FOR IMMEDIATE RELEASE

GigaIO Expands Family of Engineered Solutions Powered by AMD, Launches Purpose-Built Editions for Life Sciences & Manufacturing at SC22

Composability solutions reduce analysis time by boosting accelerator utilization while reducing costs.

San Diego, California, November 3, 2022 – GigaIO, the leading provider of workload-defined infrastructure without compromise for HPC + AI workflows, today announced the expansion of a new architecture that promises to dramatically improve computational resource utilization for both genomic and CAE analysis with the launch of its two newest Engineered Solutions. The GigaIO Engineered Solutions: Life Sciences and Manufacturing Editions, powered by AMD EPYC™ processors and AMD Instinct™ accelerators, are flexible platforms for heterogeneous compute that can easily disaggregate GPU and FPGA accelerator resources previously locked in physical servers, enabling the provisioning of genomics and CAE analysis workloads with the optimal type and amount of accelerator resources to speed analysis, increase resource utilization, and lower costs.

"AMD processors excel at running highly diverse applications and workloads," said Kumaran Siva, corporate vice president, Software & Systems Business Development, AMD. "A composable solution like GigaIO's provides increased flexibility for our processors to be leveraged in multiple ways — from a simulation environment with single-thread performance to throughput computing."

"With the expansion of these appliances, we are bringing easy-to-use infrastructure to Life Sciences and Manufacturing, where composability can provide a wide array of flexible accelerator technologies that shorten analysis times," said Alan Benjamin, CEO of GigaIO. "AMD is the perfect technology partner for this venture because we share a commitment to creating an open, industry standards-based platform. We are keen to make it easy for people to deploy this new technology through our family of products powered by AMD, building on our successful joint deployments at TACC and SDSC."

The GigaIO Engineered Solution: Life Sciences Edition was built to address many genomic analysis challenges. Modern high-throughput sequence alignment is very computationally intense, and will only get more demanding with future sequencer generations. The Solution allows researchers to speed analysis by applying the optimum mix of HPC clusters, with both GPU and FPGA accelerators, all connected by low-latency fabrics. This is made possible by FabreX™, the world's highest performance, lowest latency dynamic memory fabric, which allows differently programmed FPGAs to process different parts of the genomics analysis pipeline — a capability long sought after in modern sequencing. FabreX also liberates all of the stranded power in CPUs, FPGAs, and GPUs in order to speed genomic secondary analysis.

FabreX enables an entire server rack to be treated as a single compute resource, handling all compute communication, including server-to-server traffic (such as MPI and NVMe-oF). 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 and servers continue to communicate over a native PCIe memory fabric for the lowest possible latency and highest possible bandwidth performance, just as they would if they were still plugged into the server motherboard.

The GigaIO Engineered Solution: Manufacturing Edition shortens time to market and increases product quality by allowing higher resolution simulations and model fidelity. The Solution can be configured with different computational resources in order to provide CAE engineers with the right hardware profile for each workload, and vary resource usage within a single workflow. Compute and GPU resources can be configured to optimize mesh, simulation, and visualization on the fly. FabreX will easily scale to handle the abundance of GPU-accelerated hardware demands that will be required by introducing Machine Learning to CAE simulations.

For ease of use, both GigaIO Engineered Solutions: Life Sciences and Manufacturing Editions are delivered with NVIDIA Bright Cluster Manager pre-installed, combining their ability to easily build and manage clusters with GigaIO's ability to connect AMD Instinct accelerators, AMD EPYC CPU-powered servers, and other devices in a seamless dynamic fabric. All of the GigaIO Engineered Solutions — including the University Edition, which was launched at ISC earlier this year — are designed to accommodate a variety of accelerator types and brands and provide a truly vendor-agnostic environment. They are container-ready, easily composed via bare metal, and feature AMD EPYC processors and AMD Instinct MI210 accelerators. The Life Sciences and Manufacturing Editions are available now.

About GigaIO
GigaIO provides workload-defined infrastructure through its dynamic memory fabric, FabreX, which seamlessly composes rack-scale resources and integrates natively into industry-standard tools. FabreX lets customers build impossible servers for HPC + AI workflows — from storage to accelerators to memory — at a fraction of cloud TCO, by optimizing the utilization and efficiency of their existing hardware, allowing them to run more workloads faster at lower cost through higher utilization of resources and more agile deployment. Visit www.gigaio.com, or follow on Twitter and LinkedIn.

Contact
Danica Yatko
760-487-8395
danica@xandmarketing.com