



Engineered Solution for Universities, Powered by AMD

SOLUTION BRIEF



A turnkey, rack-scale appliance for Higher Education that can centralize, standardize, share, and optimize infrastructure across the varied workflows required by a university — both classroom and research settings.

Key University Challenges

Professors who teach students about real-world workloads such as CAE, CFD, and generative design have limited options when tied to the fixed infrastructure of their academic institution. Traditional IT architecture means limited configurations, poor resource utilization for varied types of tasks, and difficulty reconfiguring for different types of workloads.

Composable infrastructure is an emerging technology that eliminates the restrictions imposed by traditional, static server architectures. It delivers the agility and flexibility of cloud computing while lowering total cost of ownership.

Providing Turnkey, Multi-tenancy Workflows for Higher Education

GigaIO's Engineered Solution for Universities, Powered by AMD is a flexible environment for heterogeneous compute that can easily accommodate the different workloads required for teaching, professor research, and grad-student research. It is a complete, future-proofed composable infrastructure solution that facilitates dynamic workflows and is simple to deploy and manage.

Combine NVIDIA Bright Cluster Manager's ability to easily build and manage clusters with GigaIO's ability to connect AMD accelerators, servers, and other devices in a seamless dynamic fabric.

- Easy-to-deploy, turnkey appliance for Higher Education
- Centralize, standardize, share, and optimize infrastructure
- Handle varied types of workloads across departments

GigaIO FabreX™

GigaIO's FabreX is a universal dynamic memory fabric that transforms static infrastructure into software-defined hardware for advanced scale workflows by seamlessly composing an entire rack of resources into a unified compute platform, without proprietary architecture lock-in.

AMD Accelerators and CPUs

AMD EPYC™ processors and Instinct™ accelerators can tackle the world's most complex and critical workloads. Using the highest performing technical and general computing infrastructure means that workloads are run faster and less infrastructure, power, and cooling are needed.

NVIDIA Bright Cluster Manager

NVIDIA Bright Cluster Manager reduces the time, skills, effort, risk, and complexity of building clustered infrastructure. Quickly build and manage heterogeneous high-performance Linux clusters that host HPC, machine learning, and analytics applications from core-to-edge-to-cloud.

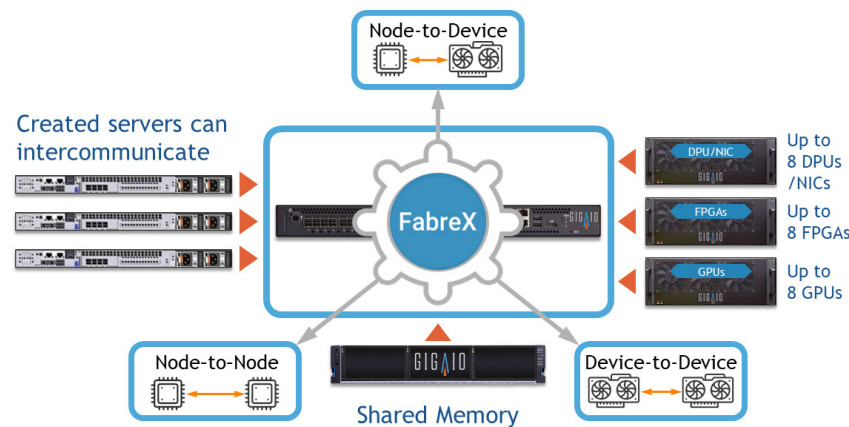
How it Works

GigaIO's Engineered Solution for Universities combines NVIDIA Bright Cluster Manager's ability to build and manage clusters with GigaIO's ability to connect AMD accelerators, servers, and other devices in a seamless fabric. Ease of use is at the forefront of this design, as its goal is to be used in a class/lab setting without requiring dedicated IT expertise.

The appliance comes pre-loaded with dynamic infrastructure management software, including schedulers and container management. Both hardware and software are preconfigured, tested, and validated in all configurations, and it is simple

to reconfigure to match existing user and group policies. Native integration with Bright Cluster Manager allows the owner to easily assign configurations prior to use, dividing hardware among students to allow them the experience of running actual simulation workloads on the same compute infrastructure they will utilize upon graduation.

Disaggregating resources outside of servers enables each department to share them. Each department can now compose systems precisely tuned to the needs of the applications they need to run. Seamless integration to burst workflows allows users to tap into the cloud as needed for added scaling and/or immediate availability.



Key Benefits

GigaIO's Engineered Solution for Universities provides a host of benefits in a higher education setting. It supplies students with access to equipment they will use in the real world, so they will be better prepared for the job market. Flexibility and composability means that systems don't remain idle while not being used for teaching — they can instead be reconfigured for actual simulation work and swapped back into teaching mode as needed. And of course the inherent benefit of using composable infrastructure is the ability to build heterogeneous configurations for varied types of workloads and trainings.

Easy to deploy and maintain: Choose from a set of fully configured, ready-to-deploy systems optimized for most working and teaching environments.

Time to value: By making underutilized and stranded resources accessible to more users, you get to research results faster.

Lower Total Cost of Ownership (TCO): Sharing resources means needing to purchase fewer accelerators and CPUs, consuming less power, and decreasing the manpower needed to administer changing needs.

Lower Research Cost of Ownership (RCO): Reduce loss of research time due to lack of availability of computing

resources and downtime while technicians physically re-able physical resources.

Efficiency: Increase utilization rate of expensive resources and share them across a greater number of users.

Future-proofing: No need to predict what your users will need in next 3-4 years, since you can add resources as needed to accommodate new accelerators and AI algorithms.

Cloud-like agility: On-premises infrastructure acquires a degree of flexibility impossible with fixed siloed architectures. Infrastructure can now change as your workloads change.

Choose Your Configuration

GigaPod

The perfect entry point and preconfigured for easy deployment.



SUPPORTS UP TO:



**640
CPU Cores**



**4.6TB
Memory**



**307TB
Storage**




**16
Accelerators**

- 1 GigaPod kit — includes FabreX™ composable switches, network adapter cards, and cables
- Up to 2 GigaIO Accelerator Pooling Appliances — each up to 8 mix-and-match GPUs
- Up to 2 storage servers — each up to 154TB
- Up to 4 compute nodes — pick the processor — dual AMD EPYC™ 7713, 75F3, or 7543 — each with up to 1TB memory and 128 cores
- Bright Cluster Manager supporting High Availability (HA) storage and Bright for Data Science

GigaCluster

Expand your system to increase compute and storage performance.



SUPPORTS UP TO:



**1,280
CPU Cores**



**9.2TB
Memory**



**614TB
Storage**



**32
Accelerators**

- 1 GigaCluster kit — includes FabreX™ composable switches, network cards, cables, and expansion pack
- Up to 4 GigaIO Accelerator Pooling Appliance — each up to 8 mix-and-match GPUs
- Up to 4 storage servers — each up to 154TB
- Up to 8 compute nodes — pick the processor — dual AMD EPYC™ 7713, 75F3, or 7543 — each with up to 1TB memory and 128 cores
- Bright Cluster Manager supporting High Availability (HA) storage and Bright for Data Science

3Pod GigaCluster

When you need the ultimate in compute power.



SUPPORTS UP TO:



**1,920
CPU Cores**



**13.8TB
Memory**



**922TB
Storage**



**48
Accelerators**

- 3Pod GigaCluster kit — includes FabreX™ composable switches, network adapter cards, cables, and expansion pack
- Up to 6 GigaIO Accelerator Pooling Appliance — each up to 8 mix-and-match GPUs
- Up to 6 storage servers — each up to 154TB
- Up to 12 compute nodes — pick the processor — dual AMD EPYC™ 7713, 75F3, or 7543 - each with up to 1TB memory and 128 cores
- Bright Cluster Manager supporting High Availability (HA) storage and Bright for Data Science

The GigaIO Engineered Solution for Universities, Powered by AMD is a turnkey, rack-scale appliance for Higher Education that can centralize, standardize, share, and optimize infrastructure across the varied workflows required by a university — both classroom and research settings. Learn more at gigaio.com/amd-higher-ed

AMD EPYC™ is a trademark of AMD. All trademarks, logos and brand names are the property of their respective owners.