




WHITE PAPER

---

## The GigalIO™ FabreX™ Network

### New Frontiers in Networking For Big Data

---



**CONTENTS**

**BIG DATA AND ORGANIZATIONAL SUCCESS ..... 3**  
**GIGAIO FABREX ..... 4**  
**THE GIGAIO FABREX OPPORTUNITY ..... 6**  
**ABOUT GIGAIO ..... 7**

**This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.**

© GigalO Networks, Inc. All rights reserved. GigalO and its affiliates cannot be responsible for errors or omissions in typography or photography. GigalO, the GigalO logo, and FabreX are trademarks of GigalO Networks Inc. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. GigalO disclaims proprietary interest in the marks and names of others.

March 2019| Rev 1.0

## Big Data and Organizational Success

The opportunity of big data has arrived. The concept of *big data* was iconically defined by Gartner as high-volume, high-velocity and high-variety information assets that demand cost-effective, innovative forms of information processing to facilitate enhanced insight, decision making, and process automation.<sup>1</sup> How organizations manage these torrential data flows with innovative new business processes will likely determine their persistence and profitability in the years ahead. Importantly, organizations will rise and fall based on their capability to transform this tsunami of data into meaningful business action and their IT systems' ability to adapt to the radically altered global data landscape.

*Big data volume* refers to the nonstop, sheer quantity of data being generated every day. For example, as of 12/31/2018, Facebook had more than 2.3 billion monthly active users, and Facebook users uploaded more than 300 million photos per day.<sup>2</sup> Every 60 seconds there are more than 300,000 Facebook status updates. In that same 1-minute time span, there are approximately 450,000 Twitter tweets, 65,000 Instagram posts, and 29 million WhatsApp messages sent.<sup>3</sup> Yet the universe of big data is about much more than social media interactions. The Internet of Things (IoT), referring to the global network of connected devices, includes an ever increasing number and type of sensors such as home appliances, private and commercial vehicles, security cameras, agricultural crop sensors, and sensors deployed across entire cities for healthcare, energy, and national security industries. In 2017, Gartner predicted there would be 20 billion IoT units installed worldwide in 2020.<sup>4</sup>

*Big data velocity* refers to the speed at which data are created and the speed at which data arrive at collection points. For example, a raw Twitter feed is referred to as a “firehose”, owing to the massive stream of Twitter posts being created in real time. Similarly, IoT sensor data are generated continuously at a near-constant rate. As well, big data streams include commercial transactions with high refresh rates. Some data can be analyzed later offline, other data needs to be analyzed in real-time. The sheer quantity of global commercial transactions requires a shift from batch processing to real-time streaming. Overall, the velocity of big data streams necessitates the computational capability of rapidly adaptable real-time analytics.

---

<sup>1</sup>Gartner IT Glossary: Big Data <https://www.gartner.com/it-glossary/big-data/>

<sup>2</sup>Zephoris: The Top 20 Valuable Facebook Statistics – Updated January 2019  
<https://zephoris.com/top-15-valuable-facebook-statistics/>

<sup>3</sup>Osman M: 28 Powerful Facebook Stats Your Brand Can't Ignore in 2018 – 2/15/2018  
<https://sproutsocial.com/insights/facebook-stats-for-marketers>

<sup>4</sup>Tung L: IoT devices will outnumber the world's population this year for the first time – 2/7/2017  
<https://www.zdnet.com/article/iot-devices-will-outnumber-the-worlds-population-this-year-for-the-first-time/>

*Big data variety* describes the numerous types and formats by which these data streams are represented. Big data may be structured as in traditional databases, semi-structured as in XML data, and unstructured as in text, video, audio files, radar and sonar data, time and location data from mobile applications, and photos and other images, including satellite images.

Big data streams provide the raw materials for breakthrough artificial intelligence (AI), machine learning (ML), and deep learning (DL) applications, which are fueling demand for fundamental change in the creation and utilization of compute and storage clusters. But the rapid proliferation of specialized compute accelerators such as graphics processing units (GPUs), field-programmable gate arrays (FPGAs), and custom application-specific integrated circuits are creating configuration problems and bottlenecks for interconnect systems. Traditional networks were not designed to handle the performance requirements of big data workloads and specialized devices.

In order to derive meaning from big data, via implementing the capabilities of big data analytics, and create new opportunities and new value, organizations must find a way to obtain radically increased overall system throughput. The GigaIO FabreX network makes such accelerated functionalities possible by achieving exceptionally low latency and high-bandwidth performance across an organization's entire network. GigaIO FabreX equips organizations to put big data to work by disaggregating legacy infrastructure, enabling dynamic pooling of resources for composability, and facilitating system adaptation to the workload (rather than vice versa) — essentially making it possible for organizations to scale up and scale out their computing resources. GigaIO FabreX enables computationally intense advanced scale computing applications and facilitates growth for enterprise, cloud, and edge data centers, while simultaneously keeping down capital expenditures (Capex) and operating expenses (Opex). GigaIO FabreX directly addresses problems of latency and bandwidth, and emphasizes organizational attention on artificial intelligence technologies, accelerators, and the cornucopia of benefits derived from big data analytics. Enterprise data centers involving telecommunications, transportation, utilities, and edge computing applications such as autonomous vehicles and urban smart grids for managing traffic flows are all facilitated by GigaIO FabreX.

## GigaIO FabreX

GigaIO FabreX is a fundamentally new network architecture that integrates computing, storage, and other communication I/O into a single-system cluster network, using industry standard PCIe (peripheral component interconnect express) technology. GigaIO FabreX enables server-to-server communication across PCIe and makes cluster-scale networking possible. The result is direct memory access by an individual server to system memories of all other servers in the cluster, creating the industry's first in-memory network. The new GigaIO FabreX architecture

enables a hyper-performance network with a unified, software-defined, composable infrastructure. GigaIO FabreX provides multiple benefits for composable hardware infrastructure, whether in the cloud or in private data centers —

- Native PCIe communication between hundreds of mixed processing units, minimizing latency by eliminating the entire translation layer required by other networks
- Direct memory access and sharing of all connected processors and memory with point-to-point connections between any two devices
- Reduced Opex and Capex by avoiding unnecessary duplication of hardware resources
- Establishing roadmaps for end-to-end data transfer improvements across the entire network

GigaIO FabreX achieves industry-leading latency (under 200 nanoseconds) and bandwidth (128 moving to 256 gigabytes/second). These specifications result in faster time to solution, less equipment and reduced power outlays, higher utilization, and easy incorporation of new technology, that is, no forklift upgrades. GigaIO FabreX is 100% compliant with PCIe-SIG standards, with the greatest range of support, low-risk adoption, easy integration, and long life. GigaIO FabreX software is open source all the way up the stack with Redfish application programming interfaces (APIs), with easy integration of existing management tools and software-defined networking and software-designed infrastructure (SDN/SDI) tools. Overall, GigaIO FabreX provides the optimal solution for rack scale computing, disaggregation, and edge deployments.

GigaIO FabreX provides unparalleled and unlimited flexibility, establishing open-ended communication among FPGAs, nonvolatile memory express over fabrics (NVMe-oF), message passing interface (MPI), and GPUDirect Remote Direct Memory Access (GPU Direct RDMA) modules, as well as Linux peer-to-peer configurations. Implementing GigaIO FabreX achieves ultra-high performance networks with memory semantics capability, as well as super-server creation. Additionally, GigaIO FabreX provides high-availability capabilities including PCIe hotplug. From the perspective of topology, GigaIO FabreX makes possible multiple options, including fat tree with spine and leaf switches, 2D torus, dragonfly, slim fly, and ring switch configurations.

In terms of scaling out with GigaIO FabreX, servers can now achieve custom reconfigurations, making any network-connected resource available to a specific server, including super-server configurations. When organizations need to compose resources to fit new applications, the GigaIO FabreX configurable fabric manager makes the necessary connections. Importantly, any resource on the network can be reassigned to a new server, with no physical hardware changes.

GigalO FabreX, making a true software-defined infrastructure that dynamically assigns resources to match changing workloads.

Utilizing PCIe tree-based input/output (I/O) topologies, GigalO FabreX enables cascading of peer-to-peer bridges that fan out to expand the I/O bus, increasing connectivity and network flexibility. Thus, as compared with outmoded network designs with restrictions of one I/O device per server, cascading expands the number of I/O resources available to any server on the network. Additionally, the GigalO FabreX cascading capability enables partitioning of ports, increasing or decreasing the number of lanes, i.e., data rate, available as needed. With cascading, I/O resources can be dynamically shared among servers with a high degree of granularity. The switch establishes software control over the entire configuration, enabling the user to custom design the number of servers and the number of I/O devices employed for a specific application, as well as the number of lanes within a port. Via GigalO FabreX Composition “sharing”, users can dynamically reconfigure, making possible an unlimited scope of change with dramatically increased system utility and lower Capex and Opex.

Additionally, GigalO FabreX network mode enables memory transactions from any server to any other server on the network. Via nontransparent bridging and address resolution protocol, a selected server allocates a block of memory to a remote server. The remote server then performs read/write operations, achieving application-specific information transfer across the entire network of nontransparent servers. GigalO FabreX cascading connects multiple switches that are in network mode, establishing an interswitch connection topology. With cascading, any server can read/write to any other server. Thus, GigalO FabreX cascading makes possible scaling up the number of servers and ports participating in network mode via sharing memory and read/write operations across the entire memory space.

Disaggregated configurations represent the new paradigm of resource-centric network architectures, replacing the old server-centric paradigm, where configurations were rigid. In outdated server-centric architectures, the big data analytics capabilities required for AI and ML/DL encounter high latency and other bandwidth constraints. In disaggregated resource-centric configurations facilitated via GigalO FabreX, the CPUs, GPUs/accelerators, application specific offloads, computational storage arrays, NVMe storage, and application servers all communicate and share data via the network at much reduced latency and higher bandwidth. Resource-centric architectures are open, flexible, manageable, dynamically configurable, and economical.

## The GigalO FabreX Opportunity

With low latency of under 200 nanoseconds and high bandwidth of 128 moving to 256 gigabytes/second, GigalO FabreX directly addresses the core problems of the rack interconnect i.e., Ethernet or InfiniBand latency of 2500-20,000 nanoseconds and bandwidth of 10-100

gigabytes/second. Overall, the GigaIO FabreX cascading capability delivers a substantial increase in resource utilization with flexible, software-defined, on-the-fly configurability for operators. With GigaIO FabreX enabling scale up and scale out, organizations can fundamentally redesign network infrastructure used in advanced scale computing and edge computing applications, including AI, ML, and DL. GigaIO FabreX assists organizations in making possible the business opportunities presented by big data, while simultaneously reducing Capex and Opex.

## About GigaIO

GigaIO has invented the first truly composable cloud-class software-defined infrastructure network, empowering users to accelerate workloads on-demand, using industry-standard PCI Express technology. The company's patented network technology optimizes cluster and rack system performance, and greatly reduces total cost of ownership. With the innovative GigaIO FabreX™ open architecture, data centers can scale up or scale out the performance of their systems, enabling their existing investment to flex as workloads and business change over time. For more information, contact the GigaIO team at [info@gigaio.com](mailto:info@gigaio.com) or [schedule a demo](#).