



Engineered Solution Brief

SuperNODE™

The Open Solution for Generative AI and Accelerated Computing

HIGHLIGHTS

BREATH TAKING 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

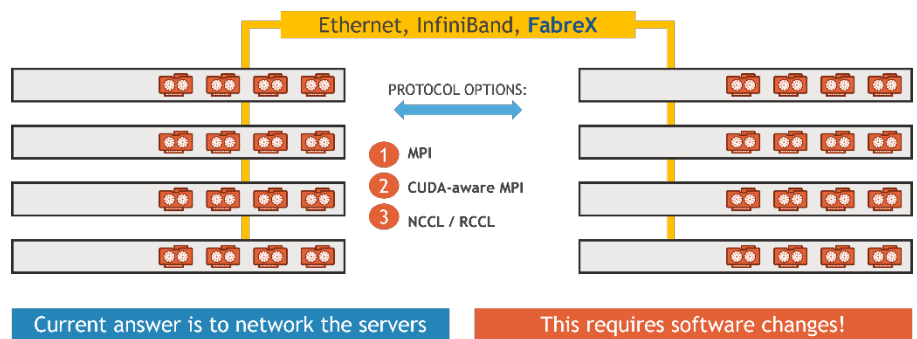
The GigaIO SuperNODE is the first single-node server with 32 accelerators and up to 1 petabyte of storage.

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 or NVIDIA accelerators.


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.

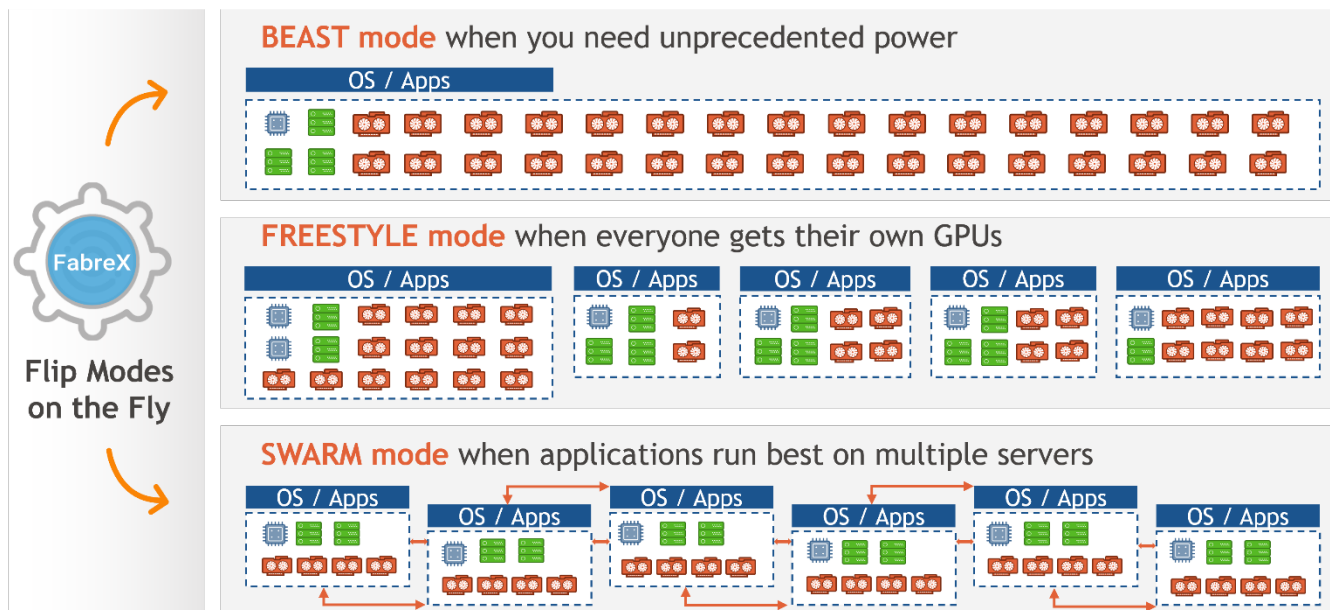


The Better Option: GigaIO's SuperNODE



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 four 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 AI Enterprise Essentials software to compose compute infrastructures that can be configured and scaled in seconds to handle the most demanding AI and HPC workloads. NVIDIA Base Command 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, 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.

HPL-MxP showed excellent scaling and reduced precision compute bandwidth, achieving 99.7% of ideal theoretical scaling.

When running training workloads such as Llama and ResNet-50, SuperNODE continues to exhibit perfect scaling.

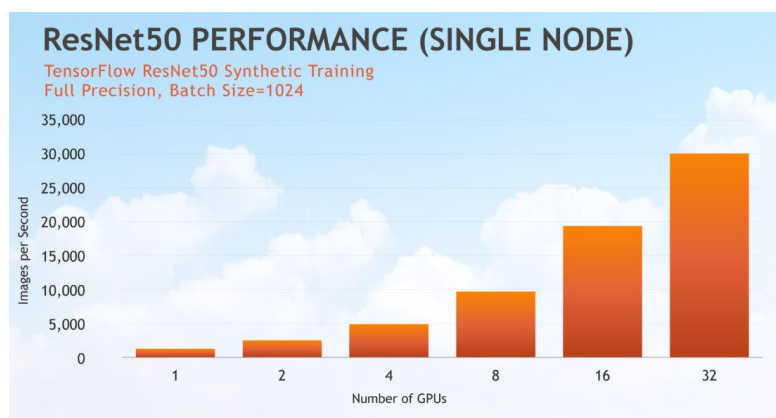
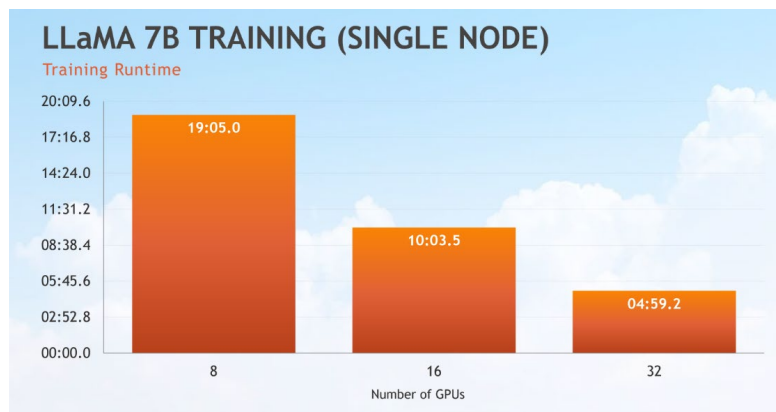

Llama 7B is part of the LLaMA (Large Language Model Meta AI) family of autoregressive large language models (LLMs), released by Meta AI starting in February 2023. “7B” refers to 7 billion parameters.

The **LLaMA** Training graph (*at right*) shows the gain in time-to-results when adding more GPUs to a single node – from 19 minutes with a standard 8-GPU server configuration to less than five minutes on a SuperNODE with 32 GPUs.

ResNet stands for Residual Network and is a specific type of convolutional neural network (CNN) introduced in the 2015 paper “Deep Residual Learning for Image Recognition” by He Kaiming, Zhang Xiangyu, Ren Shaoqing, and Sun Jian. CNNs are commonly used to power computer vision applications.

ResNet-50 is a 50-layer convolutional neural network (48 convolutional layers, one MaxPool layer, and one average pool layer). Residual neural networks are a type of artificial neural network (ANN) that forms networks by stacking residual blocks.

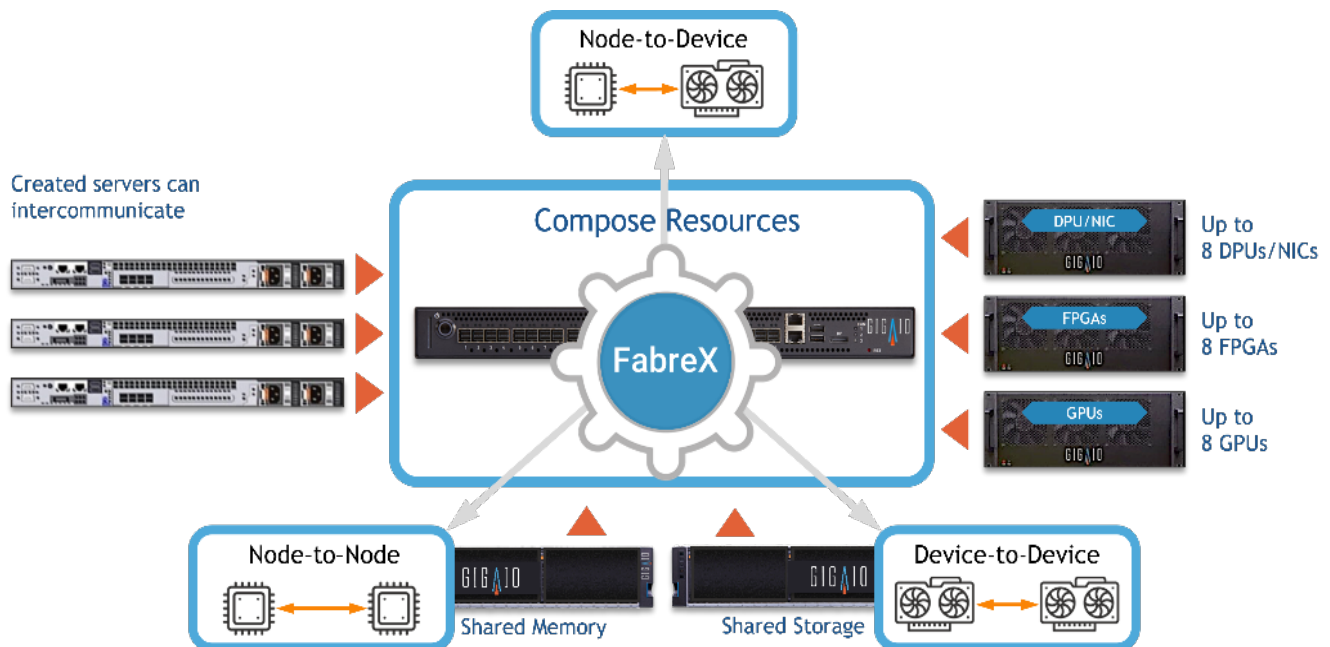
The **ResNet50** Performance graph (*above, at right*) shows the increase in images per second as GPUs are added to a SuperNODE.



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).



SuperNODE Capabilities Included


Accelerator Composability/Disaggregation - Each host can access individual GPU devices	Storage Composability/Disaggregation - 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 - NVIDIA AI Enterprise Essentials featuring Base Command Manager for Composable Cluster Management and HA Storage

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.

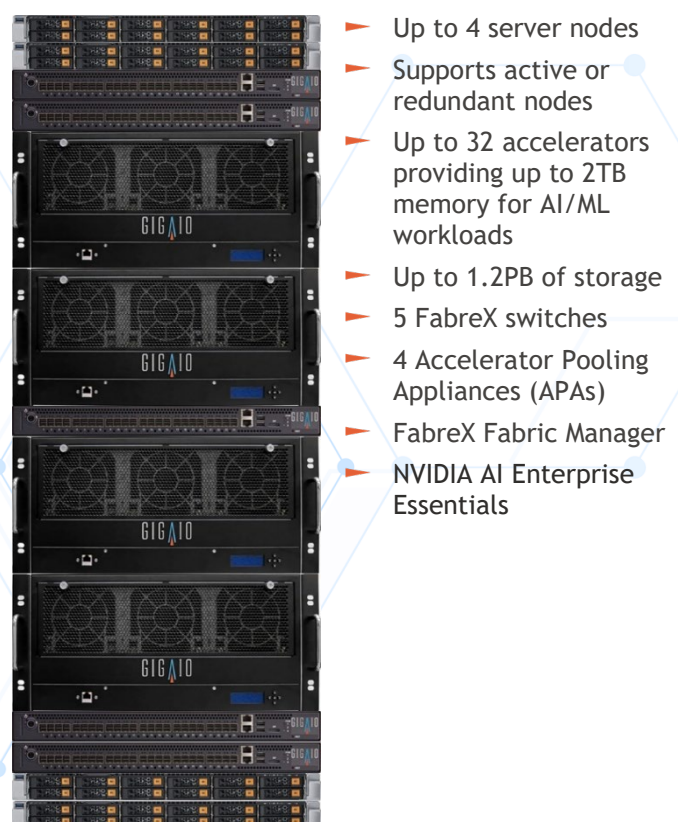
SuperNODE 256 - Single & Dual Node

Best for [BEAST mode](#) when you need unprecedented power





SuperNODE 256 - Multi Node

Best for [BEAST](#), [SWARM](#), or [FREESTYLE modes](#) for the utmost in flexibility





SuperNODE Interconnect



 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 four servers

 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](http://www.gigaio.com/supernode)