Introducing Compute Express Link™ (CXL™) 3.1: Significant Improvements in Fabric Connectivity, Memory RAS, Security and more!

By Debendra Das Sharma and Mahesh Wagh, CXL Consortium Technical Task Force Co-Chairs

The CXL™ Consortium released the CXL 3.1 specification in November 2023, which introduced enhancements to fabric capability, fabric manager API definition for PBR switch, inter-host communication using global integrated memory (GIM), trusted execution environment (TEE) security protocol, and enhancements to memory expander (e.g., up to 34-bit meta data, RAS capability enhancements). These changes will result in CXL being deployed as a composable fabric for disaggregation, pooling, and distributed processing with high reliability and security.

Fabric Enhancements with CXL 3.1:

CXL 3.1 defines GIM for communication across multiple hosts, as shown in Figure 1. The GIM domains are protected. It comprehends the fabric decode and port based routing (PBR) requirements, including the fabric manager API definition for PBR switches. Direct peer-to-peer .mem support through PBR switches is supported through symmetric link layer definition. Cross-domain peer DMA traffic is also supported. CXL 3.1 also enables direct caching of memory (HDM-DB) for an accelerator.

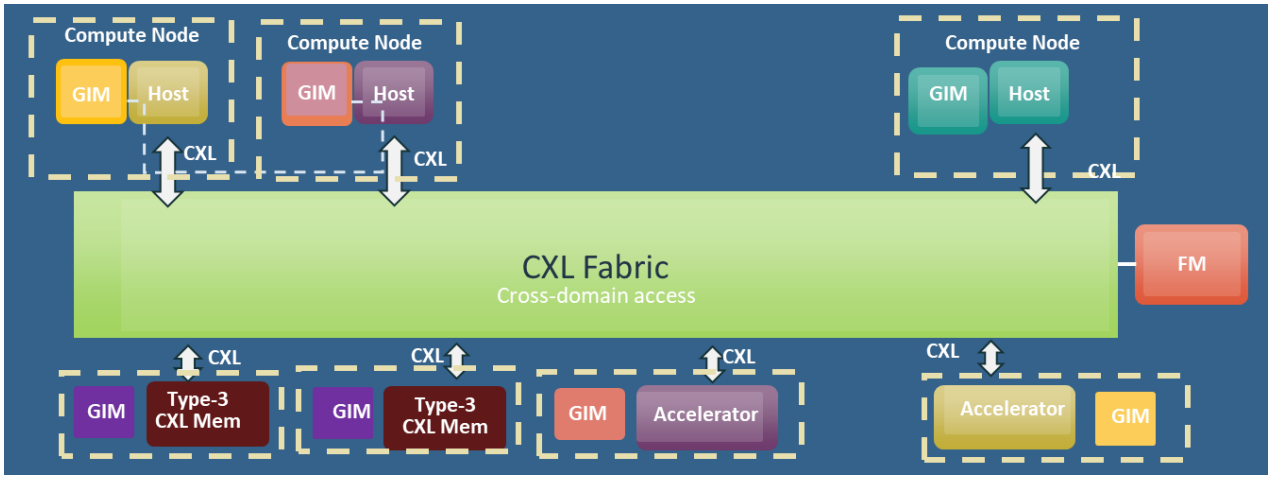

Figure 1: Fabric Enhancements with CXL 3.1

Each node has its view of the address map that needs to comprehend GIM, as shown in Figure 2. Hosts rely on the switches to do the appropriate mapping. The GIM address range is mapped within the fabric address space in the host physical address. This mechanism is used to communicate between hosts GIM, each with its own address map.

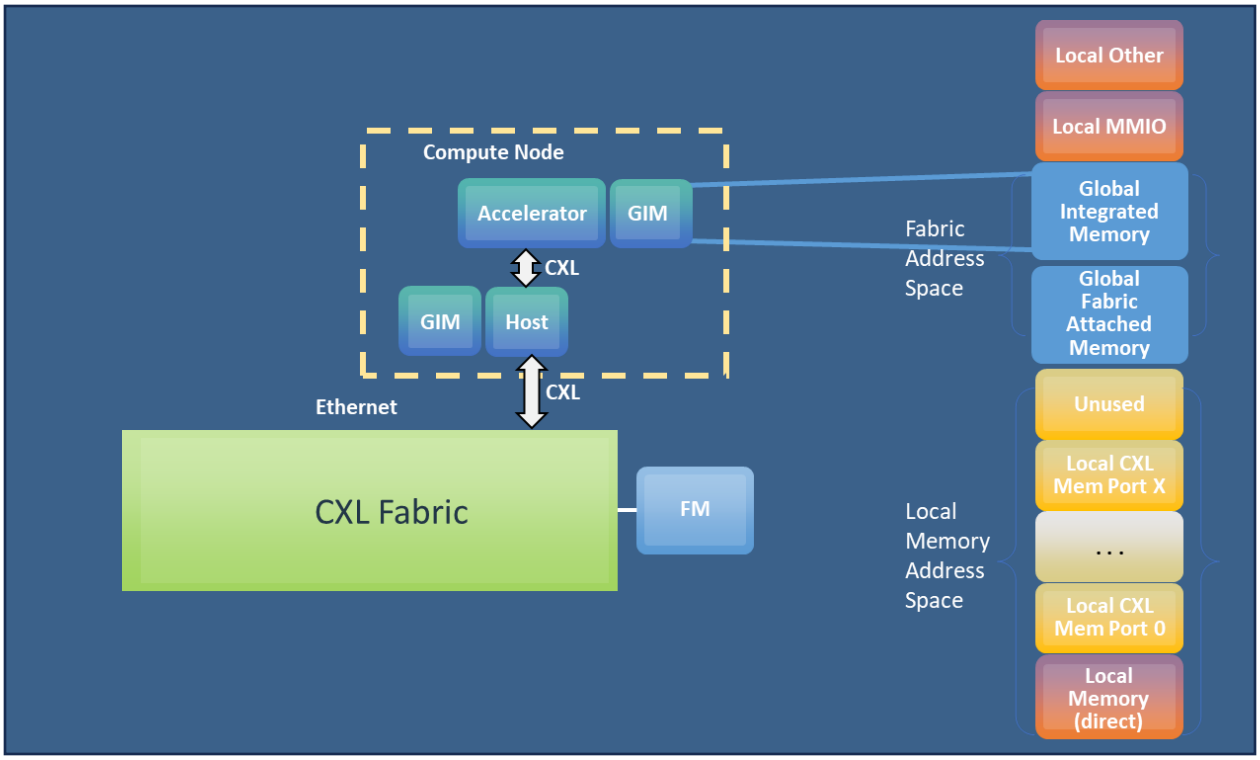

Figure 2: Address map: Node view of memory, including GIM

Accelerators can directly access (peer-to-peer, without going through host) CXL.Mem, using PBR transactions, as shown in Figure 3. Data consistency is ensured by the Type-3 devices using the back-invalidate semantics, if needed, which was introduced in CXL 3.0.

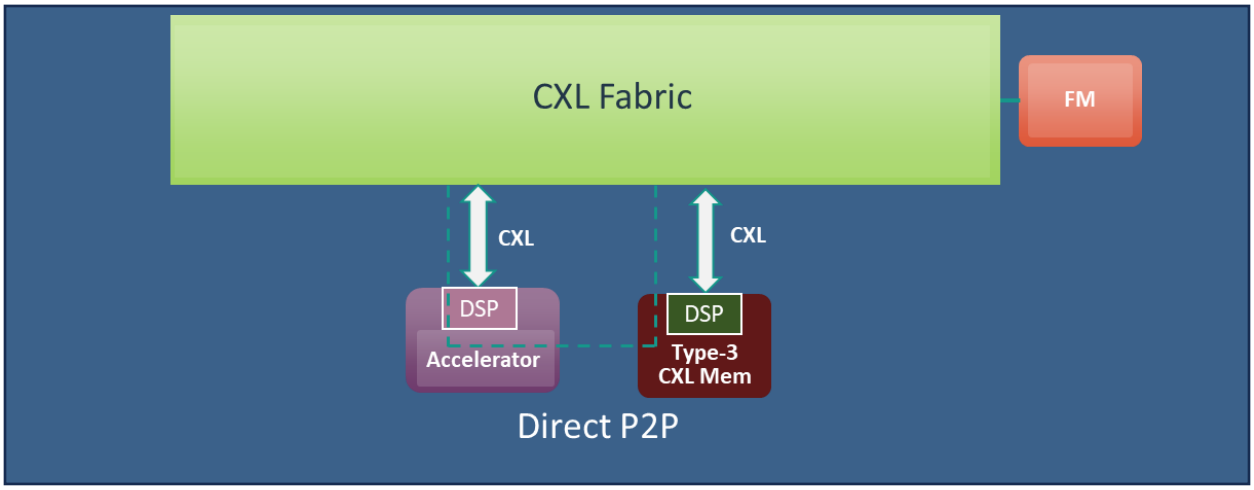

Figure 3: Direct P2P access of Accelerator to Type-3 Memory connected directly across a PBR switch

Security Enhancements with CXL 3.1:

CXL 3.1 enhances the security architecture to enable confidential compute capabilities to direct attached CXL memory expander devices. The newly defined CXL TEE (Trusted Execution Environment) Security Protocol (CXL-TSP), defines mechanisms to include direct attached CXL Memory Devices within the TVM trust boundary for confidential compute use cases. The reference architecture covers the security requirements and behaviors that are needed to support confidential computing use cases and covers the architectural scope, detecting TSP support, CMA/SPDM, attestation and authentication, memory encryption, transport security, access control, configuration, and Dynamic Capacity.

Memory Expander Enhancements with CXL 3.1:

CXL 3.1 enhances the 2-bit metadata to describe the state of the cache line (modified, shared, exclusive, or invalid) with an additional 32 bits for metadata. This metadata may be used for access control, data type tagging, memory tiering usage, etc. and this is aligned with the proposed DDR6 feature that adds 16-32 bits of metadata for each 64B of data. This extended metadata may be used with host coherent as well as device coherent memory. Further, CXL 3.1 has made RAS enhancements such as additional information on correctable error limits, source of errors including transaction during error, memory sparing, scrubbing, and capacity or performance degradation.

Interested to contribute?

Join the CXL Consortium to participate in the technical working groups and influence the direction of the CXL specification. Learn more about the CXL Consortium membership here.