li> </ul> |  <b>Data Management</b> <ul style="list-style-type: none"> <li>BeeOND for burst buffer on compute nodes</li> <li>Temporary, job-specific scratch space</li> <li>Fine-grained control with QoS and striping</li> </ul> |  <b>Architecture Agnostic</b> <ul style="list-style-type: none"> <li>All kinds of Linux distributions</li> <li>Multiple network technologies</li> <li>Multitude of CPU/GPU architectures</li> <li>All types of storage technologies</li> </ul> |


## Eliminating I/O Bottlenecks in HPC/AI Workloads

|                                            |                                                                                                                                                                                      |
|--------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>I/O Throughput &amp; Bandwidth</b>      | 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.                      |
| <b>Metadata Performance</b>                | Eliminates metadata bottlenecks by distributing operations across nodes—ideal for small file workloads.                                                                              |
| <b>Scalability</b>                         | Add servers or disks to scale linearly. No downtime, No disruption.                                                                                                                  |
| <b>Concurrency &amp; Contention</b>        | Handles many clients accessing the same data simultaneously, with no single point of contention.                                                                                     |
| <b>Latency</b>                             | Bypasses kernel and enables direct memory transfers, reducing latency and CPU overhead.                                                                                              |
| <b>Cost-Effectiveness &amp; Management</b> | Runs on standard Linux servers with HDDs, SSDs, or NVMe, flexible and cost-effective.<br>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
- ✓ High-performance for both large and small I/O workloads
- ✓ A trusted technology used in leading HPC sites worldwide
- ✓ 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                  |                  | ✓                      |



**BeeGFS<sup>®</sup>**

**[www.beegfs.io](http://www.beegfs.io)**