DUG'23
Micron CXL update

Andrey Kudryavtsev
CXL Business Development, Micron
November 2023

The Industry's Fully Composable, Scalable Vision

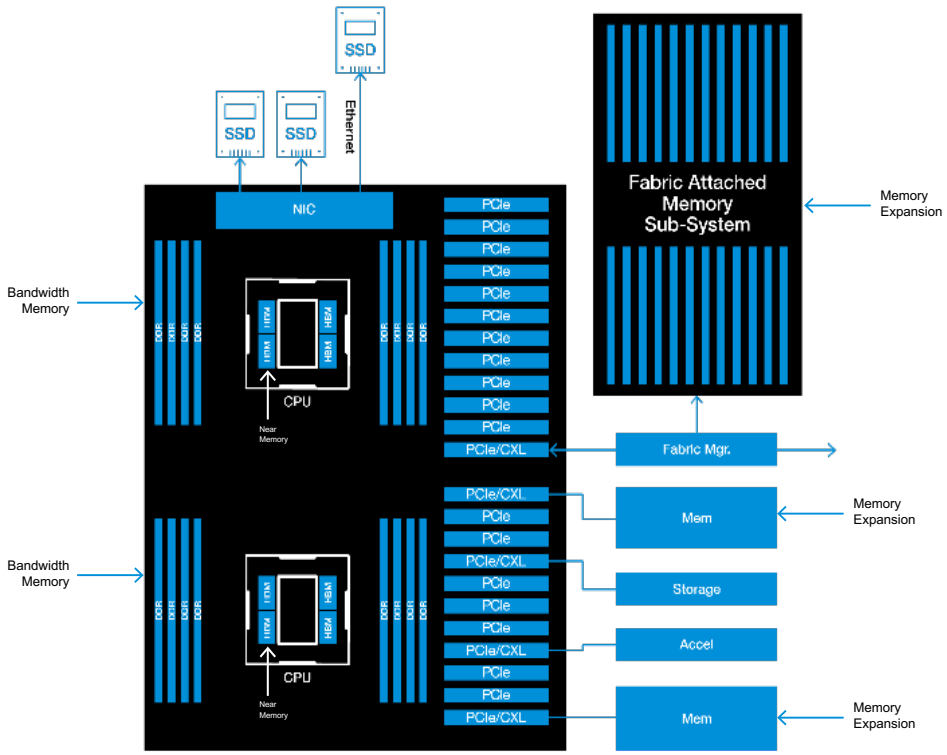

- Compute Express Link™ (CXL™) is an industry-supported Cache-Coherent Interconnect for Processors, Memory Expansion and Accelerators.
- The CXL Consortium is an open industry standard group formed to develop technical specifications that facilitate breakthrough performance for emerging usage models while supporting an open ecosystem for data center accelerators and other high-speed enhancements.

Micron

CXL Consortium / Specification

CXL Adoption Timeline

	2023	2024	2026
Features	**CXL 1.0 / 1.1**	**CXL 2.0**	**CXL 3.0**
Release date	2019	2020	1H 2022
Max link rate	32GTs	32GTs	64GTs
Flit 68 byte (up to 32 GTs)	✓	✓	✓
Flit 256 byte (up to 64 GTs)			✓
Type 1, Type 2 and Type 3 Devices	✓	✓	✓
Memory Pooling w/ MLDs		✓	✓
Global Persistent Flush		✓	✓
CXL IDE		✓	✓
Switching (Single-level)		✓	✓
Switching (Multi-level)			✓
Direct memory access for peer-to-peer			✓
Symmetric coherency (256 byte flit)			✓
Memory sharing (256 byte flit)			✓
Multiple Type 1/Type 2 devices per root port			✓
Fabrics (256 byte flit)			✓

Not supported

✓ Supported

CXL Memory Module (Type-3)

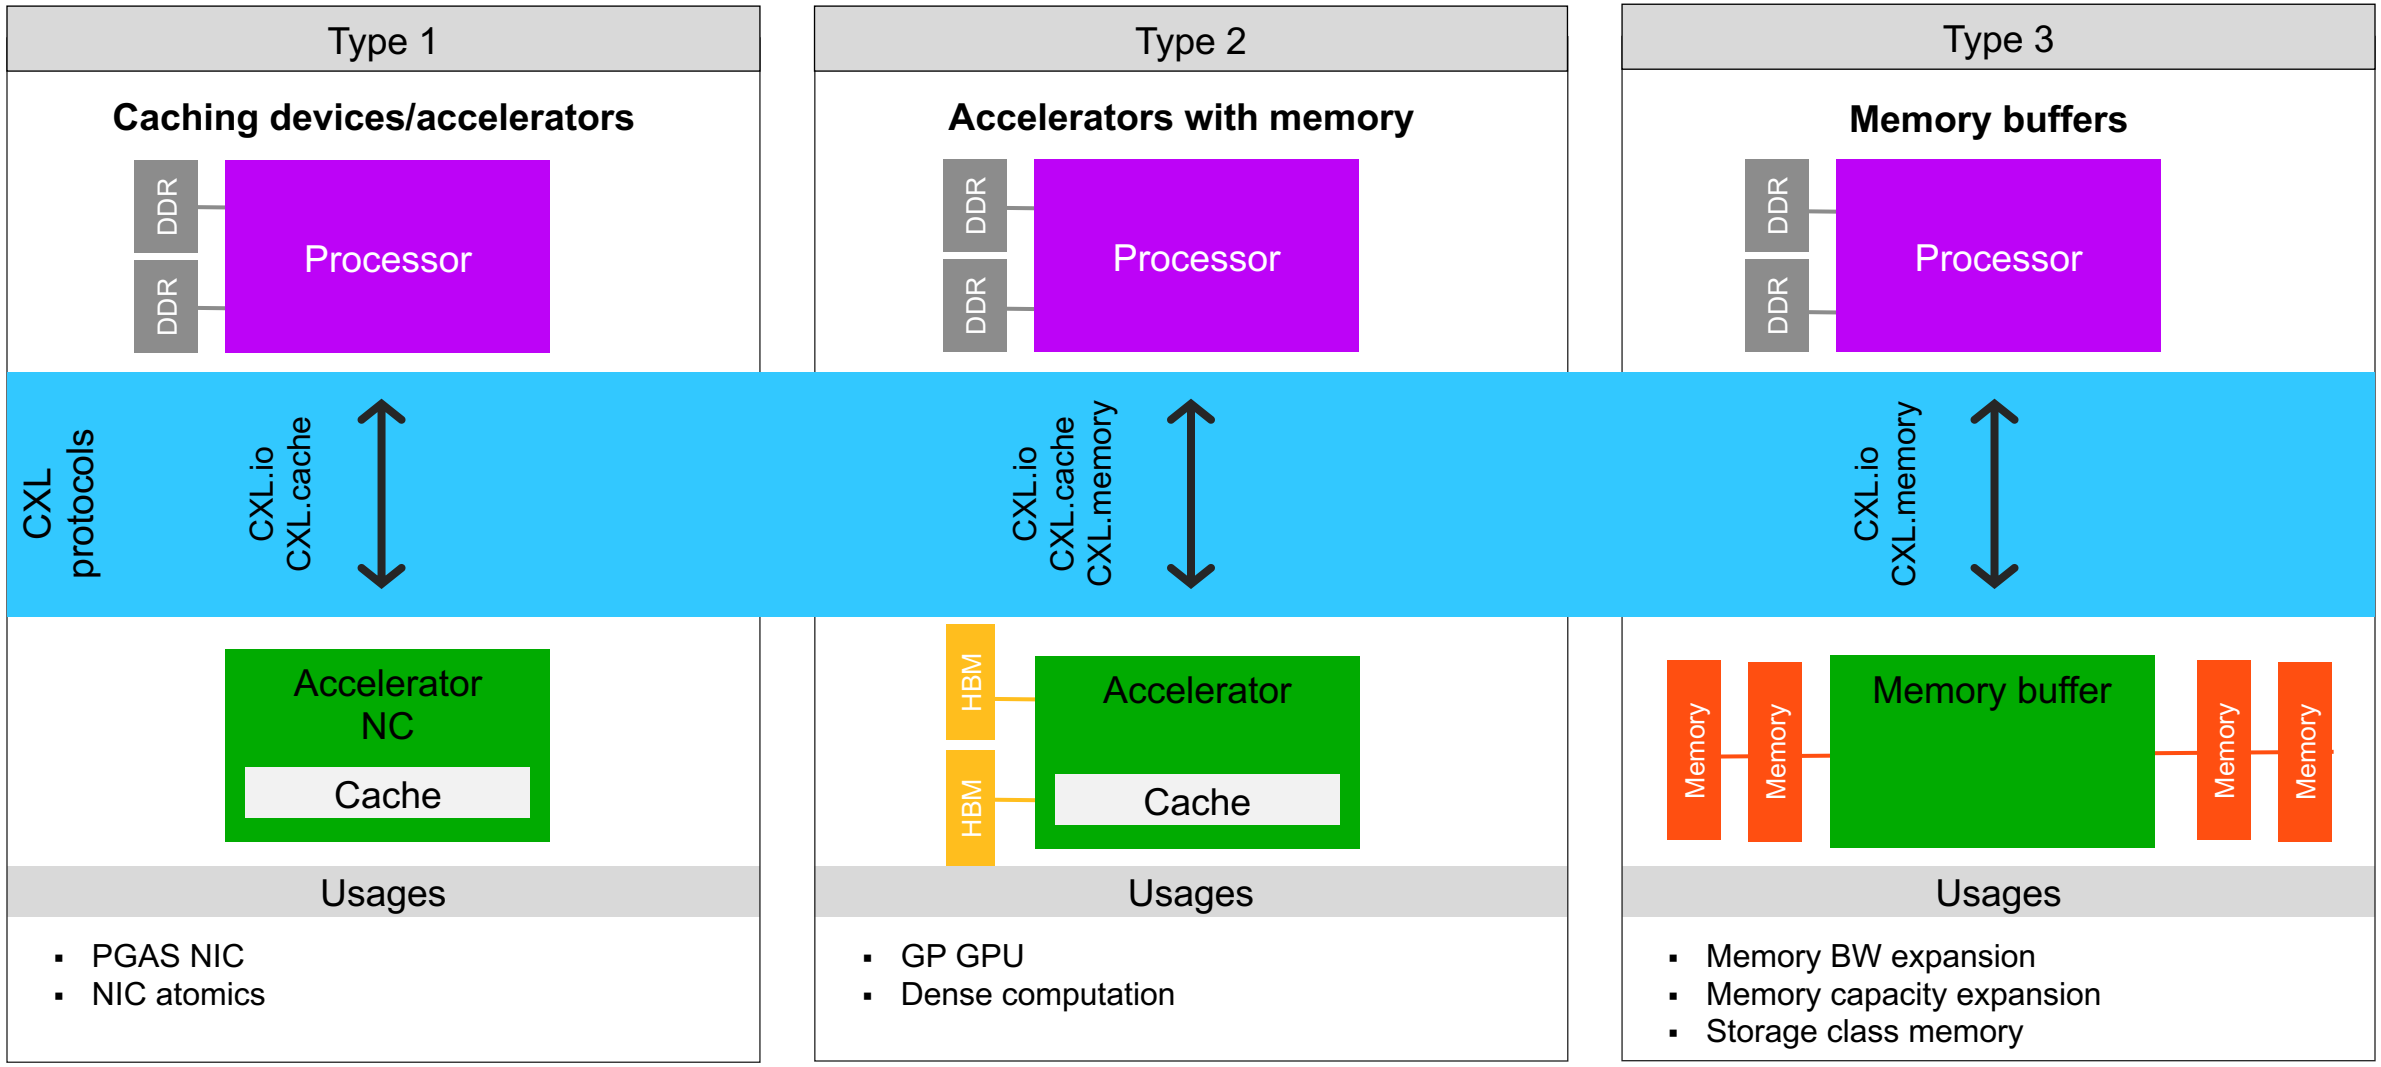

Reference: SNIA EDSFF (E3) Spec

SFF-TA-1008

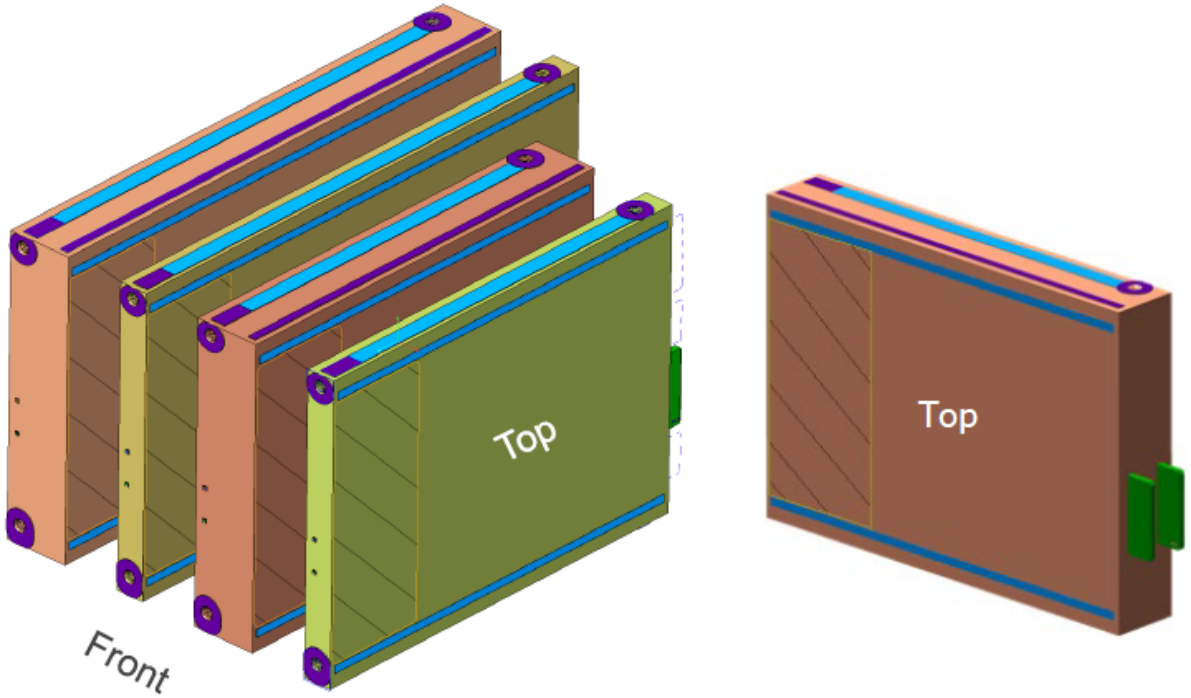

Figure 4-1. Example Device 3D Views

The example above includes the following device implementations starting from left to right.

- E3 long double thickness device (E3.L 2T)
- E3 long single thickness device (E3.L)
- E3 short double thickness device (E3.S 2T)
- E3 short single thickness device (E3.S)
- E3 short double thickness 2x1C device (E3.S 2T 2x1C)

Table 5-1 Nominal Device Form Factor Dimensions

Device Variation	Width	Length	Thickness
E3.S	76mm	112.75mm	7.5mm
E3.S 2T	76mm	112.75mm	16.8mm
E3.L	76mm	142.2mm	7.5mm
E3.L 2T	76mm	142.2mm	16.8mm

Introducing Micron CZ120 Memory Expansion Modules

Delivering capacity, bandwidth, flexibility

128GB / 256GB

Up to 2TB incremental server capacity[1] supporting CXL 2.0

Up to 36GB/s[2]

Up to 24% increased server memory read/write bandwidth[3]

E3.S 2T PCIe Gen5 x8

Industry-standard form factor for broad deployment

1. By adding 8x256GB CZ120s, system limitations may apply.
2. Measured by running MLC workload with 2:1 read/write ratio on a single CZ120 module.
3. MLC bandwidth using 12-channel 4800MT/s RDIMM + 4x256GB CZ120 vs. RDIMM only.

CZ120 product overview
Performant, secure, reliable

Product highlights

- Leverages high-volume DRAM production process
- Unique dual-channel memory architecture for higher module bandwidth
- Capacity expansion - up to 2TB[1] of incremental memory per CPU in a E3.S 2T form factor

Key features

- Secure root of trust and secure boot
- Sideband device management
- Data center RAS
  - SECDED, SDDC ECC
  - Reed-Solomon-based DRAM device error correction
  - Post package repair management

CZ120 memory expansion module

Capacity	128GB / 256GB
CXL	2.0
Form factor	E3.S 2T
Host interface	PCIe Gen5 x8
Mean Time Between Failures (MTBF)	3 Million hours
Power (typical)	27W / 31W
Module bandwidth	Up to 36GB/s

1. By adding 8x256GB CZ120s, system limitations may apply

Methods to use CXL for capacity & BW

Capacity Expansion

Direct Attached Memory Tiering

- Application Transparent
  - OS Managed, User Space Library, 2LM Mode
- Application Managed
  - Application Aware (ex: libnuma)
  - Modified (ex: libmemkind)

CXL Switch / Fabric attached Memory Tiering

- Another Memory tier added to the system with higher latencies
- Micron taking a leadership role in enabling the AI and data science toolchains to share data in fabric-attached memory

Bandwidth Expansion

- CXL Heterogenous interleave solutions
  1. Hardware based interleave
  2. Software and HW heterogenous interleave
  3. Software based NUMA interleave

  Micron actively collaborating on introducing ratio-based memory interleaving capabilities for the Linux kernel

CXL Use Cases

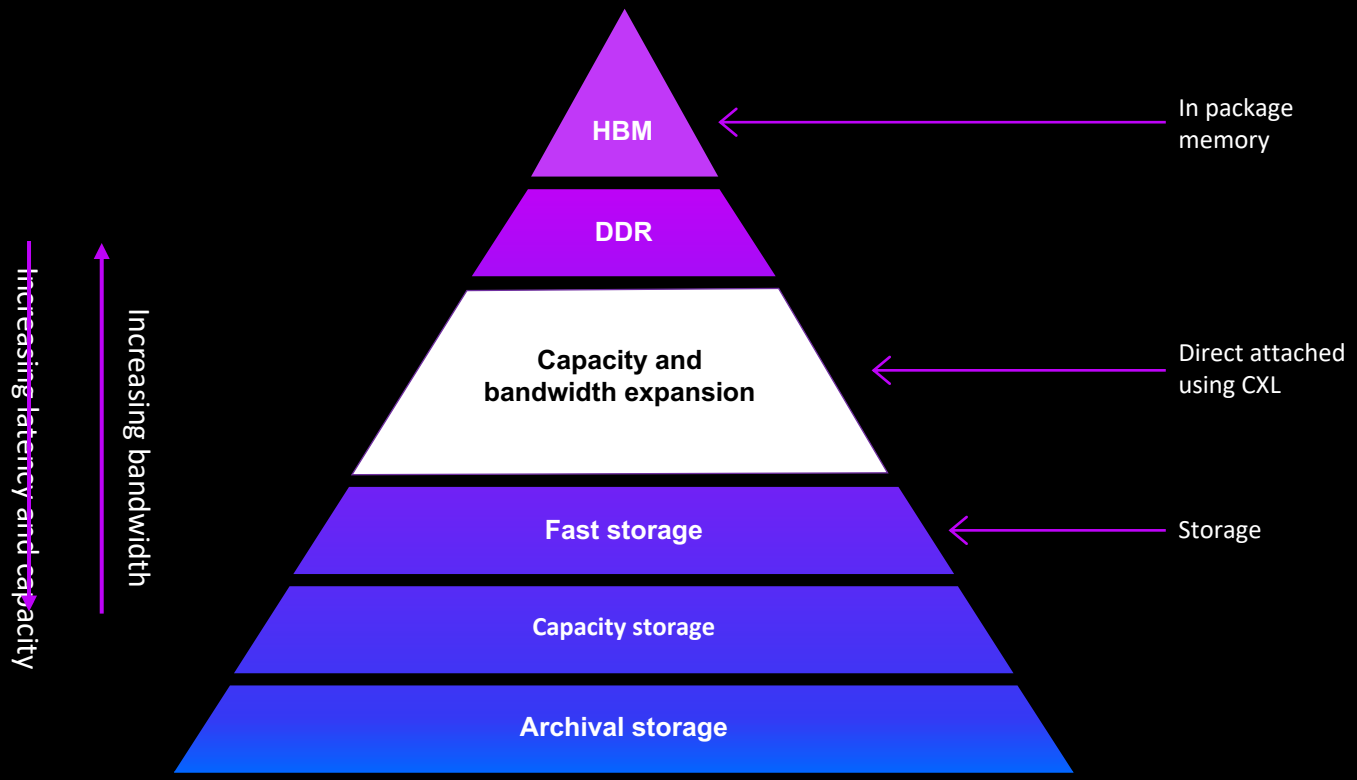

Use cases for CXL-based memory expansion

Data center workloads	Capacity expansion	Bandwidth expansion
AI/ML	●	●
In-memory database	●	
Data analytics	●	●
General purpose compute	●	
High performance compute	●	●

Enriching the CXL Ecosystem with CZ120

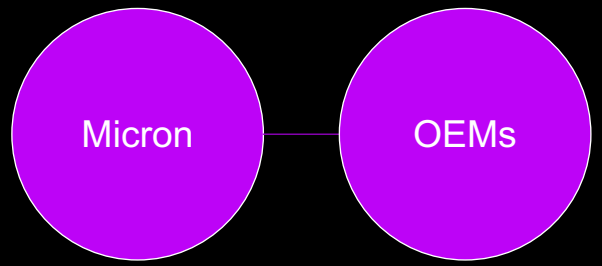

Activities:

Validation, Compatibility, Power, Thermal, EVT, Optimization CXL Configuration for different use cases, and Joint Qualification

Objective:

Achieve Timely Qualification by Systems Launch

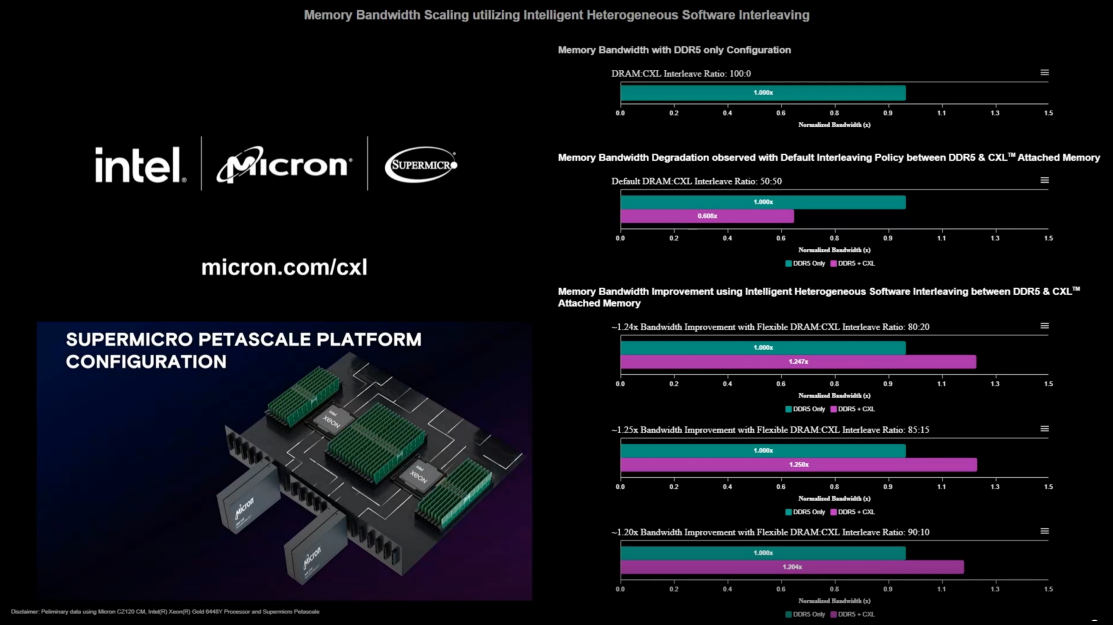

Memory solutions to meet data center infrastructure demands

Not all Products are POR and shown for conceptual discussion purposes
Left edge of bar denotes approximate availability, right edge does NOT represent End of Life

Micron Technology Enablement Program (TEP))

Cloud Service Providers, Original Equipment Manufacturers and Original Design Manufacturers — Qualify our CZ120 into your server platforms by enrolling with Micron TEP

Hands-on support to aid in the development of CXL™-enabled designs

- Technical resources including data sheets, electrical and thermal models to aid in product development and evaluation, and engineering consultation related to signal integrity and other technical support topics
- Access to other ecosystem partners who can aid in system-level design

Learn more micron.com/CXL