SMART®
Modular Technologies
A PENGUIN SOLUTIONS® BRAND

White Paper M-WP011

Introduction to Compute Express Link (CXL®) Memory Modules

September 2023

Table of Contents

Introduction

Compute Express Links (CXL®) is a high-speed interconnect standard that enables an efficient link between the CPU and platform subsystems. It builds upon PCI Express® infrastructure, leveraging the PCIe® 5.0 physical and electrical interfaces. CXL encourages heterogeneous and distributed compute architecture by enabling hardware based cache-coherency for all types of compute engines like CPU, GPU, and accelerators xPUs (TPU, DPU, IPU, etc.).

The biggest advantage of CXL comes with the extension of the memory attachment to serial interface. This fills the gap for data-intensive applications as they drive the requirements for high bandwidth and provide low-latency and sharing across multiple devices in a system. Serializing the memory interface also opens up opportunities for connecting memory in different form-factors like Add-in-Cards (AIC) and EDSFF modules, which simplifies system design, and lowers the TCO by allowing memory modules to be mechanically and electrically compatible to SSDs.

The CXL standard specifies three protocols:

- CXL.io: CXL.io is functionally equivalent to the PCIe interface. This sub-protocol is used for configuration, DMA and Interrupt handling between device and Host.
- CXL.mem: CXL.mem enables a host, such as a processor, to access the device-attached memory using load/store commands.
- CXL.cache: CXL.cache specifies semantics and rules for accelerators trying to access directly attached memory connected to the CPU's DDR bus. This enables accelerators to efficiently access and cache host memory for optimized performance.

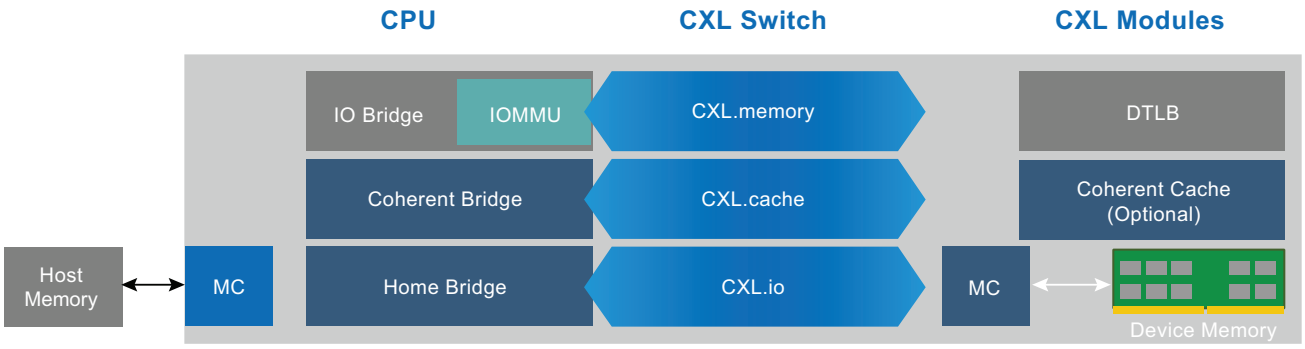

Figure 1: CXL Protocol Overview

Using these three protocols, the CXL consortium defines three types of devices or use cases:

- Type 1 devices or use case: These devices support both CXL.io and CXL.cache sub-protocols, but do not support CXL.mem. Therefore, we can infer that Type 1 CXL devices do not contain any memory which is available for host consumption, for example, Network Interface Cards (NICs).

- Type-2 devices or use case: These devices support all sub-protocols CXL.io + CXL.cache + CXL.mem. A Type-2 device would support an on-board memory which is accessible to the host address-map, and an on-board accelerator or compute function which requires frequent access to memory of other hosts or compute functions.
- Type-3 devices or use-case: These devices support only CXL.io and CXL.mem, and are targeted towards memory capacity and memory bandwidth expansion use-cases.

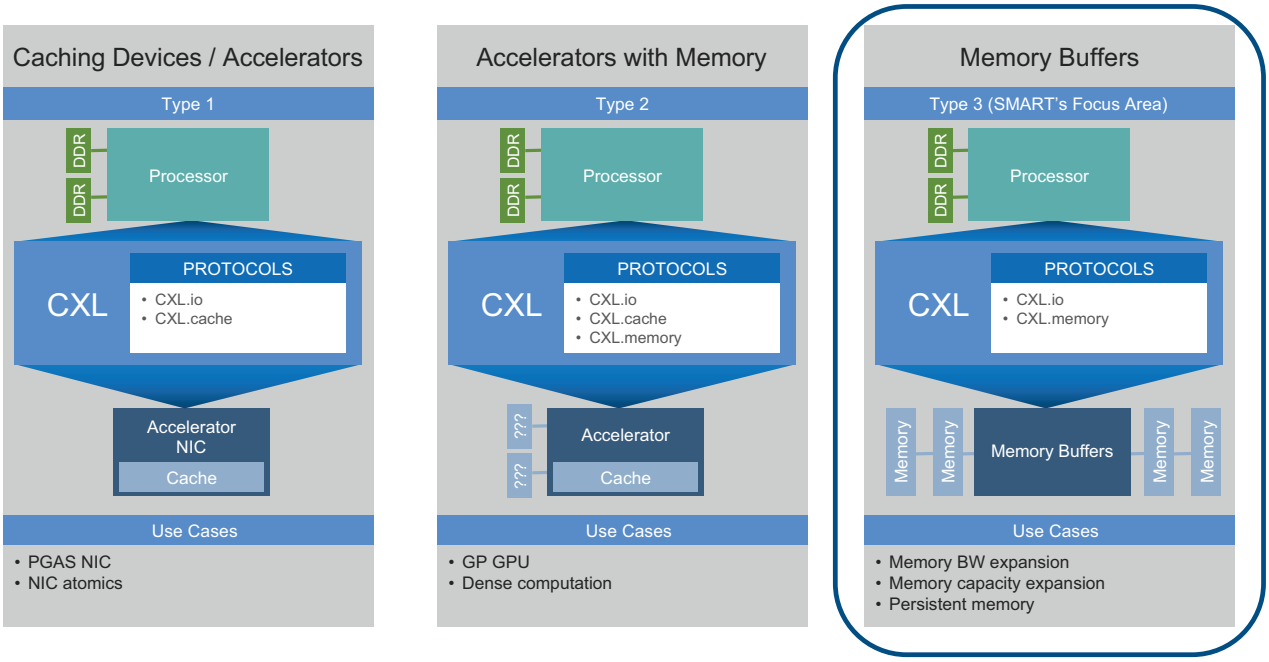

Figure 2: CXL Use cases (source: CXL Consortium 2021)

This paper focuses on Type-3 devices and explores the benefit of this use-case.

SMART Modular Technologies (SMART) Type 3 Memory Modules (CMMs), utilizing the CXL standard, which attach over a CXL link and serve as additional system memory, providing either increased memory bandwidth or increased memory expansion. The following figure shows a system deployment that implements a CXL-memory subsystem to add capacity and bandwidth to the standard DDR memory subsystem, while maintaining data coherency.

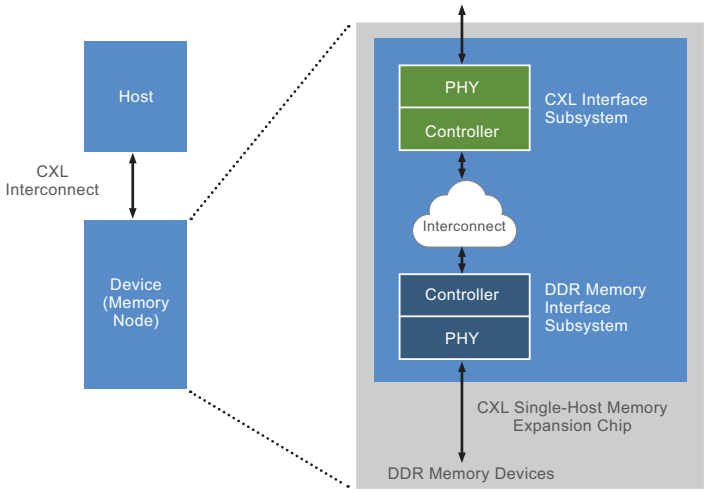

Figure 3: CXL Deployment Example

Disaggregated Memory Implementation

CXL's high-speed serial interconnect standard is targeted towards latency reduction to enable disaggregated memory. Creating shared memory pools with efficient, coherent, low-latency access mechanisms is in line with the goal of heterogeneous resource sharing. Support for Type 3 memory products in CXL provides an opportunity to separate the memory controllers. As data centers interchange with a wide range of use cases, Type 3 memory products allow for the use of different types of media. CXL controllers can be designed to support different media, including DDR4 and DDR5, as well as persistent memory. Each media type would have different performance characteristics. For example, a slower memory tier can be completely isolated from the main tier with no interference to main memory direct-attached DIMMs. Using CXL, developers can dial in memory bandwidth that is ideal for their application, use persistent memory options, and mix-and-match as the application requires.

Memory Pooling

Memory pooling is the concept of allowing multiple hosts to access a common pool of memory. This enables efficient utilization of memory resources as hosts release memory back to the pool once it's no longer needed. Hosts can also dynamically request additional memory, when required, which avoids over-provisioning the memory for workload.

CXL2.0 compliant memory products enable the pooling of memory. Memory systems can be partitioned into Multiple Logical Devices (MLD), each simultaneously accessible by a different host. The diagram below shows an active host (H1) using half the memory in one device (D1) and a quarter in another (D2) based on its workload. Remaining memory capacity in CXL devices D1 and D2 is available for use by other hosts (H2 – H16) in the system.

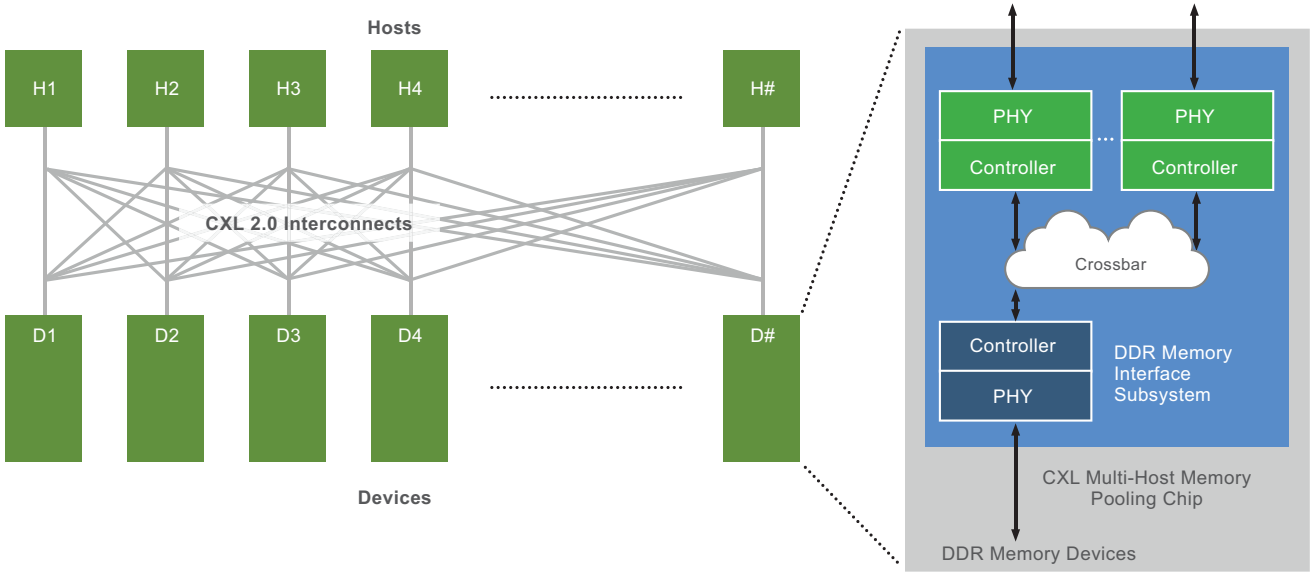

Figure 4: Memory Pooling Diagram (Source:Rambus.com)

Switching

In most data center deployments, owners prefer sharing resources for as close to 100% utilization as possible to dedicate resources where they could be underutilized. The switching implementation makes memory available to multiple hosts on an as-needed basis. Servers, provisioned for nominal workloads, can access the pool as needed for high-capacity workloads.

Integrity and Data Encryption (IDE)

CXL Type 3 memory products have built-in IDE protocols to provide confidentiality, integrity, and replay protection over the entire attack surface. Hardware-implemented IDE engines in the host and the storage devices meet the high-speed data rate requirements of data centers without introducing additional latency. CXL controllers and systems themselves require safeguards against tampering and cyberattacks. A hardware root of trust implemented in the CXL chips can provide this basis for security and support requirements for secure boot and secure firmware download.

CXL Type 3 Memory Products

SMART's memory products, utilizing the CXL standard, are targeted for memory capacity expansion and memory bandwidth expansion in server and storage systems. CXL modules can be dynamically allocated behind the CXL interface for those workloads that need it. SMART is utilizing its experience in enabling new technologies and new interconnect standards to fully support the adoption of CXL memory.

CMM E3.S Module

The E3.S DDR5 CMM incorporates DDR5 DRAM technology in the E3.S form factor and is earmarked for 2U servers. The DDR5 E3.S CMM is 64GB and CXL 2.0 compliant and can be used as a tiered cache memory with a round trip latency of less than 80ns. The throughput rate is 64GB/s with 66% reads to DDR5 over CXL x16. Some of the notable RAS features are data-path integrity, poisoning and error injection, memory ECC, chip-kill and scrubbing. It also enhances memory interleaving, performance tuning and supports low-power modes.

CXL Add in Cards

CXL DDR5 AICs enable the use of standard DDR4 or DDR5 RDIMMs. They provide maximum memory configuration flexibility for low, medium, high-end server and storage SKUs. They are ideal for data center and hyper-scale companies for memory capacity and bandwidth expansion options. Multiple types of AICs are being designed and to align with CXL industry adoption.

SMART's Value Proposition

CXL memory expander devices can be built with media-specific controllers. A system could support a variety of different memory types, including DDR3, DDR4 or DDR5, as well as persistent memory, low-power DRAM and so on, each having a media-specific controller supporting asymmetric or non-deterministic timing and error handling. A slower memory tier can be completely isolated from the main tier, with minimal interference to direct-attached DRAM dual in-line memory modules (DIMMs).

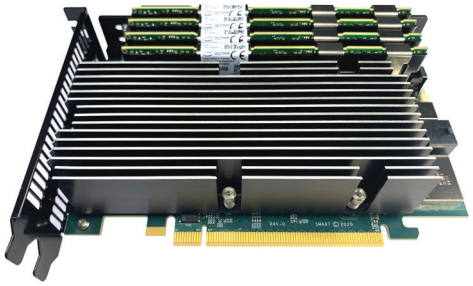

CXL is enabling new form-factors and architectures across Memory and Storage domain. SMART is on the forefront of this technology disruption.

Watch following webinars from SMART to learn more how CXL is changing Memory expansion and Persistent Memory paradigm and introducing new concepts like Computational Memory.

Title	URL
Scaling NVDIMM-N architecture for System Acceleration in DDR5 and CXL Applications	https://youtu.be/pUfCMC_fcFM
Future of Persistent Memory in Form Factors Architectures with CXL	https://youtu.be/dD1O6HmMAYU
EDSFF for Storage, Memory and Acceleration	https://youtu.be/J8FN4SxqQd8
Computational Memory: Moving Compute near Data	https://youtu.be/UJZJVM6CTyM

SMART® Modular Technologies
A PENGUIN SOLUTIONS® BRAND

For more information, please visit: **www.smartm.com**

Headquarters/North America
T: (+1) 800-956-7627 • T: (+1) 510-623-1231
F: (+1) 510-623-1434 • E: info@smartm.com

Latin America
T: (+55) 11 4417-7200 • E: sales.br@smartm.com

EMEA
T: (+44) 0 7826-064-745 • E: sales.euro@smartm.com

Asia/Pacific
T: (+65) 6678-7670 • E: sales.asia@smartm.com

Customer Service
T: (+1) 978-303-8500 • E: customers@smartm.com