BeeGFS

Made in Europe, deployed Globally

Davide Villa – VP Sales
Davide@beegfs.io

thinkparQ

Trusted by thousands of organizations worldwide

Why Choosing BeeGFS?

- **Performance:**
  Well Balanced from small to large files.

- **Scalability:**
  Scale capacity and performance from one server node to thousands

- **Ease of Use:**
  Easy to deploy and integrate with existing infrastructure.
  No HW lock-in

- **Robust:**
  High Availability design enabling continuous operations.

- **Data Management:**
  High Availability, Quota, Tiering between NVMe, HDD and Tapes, S3 offload, etc.

MADE IN EUROPE

Sovereign by design

Data Sovereign:

- ✓ You own **your data**, where it lives, how it moves, and who can access it. Always.
- ✓ **No lock-in**.

Technology Sovereign:

- ✓ **Deploy anywhere**: on-prem, private cloud, hybrid.
- ✓ Seamless compatibility with modern kernels and software ecosystems

Hardware Sovereign:

- ✓ **Your hardware**. Your choice.
- ✓ Any tier: **Flash, HDD, Tape**

Operations Sovereign:

- ✓ **Easy** to deploy, easy to use, easy to maintain,
- ✓ Easy to Scale

Compliance Sovereign:

- ✓ **Source-code available**.
- ✓ Developed in **Europe** and deployable entirely under your governance

BeeOND: On-Demand Storage for Every Job

BeeGFS On Demand

Dynamically pools local NVMe across compute nodes into your job, to build a shared on-demand parallel filesystem.

Spins up in under 60 seconds. Tears down when the job ends.

No new hardware.

No changes to global storage.

No disruption.

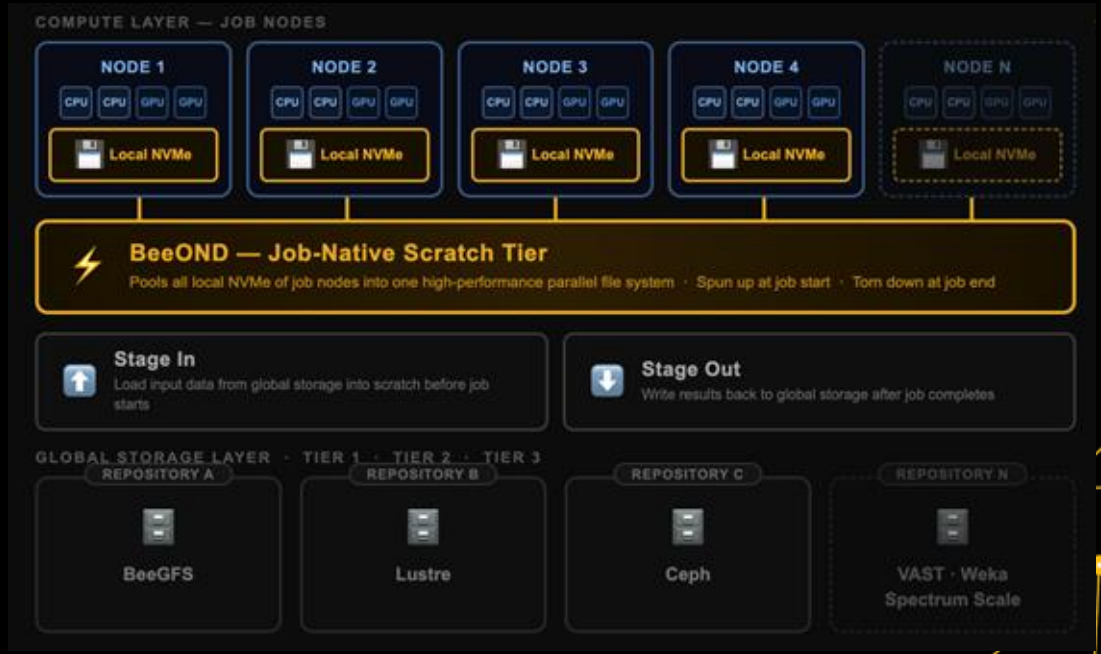

Why AI and HPC Teams Choose BeeOND

It's fast.
Aggregates raw bandwidth of every node's local drives. Linear scaling — more nodes means more performance. Proven at TB/s throughput.

It's simple.
Integrates directly with every scheduler. A constraint flag and a Prolog/Epilog hook. No agents, no management plane, no persistent services.

It's isolated.
Each job gets its own BeeOND instance. No noisy neighbors. No contention with the repository filesystem(s). Predictable performance, every time.

It's complementary.
Works alongside every repository filesystem — whatever you run today. Your existing investment stays intact.

Proven in Production. Not a Prototype.

TB/s

Throughput

World's first TB/s on-demand filesystem

10,000+

Compute Nodes

Operational at scales most products never reach

100s

Organizations

National labs to enterprise AI, running in production worldwide

<60s

Deployment

Job submission to operational parallel filesystem

thinkparq

www.thinkparq.com

www.beegfs.io