



The GigaIO Engineered Solution for Manufacturing, Powered by AMD

SOLUTION BRIEF



Speeds CAE analysis by boosting utilization while lowering costs.

Key Challenges

- Imperative to shorten design cycle times
- Underutilized use of critical, expensive GPU and CPU resources
- Controlling CAE software license costs

Solution

- The GigaIO Engineered Solution for Manufacturing, Powered by AMD
- Runs all CAE software with no changes
- Works with your existing HPC storage

Benefits

- Cloud-like agility, on-prem economics
- Scale resources independently
- Speed and simplify CAE analysis
- Reduce design cycle time in half
- Boost GPU/CPU processor utilization by more than 100%
- Increase product quality with higher resolution simulations and model fidelity
- Significantly reduced CAPEX and OPEX
- Eliminate HPC network complexity
- Immediate availability with no supply chain delay.

Dynamically Composable to Adapt to Different CAE Workloads

FabreX™ can vary computational resources in order to provide CAE engineers with the right hardware profile for each workload, and vary resources usage within a single workflow. Compute and GPU resources can be configured to optimize mesh, simulation, and visualization on the fly. What's more, FabreX will easily scale to handle the abundance of GPU-accelerated hardware demands that will be required by introducing Machine Learning to CAE simulations.

Shrink Design Cycles by Unlocking All Available Compute Power

Higher resolution simulations involving ever increasingly complex multi-domain physics scenarios are challenging CAE organizations tasked to shorten design cycles while also keeping HPC hardware and simulation costs under control. GigaIO addresses these CAE simulation challenges head-on with FabreX, the highest performance, lowest latency rack-scale server dynamic memory fabric on the planet. FabreX liberates all the stranded power in your infrastructure's GPU accelerators – while also providing engineers with the right hardware configuration at the right time to optimize simulation and visualization results.

Breakthrough Performance and Compelling Economics

Because it disaggregates GPUs from the servers into pools of resources available throughout the cluster, FabreX improves GPU utilization by more than 100%. In one side-by-side test shown below, a rack built with FabreX ran the same customer weekly workload with 16 GPUs instead of the 32 initially deployed on InfiniBand. The result is a system that costs 30% less to purchase, with similar savings on CAE software costs.

FabreX Infrastructure

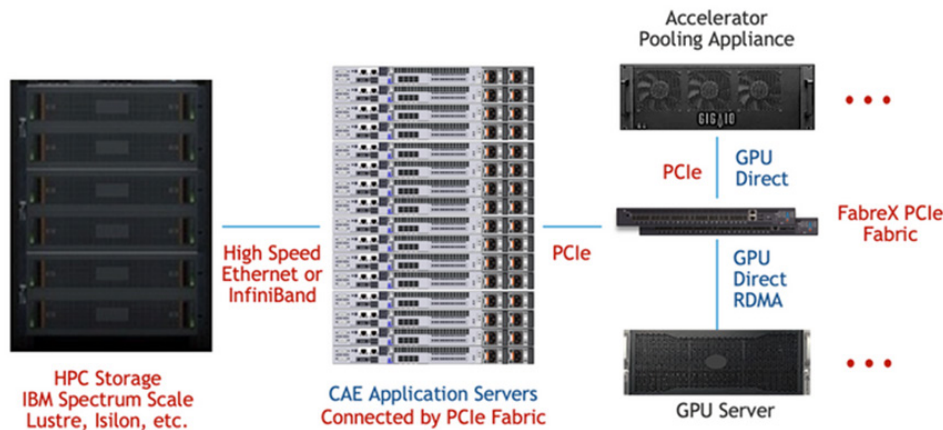
INFRASTRUCTURE INPUTS		
Number of GPU servers	0	8 GPUs/server
Number of App servers	8	CPU-only app servers
Number of JBOGs	2	8 GPUs/JBOG
Number of GPUs	16	
GPU model	V100-32	
Network b/w per port (Gb/s)	128	
Cluster Network type	FabreX	
WORKLOAD INPUTS		
Number of jobs	150	
% of jobs that need GPUs	50%	
OUTPUTS		
Elapsed Run Time	151	hours

System Purchase Price: **290K**

InfiniBand Infrastructure

INFRASTRUCTURE INPUTS		
Number of GPU servers	4	8 GPUs/server
Number of App servers	4	CPU-only app servers
Number of JBOGs	0	8 GPUs/JBOG
Number of GPUs	32	
GPU model	V100-32	
Network b/w per port (Gb/s)	100	
Cluster Network type	InfiniBand	
WORKLOAD INPUTS		
Number of jobs	150	
% of jobs that need GPUs	50%	
OUTPUTS		
Elapsed Run Time	156	hours

System Purchase Price: **412K**



The GigaIO Engineered Solution for Manufacturing, Powered By AMD

Under the Hood, What Makes GigaIO's Composable Infrastructure So Performant?

GigaIO's software-first composable platform increases resource utilization and lowers cost of ownership by allowing data center managers to compose all individual resources as needed and on-the-fly to adapt to changing workflows, and to burst to the cloud as needed, as organizations pivot to container-native architectures.

As shown above, the GigaIO Engineered Solution for HPC in Manufacturing delivers the industry's lowest latency and the highest effective bandwidth by integrating CAE compute and GPU acceleration I/O into a single-system cluster fabric, using industry-standard PCI Express (PCIe) technology. GigaIO enables true server-to-server communication across PCIe and makes cluster scale computers possible, with direct memory access by an individual server to system memories of all other servers in the cluster fabric, for the industry's first in-memory network.

Patented GigaIO technology strips away unnecessary conversion, software layers, and overheads that add latency to legacy interconnects. HPC storage is easily integrated via the CAE application servers as they normally would in conventional infrastructures.

Ultimate Power and Flexibility from only a Few Composable FabreX Building Blocks

A complete Engineered Solution for HPC in Manufacturing can be built out of a few simple FabreX Building Blocks. Lower CapEx and OpEx through less hardware, higher utilization of resources, quicker adoption of new technology, lower power consumption, and less cooling. Avoid over-provisioning and add just the components you need. Maximize utilization of the footprint of your data center and contribute to your bottom line.

The FabreX Fabric Manager (FM) is the central building block and the software engine that drives the performance and dynamic composability of GigaIO's software-defined hardware. This Linux-based, resource-efficient software layers onto FabreX hardware for easy-to-use composing of computing clusters on-the-fly. A classic Top-Of-Rack network switch, the FabreX Switch communicates with FabreX server drivers to identify and coordinate resources required by the servers.

The FabreX Network Adapter card is the high-performance, cabled interface to cluster subsystems across FabreX. The card includes both server and target (for PCIe I/O) modes and is FPGA-powered for configuration flexibility. Applications can access remote PCIe devices as if they were attached to the local system. The GigaIO Accelerator Pooling Appliance is the industry's highest performing PCIe accelerator appliance with up to 1Tb/sec bandwidth into and out of the box. It delivers advanced provisioning and monitoring of accelerators supporting up to 8 double-width PCIe Gen 4.0 x16 accelerator cards and 2 PCIe Gen 4.0 x16 low-profile slots. Industry-standard HPC cluster servers and GPU servers are also integrated as shown.

Outstanding Simplicity and Easy Management

Accessing all the power of GigaIO FabreX is as simple as using the same cluster management and workload scheduling managers as you always have. What's more, there are no problematic NIC or OFED driver bugs, incompatibilities, etc., and there are no more supply chain delays. And down the road, you can add compute, storage, and application accelerators at the component level that plug-n-play with your environment. Every major subsystem can now operate on its own upgrade cycle. And the total cost of the system is optimized over its lifecycle as FabreX drives much higher utilization of all resources.

