BeeGFS

HPC User Forum Santa Fe

www.beegfs.io

2019

Frank Herold CEO

About ThinkParQ

- ThinkParQ strives to create and develop the fastest, most flexible and most stable, turn-key solutions for every performance-oriented environment.
- Established in 2014 as a spinoff from the Fraunhofer Center for High-Performance Computing, with a strong focus on R&D (70% of the team)
- 5 rankings in the top 20 on the IO-500 list.
- Awarded the HPCwire 2018 Best Storage Product or Technology Award

BeeGFS – The Leading Parallel Cluster File System

Performance
Well balanced from small to large files

Scalability
Increase file system performance and capacity, seamlessly and nondisruptively

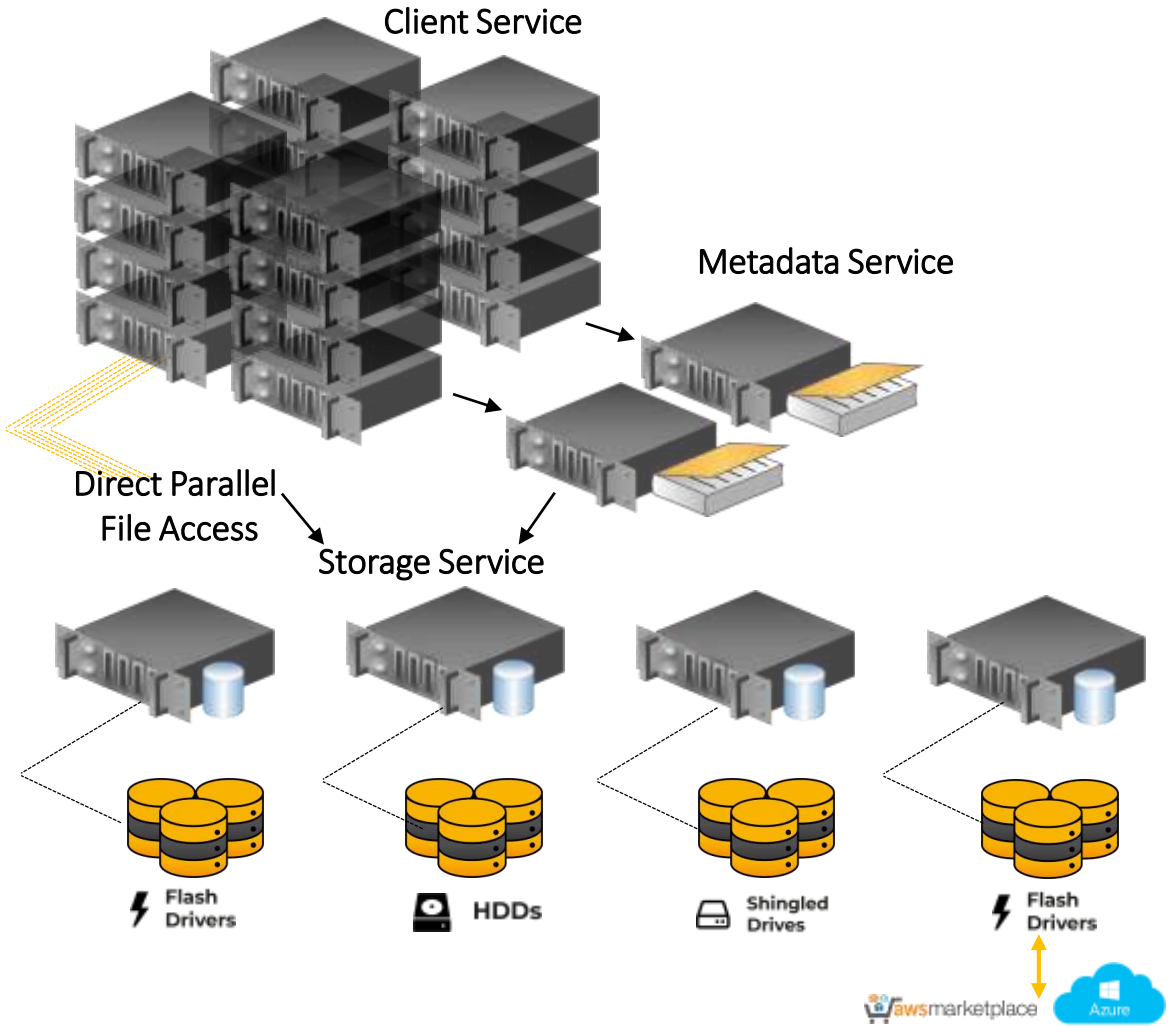

Ease of Use
Easy to deploy and integrate with existing infrastructure

Robust
High availability design enabling continuous operations

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

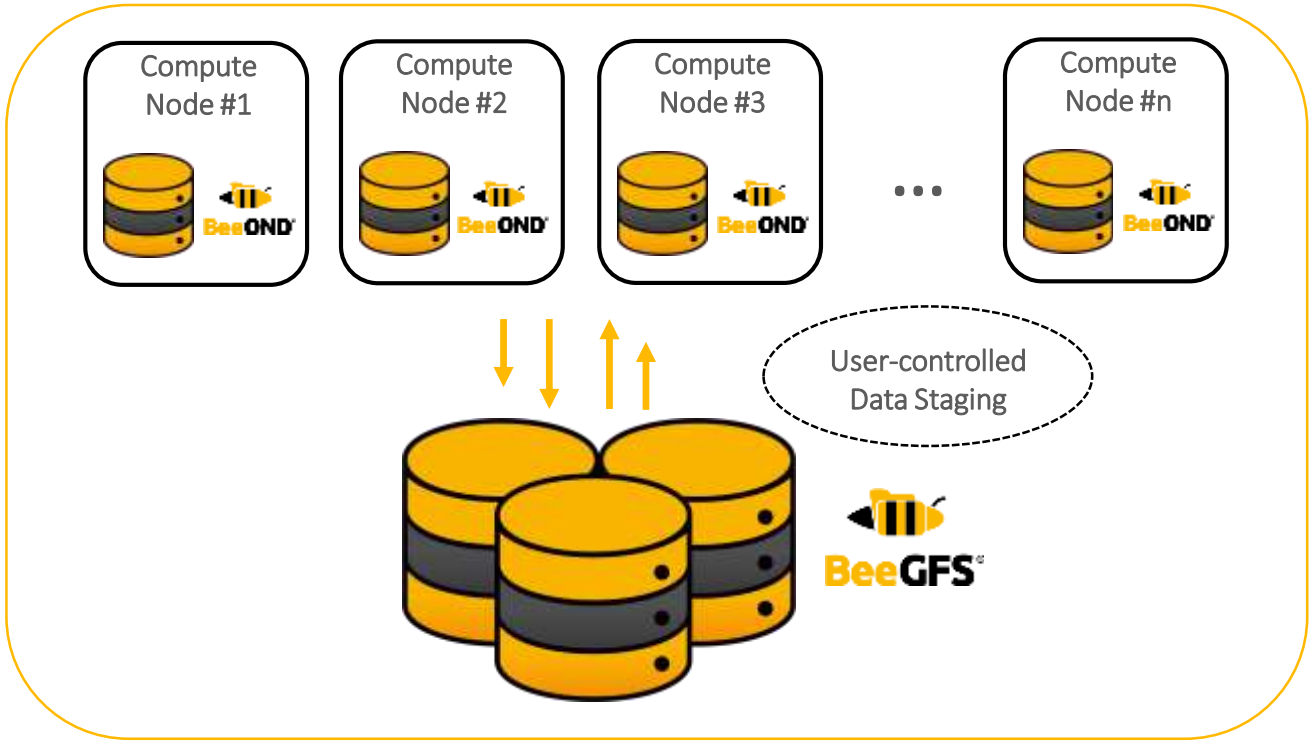

Quick Facts: BeeGFS

- A hardware-independent parallel file system (aka Software-defined Parallel Storage)
  - Runs on various platforms: X86, ARM, OpenPower, AMD …
  - Multiple networks (InfiniBand, OmniPath, Ethernet...)
- Open Source
- Runs on various Linux distros: RHEL, SLES, Ubuntu…
- NFS, CIFS, Hadoop enabled

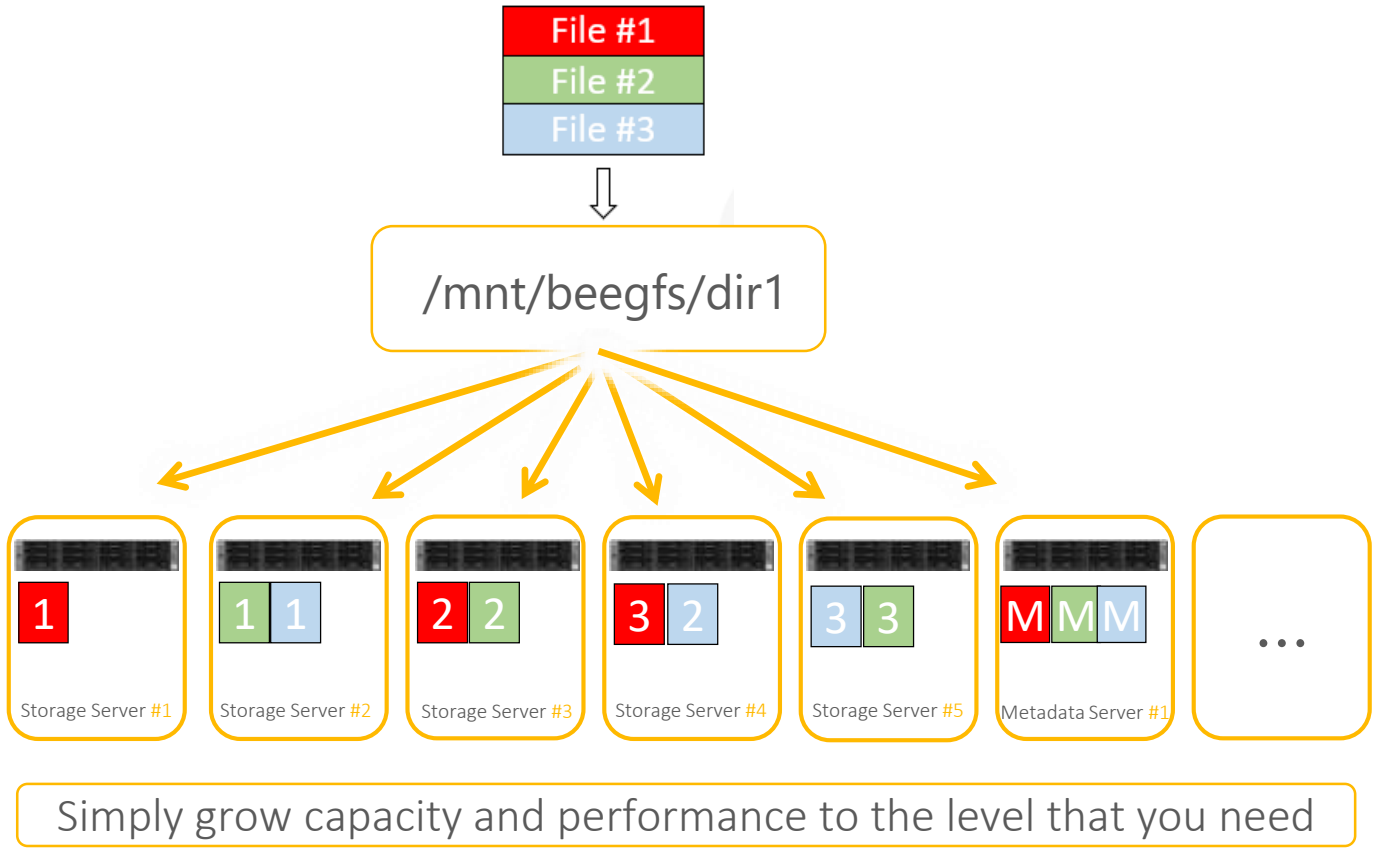

BeeGFS Use Cases

CSIRO

- The Commonwealth Scientific and Industrial Research Organisation (CSIRO) has adopted BeeGFS file system for their 2PB all NVMe storage in Australia, making it one of the largest NVMe storage systems in the world.

Overview:

- 4 x Metadata Server
- 32 x Storage Server
- 2 PiB usable capacity DELL all NVMe
- **Look forward to ISC to see what the beast can do!**
- [Further details: http://www.pacificteck.com/?p=437](http://www.pacificteck.com/?p=437)

Metadata
x 4

Storage
x 32

3.2 TB NVMe
x 24
per server

Alfred Wegener Institute for Polar and Marine Research

- Institute was founded in 1980 and is named after meteorologist, climatologist and geologist Alfred Wegener.
- Government funded
- Conducts research in the Arctic, in the Antarctic and in the high and mid latitude oceans
- Additional research topics are:
  - North Sea research
  - Marine biological monitoring
  - Technical marine developments
- **Actual mission:** In September 2019 the icebreaker Polarstern will drift through the Arctic Ocean for 1 year with 600 team members from 17 countries & use the data gathered to take climate and ecosystem research to the next level.

ALFRED-WEGENER-INSTITUT
HELMHOLTZ-ZENTRUM FÜR POLAR- UND MEERESFORSCHUNG

Day to day HPC operations @AWI

- CS400
  - 11,548 Cores
  - 316 Nodes:
    - 2x Intel Xeon Broadwell 18-Core CPUs
    - 64GB RAM (DDR4 2400MHz)
    - 400GB SSD
  - 4 fat compute nodes, as above, but 512GB RAM
  - 1 very fat node, 2x Intel Broadwell 14-Core CPUs, 1.5TB RAM
  - Intel Omnipath network
  - 1024TB fast parallel file system (BeeGFS)
  - 128TB home and software file system

ALFRED-WEGENER-INSTITUT
HELMHOLTZ-ZENTRUM FÜR POLAR- UND MEERESFORSCHUNG

Do you remember BeeOND?

- Global BeeGFS storage on spinning disks
  - 1PB of scratch fs providing 80GB/s
- 316 compute nodes
  - Each equipped with 400MB SSD each
- 316x500MB/s per SSD equals 150GB/s aggregate BeeOND burst “for free”

ALFRED-WEGENER-INSTITUT
HELMHOLTZ-ZENTRUM FÜR POLAR- UND MEERESFORSCHUNG

“Robust and stable, even in a case of unexpected power failure.“

Dr. Malte Thoma
Alfred Wegener Institute, Helmholtz Centre for Polar and Marine Research - (Bremerhaven, Germany)

BeeGFS®

Thank You

Follow BeeGFS: