



## For Immediate Release

### **GigalIO FabreX for Composable Infrastructure Now Supported Natively in NVIDIA Bright Cluster Manager 9.2**

*Integration combines ability to build and manage clusters in Bright Cluster Manager with GigalIO's ability to disaggregate and compose servers, accelerators, and devices within seamless memory fabric.*

**San Diego, California, May 19, 2022** – GigalIO, provider of the world's only open rack-scale computing platform for advanced scale workflows, today announced that GigalIO FabreX™ for composable infrastructure is now natively supported in [NVIDIA Bright Cluster Manager 9.2](#). The integration, led by NVIDIA in collaboration with GigalIO, ensures customers can build easy-to-manage, platform-independent compute clusters that scale in minutes to handle the most demanding AI and HPC workloads.

This new integration is an example of GigalIO's strategy to deliver an open platform that allows customers to access the benefits of composable infrastructure via the enterprise-class tools they already use. "With our strategy of native integration into leading software tools such as NVIDIA Bright Cluster Manager, our goal is to be invisible to data center managers so that their users can seamlessly submit jobs and not even need to know about the magic of our software-defined hardware reconfiguring resources on the fly," said Alan Benjamin, CEO of GigalIO.

"Enterprises building AI and HPC computing infrastructure are seeking solutions that provide performance, flexibility, and efficiency," said Charlie Boyle, vice president of DGX systems at NVIDIA. "With native support for Bright Cluster Manager 9.2, GigalIO FabreX customers can now compose and manage their compute systems to suit the needs of unique workloads from a single management interface."

GigalIO's universal dynamic memory fabric, FabreX, enables an entire server rack to be treated as a single compute resource. Resources normally located inside of a server, including accelerators, storage, and even memory, can now be pooled in accelerator or storage enclosures, where they are available to all of the servers in a rack. These resources continue to communicate over a native PCIe memory fabric just as they would if they were plugged into the server motherboard.

NVIDIA Bright Cluster Manager is an enterprise-class software solution that simplifies building and managing HPC clusters from edge to core to cloud, transparently to the customer, by combining provisioning, monitoring, and management capabilities in a single tool. Version 9.2 extends the goals of eliminating complexity and enabling flexibility by adding built-in support for composable infrastructure using GigalIO FabreX, where nodes can now be composed using Bright Cluster Manager Shell or BrightView.

Auto-scaling in Bright Cluster Manager creates a dynamic, multi-purpose infrastructure, and FabreX extends that agility to each hardware element in a rack by enabling the creation of composable GigaIO GigaPods and GigaClusters with cascaded and interlinked switches. PCIe devices can be monitored and health-checked, as well as pooled and assigned to nodes within a cluster using Bright Cluster Manager. Clusters may have several different fabrics defined, and Bright Cluster Manager streamlines the fabric configuration process.

Support for FabreX with NVIDIA Bright Cluster Manager allows users to handle more workloads while maximizing resource utilization, minimizing cost, and managing everything from a single state-of-the-art user interface. Cloud-like agility is now easier for on-prem infrastructure, allowing cloud bursting as needed within a single interface.

GigaIO's dedication to native integration with software tools like Bright Cluster Manager provides a best-in-class experience for customers who can continue to use their favorite tools without needing to alter their software stack. Native integration frees customers from having to rely on plug-ins with limited capabilities, learn new software, manage yet another pane of glass, and pay additional per-node license fees.

This development expands GigaIO's close collaboration with NVIDIA. GigaIO has been a member of the [NVIDIA Partner Network](#) (NPN) since 2020 and was accepted last month into [NVIDIA Inception](#), a program designed to nurture cutting-edge startups. GigaIO shares many common customers with NVIDIA, including the Texas Advanced Computing Center and the San Diego Supercomputer Center.

Dr. Frank Würthwein, Director of the San Diego Supercomputer Center, is a beneficiary of this collaboration. "Our research requires that we aggregate disparate computational elements, such as GPUs, x86 processors, and storage systems into highly usable and reconfigurable systems," he said. "GigaIO's FabreX technology combined with NVIDIA Bright Cluster Manager makes it possible to dynamically bring these elements together in a very low-latency, high-performance interconnect while allowing for distinct, non-interfering workflows to co-exist on the same infrastructure."

[Learn more](#) about the native integration of FabreX for composable infrastructure in NVIDIA Bright Cluster Manager 9.2, available now.

### **About GigaIO**

Headquartered in Carlsbad, California, GigaIO provides the world's only open rack-scale computing platform, delivering the elasticity of the cloud at a fraction of the TCO (Total Cost of Ownership). With its universal dynamic memory fabric, FabreX™, and its innovative open architecture using industry-standard PCI Express/soon CXL technology, GigaIO breaks the constraints of the server box, liberating resources to shorten time to results. Contact [info@gigaio.com](mailto:info@gigaio.com), visit [www.gigaio.com](http://www.gigaio.com), or follow on [Twitter](#) and [LinkedIn](#).

### **Contact**

Danica Yatko  
[danica@xandmarketing.com](mailto:danica@xandmarketing.com)