SNIA | Compute, Memory, and Storage

Unlocking CXL's Potential: Revolutionizing Server Memory and Performance

Live Webinar

April 2, 2025

10:00 am PT/ 1:00 pm ET

Our Speakers

Arthur Sainio

Co-Chair

SNIA Persistent Memory Special Interest Group

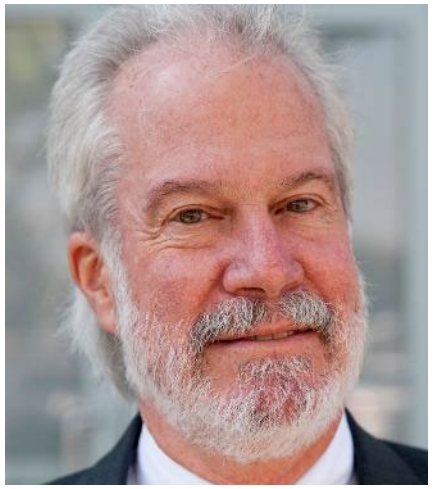

Jim Handy

General Director

Objective Analysis

Mahesh Natu

Systems and Software Working Group Co-Chair

CXL Consortium

Torry Steed

Senior Product Marketing Manager

SMART Modular Technologies

SNIA | CMS A SNIA Community

Agenda

- CXL is Exciting, but Where Is It Headed?
- CXL Consortium Update
- CXL Benefits, Implementation, and Ecosystem
- Final Thoughts and Q&A

The SNIA Community

200
Corporations, universities, startups, and individuals

2,500
Active contributing members

50,000
Worldwide IT end users and professionals

SNIA | CMS A SNIA Community

SNIA | Compute, Memory, and Storage

What We Do

- Engage technology users
- Educate on compute, memory, and storage technologies
- Accelerate SNIA standards
- Propel technology adoption

How We Do It

- Demonstrate at industry events
- Host Programming Workshops and Hackathons
- Present educational webinars, podcasts, and blogs

Learn more at www.snia.org/groups/cms

SNIA Legal Notice

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be, construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

SNIA | CMS
A SNIA Community

CXL is Exciting, but Where Is It Headed?

Jim Handy

CXL Can Do So Many Things, But What is it Really For?

- Maintaining coherency?
- Eliminating stranded memory?
- Expanding memory size?
- Increasing memory bandwidth?
- Supporting persistent memory?
- Hiding DDR4/DDR5/DDR6 differences?
- Passing messages between xPUs?

CXL Technology Basics

- Memory-speed access over PCIe physical layer
- Supports new architectures:
  - Disaggregated memory
  - Pooled memory
  - Switches for memory fabrics
  - Shared memory
  - Persistent memory

CXL DRAM Is A Lot Like An SSD

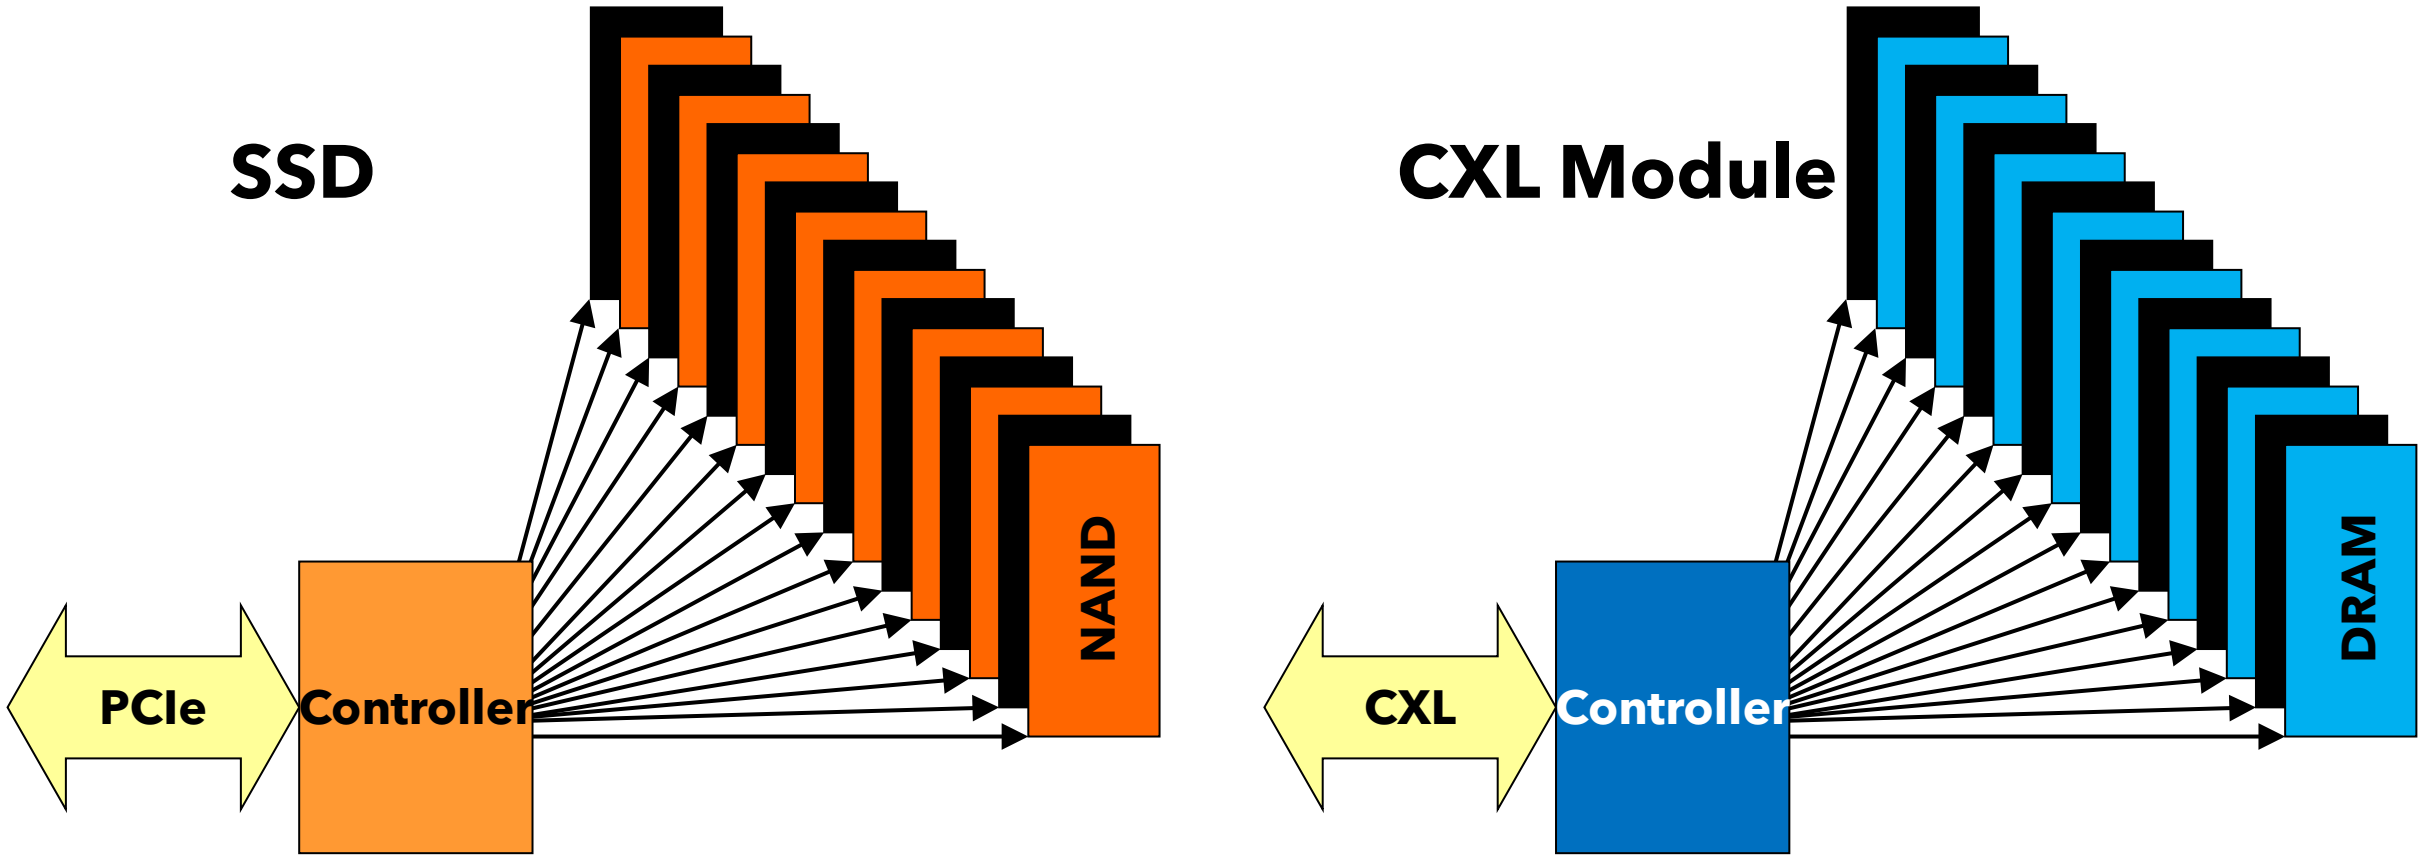

CXL DRAM Is A Lot Like An SSD

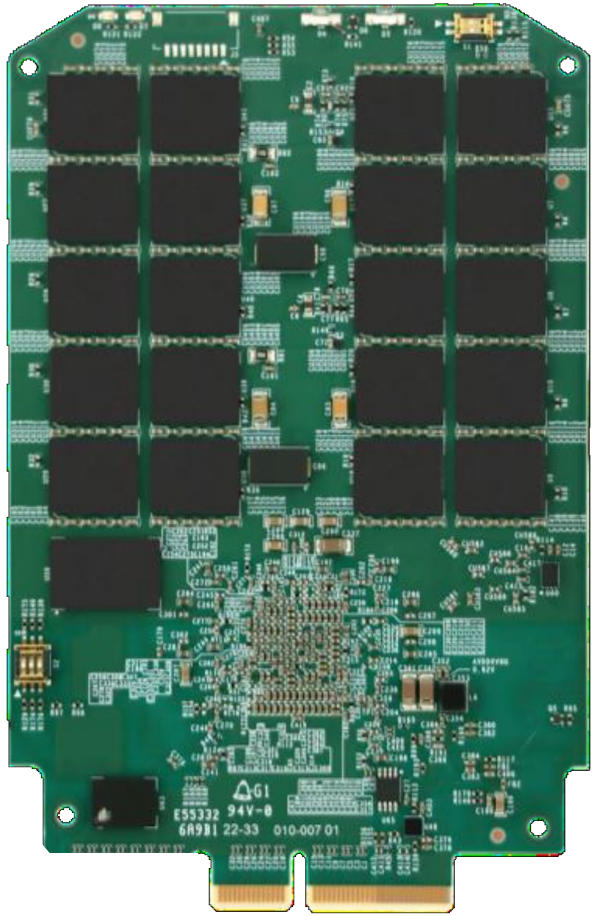

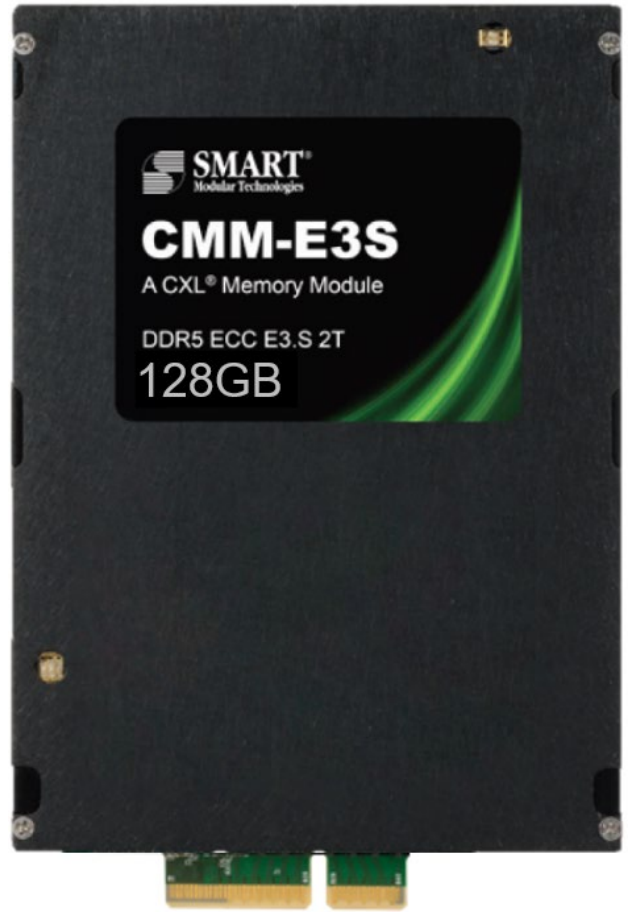

SNIA | CMS A SNIA Community

What Do Potential Users Say?

- **Google:** Stranded memory is not important
- **IBM/Georgia Tech:** DDR is a poor answer
- **AI Providers:** We need enormous memories
  - Also fast loads of GPU HBM
- **Hyperscalers:** "Any-to-Any" xPU connections
- **PC OEMs:** CXL is not immediately useful

CXL Forecast: Slow Steady Growth

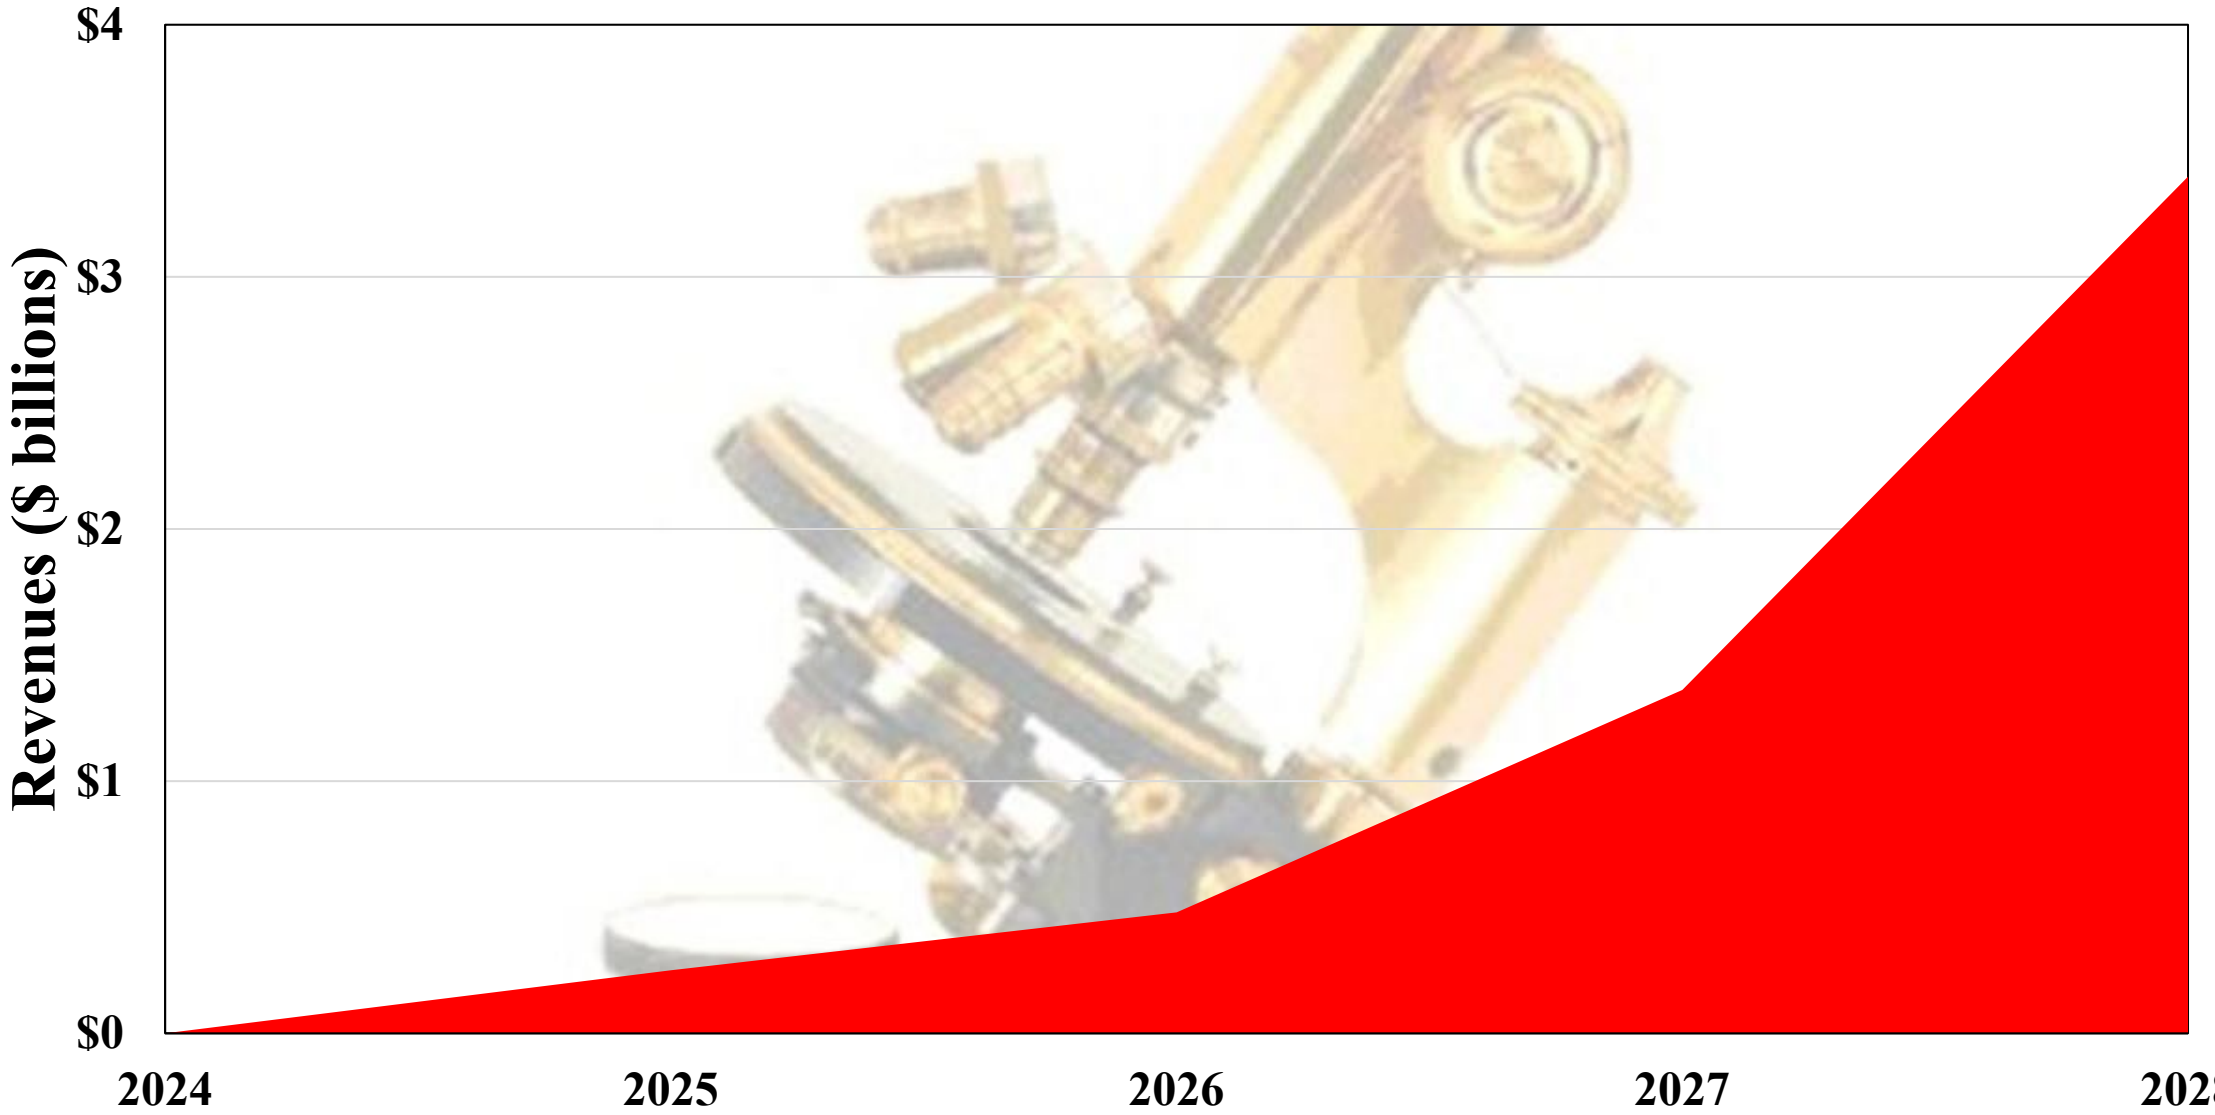

SNIA | CMS A SNIA Community

Long-Term Impact

- Re-thinking system architecture
  - Disaggregated memory
  - Processor arrays with mesh networks
  - Memory agnostic
- Better memory bandwidth & size vs. worse latency
  - Design-arounds will optimize for this
    - CXL Hot-Range Monitoring Unit (CHMU)

CXL Looks for the Perfect Home

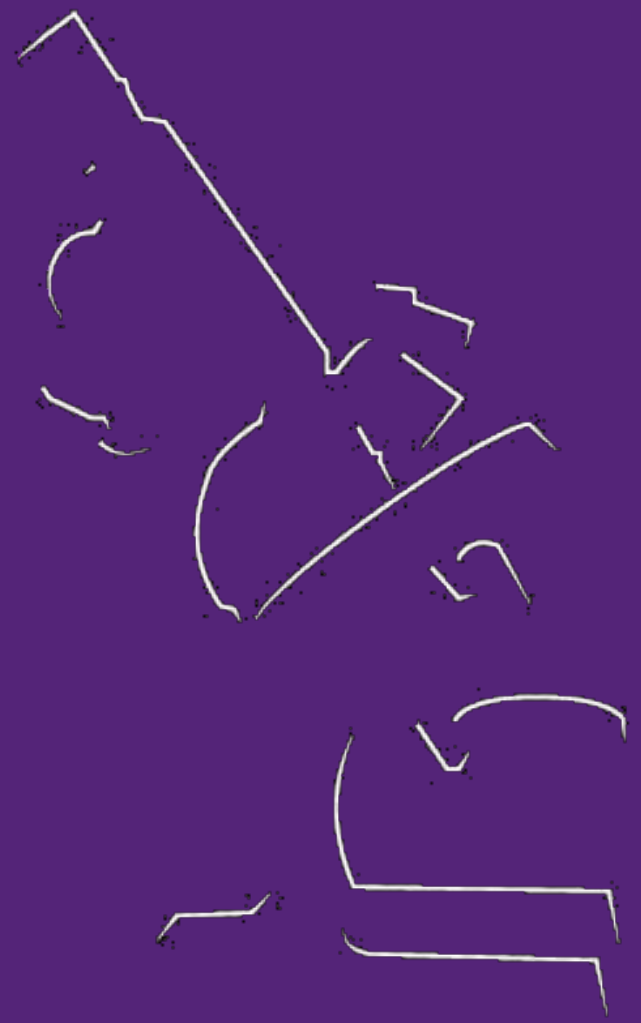

New report from Objective Analysis

- Covers all perspectives
  - Where CXL is useful, and where it isn't
  - Demand drivers for CXL DRAM modules
  - Opportunities outside of DRAM
  - Forecast (Revenues, units, ASP)
- Available for immediate download:
  - Objective-Analysis.com/reports

SNIA | CMS
A SNIA Community

CXL Consortium Update

CXL Compute Express Link®

Mahesh Natu

CXL Specification Release Timeline

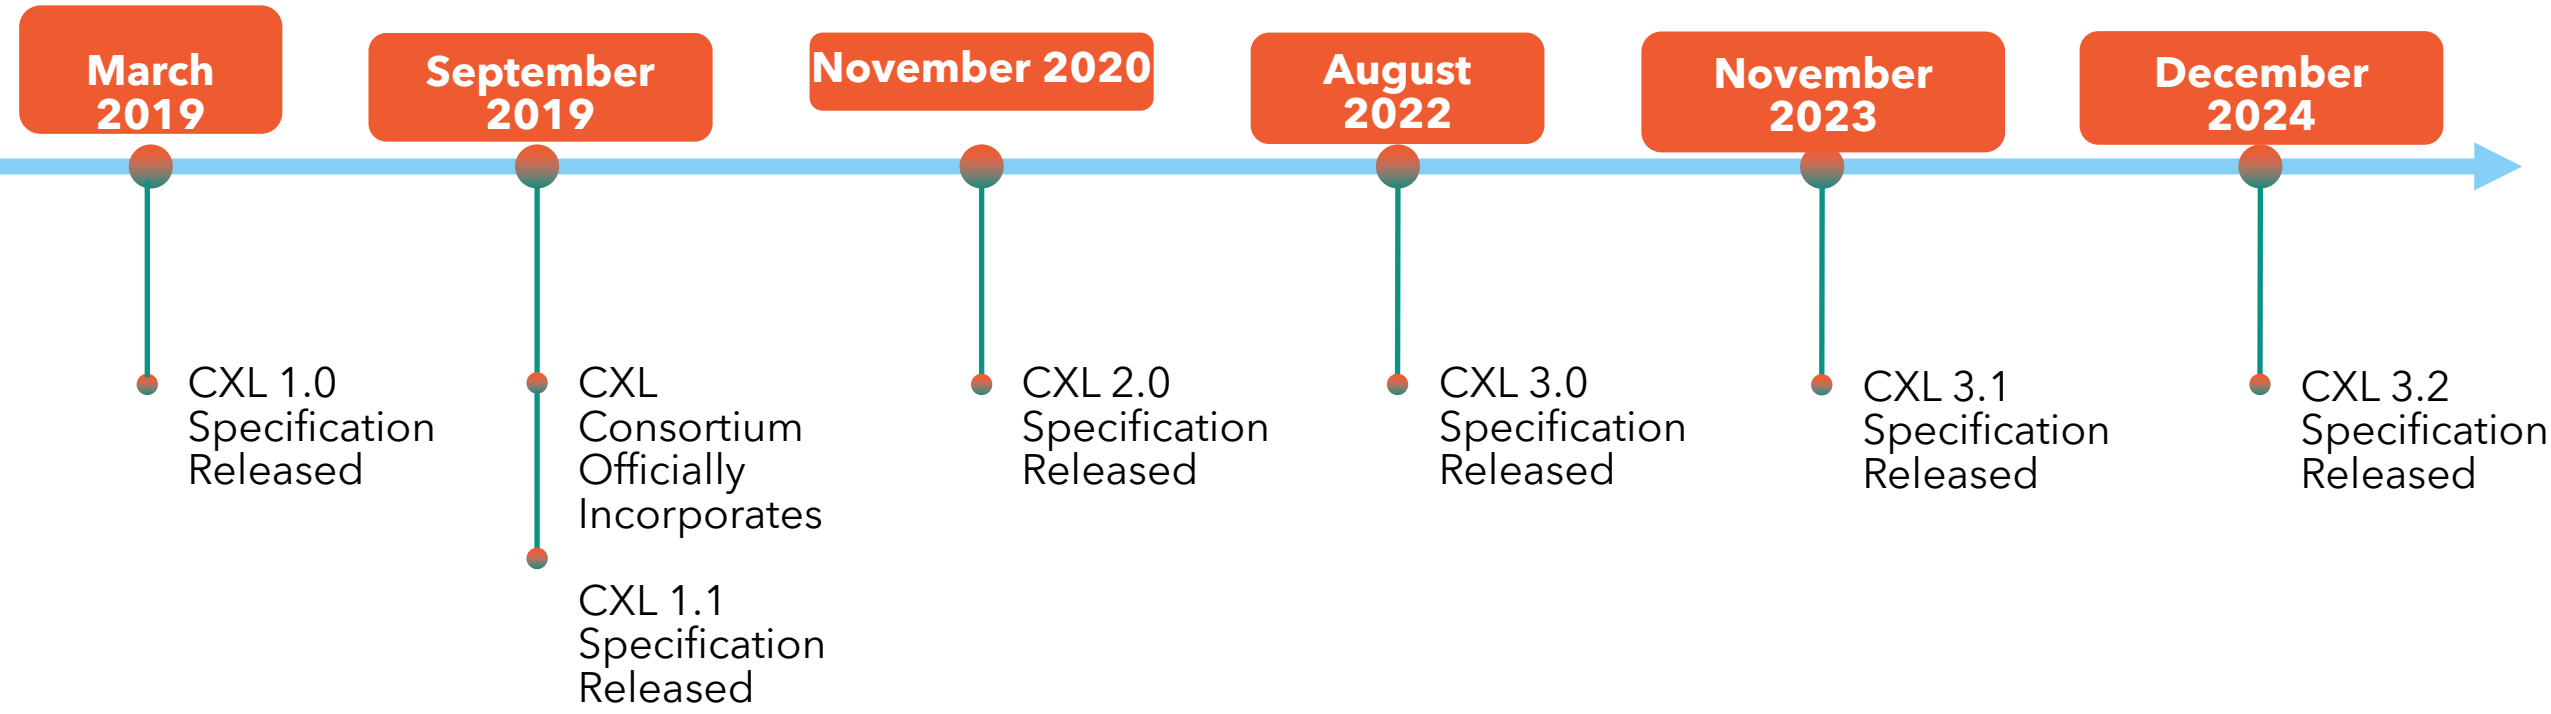

CXL 3.x Themes

- Scaling
- Memory Hierarchy
- Security

CXL: Scaling new heights

CXL 3.X

- Doubled the bit rate without latency impact
- Improved security
- Composable Fabric growth for disaggregation of memory & accelerators
- Memory sharing with back invalidate

CXL 2.0

- Multiple nodes inside a Rack/Chassis supporting pooling of resources
- Switch-based memory pooling
- Global Persistent Flush (GPF)
- Link-level encryption

CXL 1.1

- Single Node coherent interconnect
- Load/store semantics
- Hardware enforced cache coherence
- Low latency

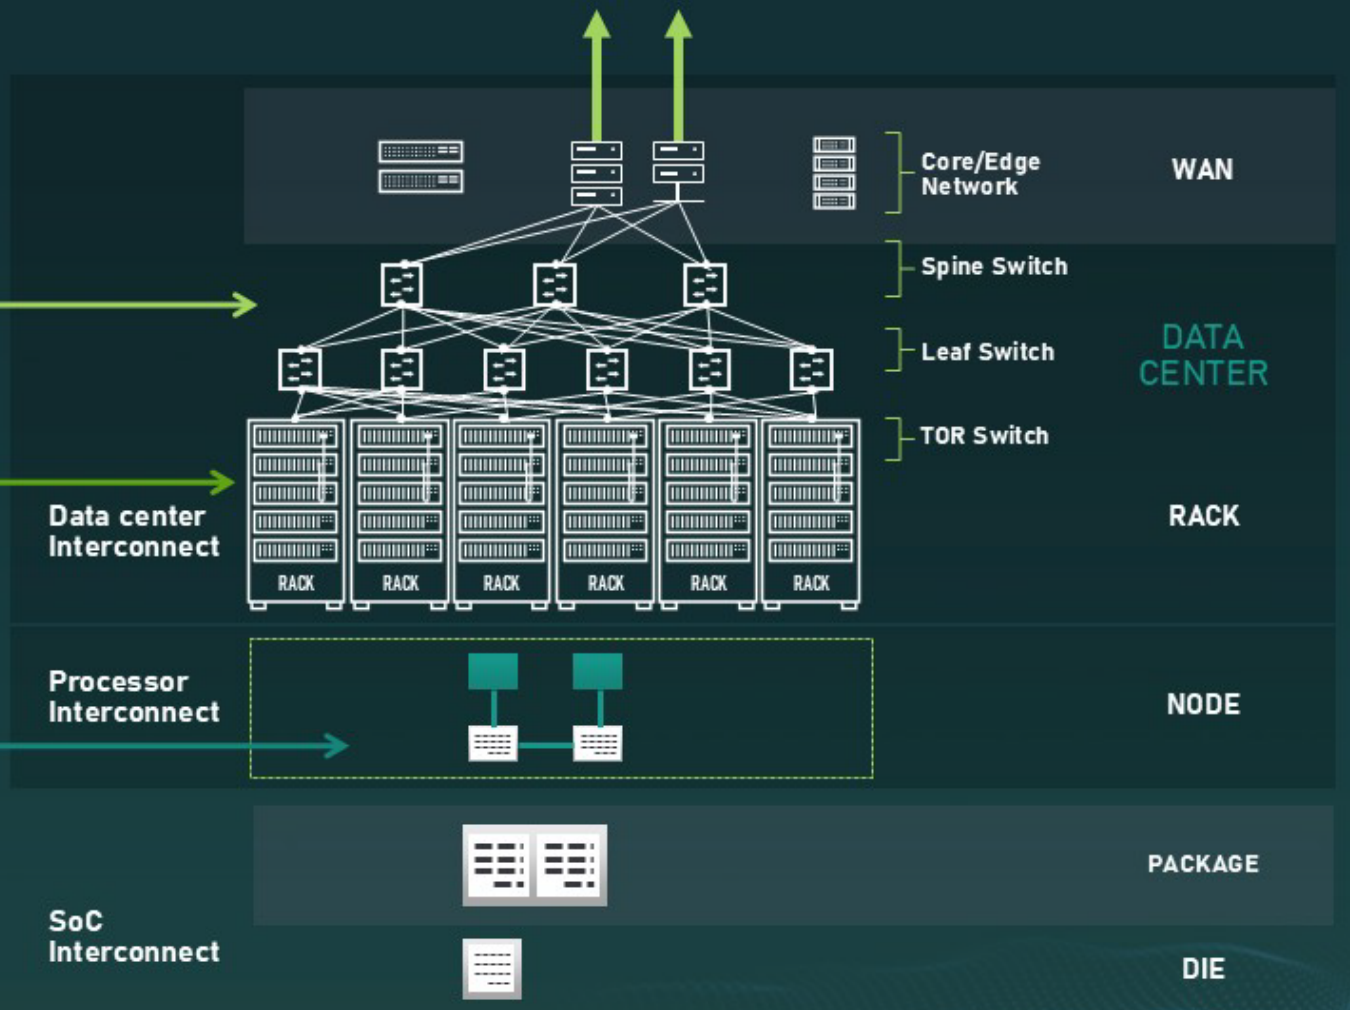

SNIA CMS A SNIA Community

CXL Memory Introduces New Tiers

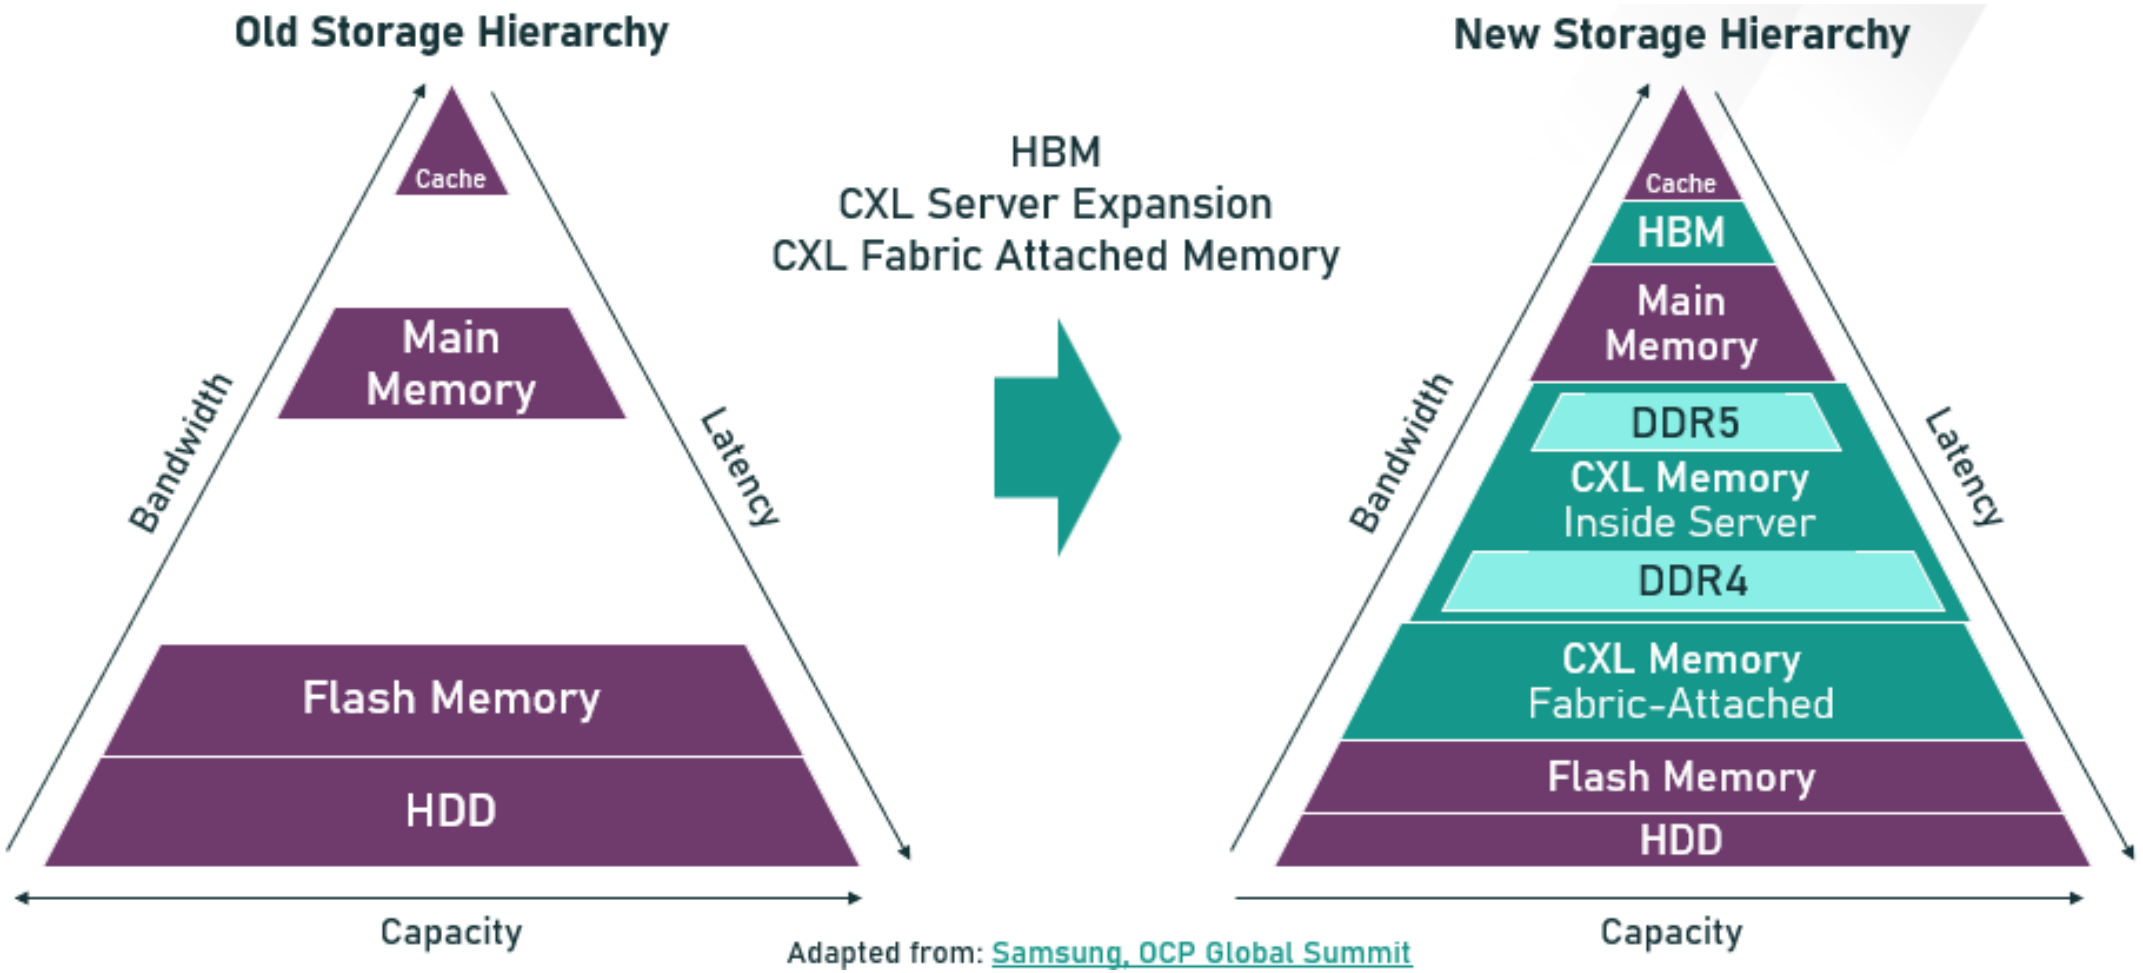

Adapted from: Samsung, OCP Global Summit

CXL Hot-Range Monitoring Unit (CHMU) for Tiering

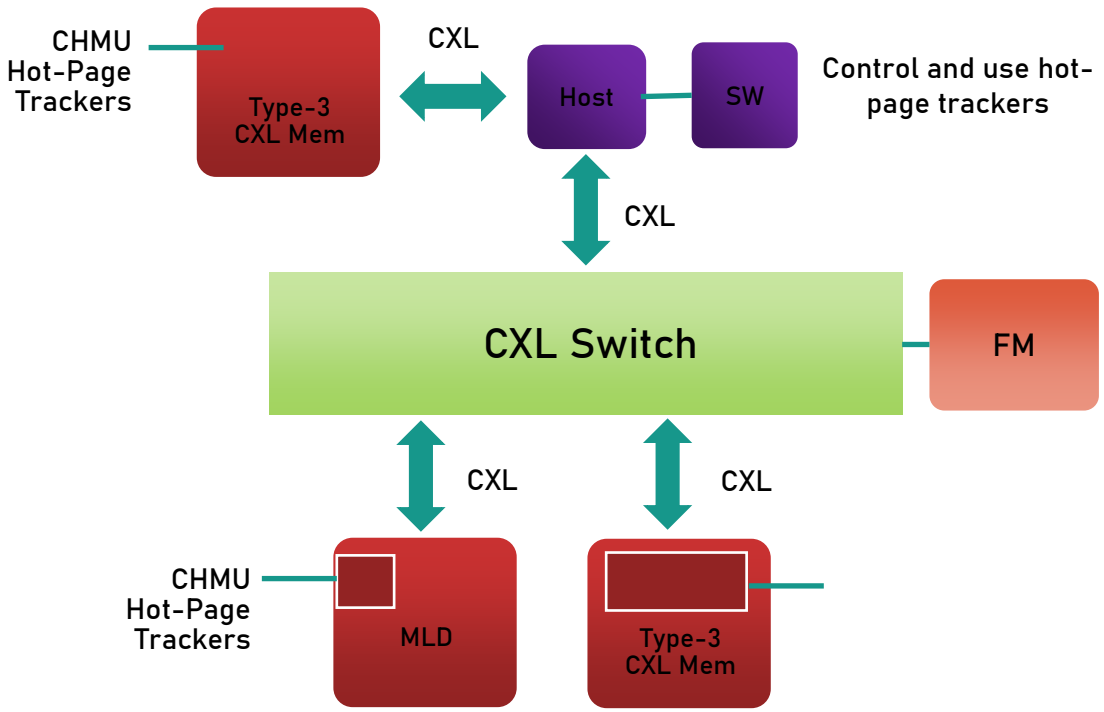

More efficient SW Memory Tiering
Better Perf, lower TCO

Challenges faced by the today's SW tiering solutions

- Requires application changes
- Must trade-off accuracy against perf overhead
- CPU vendor specific

CHMU addresses these problems

- Enables generic OS based solutions
- Works for simple and pooling memory devices
- Tracking offloaded to CXL memory device
- Configurable and extensible

CXL Trusted Security Protocol (TSP)

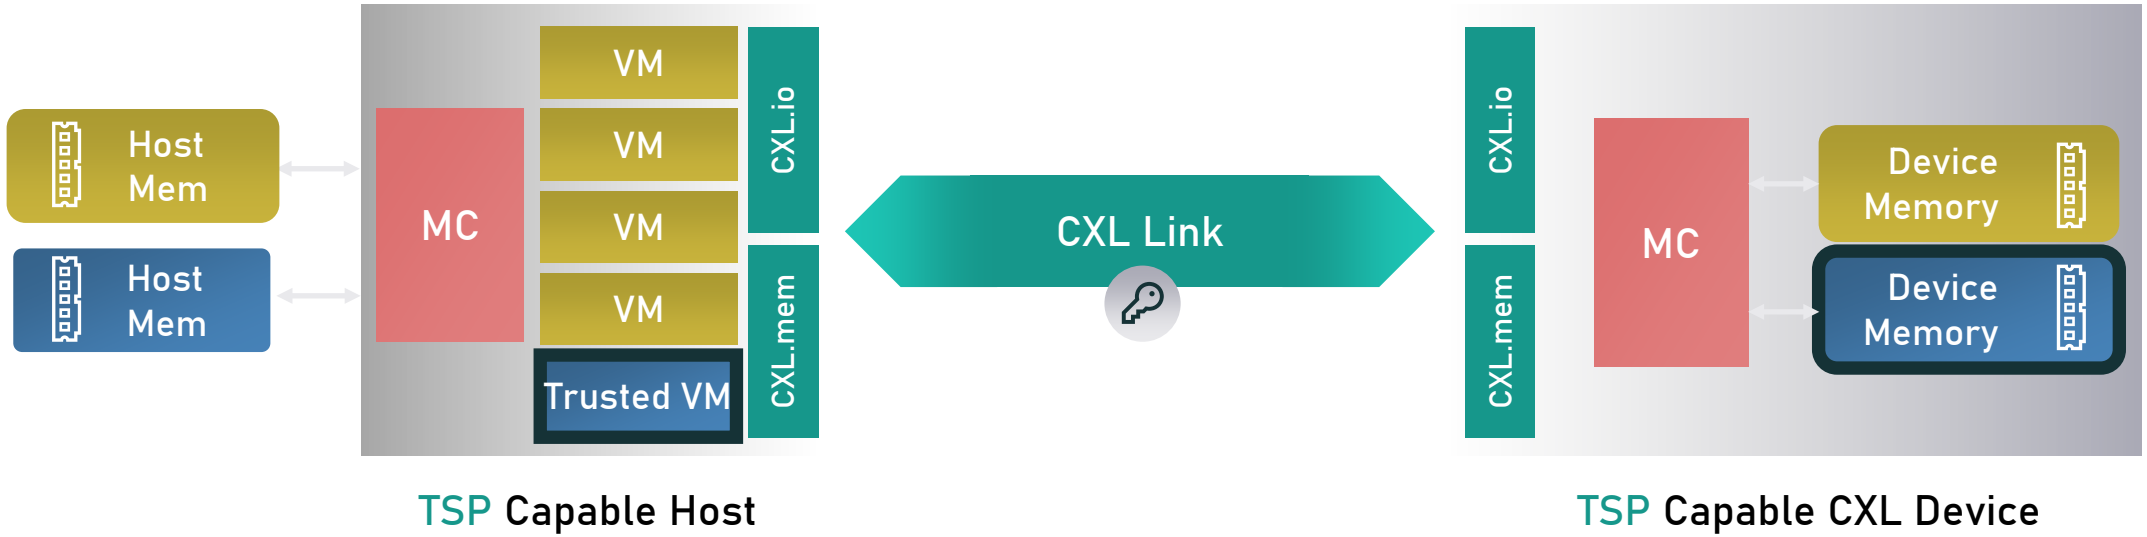

Confidential computing is a pre-requisite to migrating sensitive WLs to Cloud

TSP extensions enable CXL devices to participate in Confidential Computing

- CXL memory device can hold proprietary model parameters

2025 CXL Events

- April
  - CXL DevCon 2025, April 29-30
- August
  - FMS: The Future of Memory and Storage, August 5-7 at the Santa Clara Convention Center
    - Visit CXL kiosk at Booth No. 725
    - Attend CXL presentation on **Wednesday, August 6 at 9:45 am PT** to learn about ready to deploy CXL solutions!
- September
  - AI Infra Summit, September 9-11 at the Santa Clara Convention Center
- November
  - SC'25 (Supercomputing 2025), November 16-21 in St. Louis, MO
    - Visit the CXL Pavilion at Booth No. 817.
- Other events will be added as opportunities arise
  - Follow the CXL Consortium via LinkedIn (@CXL Consortium) for updates!

Learn how to program CXL Memory Modules in SNIA's **CXL.mem Programming Workshop**

Live online and at CXL DevCon!

https://www.snia.org/pm-summit/hackathon

SNIA | CMS A SNIA Community

SNIA | CMS
A SNIA Community

CXL Benefits, Implementation, and Ecosystem

Torry Steed

CXL® Memory Expansion Advantages

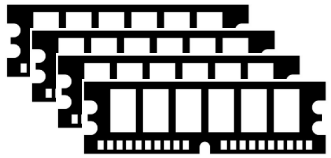

Increased Memory Capacity

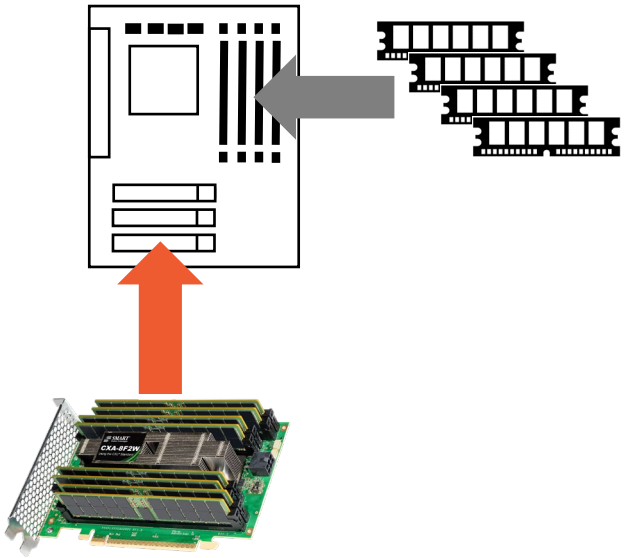

Reduced Memory Cost

DDR5 DIMM Price

Capacity	Price/Gb
64GB	1x
96GB	1.1x
128GB	1.4x
256GB	2.0x

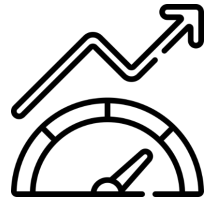

Increased Performance

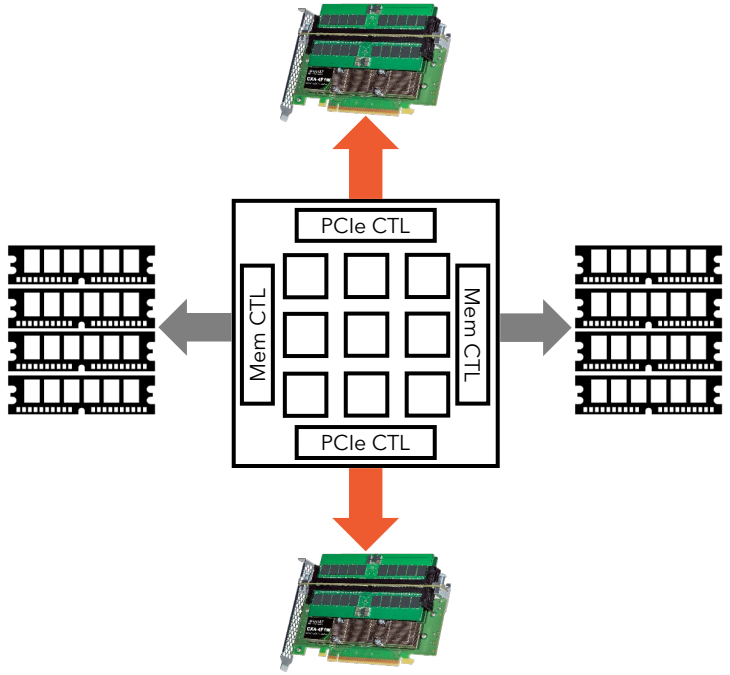

SNIA | CMS A SNIA Community

Applications Benefiting from Large Memory

Real Time Data Analytics Retrieval and Processing

In-Memory Databases

- Keep the entire dataset in memory for real time/fast processing
- Examples: SAP-Hana, Redis
- Real time Feature Stores for Edge Inference applications

HPC

- Climate modeling
- Genomics research
- Fluid dynamics
- Particle Physics

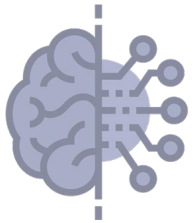

Big Data Analytics, AI and Deep Learning

- Newer analytics and machine learning utilizing very large datasets
- Enables rapid query and model training

Financial Modeling

- High frequency trading complex risk analysis
- Processing, ingesting vast amounts of market data in real-time

CXL® Generational Improvements

CXL 1.1
In-server memory expansion

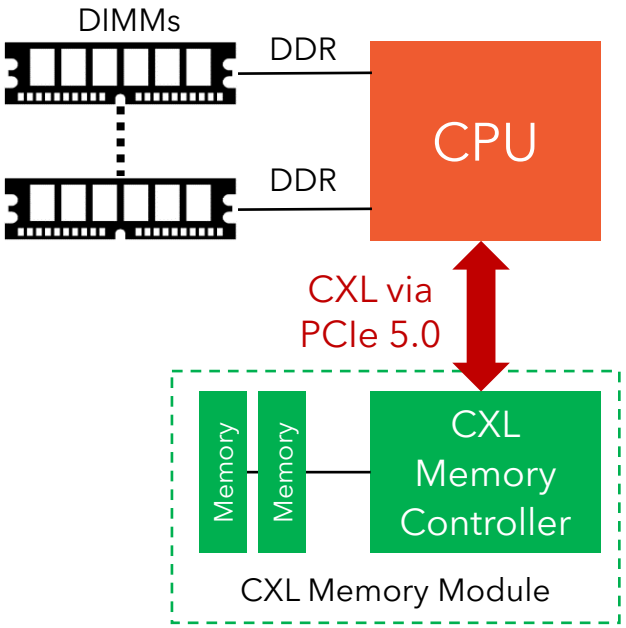

CXL 2.0
In-rack memory expansion

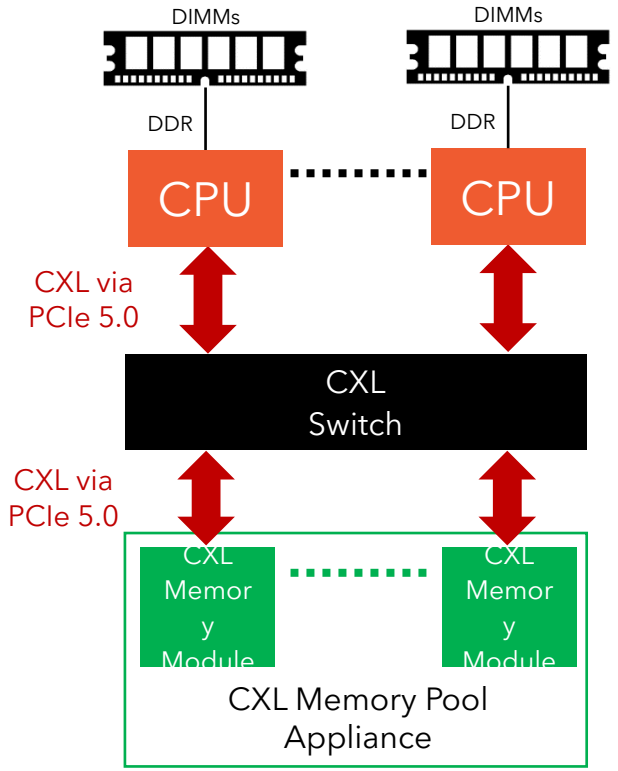

NOTE: Also added support for NV-CMM

CXL 3.0
Fully disaggregated memory pool

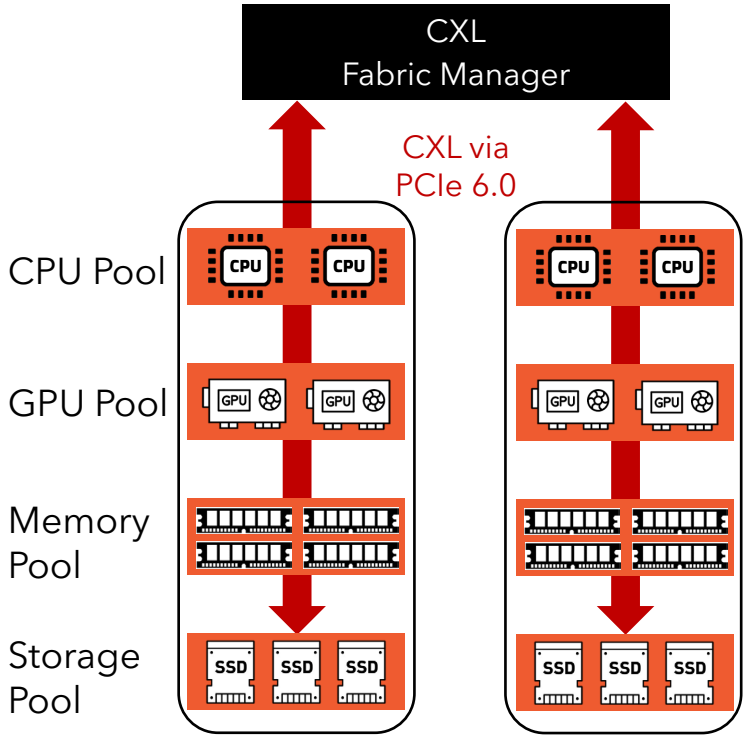

SNIA | CMS A SNIA Community

CXL® Memory Adoption

In-Server Memory Expansion

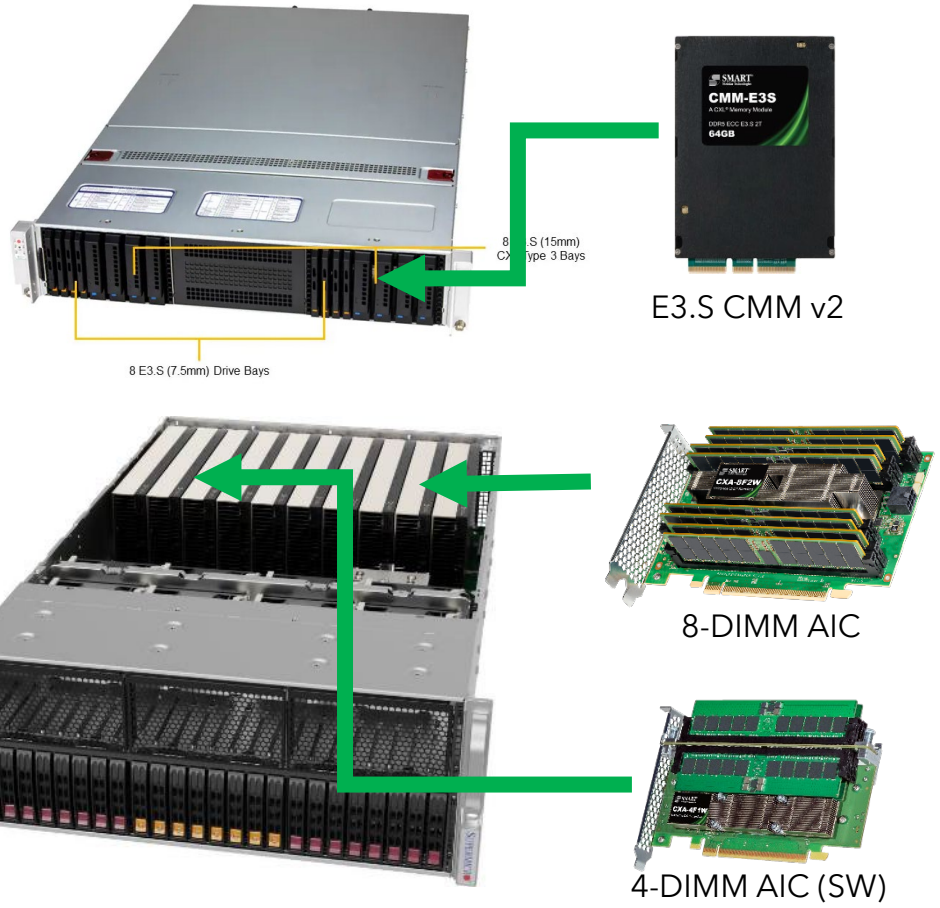

In-rack Memory Expansion
CXL 2.0 Switching

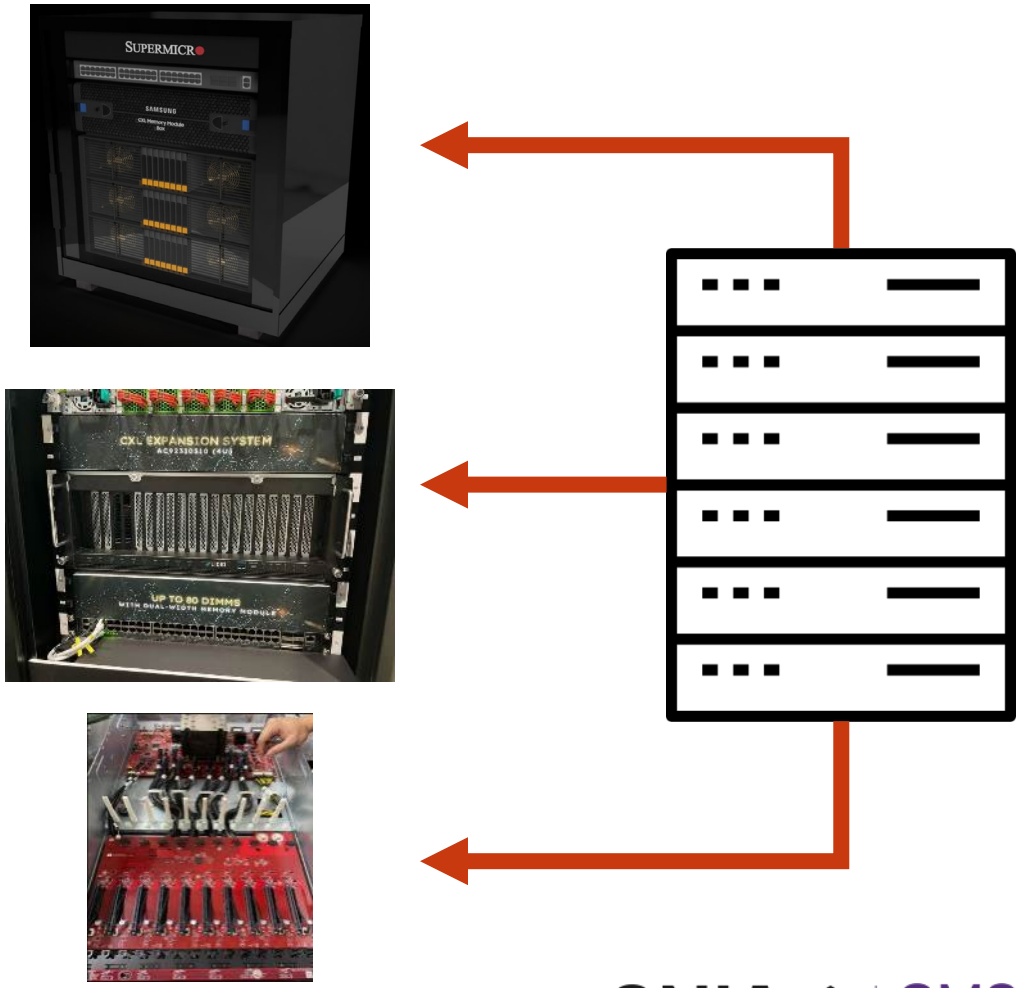

SNIA | CMS A SNIA Community

CXL® and DRAM Latency

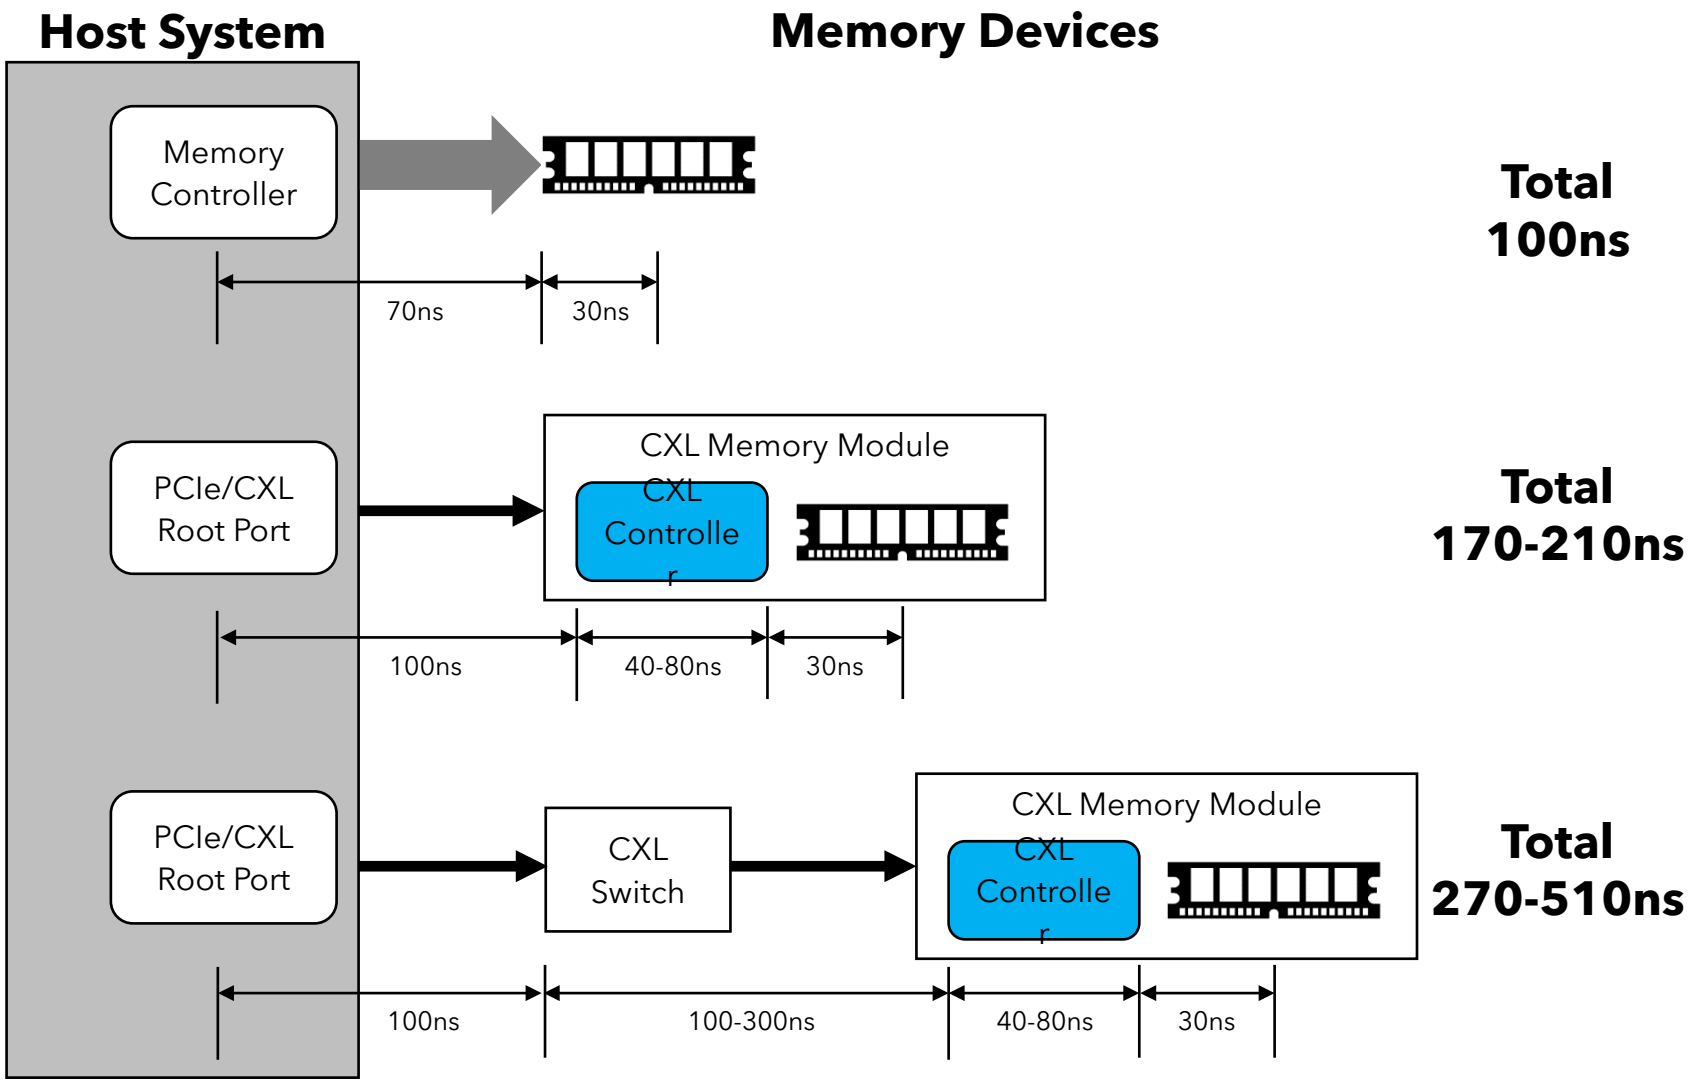

Total 100ns

Total 170-210ns

Total 270-510ns

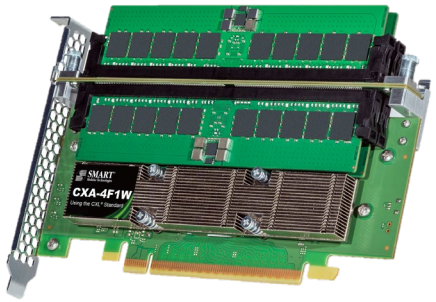

Measured Round Trip Latency

205ns

SNIA | CMS A SNIA Community

CXL® and DRAM Bandwidth

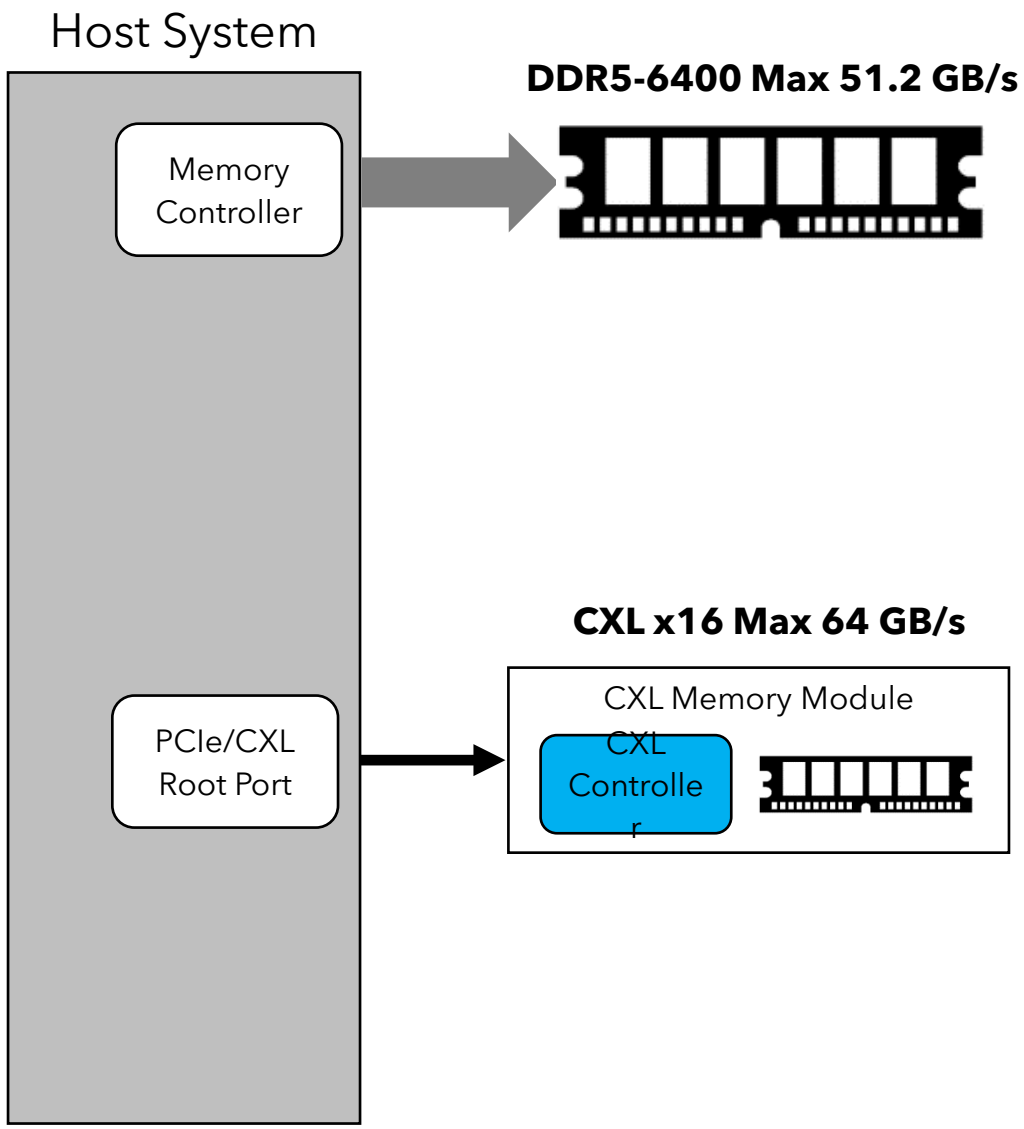

Generation	Common Name	MT/s	GB/s
DDR5	DDR5-4800	4800	38.4
DDR5	DDR5-5600	5600	44.8
DDR5	DDR5-6400	6400	51.2

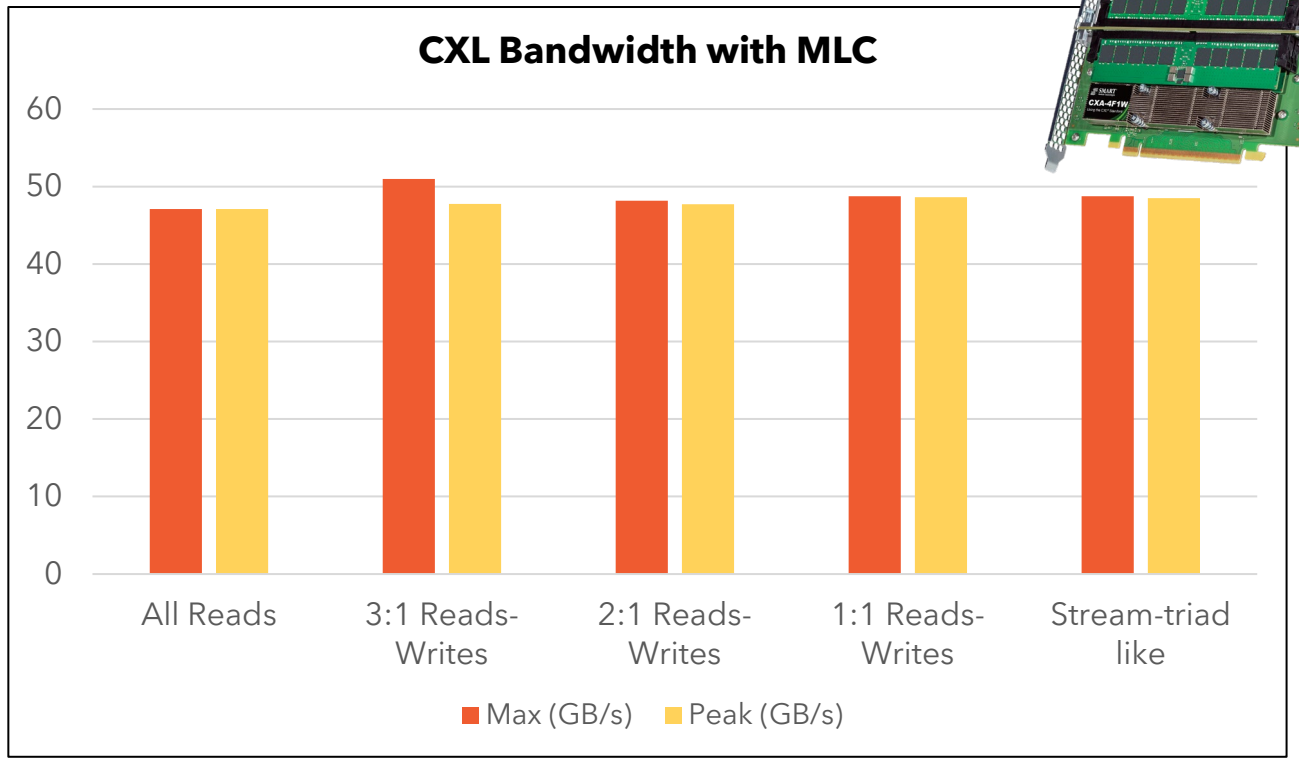

SNIA | CMS A SNIA Community

Real World Results

In-Memory Database (MS SQL + TPC-H)	**Machine Learning** (Apache Spark™ SVM)	**High Performance Computing** (CloverLeaf)
	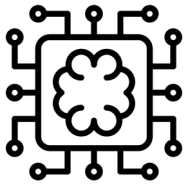	
• Consumes only 50% of memory bandwidth • Capacity limited	• Memory bound workload • Sensitive to latency & capacity	• Bandwidth intensive workload • 80% mapped to DRAM • 20% mapped to CXL
• 44%-88% reduction in SSD paging I/Os • 23% performance improvement	• 2.21x performance improvement compared to DRAM only	• CXL increases memory bandwidth by 33% • 17% performance improvement

Source: CXL Memory Expansion: A Closer Look on Actual Platform by Micron and AMD

NYCU Taiwan

Study on CXL Affects on LLM Training with CPU Offloading

National Yang Ming Chiao Tung University

Example Configuration

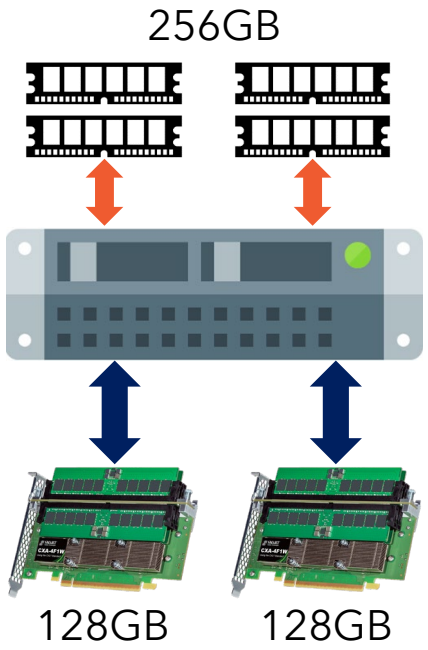

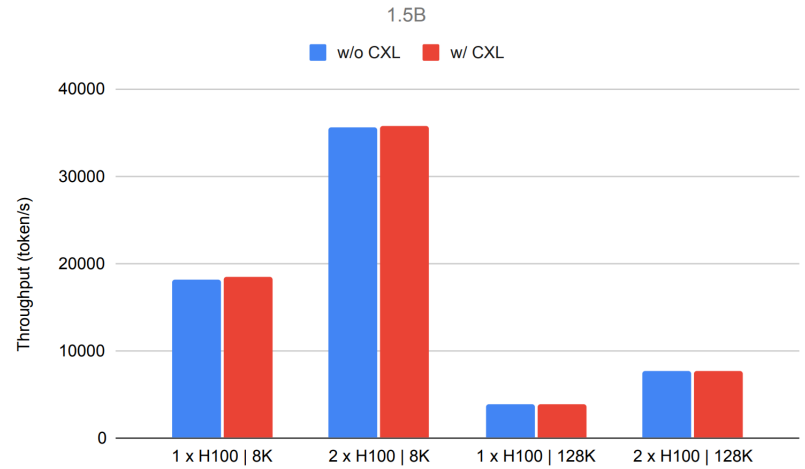

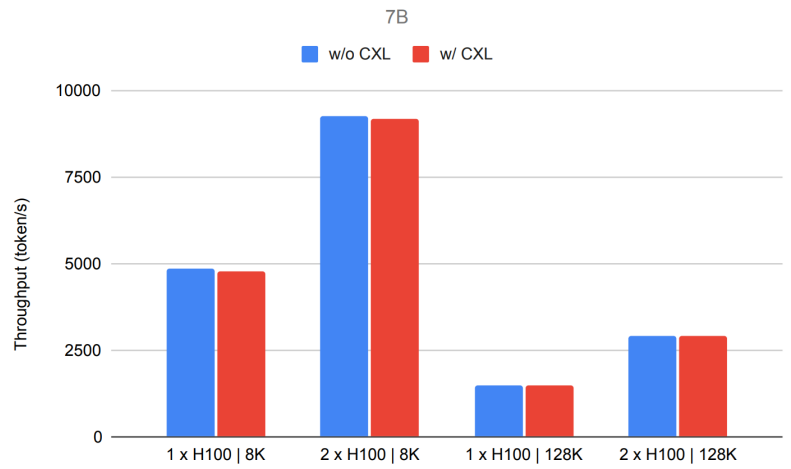

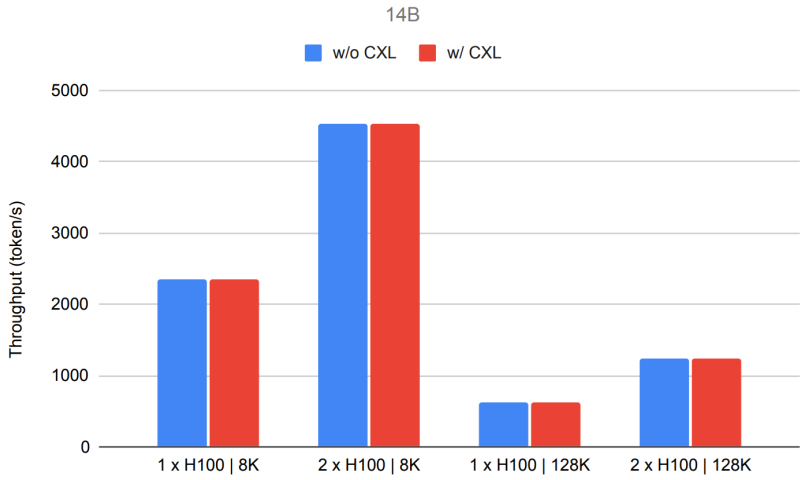

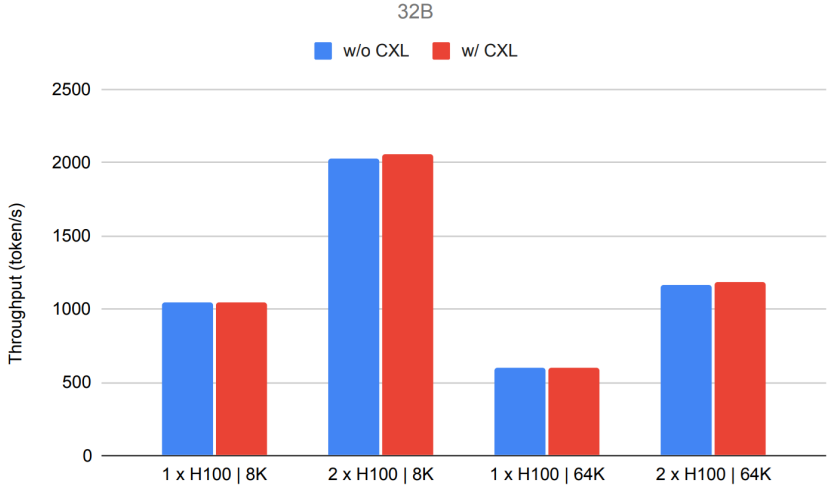

SNIA | CMS A SNIA Community

Level1Techs Real World Demo

"Benchmarking the Xeon 6 6787P in the Supermicro 222H-TN"

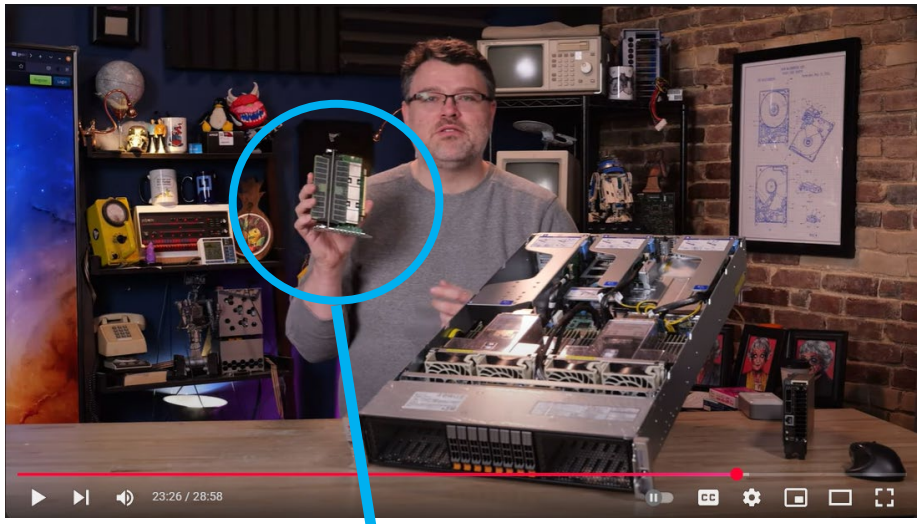

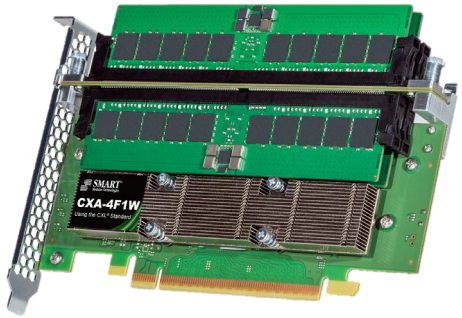

SMART 4-DIMM AIC

```
root@w-SYS-222H-TN:~/Linux# ./mlc
Intel(R) Memory Latency Checker - v3.11b
Measuring idle latencies for sequential access (in ns)...
                Numa node
Numa node            0       1       2       3       4
       0         125.3   148.5   402.0   422.8   650.6
       1             ^C
Exiting...


root@w-SYS-222H-TN:~/Linux# ./mlc
Intel(R) Memory Latency Checker - v3.11b
Measuring idle latencies for sequential access (in ns)...
                Numa node
Numa node            0       1       2       3       4
       0         125.9   148.3   403.7   429.6   651.7
       1         147.4   123.6   423.2   397.5   647.0
       2         400.3   430.2   124.7   148.8   320.1
       3         424.3   399.7   140.8   123.7   323.5

Measuring Peak Injection Memory Bandwidths for the system
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using traffic with the following read-write ratios
ALL Reads        :      560754.1
3:1 Reads-Writes :      515094.5
2:1 Reads-Writes :      499505.9
1:1 Reads-Writes :      495507.8
Stream-triad like:      493355.5

Measuring Memory Bandwidths between nodes within system
Bandwidths are in MB/sec (1 MB/sec = 1,000,000 Bytes/sec)
Using all the threads from each core if Hyper-threading is enabled
Using Read-only traffic type
                Numa node
Numa node            0       1       2       3       4
       0        140119.6        141680.2        89727.9 92632.[illegible] 52427.9
       1        140570.6        141230.4        92619.7 90423.[illegible] 48792.7
       2        90309.1 94558.2 141052.9        141918.1        49862.1
       3        89690.0 89595.8 141793.1        141231.4        46754.6


Measuring Loaded Latencies for the system
Using all the threads from each core if Hyper-threading is enabled
Using Read-only traffic type
Inject  Latency Bandwidth
Delay   (ns)    MB/sec
==========================
 00000  516.58   560865.4
 00002  482.80   560055.2
 00008  497.52   559511.0
 00015  494.33   560060.2
 00050  453.51   560140.8
```

Latency
~320ns

Bandwidth
46 - 52GB/s

Source: Benchmarking the Xeon 6 6787P in the Supermicro 222H-TN

CXL® Ecosystem "Connecting the Dots"

AMD

Year	AMD EPYC	CXL
2021	Milan	No
2022	Genoa	CXL 1.1
2023	Bergamo	CXL 1.1
2024	Turin	CXL 2.0
2025	Venice	CXL 3.0

intel

Year	Intel Xeon	CXL
2021	Ice Lake	No
2023	Sapphire Rapids	CXL 1.1
2023	Emerald Rapids	CXL 1.1
2024	Granite Rapids	CXL 2.0
2025	Diamond Rapids	CXL 3.0

BIOS Support

insyde

Full support for CXL

intel

Intel® Server M50FCP Family
CXL BIOS version v0.1.02.0002

ami

AMIBIOS supports CXL

phoenix technologies

Pending

OS Support

Kernel v6.5 and
above recommended

Windows Server
Beta testing CXL

vmware

VMWare
supports CXL

SNIA | CMS A SNIA Community

Where does that leave CXL®?

SMART® Modular Technologies
A PENGUIN SOLUTIONS BRAND

- CXL memory expansion is available now
- "Off the shelf" systems are here and working
- CXL benefits are real
- CXL rack pooling appliances are coming soon
- CXL rack to rack sharing is on the horizon

SNIA | CMS A SNIA Community

SNIA | CMS
A SNIA Community

Questions?

Next Steps. . .

- ✓ Please rate this webinar and provide us with feedback
- ✓ Look for a Blog answering questions from this webinar at www.sniablog.org
- ✓ Get more SNIA education!
  - **Live**
    - FMS Future of Memory and Storage, August 4-7, 2025 Santa Clara CA fms.com
    - SNIA Developer Conference, September 15-17, 2025 Santa Clara CA sniadeveloper.org
    - SC25, November 16-21, St. Louis MO sc25.supercomputing.org
  - **Online**
    - This webinar and many other videos and presentations on today's topics are in the **SNIA Educational Library** https://snia.org/educational-library
- ✓ Join SNIA and the Persistent Memory Special Interest Group
  - www.snia.org/join
  - https://www.snia.org/groups/cms

SNIA | CMS
A SNIA Community

Thank You