Session: CXLT-201-1: CXL Form Factors

Getting Started with CXL® in 2024

Presenter:

Andy Mills

SMART Modular Technologies

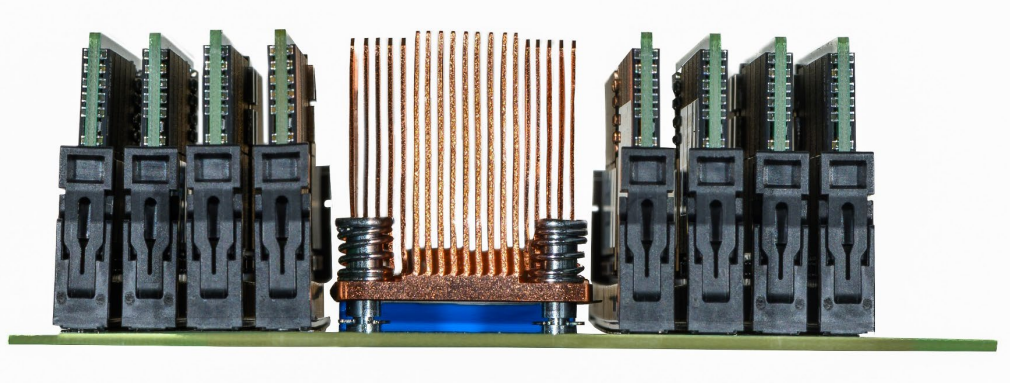

CXL Market Outlook

Percent of Servers with CXL

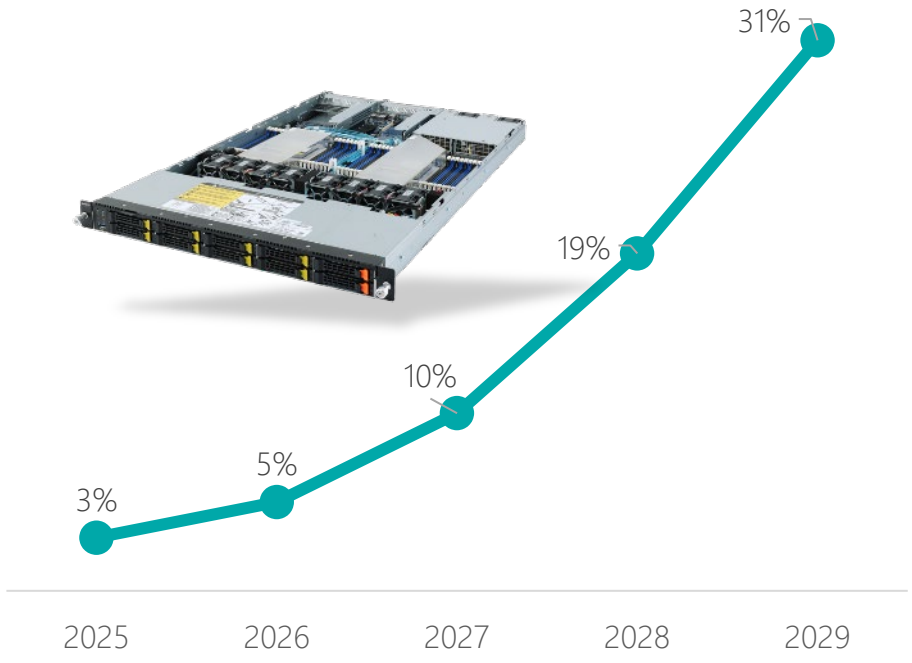

TechInsights, March 2024

CXL DRAM Revenue Forecast

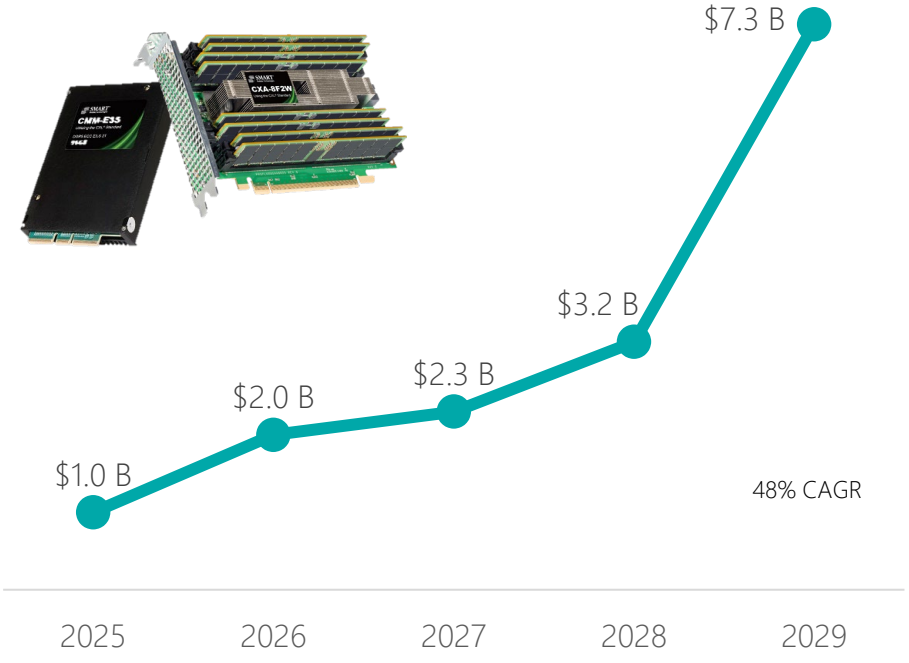

TechInsights, March 2024

Why We Need CXL Memory

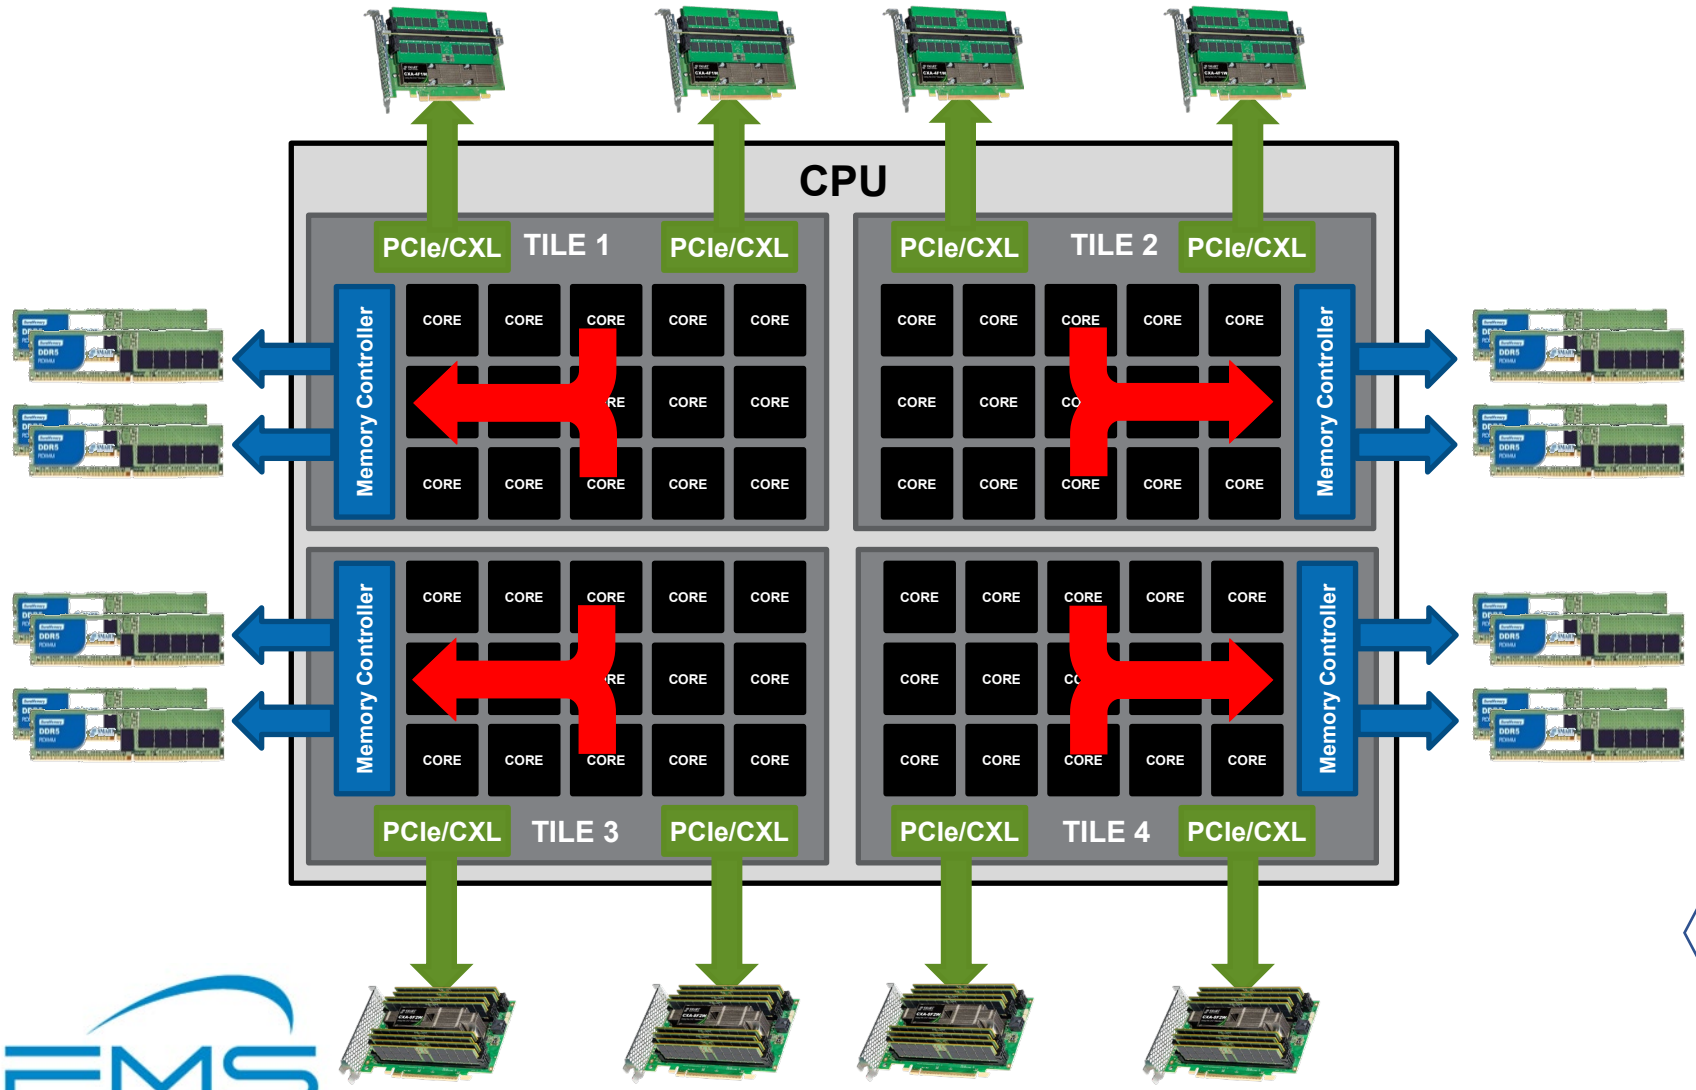

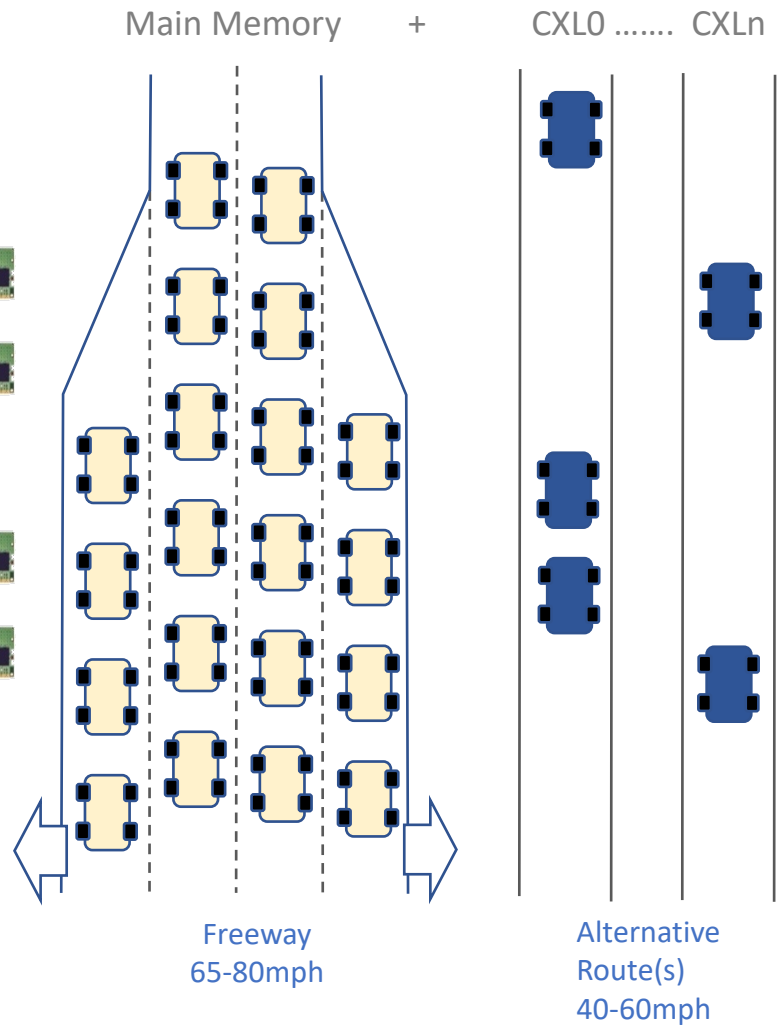

Source: A Case for CXL-Centric Server Processors

Server Memory Expansion

SMART® Modular Technologies

Without CXL – "Scale out"

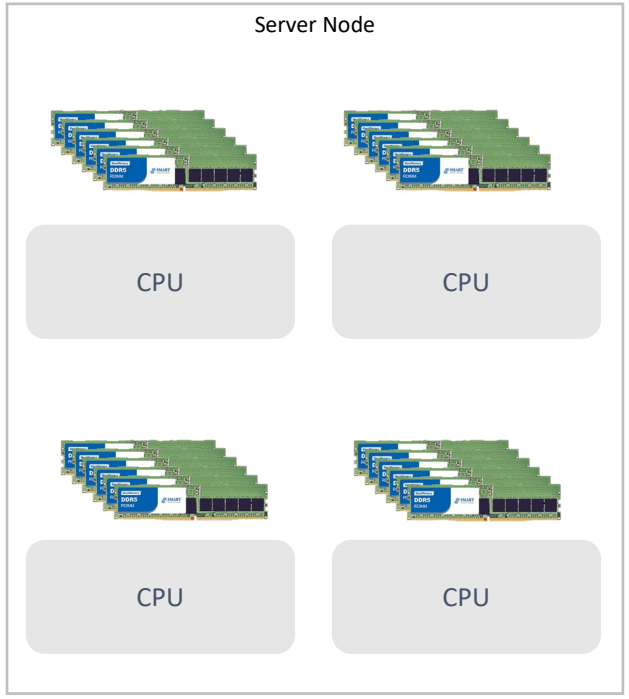

- Add more memory by adding more CPUs+memory
- Adding more memory requires more CPUs

With CXL – "Scale-up"

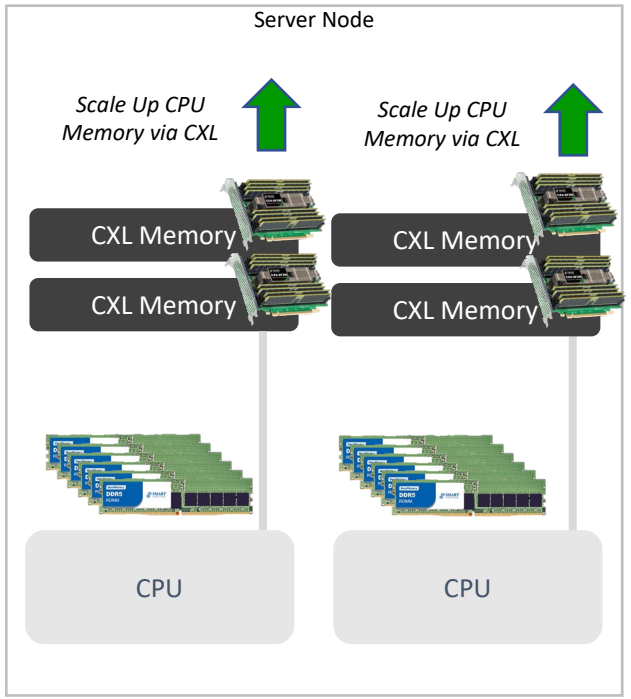

- Memory expansion no longer limited to CPU additions
- CXL allows addition of more memory to the existing CPUs

✓ Ease of CXL Adoption
- Add-in-cards CXL memory leveraging low cost RDIMMs (existing server chassis)
- E3.S modules (newer server chassis)

✓ Added Benefits using CXL
- Lower power and cost
- Enhanced RAS functions
- Enhanced telemetry
- Ease of expansion

✓ Up to 50-60% savings in system cost

FMS

CXL Rollout Expectations

2024/2025

Memory Expansion
CXL 1.1/2.0 Enabled Servers

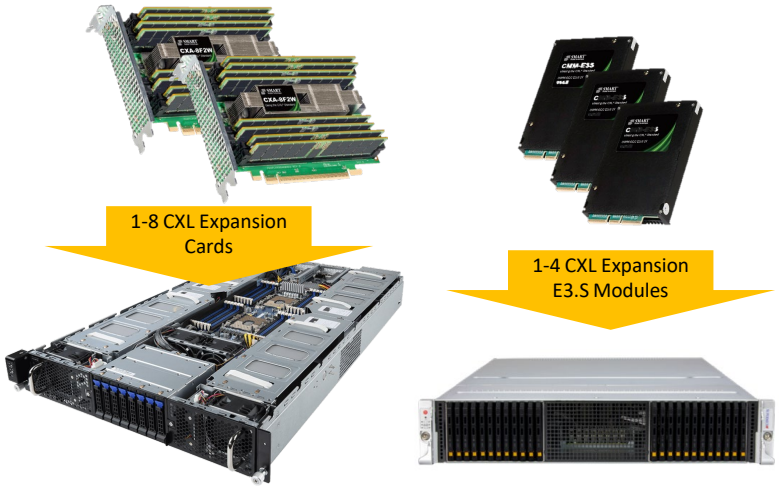

2025/26

Memory Pooling
CXL 1.1/2.0 Direct or Switched

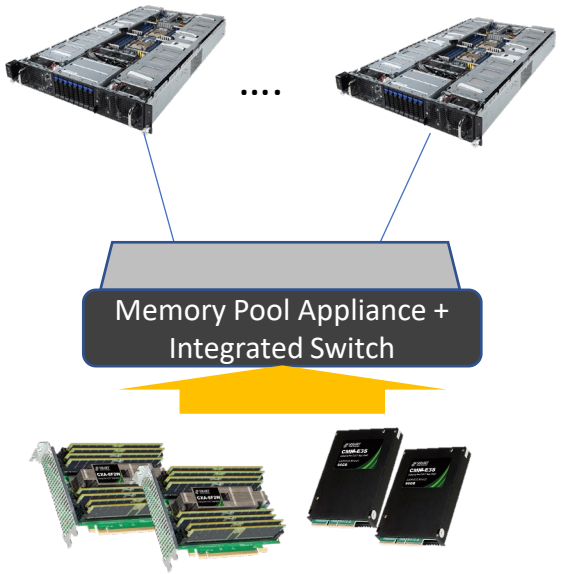

2026/27

Memory Sharing
CXL 3.1 Direct and Switched

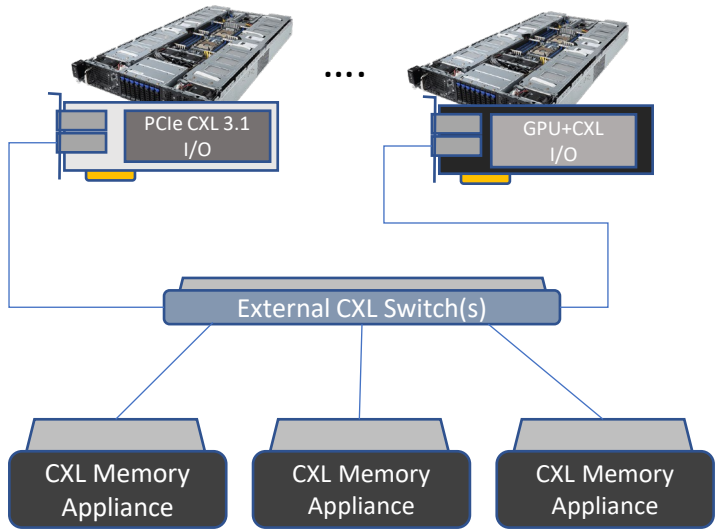

Images courtesy of Gigabyte and SuperMicro

Enabling Use of RDIMMs with CXL

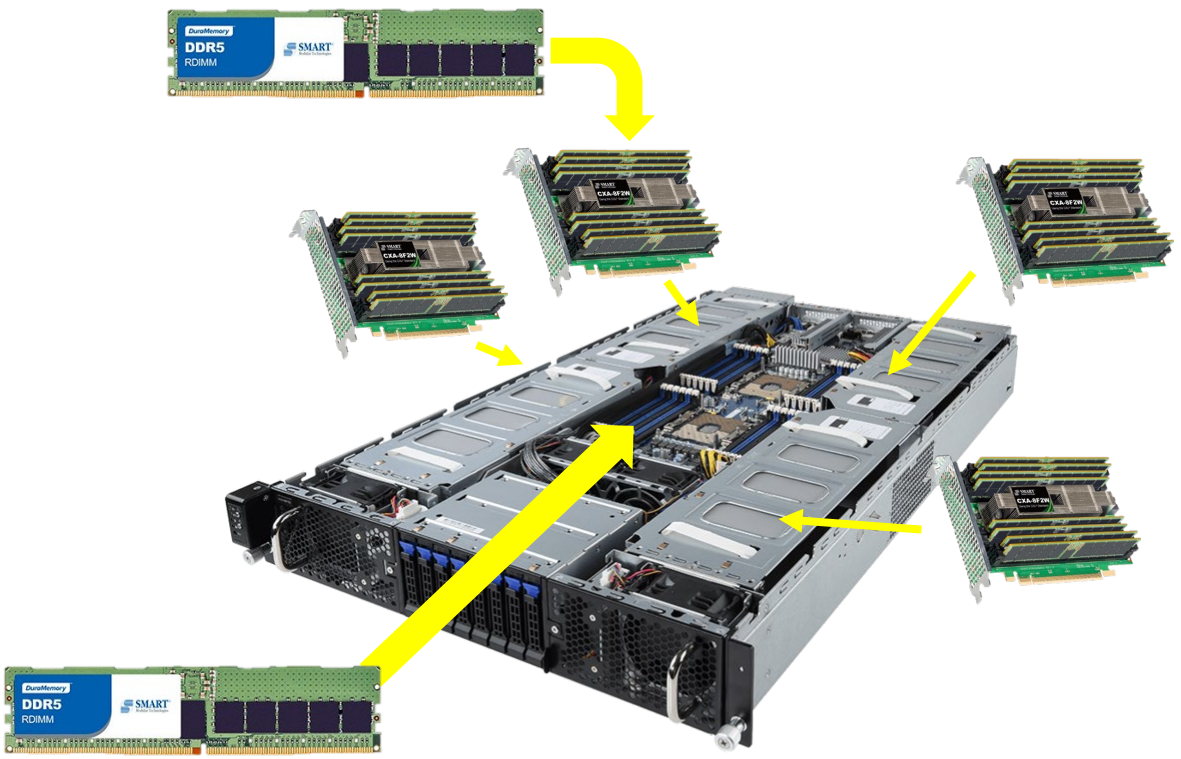

2U 4 or 8 GPU Server
Repurposed as a Memory Server

- Leverage commonly available DDR5 RDIMMs
- Octal RDIMMs require higher power (120-130W)
- Multiple existing server options available built for higher powered GPUs with AUX power connections
- Example: Increase RDIMM count from 24 to 88 in a 2U, 8-way GPU server

RDIMMs	No CXL	With CXL
96GB	2304 GB	8448 GB
128GB	3072 GB	11264 GB
256GB	6114 GB	22528 GB

Case Study: CXL-enabled 1U GPU Server

- Requirement: 20TB in 4 1U servers for large in-memory analysis workload under $200K cost
- Without CXL – only option is 256GB TSV based DIMMs
- Standard GPU server with 4x 8-DIMM CXL expansion card added 32 extra DIMMs
- Up to 66% lower system cost

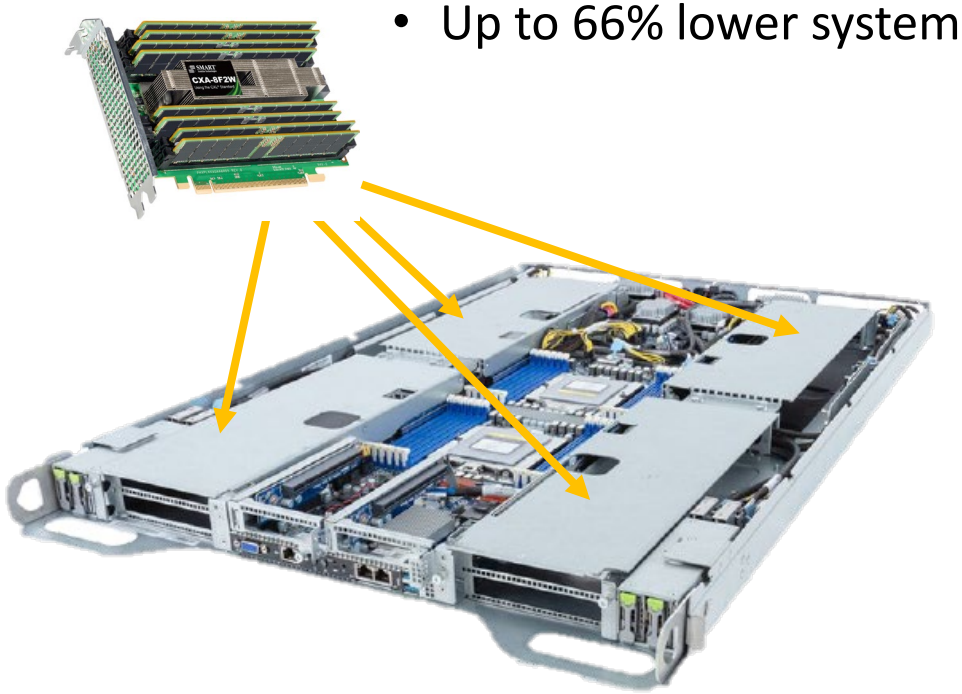

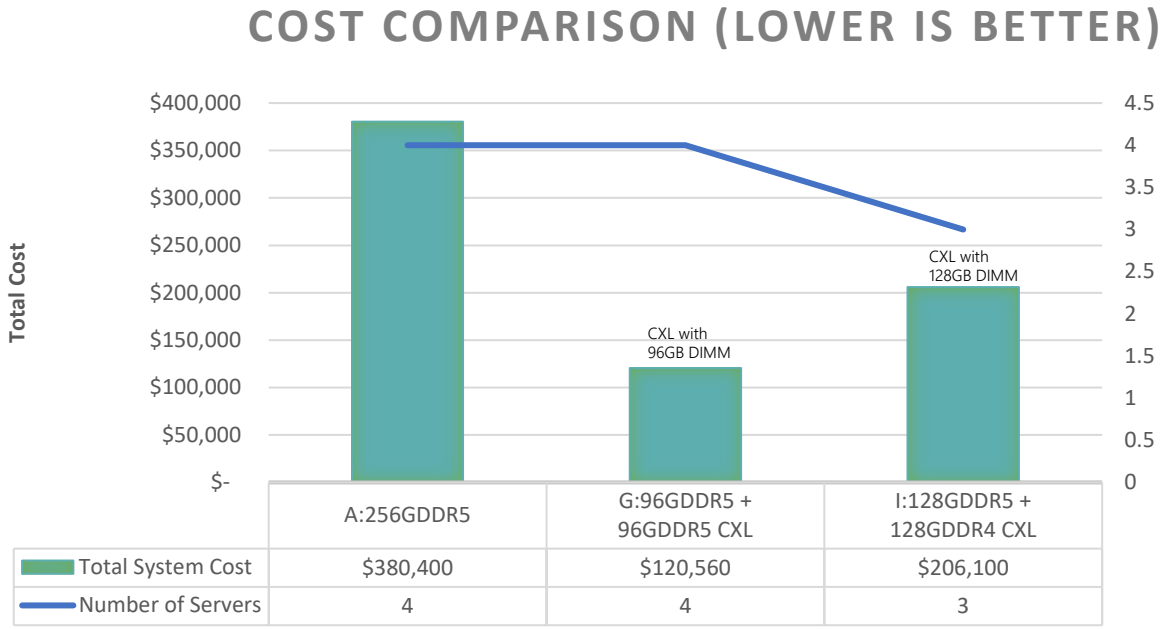

Source: SMART Modular based on estimated memory and system pricing as of April 2024

Cost Reduction Case Study – 1TB Memory

DDR5 DIMM Price

Capacity	Price[1]	Price/Gb
32GB	$120	$0.47
64GB	$230	$0.45
96GB	$415	$0.54
128GB	$1,000	$0.98
256GB	$3,300	$1.61

(1) Approximate as of May 2025

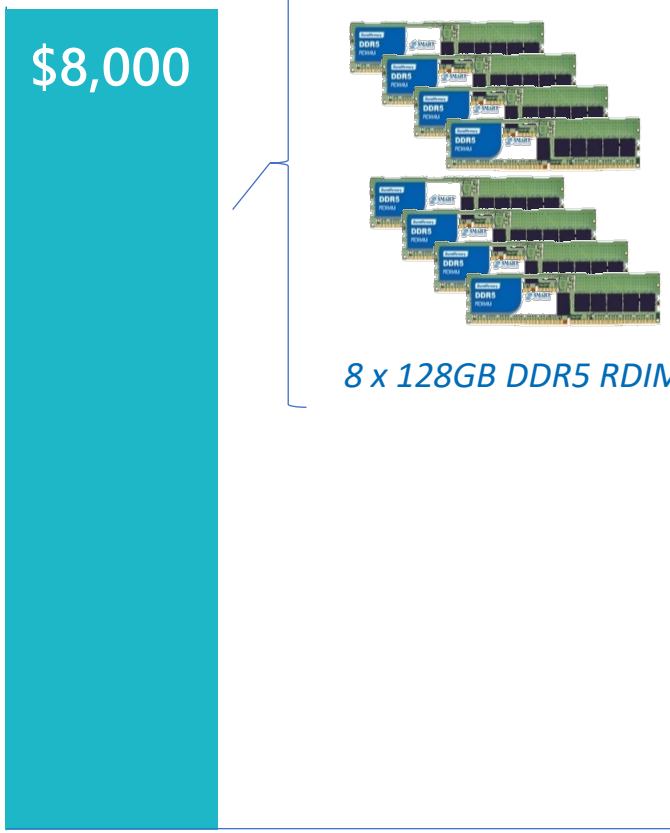

Deploying CXL in Standard 2U or 1U Servers

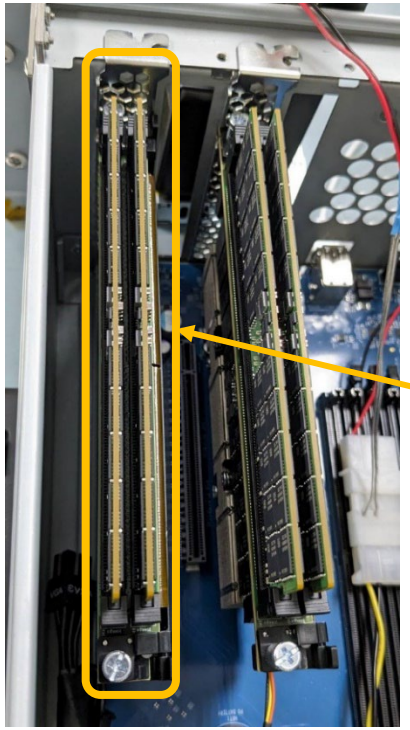

Conventional:
Takes up 2 slots and blocks off PCIe connectors for other cards

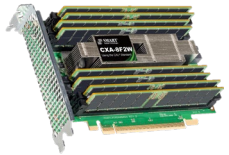

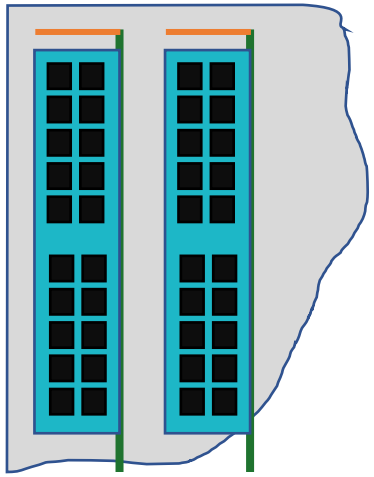

4-DIMM Slim Board:
Takes up 1 slot
=
Higher Density, utilze all PCIe slots plus works in 1U systems

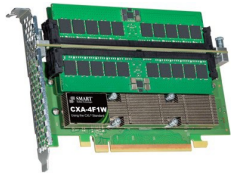

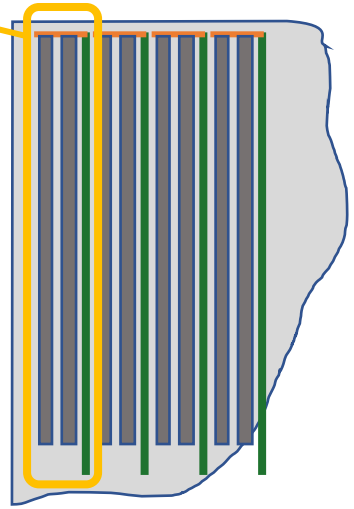

DIMMs vertically mounted on unique flex board

- Slim CXL card option available utilizing RDIMMs take up 1 slot only
- Designed for more commonly available chassis/PSUs
- Powered from PCIe slot only (no AUX power required)
- Flexible Options
  - One x16 for performance
  - Dual x8 on a single riser for capacity

E3 EDSFF CXL Options

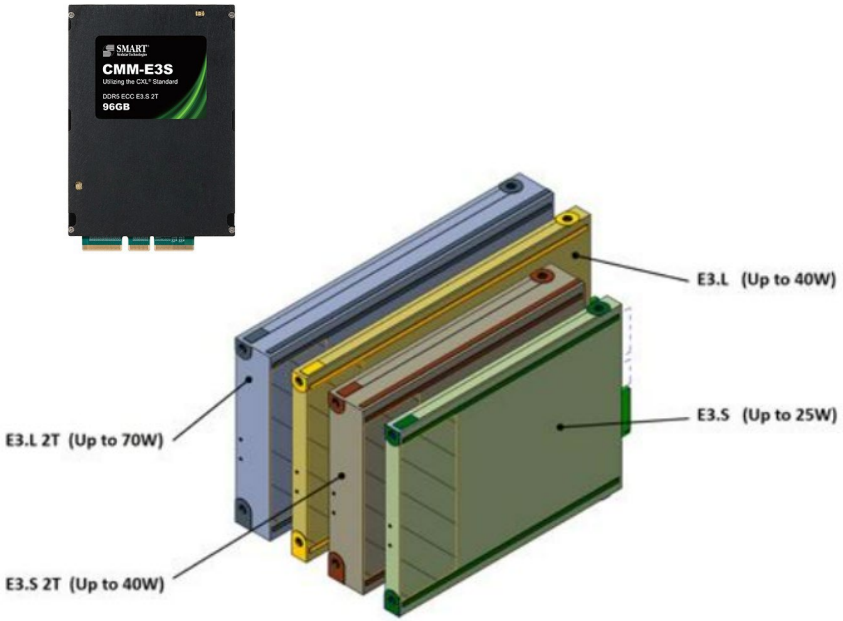

EDSFF E3 EDSFF
(Source: SNIA specifications)

- E3.S/L 2T
- Dual ported devices dual x4/x8
- High Reliability/Memory Failover
- Accelerated/Intelligent CXL Options

- E3.S/L 2T
- Single ported devices
- High Performance x8/x16
- Non-volatile/Hybrid CXL Options

- E3.S 1T
- Single ported devices
- Low-Medium Performance x4/x8
- 1T enables up to 40 per server

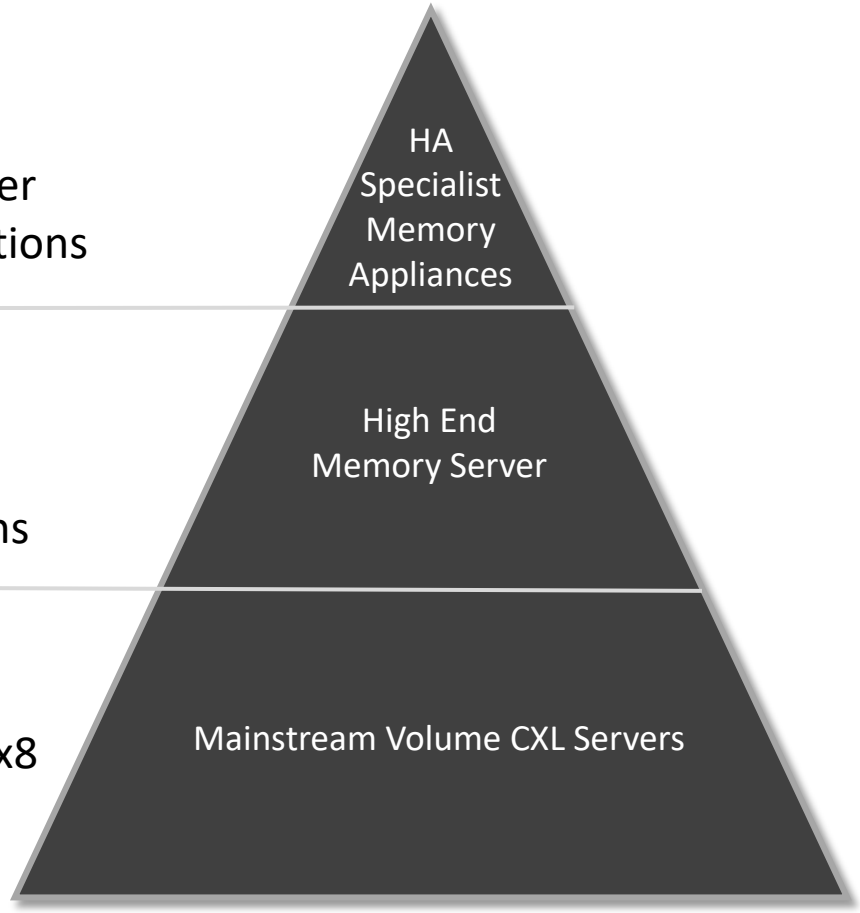

A Word on Volatility…

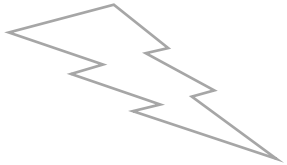

- Memory intensive applications require non-volatile options at memory speeds
  - Faster check-pointing and recovery in AI/ML
  - Saving system state at memory speeds
  - Protecting cached data at memory speeds during power losses
- E3.S/L enables a new class of non-volatile devices
  - Single port for mainstream/high end servers
  - Dual port for high end, high availability servers
  - Fully integrated unit including the energy source module
- Two approaches
  - NV-CMM: conventional JEDEC compliant NVDIMM with a CXL front end
  - Hybrid memory/flash device (aka Memory Semantics or CMM-H)

Non Volatile CXL Modules

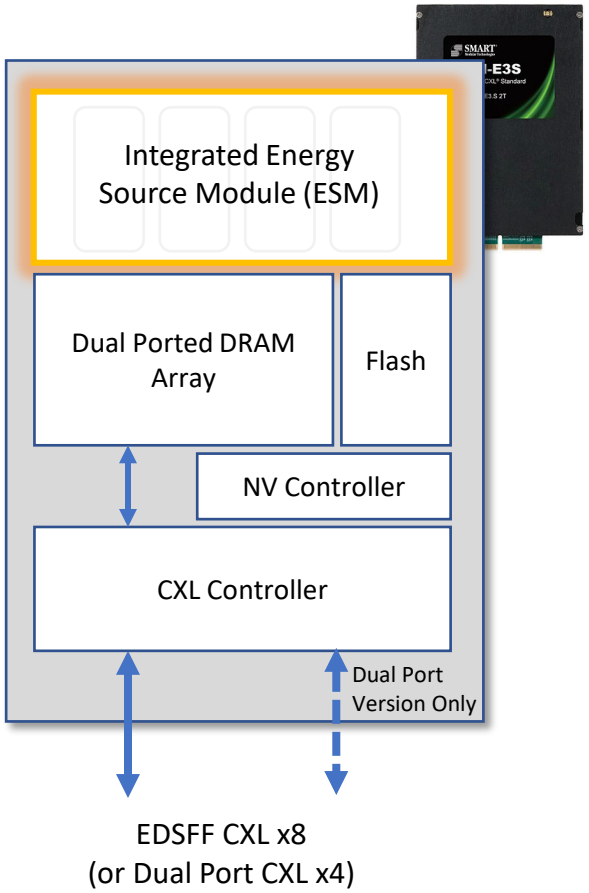

- Designed for maximum system compatibility
  - Evolution of classic NVDIMM
  - System must support CXL GPF command generation on power loss and leave sufficient time for restore
- Integrated energy source module
  - Easier to service/maintain
  - Increased reliability
- Significant improvements over classic DIMM based NVDIMM
  - Avoids DIMM balancing acts (mixing and matching of DDR speeds)
  - Zero/less dependence on Platform BIOS
  - Zero dependence on SMBus board architecture, management and configuration
  - Single FRU for memory and ESM

NV-CMM vs. NVDIMMs

	NV-CMM	NVDIMM-N	NVDIMM-P
Interface	CXL.mem compatible interface (x8, x16 PCIe Gen 5/CXL)	DDR3/4 DIMM interface [1]	DDR3/4 DIMM interface
Form factors	E3.S, E3.L or Add-in-card	Traditional DIMM form factor	Traditional DIMM form factor
Backup Energy Source	Yes, integrated or external battery or supercaps	Yes, cable to external battery or supercaps or separate system backup power source	Not needed. Media itself is persistent.
Hot Plug Capable	Yes	No	No
Backup method	Primary media is DRAM, backed up to NAND flash during power fail or CXL GPF command	Primary media is DRAM, backed up to NAND flash during power fail or hardware trigger	
Performance	Same speed and latency as other CXL volatile memory components	Same speed and latency as other DIMMs in system	Slower than RDIMM speeds
Benefits	Fully integrated energy source Common management across Intel and AMD Significantly less dependent on DDR type and speeds	Common DIMM formfactor	Large capacities
Challenges	Requires newer E3.S/L chassis to utilize.	Different management approaches for Intel vs. AMD platforms, limited server SKUs, no follow on DDR5 versions	Availability e.g. 3DXP now obsolete, Re-RAM never quite there…. Often not memory speeds.

(1) DDR5 versions developed but never brought to market in favor of emerging CXL

CXL Reality Check

- Server support is still spotty but much better since Q2 of 2024
  - Advertised support doesn’t mean it supports CXL out of the box
- Server BIOS version is important
  - Both AMD and Intel have specific BIOS minimum versions
  - Check with server vendor for latest BIOS support
  - Non volatile/GPF support
- Operating system support
  - Linux kernel 6.5 or higher, preferably 6.7 appears to be most stable
  - Windows – shows up as additional memory in 2022, but lacking generic tools

Call to Action

- Check out the SMART CXL products at the following booths:
  SNIA, Microchip, Asterlabs booths

- Check in with your server vendor to verify which models have been verified with CXL add-in cards and/or E3.S modules
  Also check for GPF and power loss support as we migrate to CXL 2.0 compliant systems

- Follow CXL progress at CXL Consortium – https://computeexpresslink.org/

Thank you!