

# Solving HPC challenges with Micron CXL attached memory products

CERN Openlab workshop track

Andrey Kudryavtsev, CXL BDM, Micron Technology

© 2023 Micron Technology, Inc. All rights reserved. Information, products, and/or specifications are subject to change without notice. Micron, the Micron logo, and all other Micron trademarks are the property of Micron Technology, Inc. All other trademarks are the property of their respective owners.



# CXL® Feature Adoption



## CXL Adoption Timeline

| Features                                     | CXL Adoption Timeline |         |         |
|--|-----------------------|---------|---------|
|  | 2023                  | 2024    | 2026    |
|  | 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



**OCP**  
GLOBAL  
SUMMIT

OCTOBER 17-19, 2023  
SAN JOSE, CA

Scaling Innovation Through Collaboration



Delivering capacity, bandwidth, flexibility

# 128GB / 256GB

Up to 2TB incremental server capacity<sup>1</sup> supporting CXL 2.0

# Up to 36GB/s<sup>2</sup>

Up to 24% increased server memory read/write bandwidth<sup>3</sup>

# 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.

## Product highlights

- Leverages high-volume DRAM production process
- Unique dual-channel memory architecture for higher module bandwidth
- Capacity expansion - up to 2TB<sup>1</sup> 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  |
| Power (typical)  | 27W / 31W     |
| Module bandwidth | Up to 36GB/s  |

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

# Memory Workloads targets

## Application aware

- Optimized for memory tiers
- Will be the most performant
- Target workloads with high ROI
- Extract the most value out of CXL

## OS/Hypervisor aware

- Some mitigation for latency
- Better performance (compared to flat memory address)
- Address broadest workloads and cold memory pages

## Flat memory address space

- Customer CXL starting point (easiest lift)
- No mitigation for latency – targeting workloads that are insensitive to latency

Challenging



Difficulty

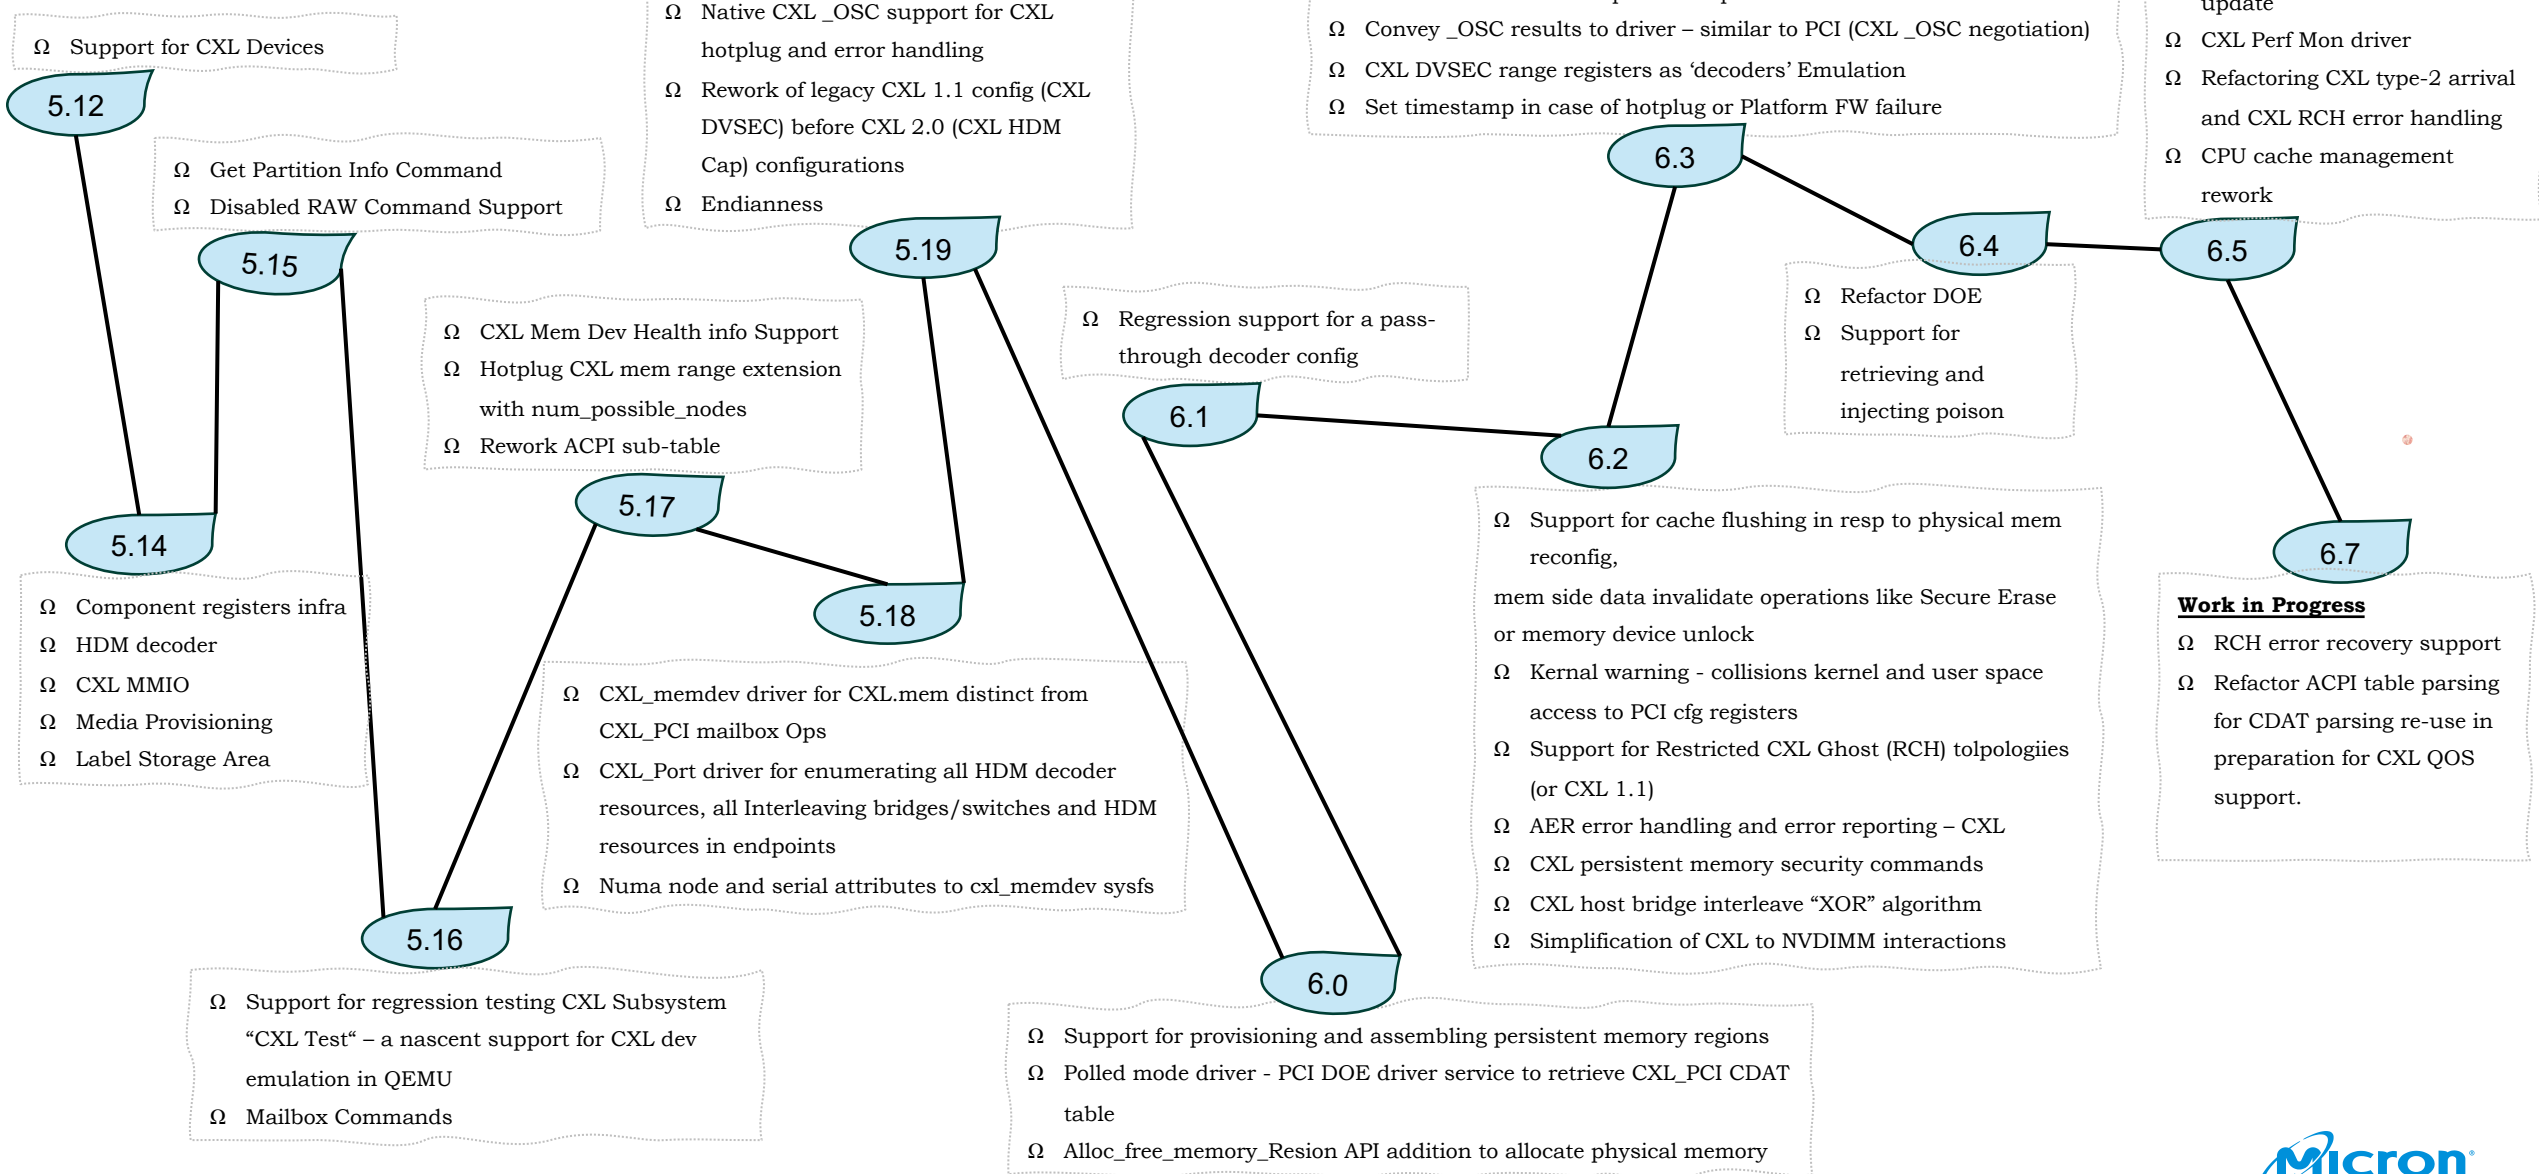
Performance

Easy

## 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 |                    | ●                   |

# Linux Kernel Updates

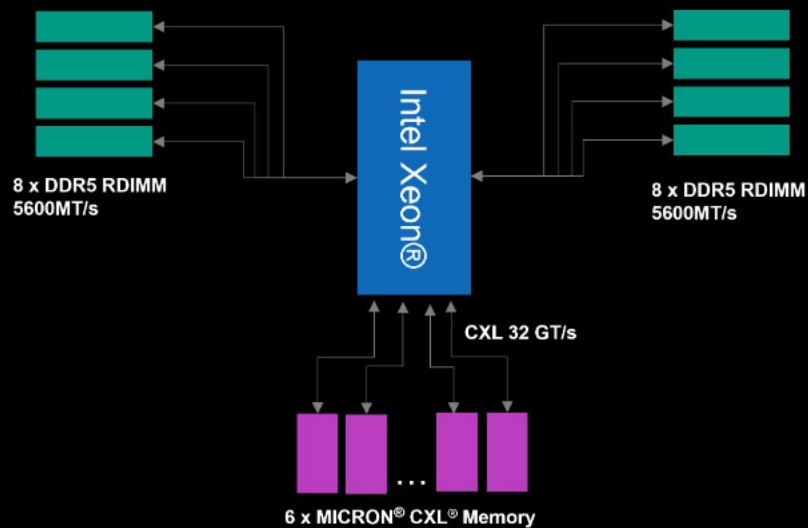


# LLM Performance Optimization with Micron's Cutting-Edge CXL Memory and Software Memory Interleaving




Unleash the Power of LLM using Micron CZ120 CXL Bandwidth Expansion and Next Gen Intel Xeon Processor


## System Configuration



Max Bandwidth using Flexible Interleave Ratios



~1.22x Runtime Improvement with LLAMA



User: Why is the sky blue?

**DRAM Only:** Well, the answer is a bit more complicated than just one reason. There are actually several factors that contribute to the blue color of the sky. Here are some of the main reasons why the sky appears blue: 1. Rayleigh scattering: This is the primary reason why the sky appears blue. When sunlight enters Earth's atmosphere, it encounters tiny molecules of gases such as nitrogen and oxygen. These molecules scatter the light in all directions, but they scatter shorter (blue) wavelengths more than longer (red) wavelengths. This is known as Rayle

9.312 tokens/second

User: Why is the sky blue?

**DRAM + CXL:** The sky appears blue because of a phenomenon called scattering. When sunlight enters Earth's atmosphere, it encounters tiny molecules of gases such as nitrogen and oxygen. These molecules scatter the light in all directions, but they scatter shorter (blue) wavelengths more than longer (red) wavelengths. This is known as Rayleigh scattering. As a result of this scattering, the blue light is distributed throughout the sky, making it appear blue. The reason why the sky appears blue during the day and not at other times is because the sun'

10.373 tokens/second

Enter your Question/Statement here ...

Run with DRAM

Run with DRAM + CXL

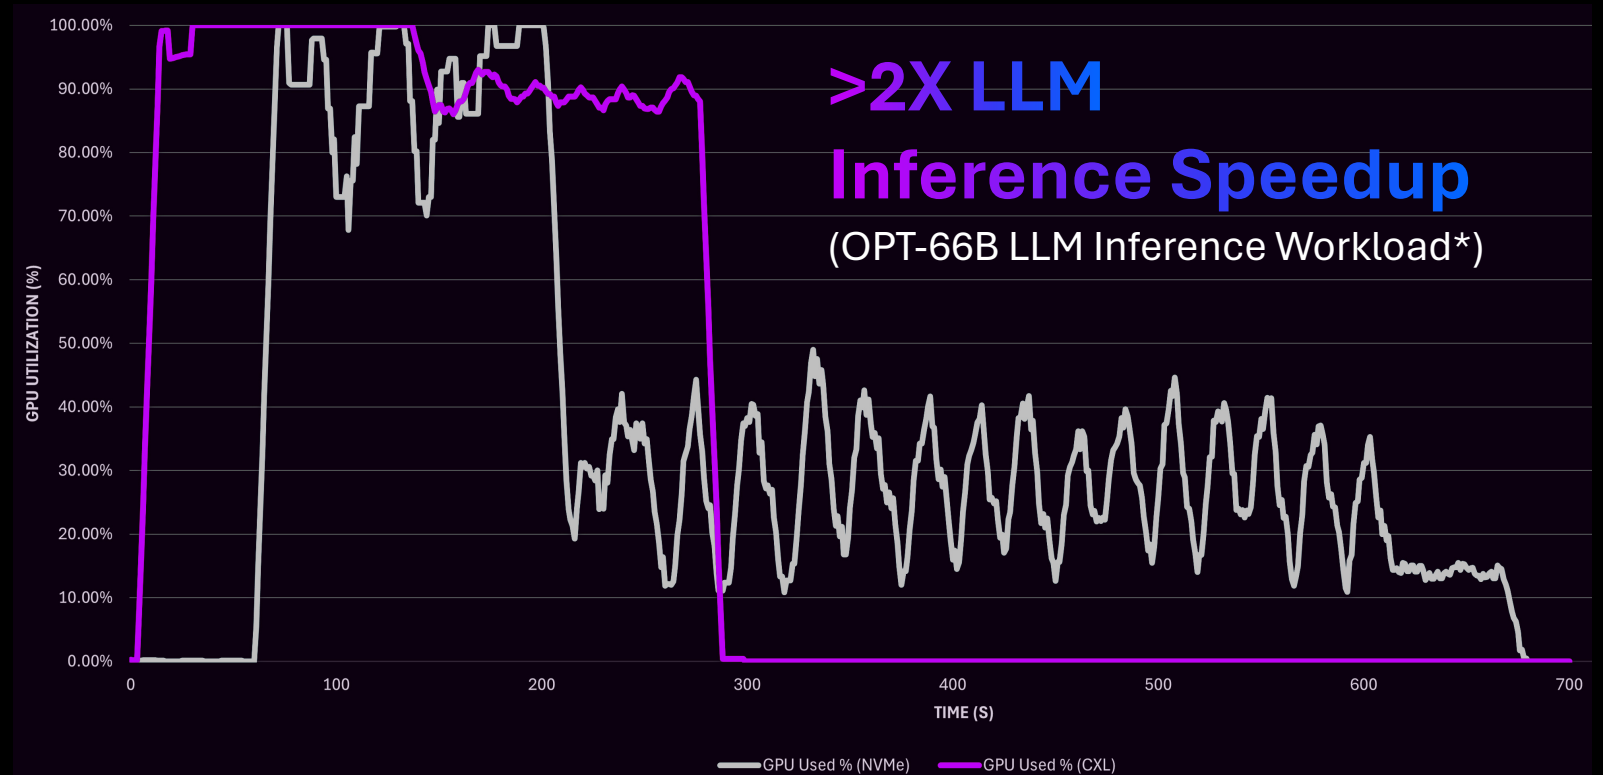
# Micron CXL enables higher GPU utilization

Higher throughput and faster task completion times for AI/ML workloads

**Config A:** Supermicro Petascale server, AMD Genoa 9634 DP/UP 84C/168T, 8 \* 32GB Micron DDR5-4800, 2 x Micron 7450 960GB M.2, Nvidia A10 GPU, Ubuntu 22.04.04 using Kernel 5.15.0, [FlexGen AI](#)

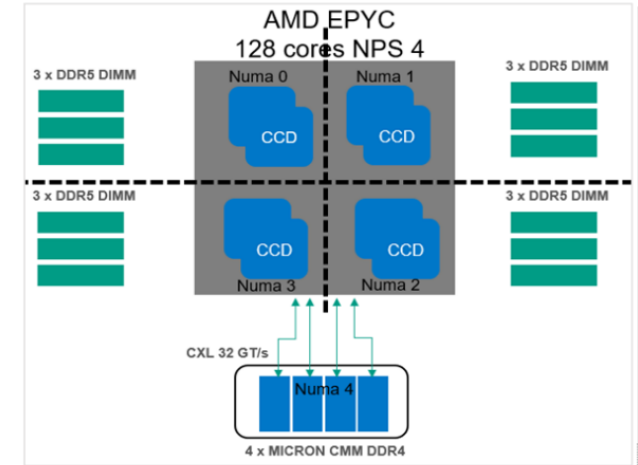
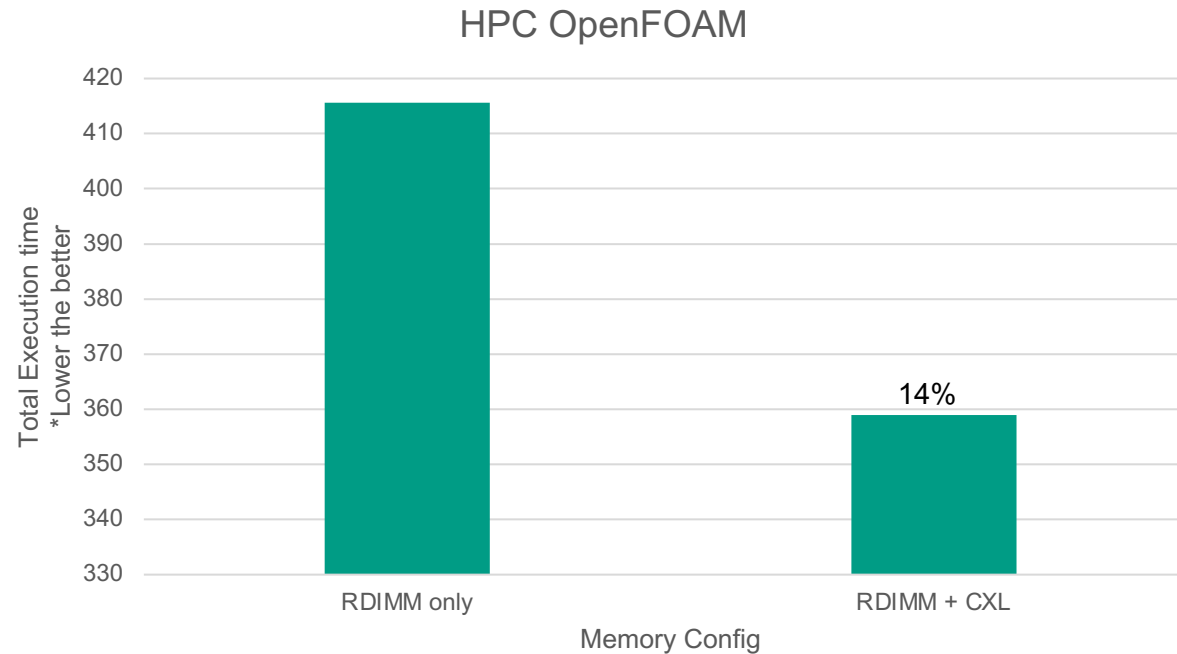
**Config B:** Supermicro Petascale server, AMD Genoa 9634 DP/UP 84C/168T, 8 \* 32GB Micron DDR5-4800, Nvidia A10 GPU, **2x Micron CZ120 256GB CXL**, Ubuntu 22.04.04 using Kernel 5.15.0, **MemVerge Memory Machine 2.5.1**, [FlexGen AI](#)

## GPU utilization NVMe SSD vs CXL memory module





# HPC- OpenFOAM



(a)


Balanced  
Interleave 4:1

## Key takeaway:

- For OpenFOAM being a bandwidth sensitive application, we observe about 14% gain in performance with ~25% increase in available bandwidth. Optimizations can lead to better performance.

CXL provides better performance with additional bandwidth

# CXL Server Example



PCIe 5.0 Slots  
Up to 2 x16 slots +  
2 AIOMs

Micron DDR5 DRAM  
Up to 24 DIMMs

Single AMD Genoa  
Processor

E3.S 1T (x4) SSD

E3.S 2T (x8)  
CXL Type 3 CMM



Redundant Power Supply  
1600W (Titanium level)

PCIe 5.0 Slots  
x16 AIOMs

PCIe 5.0 Slots  
x16 slots



# 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](https://micron.com/CXL)

<https://github.com/cxl-micron-reskit>



