SMART®
Modular Technologies
A PENGUIN SOLUTIONS® BRAND

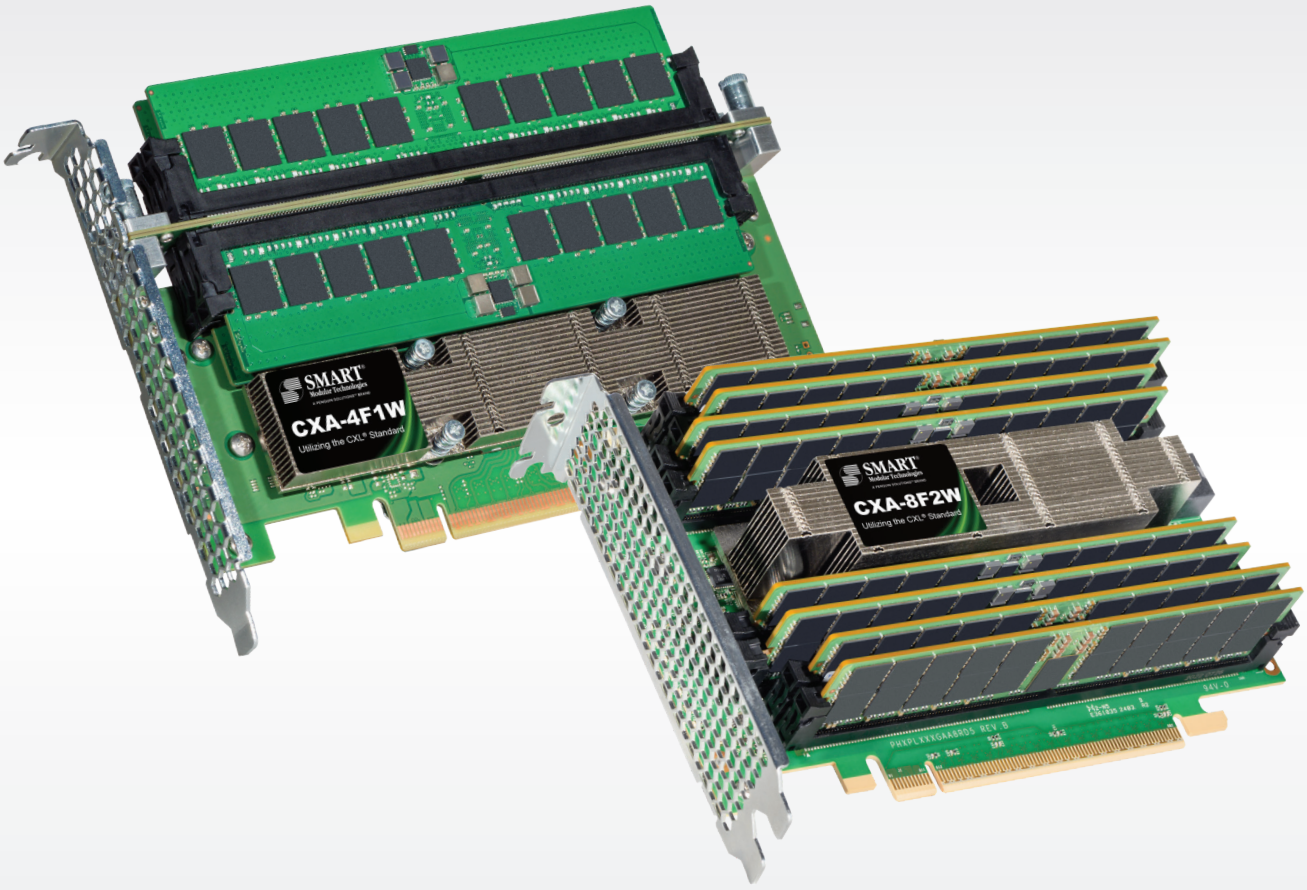

White Paper M-WP014

CXL® for Memory Expansion

April 2024

Table of Contents

Introduction

Compute Express Link or “CXL” is a new interconnect standard that operates on the PCIe physical layer. CXL implements three new protocols, CXL.io, CXL.cache, and CXL.mem. CXL.io is essentially equivalent to the PCIe protocol and is used for device discovery and enumeration. CXL.cache is a cache coherent protocol primarily intended to enable accelerator devices. CXL.mem enables load/store commands best suited for memory expansion devices. CXL has also defined three “types” of devices that each support a subset of these new protocols.

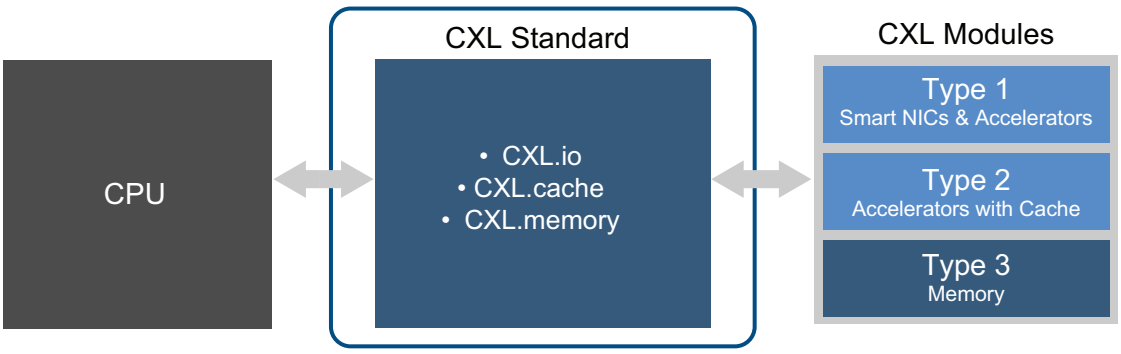

Figure 1 - CXL Protocol and Module Types

CXL Device “Type”	Supported Protocol	Example Use Case
Type 1	CXL.io and CXL.cache	NICs and Accelerators
Type 2	CXL.io, CXL.cache, and CXL.mem	GPU
Type 3	CXL.io and CXL.mem	Memory expander

The CXL standard has been released in three phases, CXL 1.0/1.1 released in 2019, CXL 2.0 released in 2020, and CXL 3.0 released in 2023. Each version implements additional features, but the primary difference is introduced in CXL 3.0 which adds memory sharing and multi-level sharing as key features. CXL 3.0 also uses the PCIe gen 6 interface, while previous versions use PCIe gen 5. As CXL technology evolves, it has the potential to revolutionize memory deployment in the industry.

Memory Expansion

As modern applications such as artificial intelligence (AI), machine learning (ML), in-memory databases, high-resolution image processing and others drive the need for ever expanding amounts of memory, system designers face a difficult challenge. Attaching more memory to a CPU with the traditional DRAM interface requires a large number of pins due to the nature of its parallel interface. CPUs often only have four to six dedicated memory controllers which limits the amount of traditional RDIMMs that can be attached. As the industry moves towards DDR5-6400 speeds, it will likely be necessary to limit memory to one DIMM per channel, which further restricts the amount of memory per processor.

With the introduction of CXL, we now have the opportunity to connect high speed, low latency memory to the CPU with an alternative interface. CXL capable PCIe interfaces can now be used to attach CXL “Type 3” memory devices, instead of storage or other peripherals. CXL essentially enables PCIe for DRAM devices, which greatly expands the amount of memory per CPU.

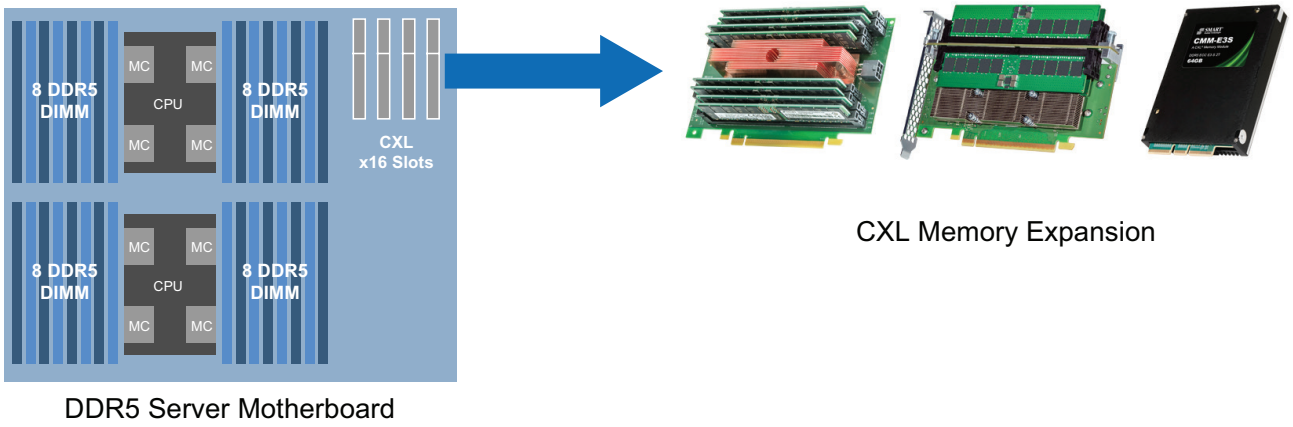

Figure 2 - CXL Memory Expansion Devices

Memory Expansion Applications

As technology advances, the demand for larger and larger amounts of system memory is obviously an advantage, but it’s worth looking at some specific examples to understand why exactly this trend continues.

Financial Technology, or Fintech, applications require large amount of memory to process thousands of transactions a day, all of which require real-time, high-speed data analytics. The capability to use completely in-memory databases, without resorting to slow, high-latency storage devices, to drive event driven analytics is key for differentiation vs. competitors.

In biotech applications, DNA sequencing has become the poster child pushing the need for large amounts of data, which in turn requires large amounts of memory. DNA Data sequencing machines have dropped dramatically in price enabling widespread adoption for numerous applications, but they still require large amounts of high performance memory to process all the data they generate on a large scale.

Additionally, there is Artificial Intelligence (AI), which is the best example of a modern application requiring large amounts of RAM. AI applications are required to process vast amounts of data when initially training to increase performance and accuracy. The faster this data can be processed, the faster each iteration can train and processing the data in RAM is orders of magnitude faster than retrieving it from storage. In addition, when the AI is executing, it is far better to store large amounts of data in RAM to allow the AI to avoid repeated calculations and to enhance overall efficiency.

A New Memory Tier

CXL memory modules require a dedicated ASIC to translate the CXL data packets into the traditional DRAM interface. These CXL controllers, by their very existence, introduce an additional latency compared to CPUs accessing the DRAM directly through a memory controller. This additional latency turns out to be about the same as one NUMA hop or an additional 100ns compared to direct attached parallel memory.

While supplementing additional latency seems problematic at first, it's necessary to remember what happens in a typical application if it runs out of memory. The data must then be placed in storage, which is orders of magnitude slower to access than CXL attached memory. So while CXL memory is slower, in terms of latency, than direct attached DRAM modules, it is significantly faster than even the fastest storage and can result in dramatic performance increases.

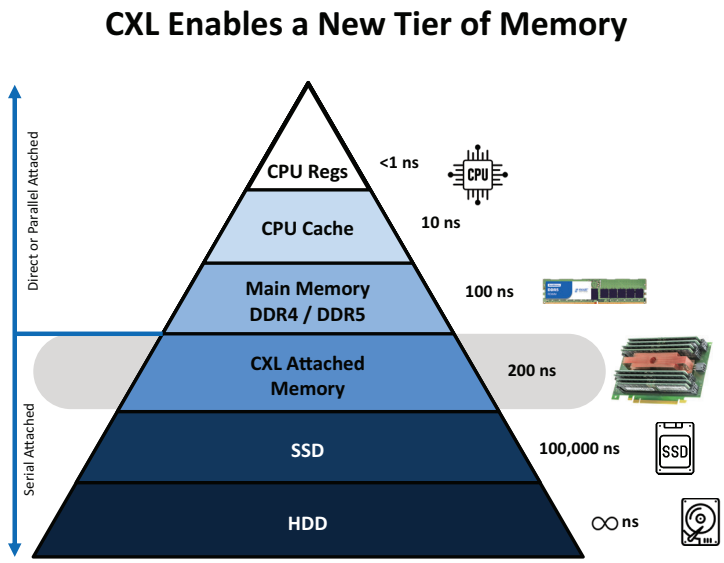

Figure 3 - Memory and Storage Tiers

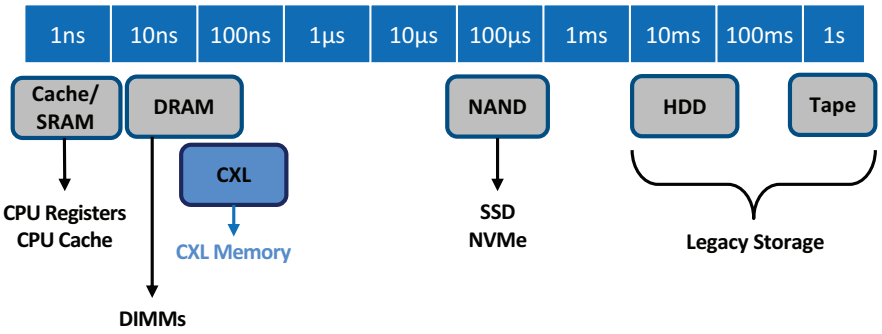

Figure 4 - CXL Latency Compared to DRAM and NAND

A detailed analysis of CXL bandwidth is beyond the scope and intention of this article, but to a first order, CXL operates at PCIe gen 5 rates so a x16 lane device can operate at speeds up to a theoretical maximum of 63GB/s (once overhead is considered). Also, PCIe and CXL protocols are bi-directional so data can be written and read simultaneously which can further increase the bandwidth depending on application loads.

CXL Form Factors

CXL modules are being developed in multiple standardized form factors. They include EDSFF (Enterprise and Datacenter Standard Form Factor) and AIC (Add-In-Card) form factors. The AIC form factor, also sometimes referred to as CEM (Card Electromechanical), has been in use for many years and has the advantage of being widely supported in many systems. Because these cards often attach directly to the motherboard, no special backplanes are required for CXL support, they can simply be installed in a compatible system today.

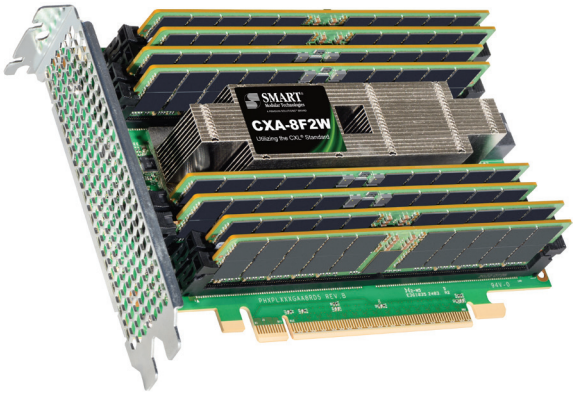

Figure 5 - SMART 8-DIMM AIC

EDSFF modules come in a variety of sizes and shapes, but the most interesting form factor for CXL attached memory modules is the E3.S for factor. E3.S supports multiple PCIe lane configurations (x4, x8, and x16) and is typically installed in the “front” of a server and can be removed without the need to open the server.

Figure 6 - SMART E3.S Memory Expansion Module

Form Factor	Capacity	Cost/GB	Power (watts)	Size (H x L) mm
E1.S	Low	High	25	31.5 x 111.49
E3.S	Moderate	Moderate	25/40	76 x 112.75
E3.L	Moderate	Moderate	40/70	76 x 142.2
AIC (FHFL)	High	Low	75+	167 x 312
AIC (HHHL)	High	Low	75+	111 x167

Summary

CXL is an exciting new protocol developed for the PCIe physical layer that enables new devices, such as memory expansion devices, to be attached to the CPU via a high-speed, low-latency interface. It can dramatically increase the amount of memory in a server with a considerable improvement to latency when compared to data on a storage device. CXL will be available in a variety of standardized form factors as is shaping up to be a transformative new technology.

Contact: Torry Steed at torry.steed@smartm.com for any sales and sampling queries.

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

10.08.24/M-WP014MemoryExpansion/Rev.4