The GigaIO Composable Appliance for Life Sciences Powered by AMD Speeds Genomic Analysis by Boosting Accelerator Utilization While Lowering Costs

Solution Brief

Challenges

- Sequencers are getting faster and faster, swamping genomics pipelines
- Secondary Analysis processing times are now the bottleneck in genomics
- Personalized Medicine is throttled by the lack of computing resources

Solution

- The GigaIO Platform for Life Sciences powered by AMD
- Accelerate analysis with FPGA and GPU resources tailored to the growing needs of expanding genomics workloads
- Works with your existing HPC storage

Benefits

- Speed Secondary Analysis
- Deploy FPGA and GPU resources independently to optimize acceleration
- Lower processing times with GigaIO's low latency PCIe fabric
- Boost FPGA and GPU utilization by more than 100%
- Easily expand accelerator resources to keep up with future sequencers
- Significantly reduced CAPEX and OPEX
- Eliminate HPC network complexity
- Immediate availability with no supply chain delay

Composable Dissagreagated Infrastructure for Genomic Workloads

A new architecture is emerging that promises to dramatically improve Genomic Secondary Analysis computational resource utilization, called Composable Disaggregated Infrastructure (CDI). CDI disaggregates FPGA and GPU accelerator resources previously locked in physical servers, enabling the provisioning of genomics analysis workloads with the optimal type and amount of accelerator resources to speed analysis, increase resource utilization, and lower costs.

Speed Analysis With Targeted FPGA and GPU Acceleration

Modern high throughput sequence alignment is considered among the most challenging computational problems of our time. And it's only going to get more demanding with subsequent sequencer generations. Researchers are looking for ways to apply the optimum mix of HPC clusters, with GPU and FPGA accelerators, all connected by low-latency fabrics to speed analysis. GigaIO addresses these genomic analysis 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 genomic analysis infrastructure's CPU, FPGA, and GPU accelerators to speed genomic secondary analysis.

Breakthrough Performance and Compelling Economics

Because it disaggregates FPGAs and GPUs from the servers into pools of resources available throughout the genomics analysis cluster, FabreX improves accelerator 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, lowering costs significantly.

FabreX Infrastructure

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**

VS.

InfiniBand Infrastructure

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**

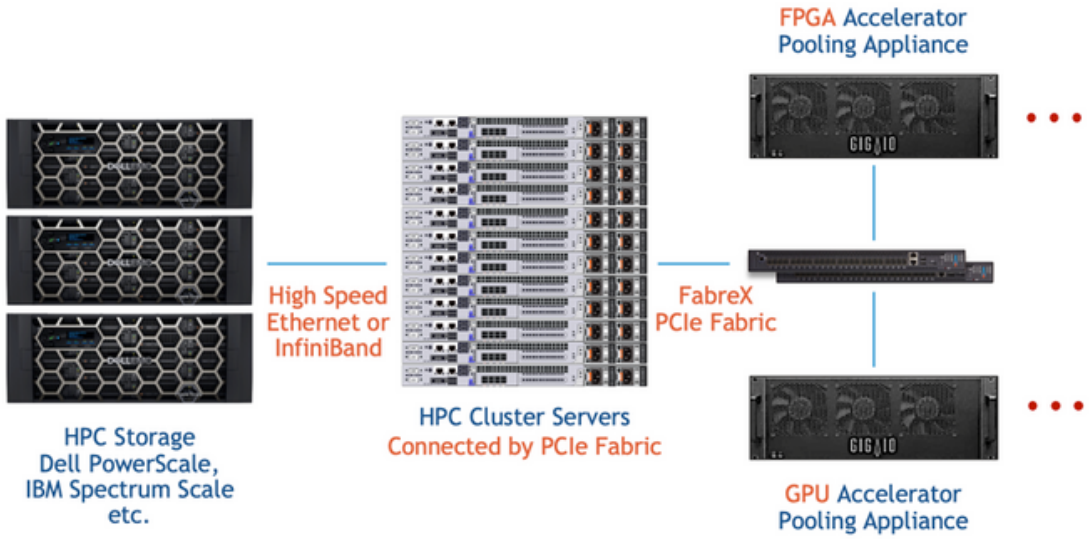

The GigaIO Platform for Life Sciences Powered By AMD

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

GigaIO's software-first composable platform increases resource utilization and lowers the cost of ownership by allowing genomics researchers to compose individual resources as needed and on the fly to adapt to changing workflows. As shown above, the AMD GigaIO Platform for Life Sciences delivers the industry's lowest latency and the highest effective bandwidth by integrating compute, FPGA, 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 HPC Cluster servers as they normally would in conventional infrastructures.

FabreX Building Blocks Compose Ultimate Power and Flexibility

A complete Platform for Life Sciences 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 device 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. Moreover, 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 primary 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.

 www.gigaio.com