GigaIO
Orchestrating High-performance Fabric with Maximum Flexibility

Data Centers today increasingly are struggling with a massive increase in the amount of data that is being collected, analyzed and stored in compliance with regulators and auditors. Unfortunately the networks and interconnects are not keeping pace with the advances in computing power (particularly with GPUs) and with the performance of storage with solid state drives with NVMe controllers. Today's networks were never designed to handle the amount of information that must be moved with the performance and latency required. Data centers, whether in the cloud or on-prem, require a high performance interconnect fabric, based on a recognized industry standard, that can deliver a completely new level of high bandwidth, low latency, and flexibility. Identifying this pressing requirement, a San Diego based firm, GigaIO, brings breakthrough interconnect performance and flexibility to advanced scale computing clusters used in AI/ML/DL, advanced analytics and high performance computing through its disruptive and patented technology—Link Express (LEX)™.


Alan Benjamin

Using PCI Express (PCIe) as the basis of the fabric, the patented Link Express fabric increases performance, scalability, and power efficiency required for advanced-scale computing clusters.

"PCI Express offers an unmatched combination of very high bandwidth and low latency. And when combined with its unique disaggregation capabilities and flexibility, we see this as a disruptive technology in the market," states Alan Benjamin, President, and CEO of GigaIO. In addition, Link Express is easily compatible with existing computing systems across verticals.

Harnessing the actual strengths of PCIe—easy accessibility, higher performance, low power consumption, and affordability, GigaIO's Link Express architecture can address multi-host communication and I/O sharing capabilities, increasing its status in high-performance cloud and on-premise data center installations. In addition, its advanced architecture enables data centers to scale performance and functionality, to enhance density and reduce power consumption while conserving the existing investment of an organization in PCIe devices.

Link Express technology runs up to 128 gigabits/secs of bandwidth per port, and delivers 768 Gb/s bandwidth through the switch. Latency numbers across the fabric, from system memory on one system to system memory on another, are less than 200 nanoseconds, based on a memory-to-memory data movement benchmark. Depending on the application, application software and the software stack that the user is running, the latency will range between 200-500 nanoseconds. This is an order of magnitude lower than any other fabric switch technology in the market, providing GigaIO a remarkable edge over its competitors.

"As a result, we save lots of time with exceptionally low latency and exceptionally high bandwidth. Coupled with our composability software, users can gain better business outcomes and reduce the total cost of ownership, because you add resource sharing and better utilization," determines Benjamin.

One early adopter of GigaIO started tests by mirroring their system at a smaller scale. A multi-server configuration was created for the client with the necessary amount of storage, along with four GPUs per server, all connected to a single fabric that combined GPUs, storage, and a management CPU.

The customer required the testing of the speed between the servers and to storage as computing performance was the organization's foremost priority. They saw considerable performance gain from having the storage not only physically close but logically close as well as a superior low latency interconnect between the GPUs and the servers. They also tested the ability to recompose the cluster GPUs and storage within the installed servers, as they wanted the ability to add more computing power on occasion to run their in-house AI codes.

By using the GigaIO's composable software, they were able to easily add as many GPUs to a given server, and determine the best ratio of server to GPU for any given application. Link Express provided the client organization with flexibility and superior performance as the applications – and entire system – operated with ease.

Established by industry visionaries with years of domain expertise in networking, data centers, high-performance computing, open source, and infrastructure management, GigaIO is now ready to help customers solve the processing power/interconnect bottleneck with the introduction of its high-performance interconnect fabric. CR