



PacificTeck
HPC and Machine Learning Experts

HPC/Machine Learning Experts in APAC

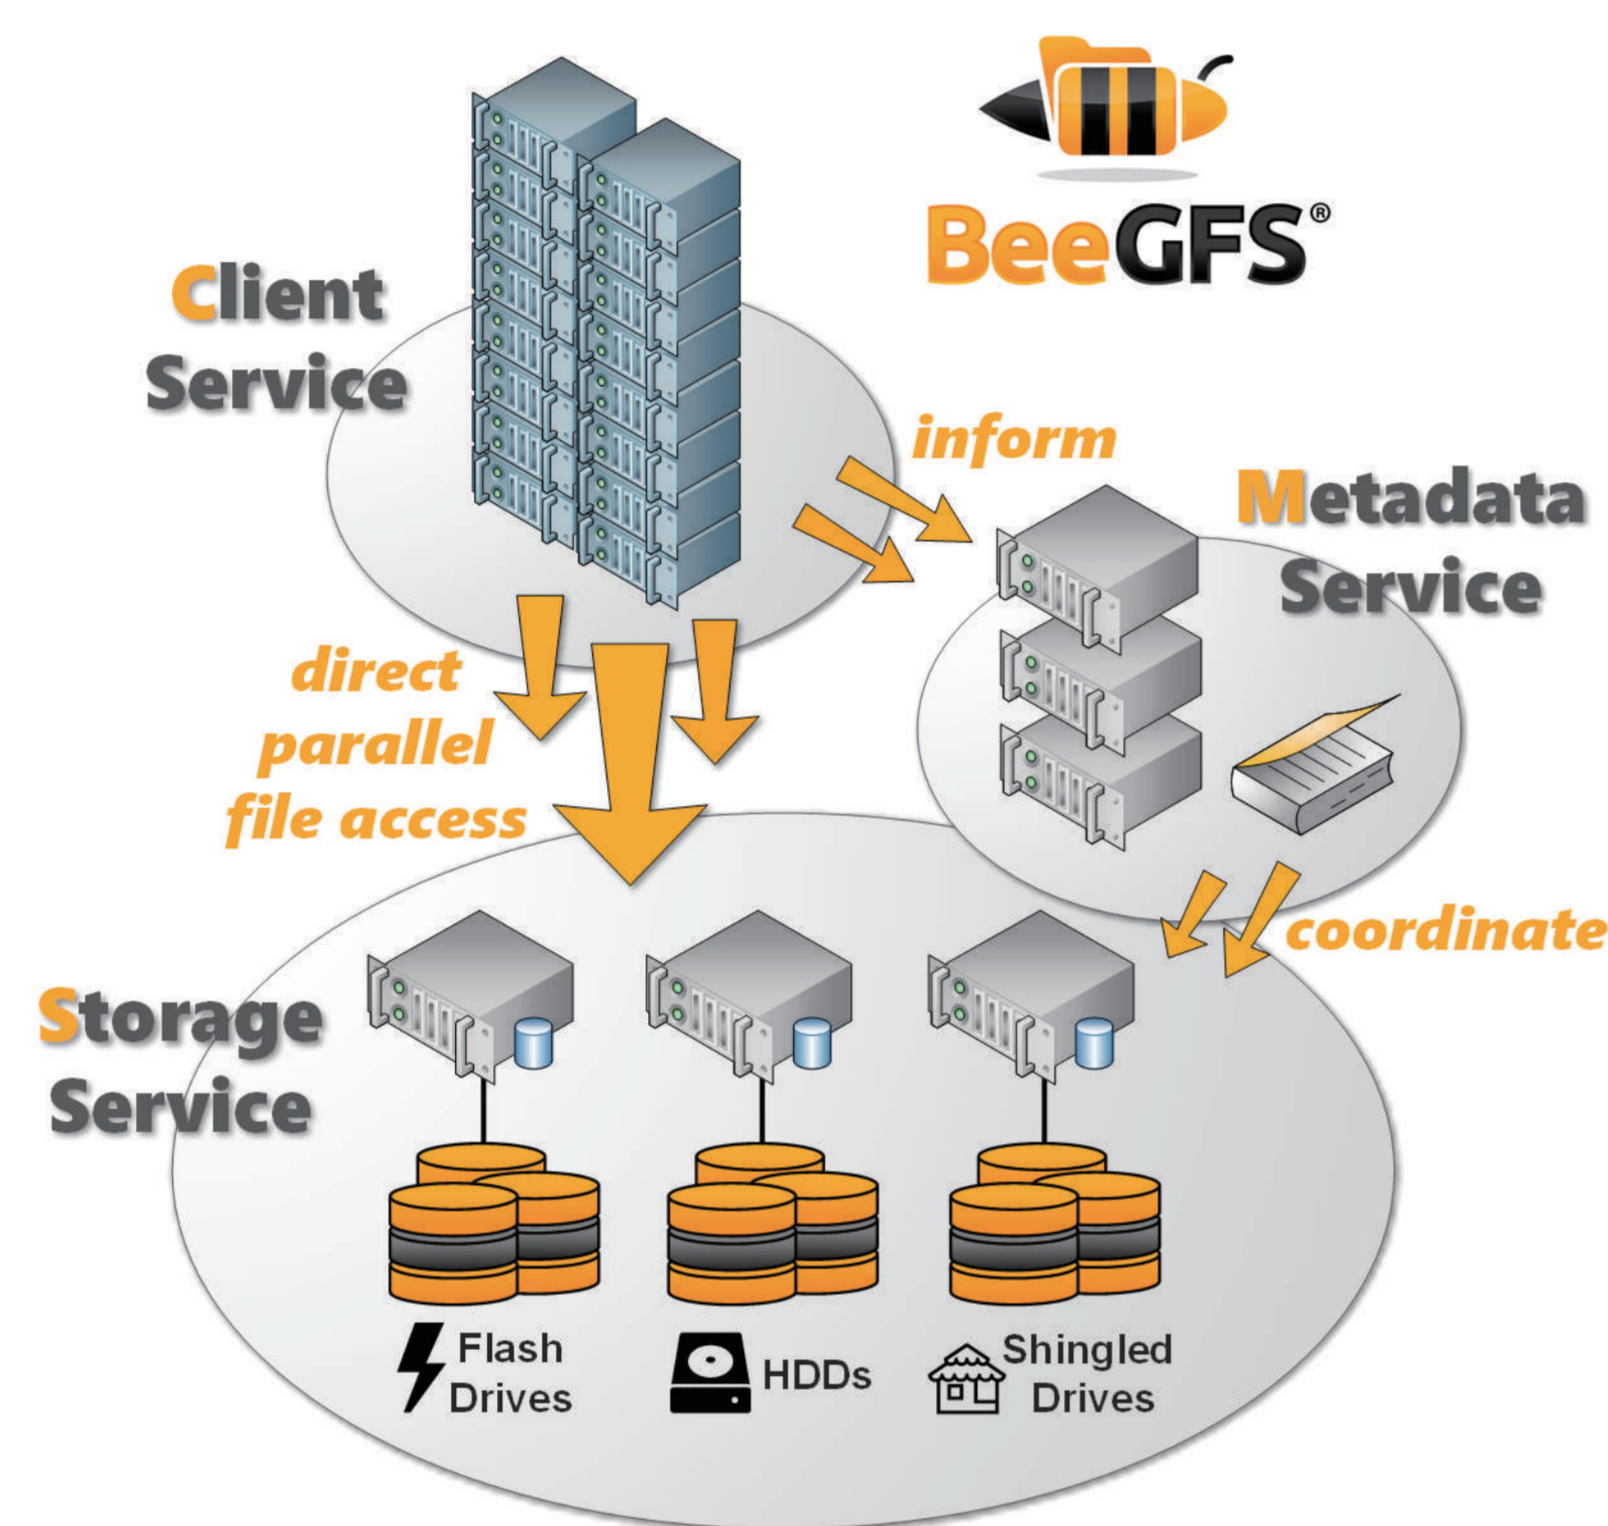
Pacific Teck brings cutting edge technology products from around the world to the Asia Pacific region. We focus on products related to building the best environments for enterprise computing, high performance computing and machine learning applications.

STORAGE



BeeGFS

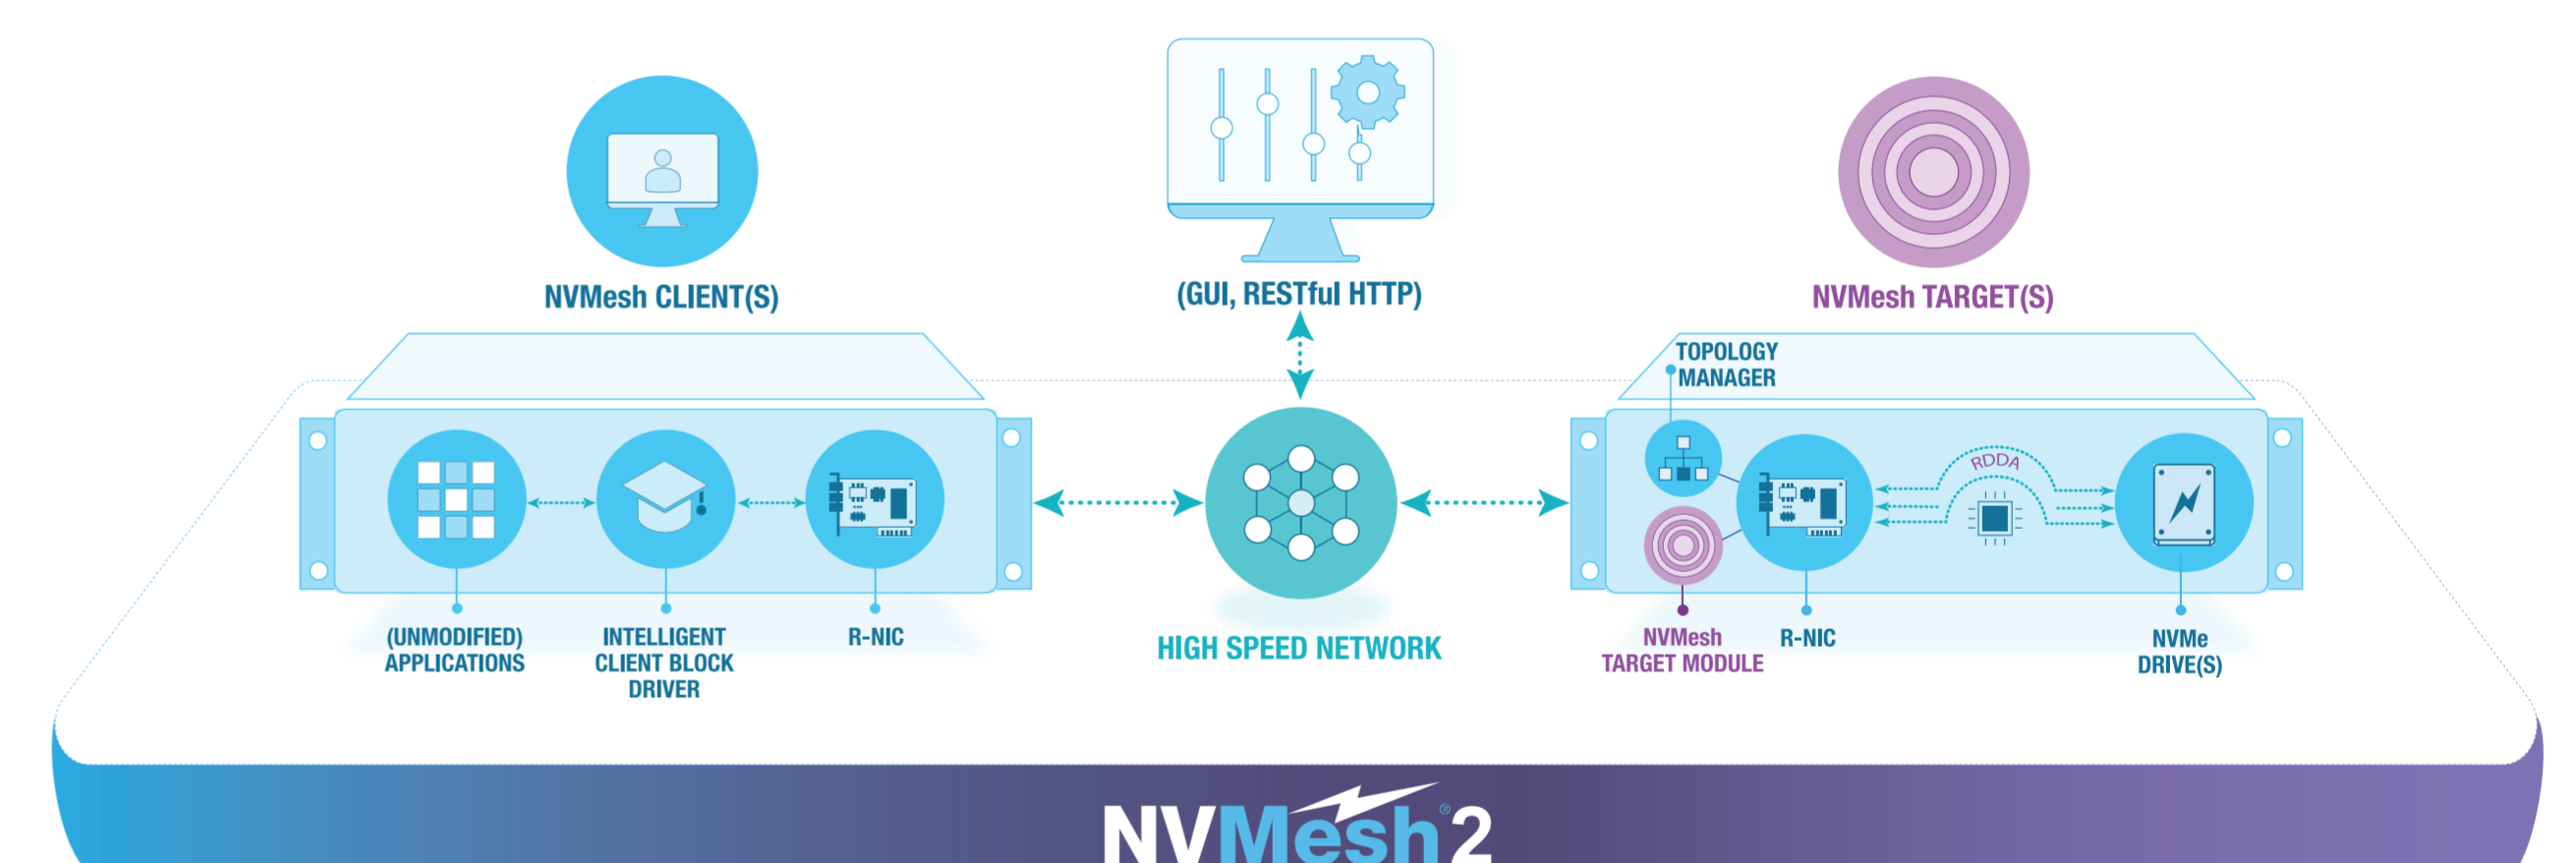
A modern, easy to install and manage parallel file system



NVMesh²

NVMesh 2

Provides high speed, low latency block level NVMe data protection



VIRTUALIZATION



SingularityPRO

Commercial version of Singularity, enabling jobs to run independent of host server OS

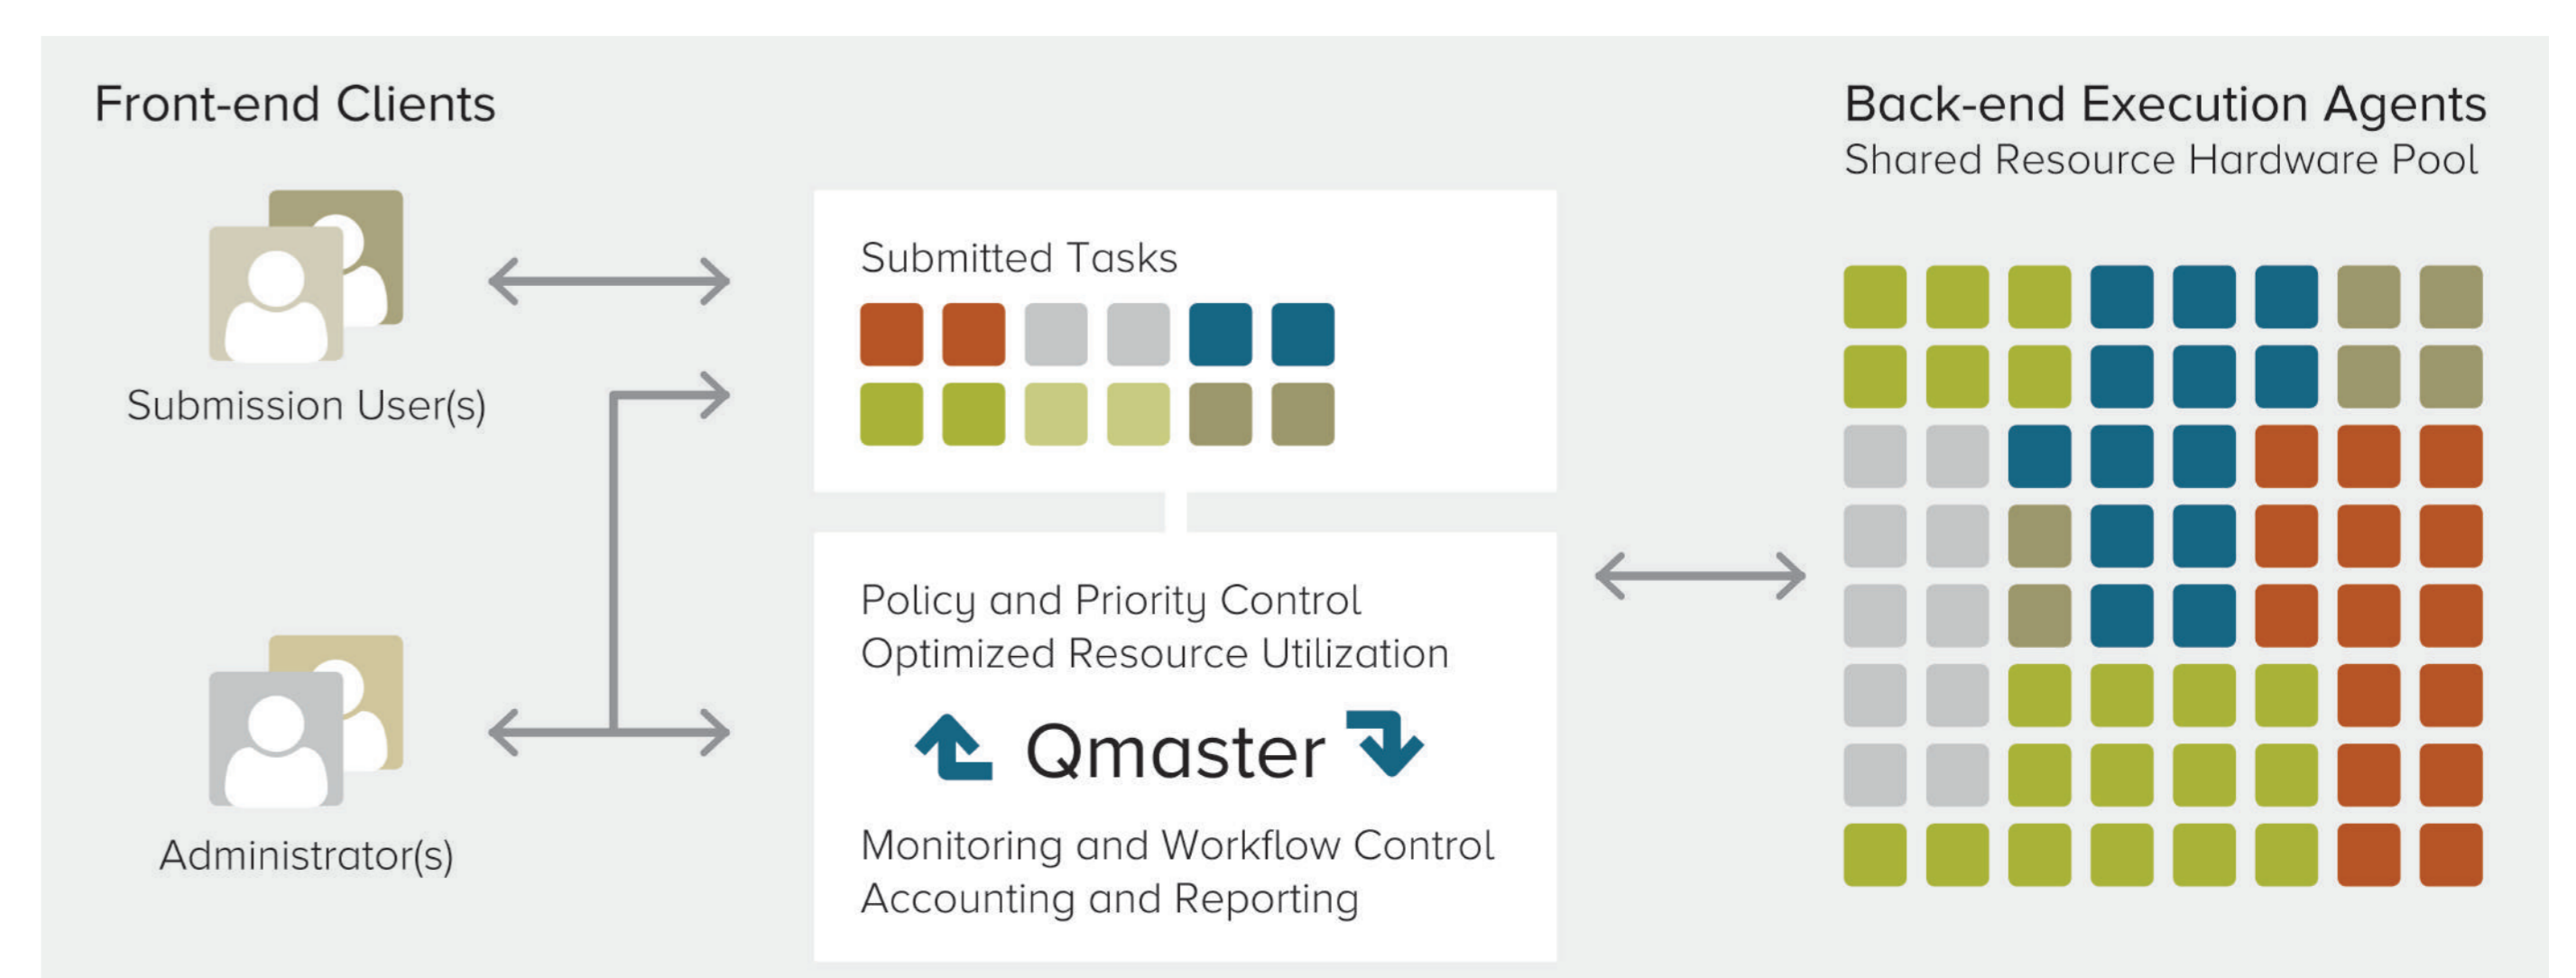
Features	Singularity Community	SingularityPRO
Supports Traditional HPC Resources	✓	✓
Resource Manager Agnostic	✓	✓
SIF: Single File Container Format	✓	✓
Cryptographically Signed and Verifiable SIF	✓	✓
Compatible with OCI Images	✓	✓
No Persistent Daemon Processes	✓	✓
Support for Non-Root (Rootless) Users Running Containers as Themselves	✓	✓
Blocks Privilege Escalation Within a Container	✓	✓
BOYE: Bring Your Own Environment Usage Model	✓	✓
Support for AI/HPC Workflows and Architectures	✓	✓
Supports GPU Natively	✓	✓
Self-Service Support Model (Open Source)	✓	✓
Code Curation		✓
Streamlined Security Updates		✓
Preferred Security Updates		✓
Signed Singularity RPM/DEB Packages		✓
Access to Sylabs Cloud KeyStore	Freemium	Premium
Access to Sylabs Cloud Container Library	Freemium	Premium
Access to Sylabs Cloud Remote Build Services	Freemium	Premium
Enterprise Support		✓
Enterprise Intergration		✓

JOB MANAGEMENT



Univa Grid Engine

Re-designed workload manager to enable multiple jobs to run on complex GPU systems



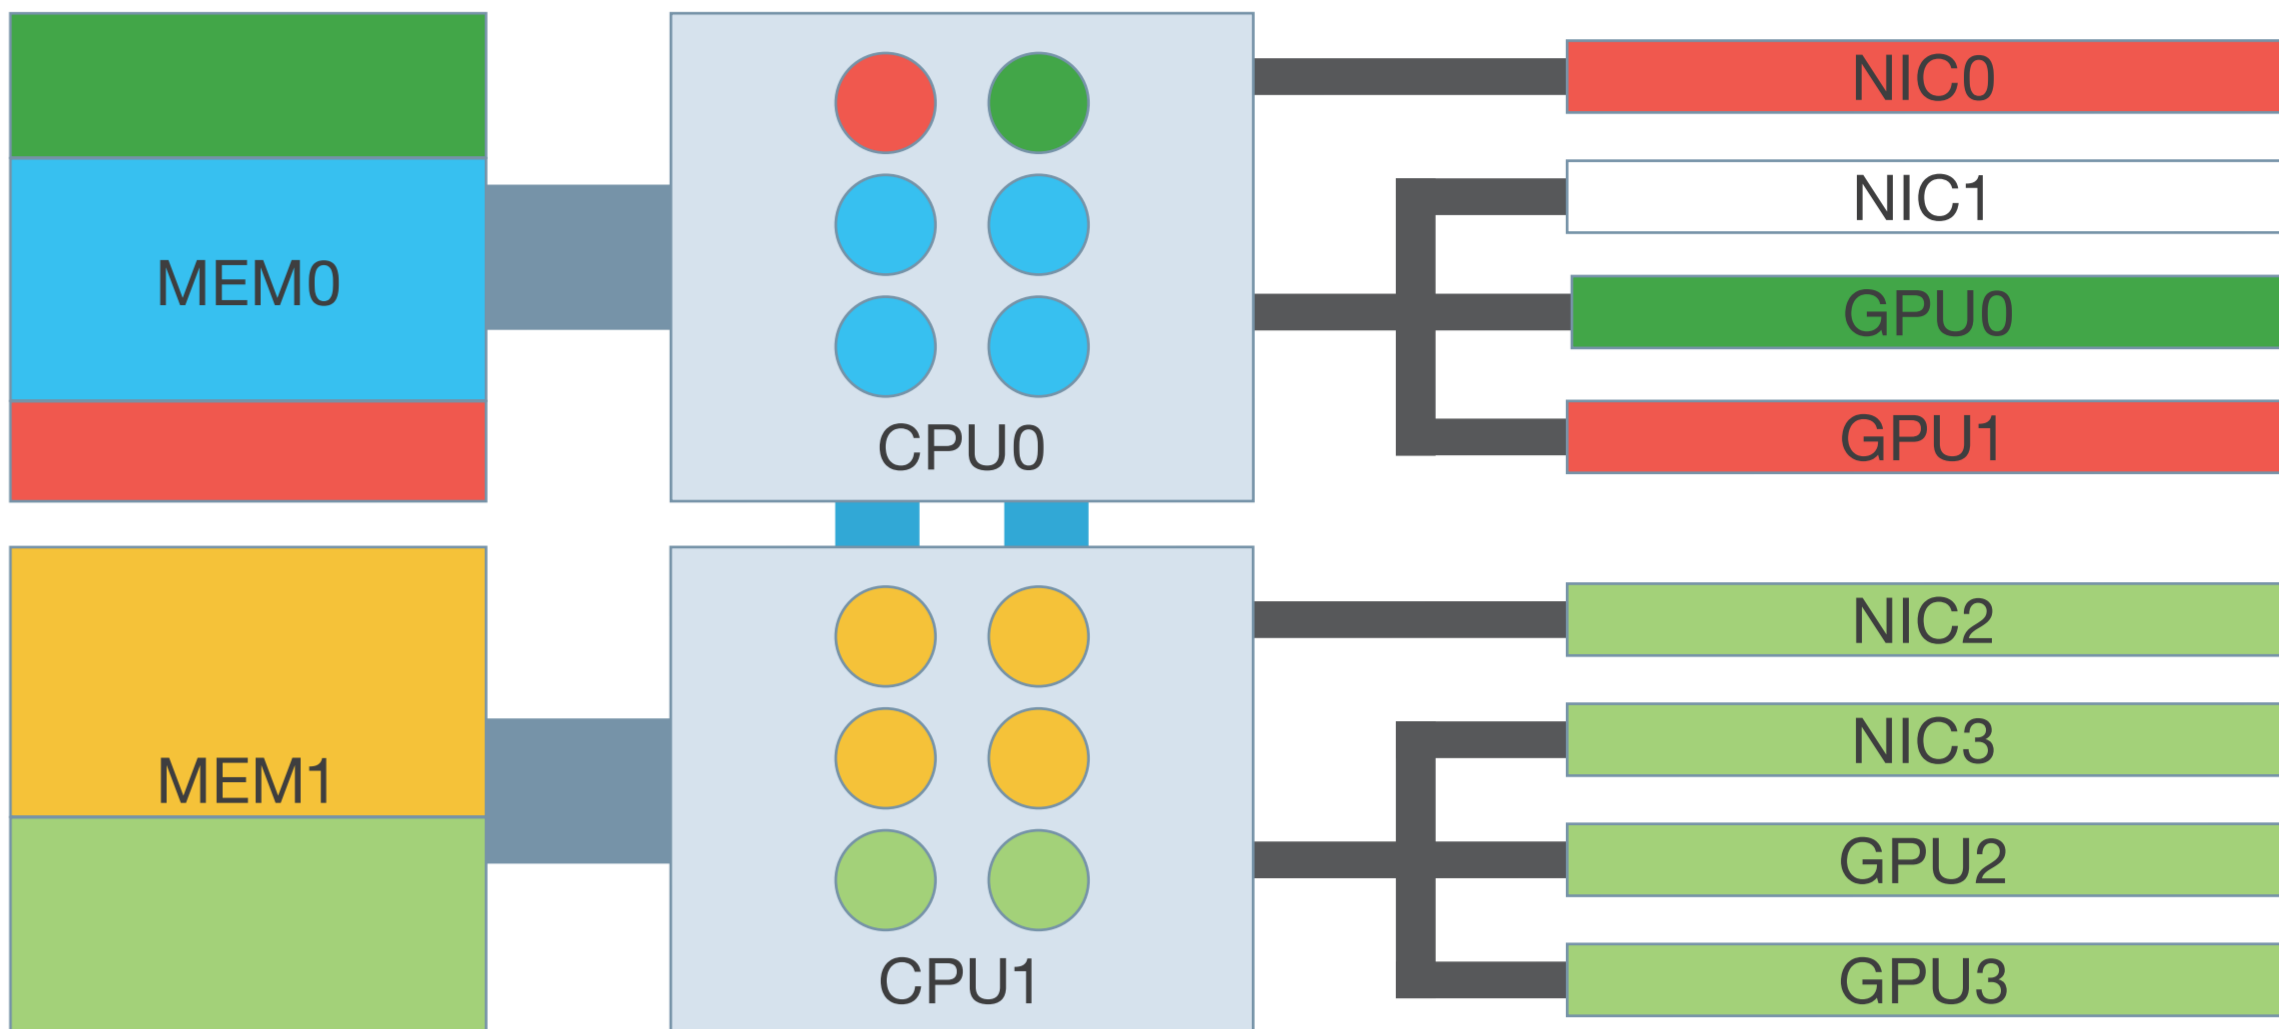
Pacific Teck References

TSUBAME 3.0 STORAGE / JOB MANAGEMENT



Tokyo Institute of Technology TSUBAME3.0 utilizes Univa Grid Engine and BeeOND in an environment with 540 nodes, each with four Nvidia Tesla P100 GPUs (2,160 total), two 14-core Intel Xeon Processor E5-2680 v4 (15,120 cores total), four Intel Omni-Path Architecture (Intel OPA) 100 Series host fabric adapters (2,160 ports total), and 2 TB of Intel SSD DC Product Family for NVMe storage devices.

TSUBAME 3.0 Container-Based Fine-grained Spatial Resource Allocations of Fat Nodes



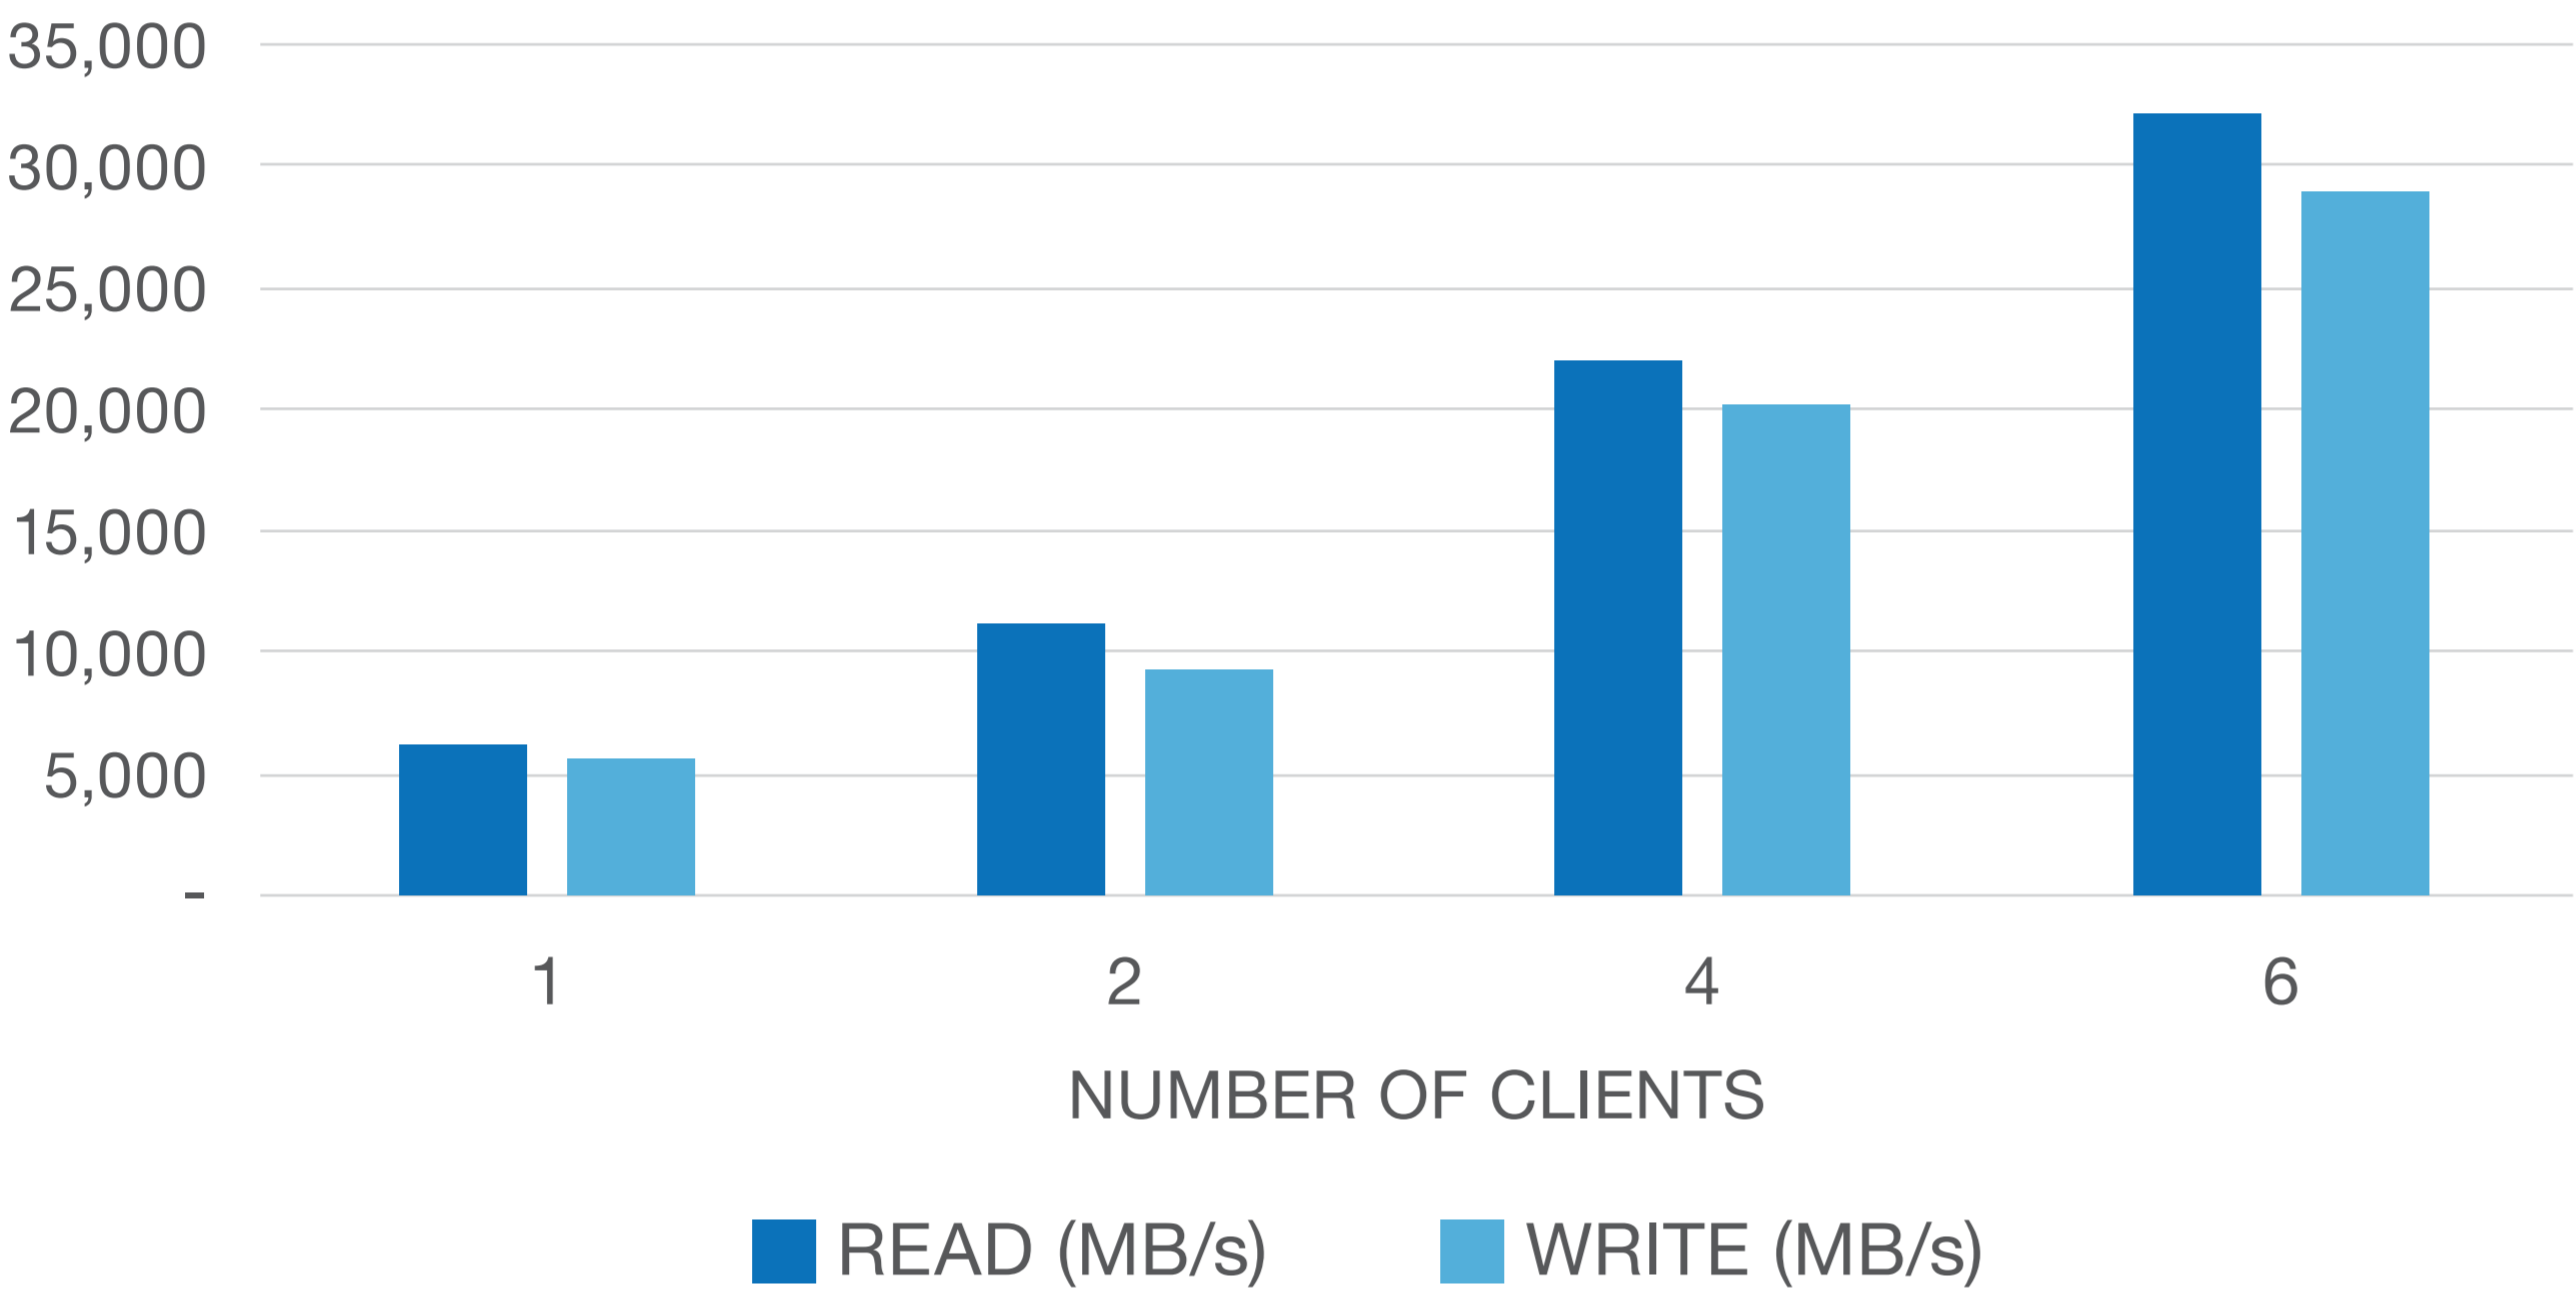
Resource Isolation via UGE Containers (future Docker etc.)

Job	Allocated Resource
1	CPU 2Cores, NIC0, GPU1, 32GB Mem
2	CPU 8Cores, 64GB Mem
3	CPU 4 Cores, GPU0, 16GB Mem
4	CPU 8 Cores, 64GB Mem
5	CPU 4 Cores, NIC2&3, GPU2&3, 48G Mem

Container configuration and deployment tied to Univa Grid Engine

A*STAR STORAGE / JOB MANAGEMENT

IOR RESULTS



The Agency for Science, Technology, and Research (A*STAR) Genome Institute of Singapore (GIS) has adopted the BeeGFS file system for their 288TB 15K SAS File System using the Dell PowerEdge R740. GIS is using InfiniBand FDR as the interconnect and 15K SAS drives. With BeeGFS A*STAR is able to fully utilize the bandwidth made available by the high speed InfiniBand interconnect.

CSIRO STORAGE

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. The storage hardware consists of 32 x Dell PowerEdge R740XD with 24 x 3.2TB NVMe per server. Using NVMe storage with the BeeGFS filesystem will remove their I/O bottleneck and provide better performance for projects such as “de-novo genomic sequence alignment and medical imaging analysis.”

Metadata: Dell PowerEdge R740XD



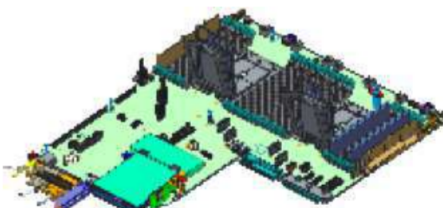
x 4

Storage: Dell PowerEdge R740XD



x 32

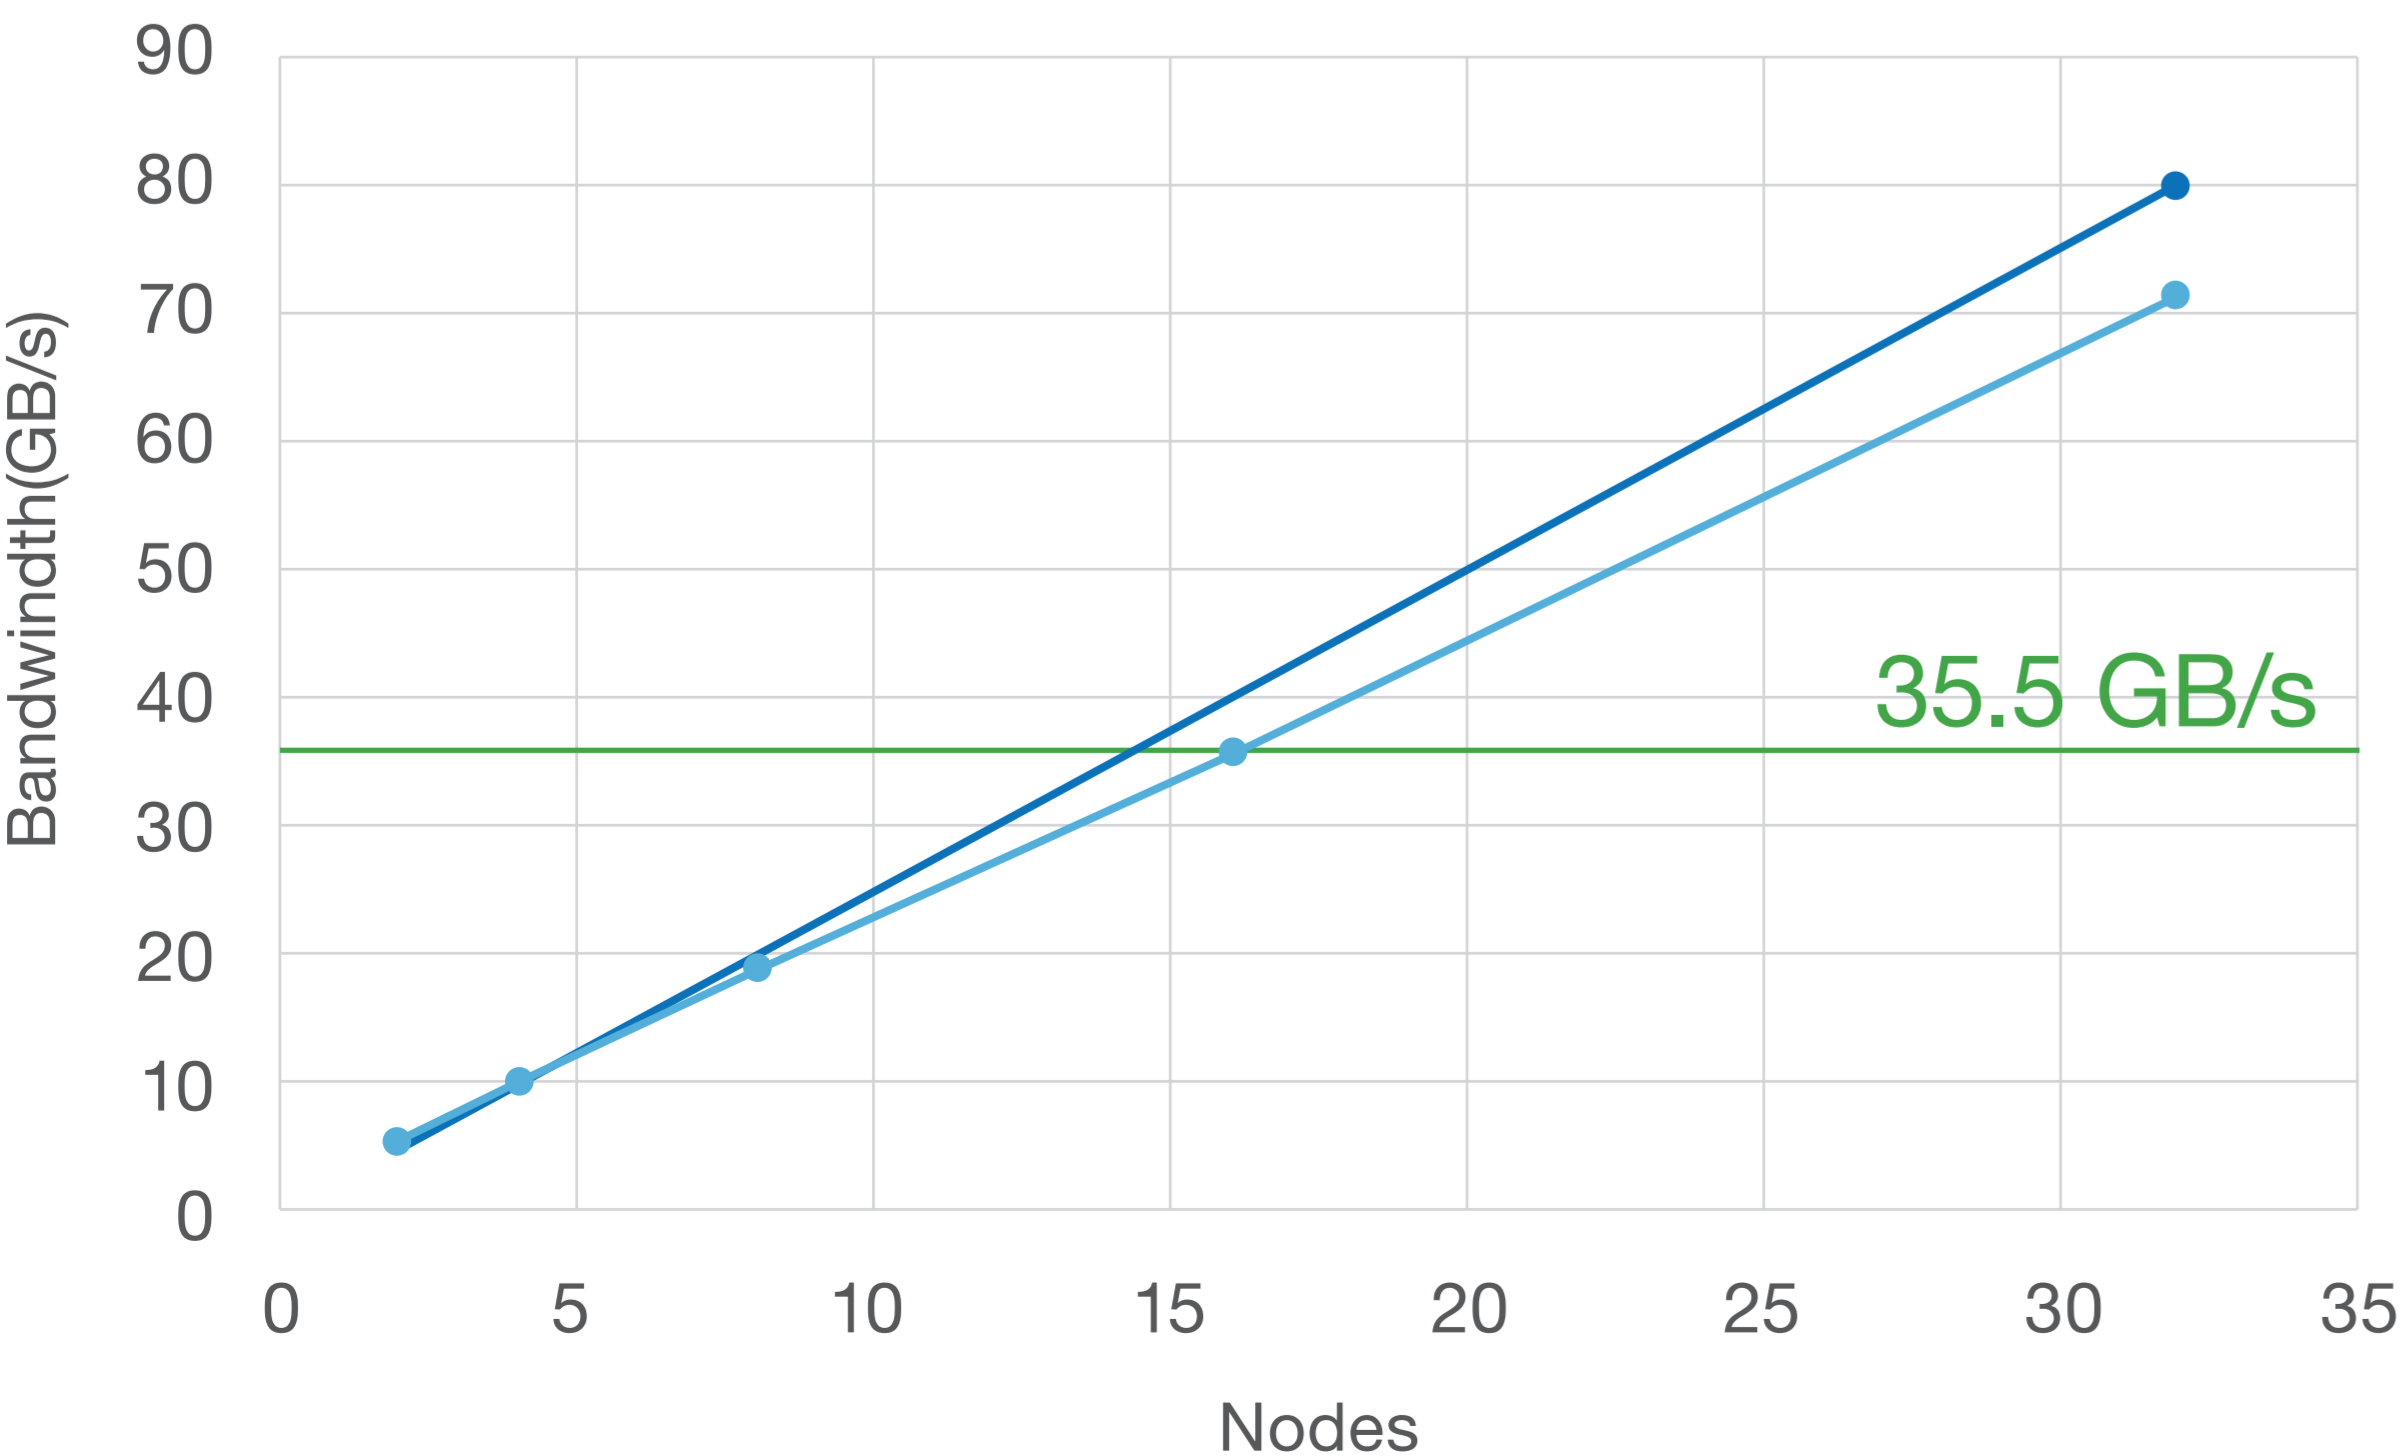
3.2 TB
NVME



x 24
per server

AIST STORAGE / VIRTUALIZATION / JOB MANAGEMENT

Sequential Read



Pacific Teck worked closely with the system integrator Fujitsu to understand and communicate the benefits of Univa Grid Engine and BeeOND with the end user. Thanks in part to Pacific Teck’s efforts, Univa Grid Engine and BeeOND were utilized in a similar project, Tsubame 3, and the end user had an interest in replicating those elements. The end user also had an interest in Singularity, for which Pacific Teck was able to make an introduction. Pacific Teck is also involved in post sales and acting as bridge between the vendors and system integrator. We also assisted with Univa Grid Engine managing Singularity containers to run MPI jobs in containers. BeeOND creates a temporary file system from NVMe in the compute nodes with a maximum capacity of about 1PB. The size utilized at a given time is determined when kicked off by Univa Grid Engine.