SuperNODE™

The Open Solution for Generative AI and Accelerated Computing

HIGHLIGHTS

BREATHTAKING PERFORMANCE
Lowest possible latency and highest effective bandwidth

EXTREME FLEXIBILITY
Easy reconfiguration of accelerators from one to several nodes as needed

OUTSTANDING SIMPLICITY AND EASY MANAGEMENT
Turnkey, easy to deploy, no software changes needed

COMPELLING ECONOMICS
Lower TCO, higher resource utilization, and less complexity, space, power, cooling, CAPEX, and OPEX

LEADERSHIP EFFICIENCY
Strips away conversion and overhead and eliminates bounce buffers with 100% PCI Express interconnect

The GigaIO SuperNODE is the first single-node server with 32 accelerators and up to 1 petabyte of storage.

The power of all these accelerators seamlessly connected by FabreX™ can now be harnessed to drastically speed up time to results. SuperNODE can connect up to 32 AMD Instinct™ MI210 accelerators and up to 24 NVIDIA A100s; other types of accelerators will be validated as testing progresses.

Workloads run faster because they utilize the power of each accelerator as if it was located in a single rack-scale server rather than in multiple constrained server nodes connected by a congested, complex network. All resources communicate over a single dynamic memory fabric.

Breathtaking Performance

Prior to SuperNODE, the only way to harness 32 GPUs would have required four servers with eight GPUs apiece, or 8x 4-GPU servers, connected using messaging protocols. While FabreX can also natively operate with the same messaging protocols, SuperNODE introduces a fourth option with all GPUs directly benefiting from native PCIe performance when communicating with either the processors or with other GPUs.
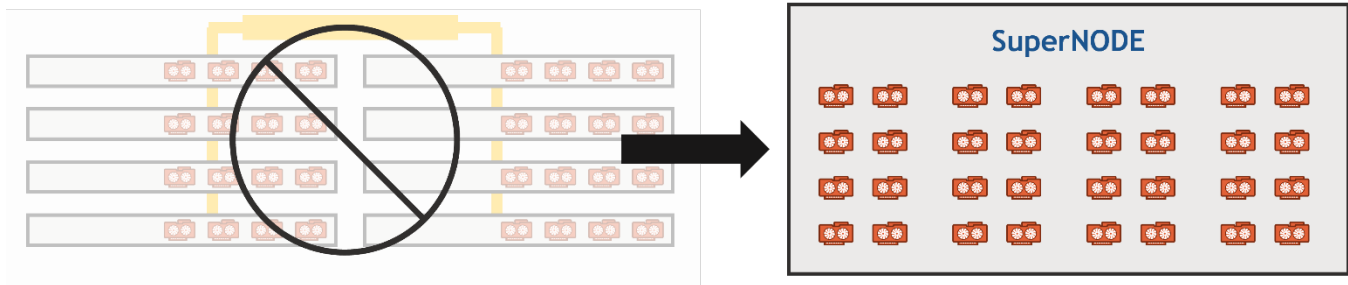

Ethernet, InfiniBand, **FabreX**

PROTOCOL OPTIONS:

1. MPI
2. CUDA-aware MPI
3. NCCL / RCCL
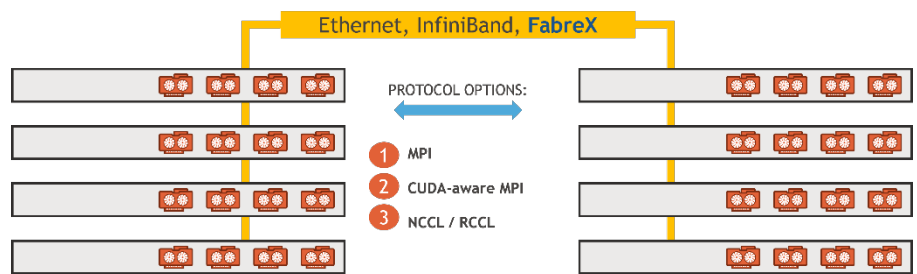
Current answer is to network the servers

This requires software changes!

The Better Option: GigaIO's SuperNODE

Extreme Flexibility Through Composability

The GigaIO SuperNODE can not only harness the power of up to 32 accelerators in a single node, but also easily distribute those same resources to up to six servers when "Beast Mode" is not required — for example, when workloads require fewer accelerators, or to accommodate multiple users.
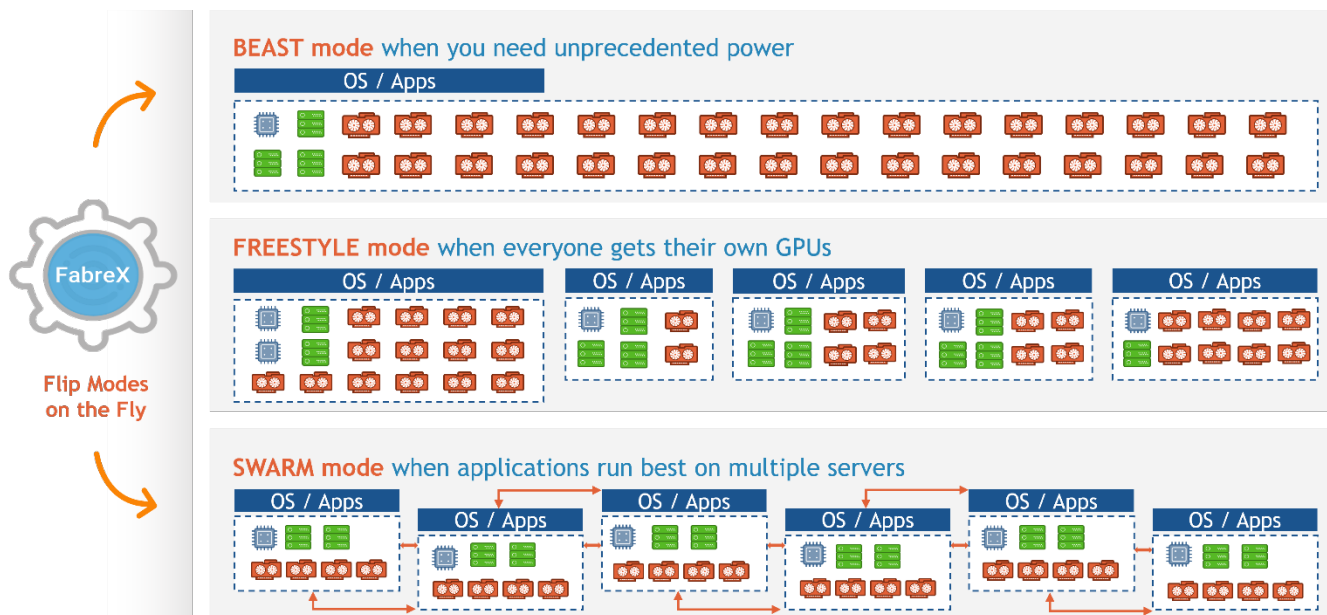
Outstanding Simplicity and Easy Management — "It Just Works"

Accessing the entire power of a SuperNODE is as simple as using the same AI frameworks you already use, such as PyTorch or TensorFlow — no changes are needed to your software. In order to provide a complete turnkey solution, all SuperNODEs include NVIDIA Bright Cluster Manager Data Science software to compose compute infrastructures that can be configured and scaled in seconds to handle the most demanding AI and HPC workloads. NVIDIA Bright Cluster Manager provides complete dynamic infrastructure management, including schedulers and container management. All hardware and software are preconfigured, tested, and validated.

An Enduring OPEN Infrastructure Built for Change

Gone are the days of having to make massive and costly forklift upgrades to replace very expensive fixed systems in order to incorporate new technologies. Upgrade or add compute,

www.gigaio.com

accelerators, and storage at the component level that plug-n-play with your environment. Every major subsystem can now operate on its own upgrade cycle — from whichever component vendor best meets your needs — with no vendor lock-in. The total cost of the system is optimized as FabreX drives much higher utilization of its resources.

Impeccable Scaling

While scaling accelerators over multiple nodes — typically done using MPI — yields as low as 50% scalability, SuperNODE delivers close to perfect scale factors. The Hashcat graph below shows GPUs scaling perfectly linearly all the way to 32.

The second graph shows the results of running ResNet50 over the same 32 GPUs. The scale factor drops to 99% for a 1% degradation per GPU, or an overall scale factor of 70%.


How it Works

The magic in FabreX is its ability to use PCIe as a routable PCI network that enables all server-to-server communication, rather than simply connecting resources to a single processor using a PCI tree. FabreX is the only PCIe-based routable network on the market, and uses Non-Transparent Bridging (NTB) to shatter the boundary between two PCIe hierarchical domains.

At its core, FabreX incorporates a construct for a memory address-based router with a built-in intelligence unit hosted in the switches that comprise the FabreX interconnect network. This routing mechanism virtualizes all hardware resources that consist of processor ecosystems and I/O devices as memory resources within a 64-bit virtual address space.

Communication between these resources consists of exclusively using the memory semantics of 'Memory Read' and 'Memory Write'. This means you can compose servers and CPUs in the exact same way as you compose end points (GPUs, FPGAs, ASICs, Smart NICs — anything with a PCIe connection).

www.gigaio.com
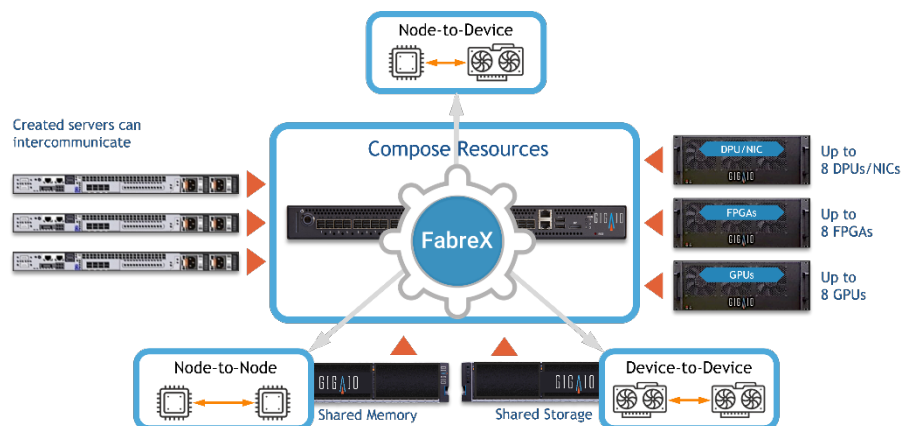
SuperNODE Capabilities Included

Accelerator Composability/Disaggregation	Storage Composability/Disaggregation
- Each host can access individual GPU devices	- Each host can access individual NVMe-oF storage devices
Management - Supports Redfish® RESTful APIs; real-time cluster topology; and utilization and performance dashboards - Predictive health and telemetry monitoring - Role-based authentication and access control	**Bundled Software** - FabreX Fabric Manager - Preloaded Linux & libraries, including NVMe-oF and GDR - Composable cloud management platform with Bright Cluster Management supporting High-Availability (HA) Storage and Bright for Data Science

SuperNODE is Delivered as One of Two Fully Engineered Solutions

The GigaIO SuperNODE comes pre-configured with FabreX composable switches, network cards, cables, management software, and accelerator pooling appliances (APAs) — everything you need to easily deploy the world's most advanced, open, and flexible accelerated computing platform.

SuperNODE 128 Gb/sec
The best value for flexibility

- Up to 32 accelerators
- Up to 1PB storage
- 1 to 6 servers
- 2 FabreX switches
- 4 APAs
- FabreX Fabric Manager
- Bright for Data Science

SuperNODE 256 Gb/sec
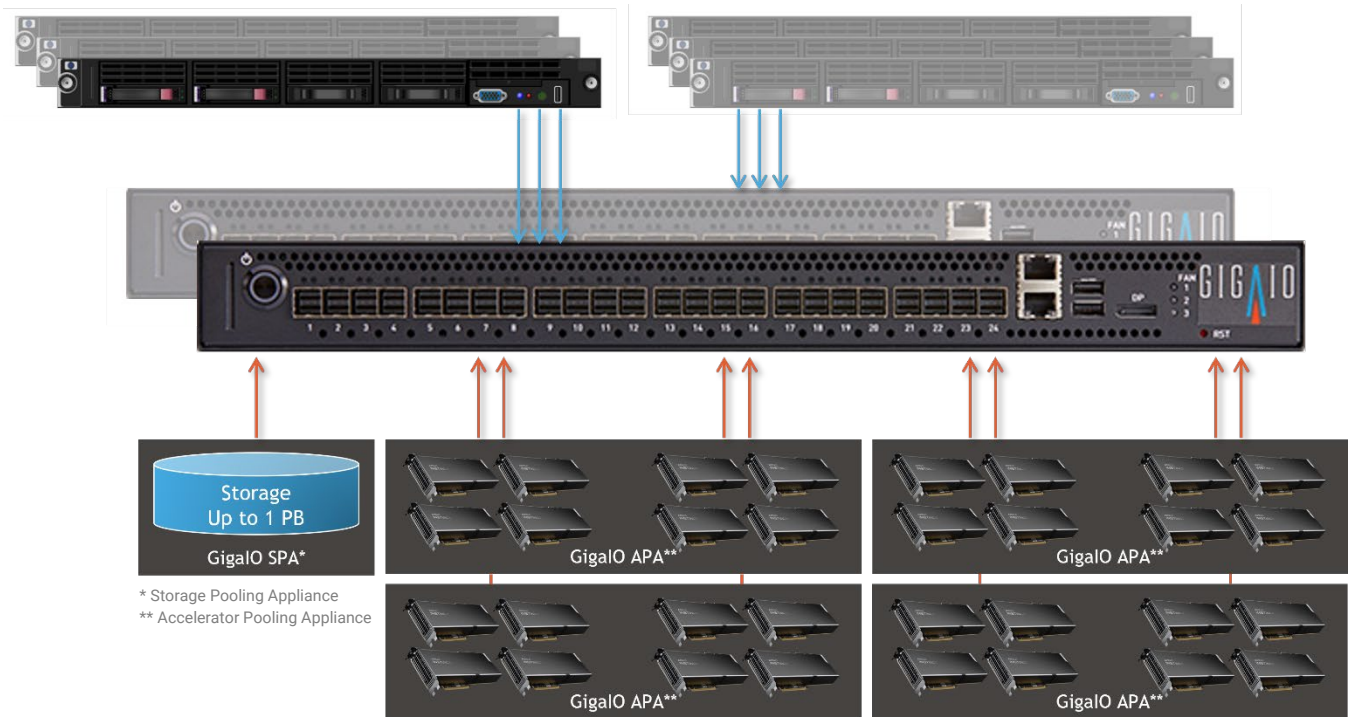*When you need the most bandwidth between GPUs*

- Up to 32 accelerators
- Up to 1PB storage
- 1 to 4 servers
- 5 FabreX switches
- 4 APAs
- FabreX Fabric Manager
- Bright for Data Science

www.gigaio.com

SuperNODE Configurations

Storage
Up to 1 PB
GigaIO SPA*

GigaIO APA**

* Storage Pooling Appliance
** Accelerator Pooling Appliance

Increase performance with native intra-server P2P, eliminating legacy network overhead

Designed for simplicity — no need to modify your existing software

Drastically reduce AI costs, in both GPUs and core licenses

Provide ultimate flexibility for any workload: use 32 accelerators on one node, or distribute the 32 across six servers

Save on power (7KW per 32-GPU deployment) and rack space (30% per 32-GPU deployment)

Effortlessly expand to increase compute and storage performance

SuperNODE is the ultimate computing machine for generative AI and demanding accelerated computing applications, delivering the leadership amount of raw GPU computing when in Beast Mode, yet easily reconfigurable to service several users sharing those accelerators.
Learn more at gigaio.com/supernode