Quick Take

CXL and Gen-Z Consortiums Combine Forces

Michael Feldman, Alex Norton, and Earl Joseph
July 2020

HYPERION RESEARCH OPINION

The recent decision by the Compute Express Link (CXL) and Gen-Z consortiums to collaborate on their respective technologies is a positive development that is likely to accelerate the acceptance of both memory fabric standards. Although not strictly competing against one another, prior to the agreement there was no assurance that Gen-Z and CXL supported devices would be compatible on future hardware. With the Memorandum of Understanding (MOU) recently agreed to by the two organizations, a path forward for interoperability and other collaborations now exists.

This is a potentially important development for high performance computing, which is increasingly limited by memory performance and capacity. Both CXL and Gen-Z seek to improve the speed and flexibility of the composable datacenter hardware, both within the server and at the rack and row level. With the advance of applications in AI/machine learning, graph analytics, and other big data workloads, demand is growing for a more disaggregated infrastructure, as well as for a more memory-centric approach to computing. Much of this shift centers on the ability to pool memory, storage, and other devices and make them universally accessible within a server or across an entire cluster.

The CXL/Gen-Z collaboration harkens back to the 1999 merger of Future I/O and Next Generation I/O, which led to the creation of InfiniBand. At the time, they were competing fabric specifications: Future I/O was backed by IBM, Hewlett-Packard, and Compaq, while Next Generation I/O was being promoted by Intel, Microsoft, and Sun Microsystems. While the original intention of InfiniBand was to create a universal fabric that would usurp Ethernet and PCI, its current use in HPC and for certain hyperscale applications demonstrates how new standards can be established in the midst of legacy technologies with entrenched markets.

Moreover, such collaborations demonstrate how a critical mass of vendors can wield their influence in the pursuit of standardization. In this case, CXL members represent nearly all of the largest processor, memory, and network providers, as well as a number of the most prominent hyperscale companies and server OEMs. The Gen-Z consortium has a similar membership, much of which overlaps with that of the CXL organization.

SITUATION OVERVIEW

In April 2020, the CXL and Gen-Z consortiums announced an MOU to “promote interoperability between the technologies, while leveraging and further developing complementary capabilities of each technology.” To accomplish that, the two organizations have committed to forming common workgroups that will define bridging technologies between the protocols.

Both CXL and Gen-Z are open-source fabric technologies that use memory semantic protocols to provide high performance interconnectivity among processors, accelerators, and memory expansion devices. CXL and Gen-Z are not mutually exclusive standards, as CXL is principally designed for memory within a server, while Gen-Z works mainly at the level of a cluster.

CXL is specifically concerned with making memory globally accessible and cache coherent within a node, whether it is system DRAM or memory associated with an attached device. The protocol rides atop PCIe, which is a fabric for locally attached devices that does not support cache coherency. Competing fabrics with the similar capability to CXL include IBM's Coherent Accelerator Interface (CAPI), Xilinx's Cache Coherence Interconnect for Accelerators (CCIX), and the AMD's Infinity Fabric. With Intel's backing, CXL has been able to attract more adherents than competing standards. With CXL as an accepted standard, component providers will be designing devices that can talk with hardware manufactured by other providers.

As such, CXL could furnish a standard bus/interconnect for heterogeneous nodes, which could help expand the market for AI-powered servers. At the same time, it could reduce the cost and increase the flexibility of HPC servers that rely on locally attached accelerators and non-volatile memory. Niche applications in the hyperscale and enterprise environments that employ heterogeneous compute within the same node could also benefit.

Gen-Z offers much the same memory-centric approach as CXL, but from a multi-node perspective. Unlike CXL, Gen-Z supports switched topologies, enabling it to extend memory accessibility access across a server rack or row. This has particular relevance in configurations that use disaggregation to pool memory, memory class storage, or accelerations on specialized servers.

The memory-centric approach of Gen-Z distinguishes itself from older system interconnects such as Ethernet and InfiniBand, which are compute-centric and, despite the incorporation of technologies like RDMA, provide limited support for pooling memory and devices. Gen-Z also promises higher bandwidth and lower latency to deliver greater performance than these legacy interconnects, even as it supports interoperability with them.

In essence, Gen-Z would build on the memory-centric approach of CXL and offer a standard way to disaggregate resources. The AI and HPC markets stand to benefit from such standardization, but the technology would also extend to less performance-demanding datacenter applications as well, such as web services and conventional enterprise computing.

FUTURE OUTLOOK

In the short term, we expect the collaboration of the CXL and Gen-Z consortiums will lead to an expansion of the standards to support interoperability between the two technologies. Building processors, accelerators, and memory components that are interoperable on the fabric is a winning formula for both vendors and users.

This is especially relevant for OEMs like HPE and IBM (both are members of each consortium) that are committed to building scalable memory-centric systems for the most performance-demanding applications. The benefits will also be conferred to second-tier OEMs and ODMs.

In the longer term, the collaboration could result in the merging of the two standards, resulting in a single open-source fabric technology. Such a development would further simplify the task of the

component and server providers and accelerate acceptance of the technologies. The likelihood of such a merger will depend upon the degree to which interoperability can be achieved, as well as initial acceptance by a critical mass of vendors and users.

About Hyperion Research, LLC

Hyperion Research provides data driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user and vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors, and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com
