

Scalable and innovative high-performance HPC solution

BeeGFS and NetApp E-Series storage

Jean-François Le Fillâtre

Senior Systems Architect, ThinkParQ GmbH

CS3, January 2020



BeeGFS[®]



www.beegfs.io


January 2020

Why Use a Parallel File System?



Better
Performance

- Simultaneous Read/Writes
- Separate Meta Data from Data

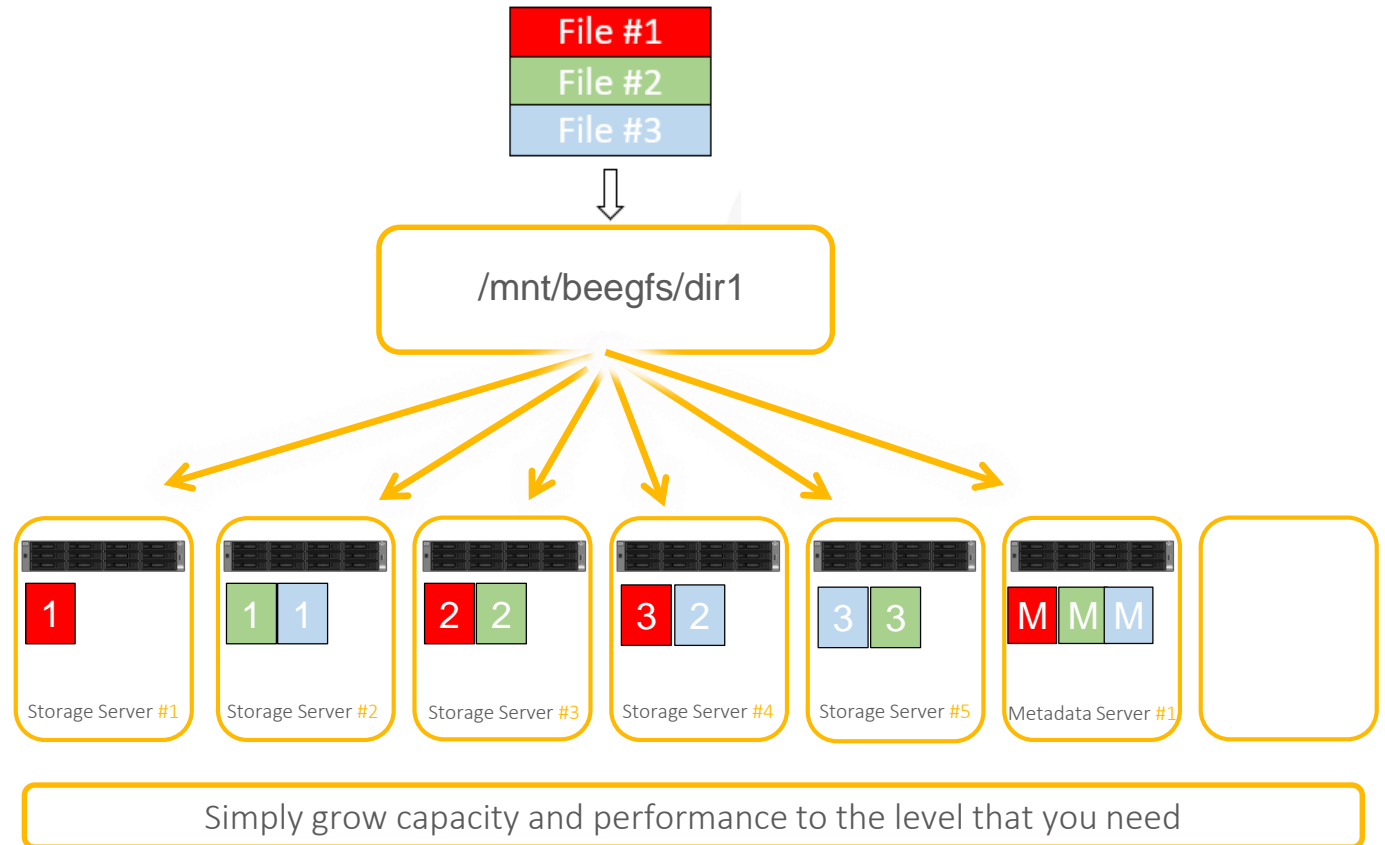


Scalable

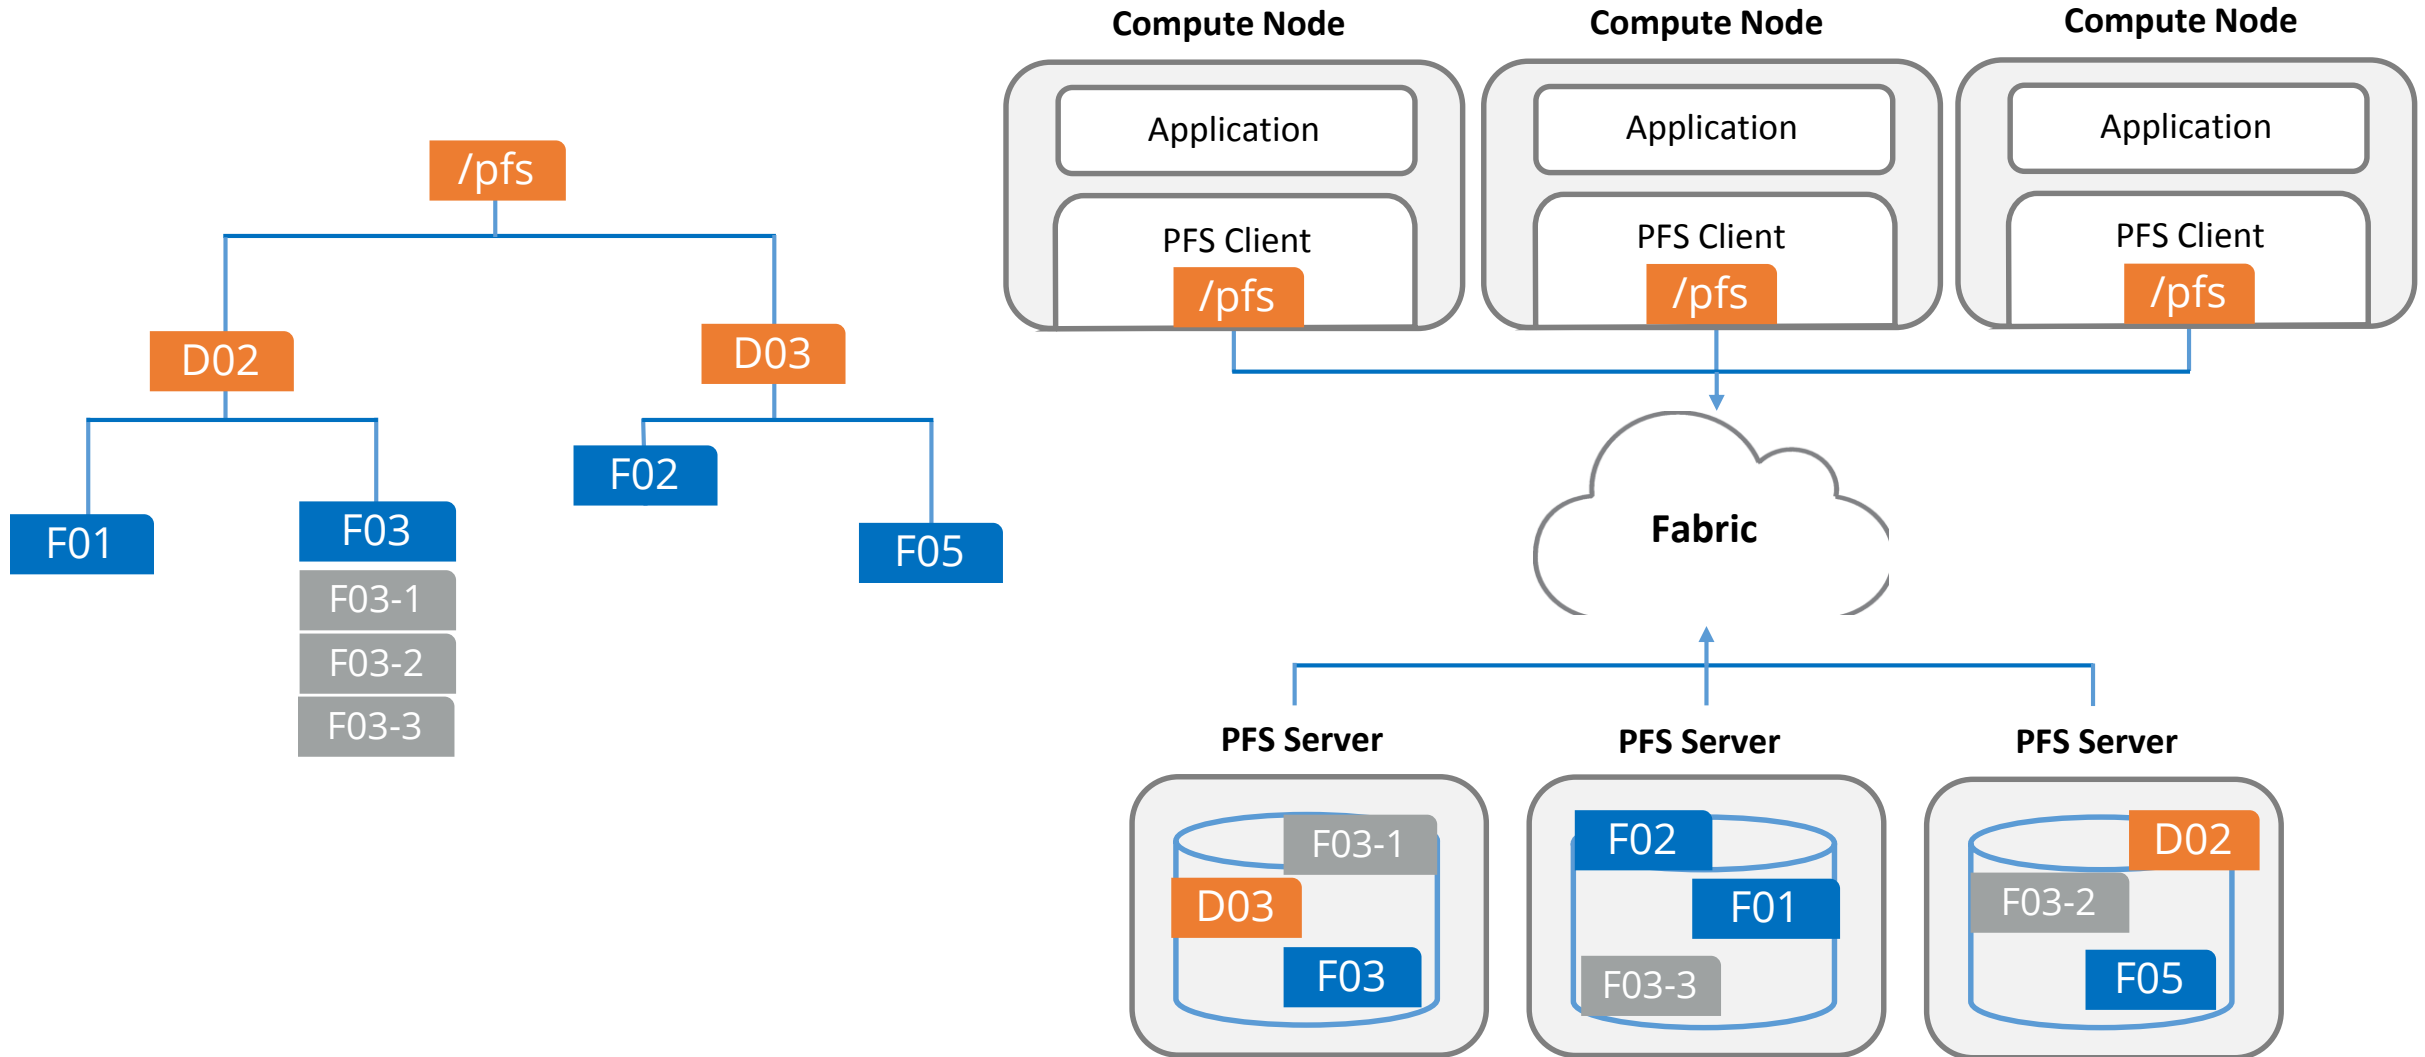
- Easy Scale-up and Scale-out
- Decouple Storage and Compute

Quick Facts: BeeGFS

- A hardware-independent parallel file system (aka Software-defined Parallel Storage)
 - Runs on various platforms: X86, ARM, OpenPower, AMD ...
 - Multiple networks (InfiniBand, OmniPath, Ethernet...)
- Available Source (Source is publicly available)
- Runs on various Linux distros: RHEL, SLES, Ubuntu...
- No Linux kernel patches
- Runs on top of EXT, XFS, ZFS, BTRFS, ..
- Scalable multithreaded architecture
- Supports RDMA / RoCE & TCP (InfiniBand, OmniPath, 200/100/40/10/1GbE, ...)
- All server-side components run in user space



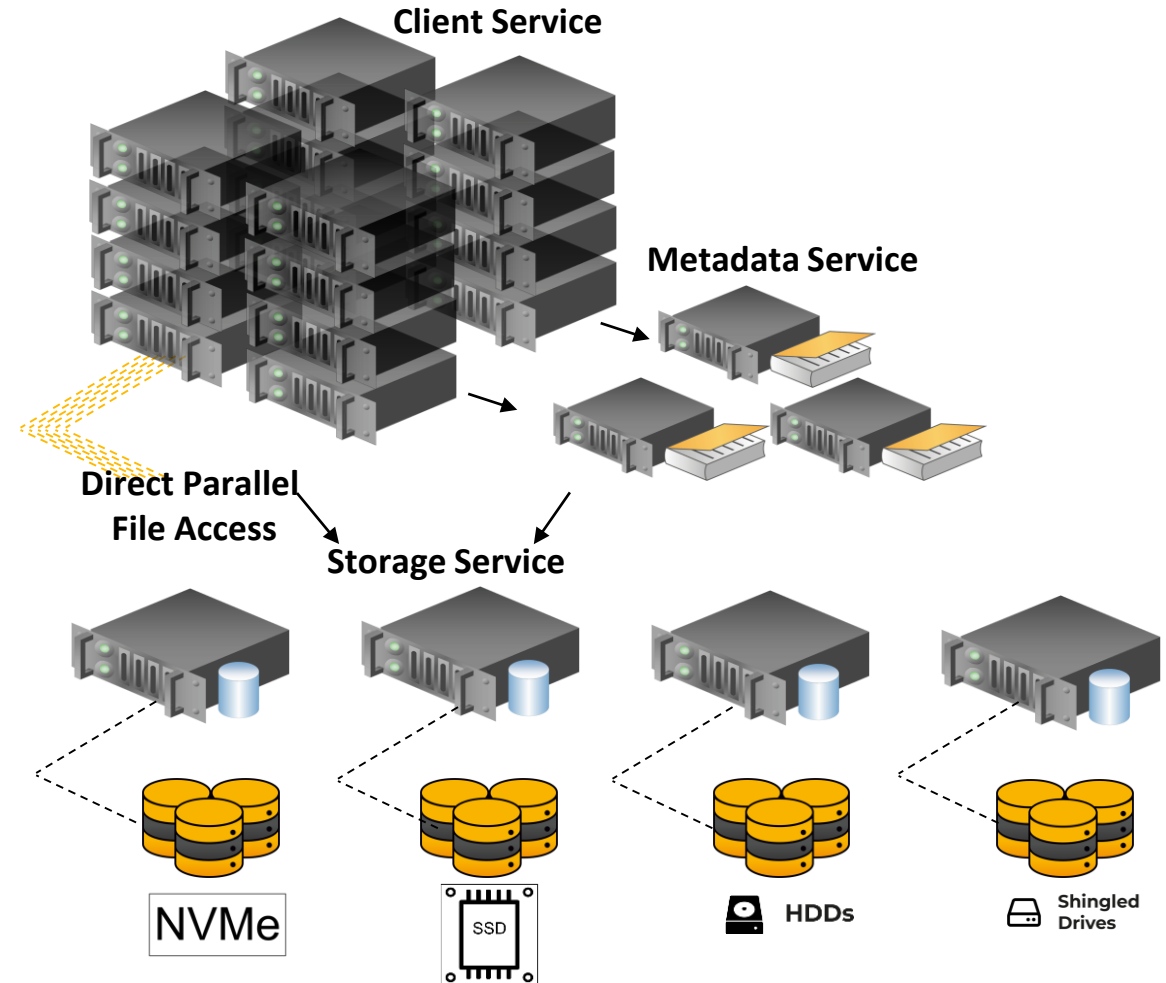
Parallel File System Architecture



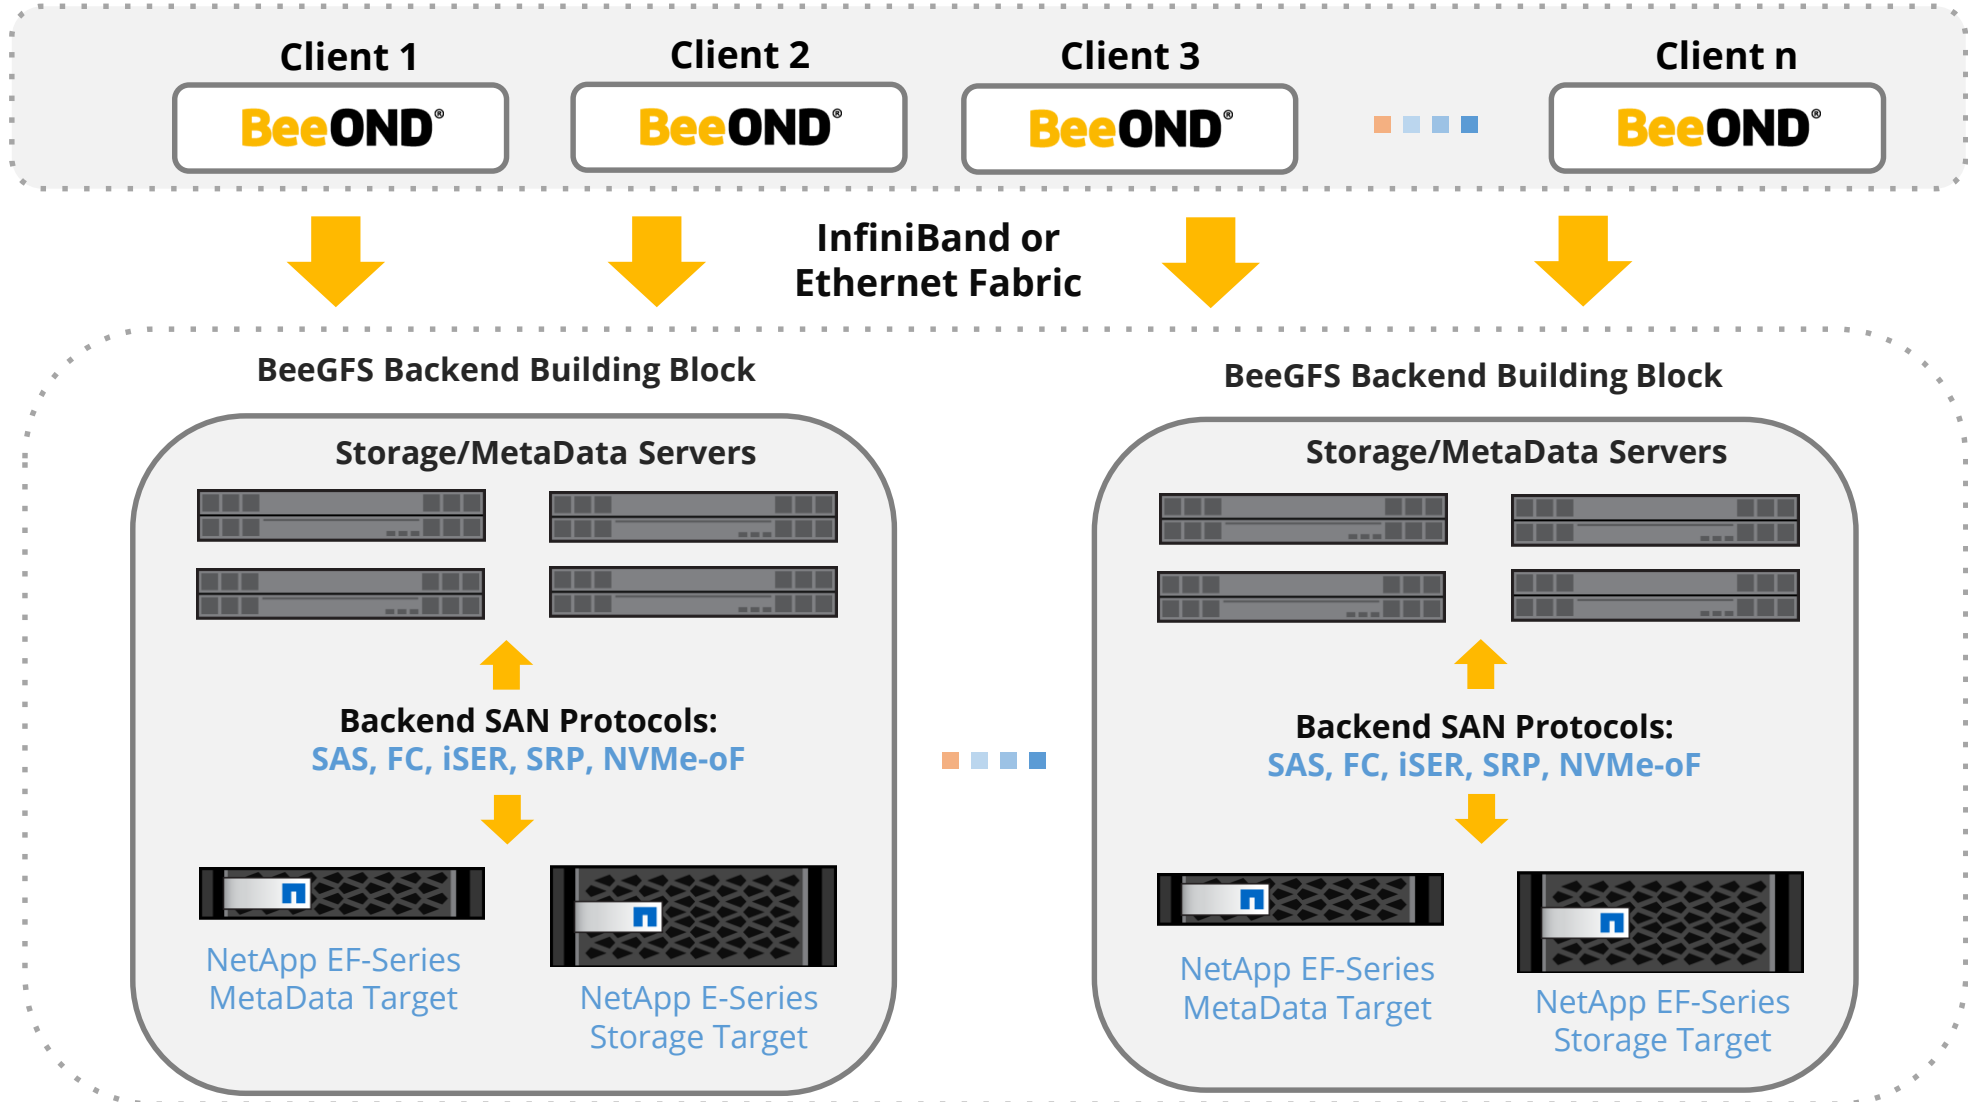
BeeGFS Architecture



- **Client Service**
 - Native Linux module to mount the file system
- **Management Service**
 - Service registry and watch dog
- **Metadata Service**
 - Maintain striping information for files
 - Not involved in data access between file open/close
- **Storage Service**
 - Store the (distributed) file contents
- **Graphical Monitoring Service**
 - WebUI based on Grafana with custom dashboards

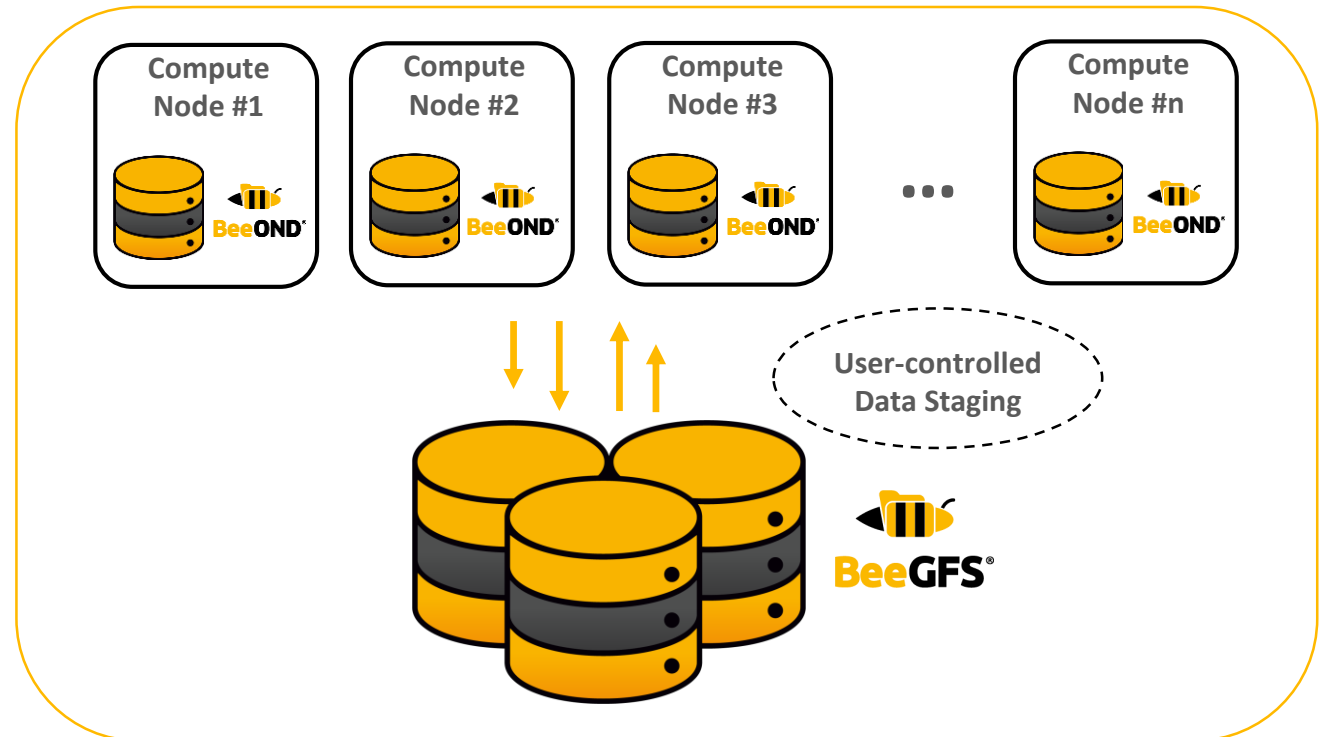


BeeGFS Cost Optimized Building Blocks on E-Series



BeeOND – BeeGFS On Demand

- Create a parallel file system instance on-the-fly
- Start/stop with one simple command
- Use cases: cloud computing, test systems, cluster compute nodes,
- Can be integrated in cluster batch system
- Common use case: per-job parallel file system
 - Aggregate the performance and capacity of local SSDs/disks in compute nodes of a job
 - Take load from global storage
 - Speed up "nasty" I/O patterns



simula



HPC SUCCESS STORY

Simula Research Lab
In Norway

- High-performance CPUs: AMD Epyc 7601, Cavium ThunderX2 CN9980 and Intel Scalable Platinum 8176M
- GPU nodes: DGX-2 system with 16 NVIDIA V100 processors internally connected via a 300 GB/s NVLink fabric
- Fabric: Mellanox HDR InfiniBand capable of 200 Gb/s
- Parallel file system: **BeeGFS**
- Parallel file system storage: 500TB of enterprise level hybrid storage (**E5760**) from **NetApp**
- Parallel file system scratch space: 85TB NVMe/SSD scratch space at the node level, globally shared using **BeeGFS's BeeOND** feature



More information at: <https://www.simula.no/news/simula-signs-contract-nextron-delivery-first-procurement-ex3-infrastructure>



Thank You

Follow

BeeGFS:





WE'RE HIRING!

Bee part of the Swarm.

Do you want to be part of a company that is a hive of activity?
An enthusiastic team powers us and we strive on the outstanding agility and understanding of customer requirements to always stay ahead. Be part of a team that creates and develops technology that empowers human development.

Open positions:

[Software Developer](#)

[Support Engineer](#)