Technology Spotlight

Gen-Z: A Memory-Centric Interconnect Fabric

Alex Norton, HPC Research Analyst
February 2019

HYPERION RESEARCH OPINION

The Gen-Z Consortium, launched in October 2016, brought together multiple technology vendors to create a new open system interconnect standard to work with the growing number of new processors and memory options available on the market now, as well as in the future:

- This fabric is designed to provide a new level of connectedness between major types of processors, including CPU, GPU, DSP, FPGA and other accelerators, supporting a larger pool of shared data regardless of the manufacturer or architecture of the processor.
- Processor access to memory is also independent of the type of memory hardware. It's targeted to support both DRAM and storage-class memory (such as flash).
- The Gen-Z fabric design allows for fast, low latency communication of large amounts of data seamlessly.

THE CONSORTIUM

The Gen-Z Consortium arose in reaction to the growing set of processor and memory options that were becoming increasingly available to system designers looking to build a processor/memory infrastructure optimally suited to their target workloads. Currently, the consortium consists of over 50 member companies, with the leadership coming primarily from processor makers AMD and Xilinx, server OEMs Cray, Dell, EMC and HPE, as well as companies such as ARM and Mellanox.

- A few major technology companies, most notably Intel and NVIDIA, are not members despite being major contributors to the HPC as well as the enterprise computing space.

The collective member companies cover most areas of the HPC food chain from systems, to storage companies to interconnect companies, ISVs and end users. The goal of the consortium is to "Create an open, industry standard for a high speed, low latency, scalable, memory centric fabric." [1]

- At SC18, the consortium presented a demonstration that connected two servers from different vendors through the Gen-Z interconnect fabric with a shared pool of data.

THE TECHNOLOGY

Gen-Z's interconnect fabric design provides high speed, low latency, and is equipped to offer high reliability while giving a wide variety of processors access to shared or mixed memory hardware. To

[1] https://www.opencompute.org/files/OCP-GenZ-March-2018-final.pdf

date, four 1.0 specifications have been released, and the consortium is targeting the next two years to start implementing the technology in products available to end users.

The fabric was designed specifically to target key bottlenecks within the workload stream, particularly the latency and throughput issues associated with current standards. The Gen-Z fabric has ultra-low latency messaging, at less than 250 nanoseconds for one-way transit.

For throughput, Gen-Z claims their fabric will be the equivalent of a multi-lane highway, compared with a traditional network, which they say is more akin to a simple two-lane road. Through low latency switches, many channels, universal protocols, and hardware-driven multi-link and multi-path technology, Gen-Z technology updates the current fabric to a much higher performance.

With the growing number of architectures and processors coming to the market, the Gen-Z fabric aspires to be the universal connector system for memory, I/O. storage, and more. Solutions are being developed for both copper and optical cable schemes. The Gen-Z fabric is designed to support a wide range of component types, as well as other existing interconnect technologies, such as PCIe.

FIGURE 1

Gen-Z Connectedness

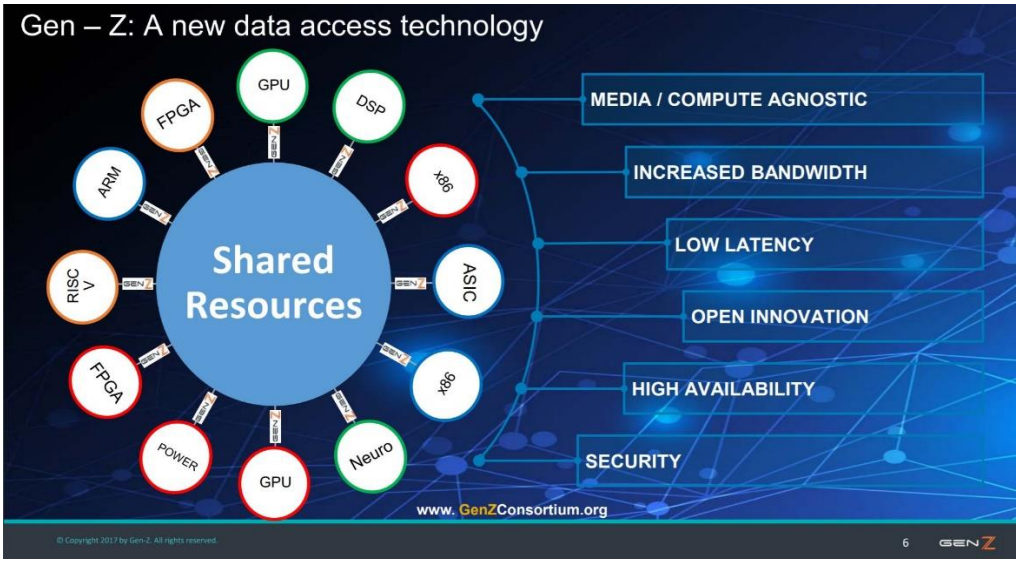

Source: Hyperion Research, 2019[2]

Figure 1 shows the wide variety of processors types that can access the same shared memory through the Gen-Z interconnect fabric, according the consortium. This is made possible through the Gen-Z point to point or switch fabric.

[2] https://www.opencompute.org/files/OCP-GenZ-March-2018-final.pdf

Furthermore, this list of different processor types continues to grow as companies including Google bring TPU technology to the table, AWS implements its Inferentia chip, and several other processors are being developed in small startups in the United States and abroad.

FUTURE OUTLOOK

As is evident from the number of companies that have joined the Gen-Z consortium, there is a strong perceived need for an interconnect fabric that allows for shared memory and composable storage, as well as a universal connector. As systems move toward being increasingly heterogeneous, interconnect fabrics will increasingly need to support fast and flexible connections between existing and emerging technologies.

Additionally, as the latency issues continue to constrain certain workloads, a technology like this that promises low latency with access to a large number of processors as well as a shared memory pool could enable a growing class of new applications to be tackled in a more efficient and timely manner.

Despite the promise of Gen-Z, no company yet has announced a serious product line incorporating the Gen-Z technology. Vendors are looking at the technology as a valid alternative to the current standards, but many companies are still in the exploration phase and are not yet willing to commit to the development of marketable products:

- HPE is probably the farthest along in developing products that utilize the Gen-Z fabric.
- Dell is chairing the efforts, so is expected to offer Gen-Z based products.
- Demonstrations at SC18 as well as other conferences have showcased the capability of Gen-Z, albeit at a relatively small scale, and companies like HPE, Dell, and AMD have already demonstrated the Gen-Z fabric will work with their products.

Gen-Z is beginning to gain some traction, but much of that traction is long-term interest in an alternative to the current standards, especially with the expansion of new HPC workloads in the AI space. However, as Gen-Z has not been installed in a system with benchmarks on real workloads at a large scale, the rush to productize and deploy Gen-Z has not come yet. The next year is crucial for Gen-Z, as they plan to send early adopters their technology to install in real systems, and the buzz continues to grow. Now, Gen-Z needs to prove they can tackle some of the emerging AI workloads with their interconnect fabric.

In summary, participation in the consortium by so many vendors is a tacit recognition of the need for HPC interconnects to embrace the trend toward heterogeneous workloads and system elements, by providing higher performance capabilities. Gen-Z is a bit ahead of the market at the moment, but it has the potential to help guide the HPC community into a higher performing, more heterogeneous future.

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 & 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
