Introducing Compute Express Link® (CXL®) 4.0: Significant Improvements in Bandwidth, Connectivity, Memory Maintenance, and Security

By Tony Benavides and Mahesh Wagh, CXL Consortium Technical Task Force Co-Chairs

The release of the Compute Express Link® (CXL®) 4.0 specification marks a pivotal advancement in high-speed interconnect technology. Building on the foundation of PCIe® 7.0, CXL 4.0 delivers significant performance and scalability for modern computing environments. CXL 4.0 doubles the data rate to 128 GT/s and introduces architectural innovations such as Bundled Ports while maintaining full backwards compatibility. These improvements are complemented with an enhanced robust memory maintenance capability framework, positioning CXL 4.0 for the next-generation data-intensive applications. From AI and machine learning to cloud and hyperscale deployments, CXL 4.0 continues to deliver double the throughput within the same low latency and power efficiency over CXL 3.0, with improved flexibility in system architecture and topology.

Increased Bandwidth and Connectivity

CXL 4.0 delivers the performance needed for next-generation workloads. It defines a major step forward in high-speed interconnect technology, doubling the data rate to 128 GT/s. This leap is achieved by doubling the Nyquist frequency while preserving PAM4 signaling, Flit-based structure with FEC and CRC from CXL 3.0, ensuring continuity, while delivering power-efficient performance with very high reliability (as measured by Failure in Time or FIT, of lower than 10^{-3} failures in a billion hours).

The specification introduces the concept of native x2 width, to support increased fan-out in the platform, and support for up to four retimers for increased channel reach beyond CXL 3.0. These features strengthen signal integrity across extended topologies. It also enables logical aggregation of multiple upstream ports, reducing latency and improving throughput in multi-host environments. These capabilities are critical for scaling performance in complex, heterogeneous systems.

Bundled Ports

The 4.0 specification introduces Bundled Ports, an architectural enhancement that aggregates multiple physical CXL device ports into a single logical entity. This approach enables a device to connect to one or more host Root Port (RPs) or switch upstream ports while still maintaining backward compatibility with existing software models. By combining links, Bundled Ports dramatically increase bandwidth and connectivity without altering the enumeration model. Importantly, bundling enables internal designs to deliver doubling of bandwidth easily by simply doubling the stacks, vs either doubling the data path widths or the internal frequency. Each bundle contains at least one full-capability port and may include additional Streamlined Ports optimized for data movement, allowing flexible scaling of performance.

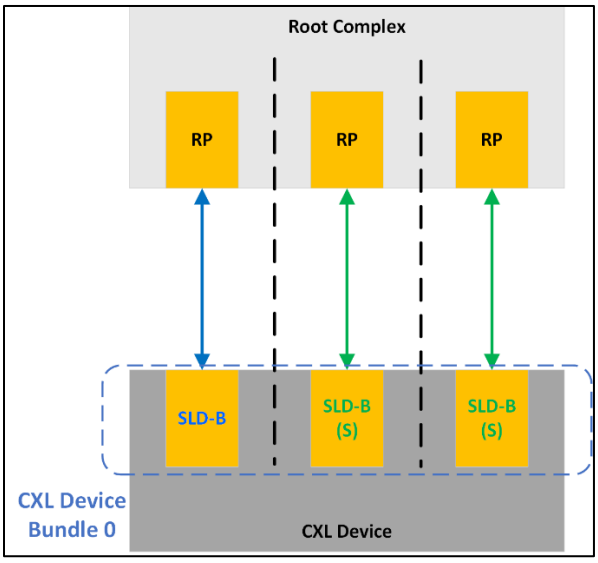

Figure 1: Bundled Port Between Host and Type 1/2 Device with CXL 4.0. Assuming a x16 link at 128 GT/s, the figure depicts 768 GB/s BW in each direction allowing for a total of 1.536 TB/s total BW between the device and CPU.

A Root Port (RP) is the host's PCIe/CXL port that anchors the hierarchy and, in CXL 4.0, can participate in a Bundled Port attachment. A Single Logical Device (SLD) represents a CXL device exposing exactly one logical device per port, binding like a standard endpoint without LD-ID multiplexing. When an SLD implements a Bundled Port (SLD-B), multiple device ports are aggregated into one logical port, allowing the system to see a single device while bandwidth scales across the bundled links.

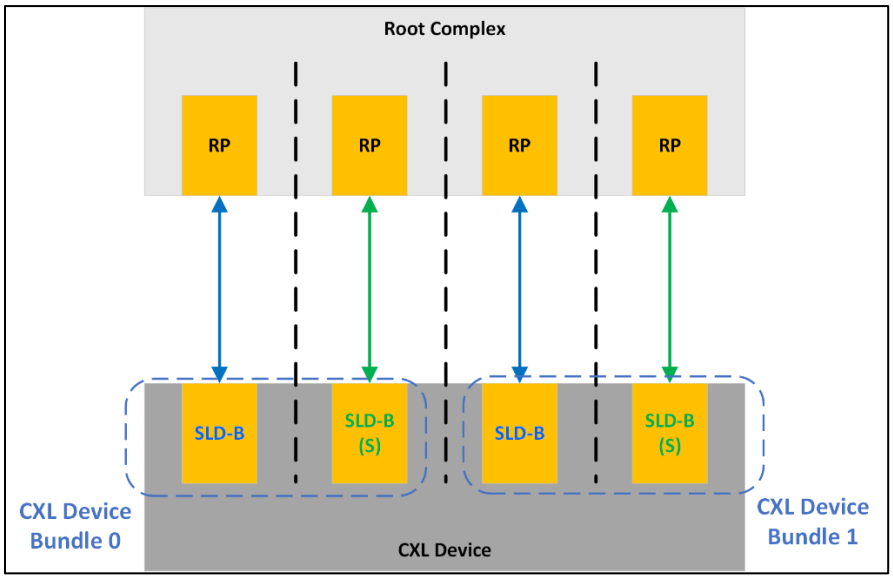

Figure 2: Multiple Bundled Ports Between Host and Single Device with CXL 4.0

This approach preserves software simplicity while enabling high-performance, multi-link connectivity for next-generation CXL deployments. Bundled Ports can be optimized for 256B Flit Mode, eliminating the overhead of legacy 68B Flit Mode and reducing hardware complexity and cost. Their primary purpose is to expand the data path and maximize throughput, making them ideal for modern deployments where efficiency and scalability are key. At least one of the bundled ports must be a 68B capable FLIT to satisfy backwards compatibility.

Memory Maintenance Enhancements

CXL 4.0 enhances memory reliability and serviceability with features designed to improve system resilience and uptime. It achieves that by introducing advanced corrected volatile memory error reporting, providing finer-grained error detection and event generation during patrol scrub cycles. The specification also enables host-initiated Post Package Repair (PPR) operations at device boot, allowing proactive maintenance before launching customer workloads. Additionally, memory sparing operations can now be performed during boot or deferred to the next cycle, offering flexibility for administrators to maintain performance without disrupting system availability.

Interested in contributing?

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.