

Compute, Memory and Storage Technology Trends for the Application Developer

SNIA Moderator: Tom Coughlin, SNIA

- Computational Storage and Memory (David McIntyre, Samsung)
- CXL and UCIe (Debendra Das Sharma, Intel)
- Emerging Persistent Memory Types (Arthur Sainio, Smart Modular)
- Bridging to the Application Layer (Arvind Jagannath, VMware)

Today's Agenda

- About SNIA
- Panel Introduction and Brief Remarks
- Discussion

What Does SNIA Do?

- SNIA is a non-profit global organization dedicated to developing standards and education programs to advance storage and information technology.

Industry Leading Organizations



180

Active Contributing Members



2,500

IT End Users & Storage Pros Worldwide



50,000

www.snia.org/cmsi
www.snia.org/memcon



What is CMSI?

- Part of SNIA, the SNIA Compute, Memory, and Storage Initiative (CMSI) is a community of storage professionals and technical experts who support:
 - The industry drive to combine processing with memory and storage
 - The creation of new compute architectures and software to analyze and exploit the explosion of data creation over the next decade
- CMSI's four Special Interest Groups –
 - Computational Storage
 - DPU
 - Persistent Memory,
 - Solid State Drives
- evangelize and educate on these technologies to the industry.

Speaker



- David McIntyre, SNIA Board of Directors/Director Product Planning and Business Enablement, Device Solutions America, Samsung
 - David McIntyre drives computational storage acceleration solutions and strategy for Samsung. He has held senior management positions with IBM, AMD and Intel along with Silicon Valley startups. He has consulted for institutional investors including Fidelity, Goldman Sachs and UBS. David is a frequent presenter at technical conferences where he strives to bridge the gap between technical solutions, application developers and the end customer experience.

Speaker



- **Debendra Das Sharma, Senior Fellow and Co-GM of Memory and I/O Intel**
 - Dr. Debendra Das Sharma is an Intel Senior Fellow and co-GM of Memory and I/O Technologies in the Data Platforms and Artificial Intelligence Group at Intel Corporation. He is a leading expert on I/O subsystem and interface architecture.
 - Das Sharma is a member of the Board of Directors for the PCI Special Interest Group (PCI-SIG) and a lead contributor to PCIe specifications since its inception. He is a co-inventor and founding member of the CXL consortium and co-leads the CXL Technical Task Force. He co-invented the chiplet interconnect standard UCIe and is the chair of the UCIe consortium.

Speaker



- **Arthur Sainio, Director Product Marketing, SMART Modular Technologies**
 - Arthur is Co-Chair of the SNIA Persistent Memory Special Interest Group, which accelerates the awareness and adoption of Persistent Memories for computing architectures. As a Director of Product Marketing at SMART, Arthur has been driving new product launch and business development activities at SMART since 1998. He has supported a wide variety of product, technology and interconnect areas including NVDIMMs, MRAM, Optane AIC, OpenCAPI™, CCIX™, Gen-Z™, and CXL™.

Speaker




- **Arvind Jagannath, Sr. Product Line Manager for vSphere, VMware**
 - Arvind Jagannath works in Product Management at VMware. With over 25 years of experience in the industry working on memory, networking, storage, embedded, and kernel development, he currently leads infrastructure and core platform enablement for vSphere, working across the VMware ecosystem of server/CPU, IO, and storage partners. Arvind most recently drove platform product management at Cohesity and NetApp. Arvind holds an MBA from the University of Chicago, Booth school of Business and a Bachelors in Computer Science and Engineering from India.

Moderator



- Tom Coughlin, SNIA Member/President, Coughlin Associates, Inc.
 - Tom Coughlin is a digital storage analyst and business/ technology consultant. He has over 40 years in the data storage industry with engineering and senior management positions. Coughlin Associates consults, publishes books and market and technology reports and puts on digital storage-oriented events. He is a regular contributor for forbes.com and M&E organization websites. He is an IEEE Fellow, 2023 IEEE President Elect and is also active with SNIA and SMPTE. For more information on Tom Coughlin go to <https://tomcoughlin.com>

New Report: Emerging Memories Enter the Next Phase

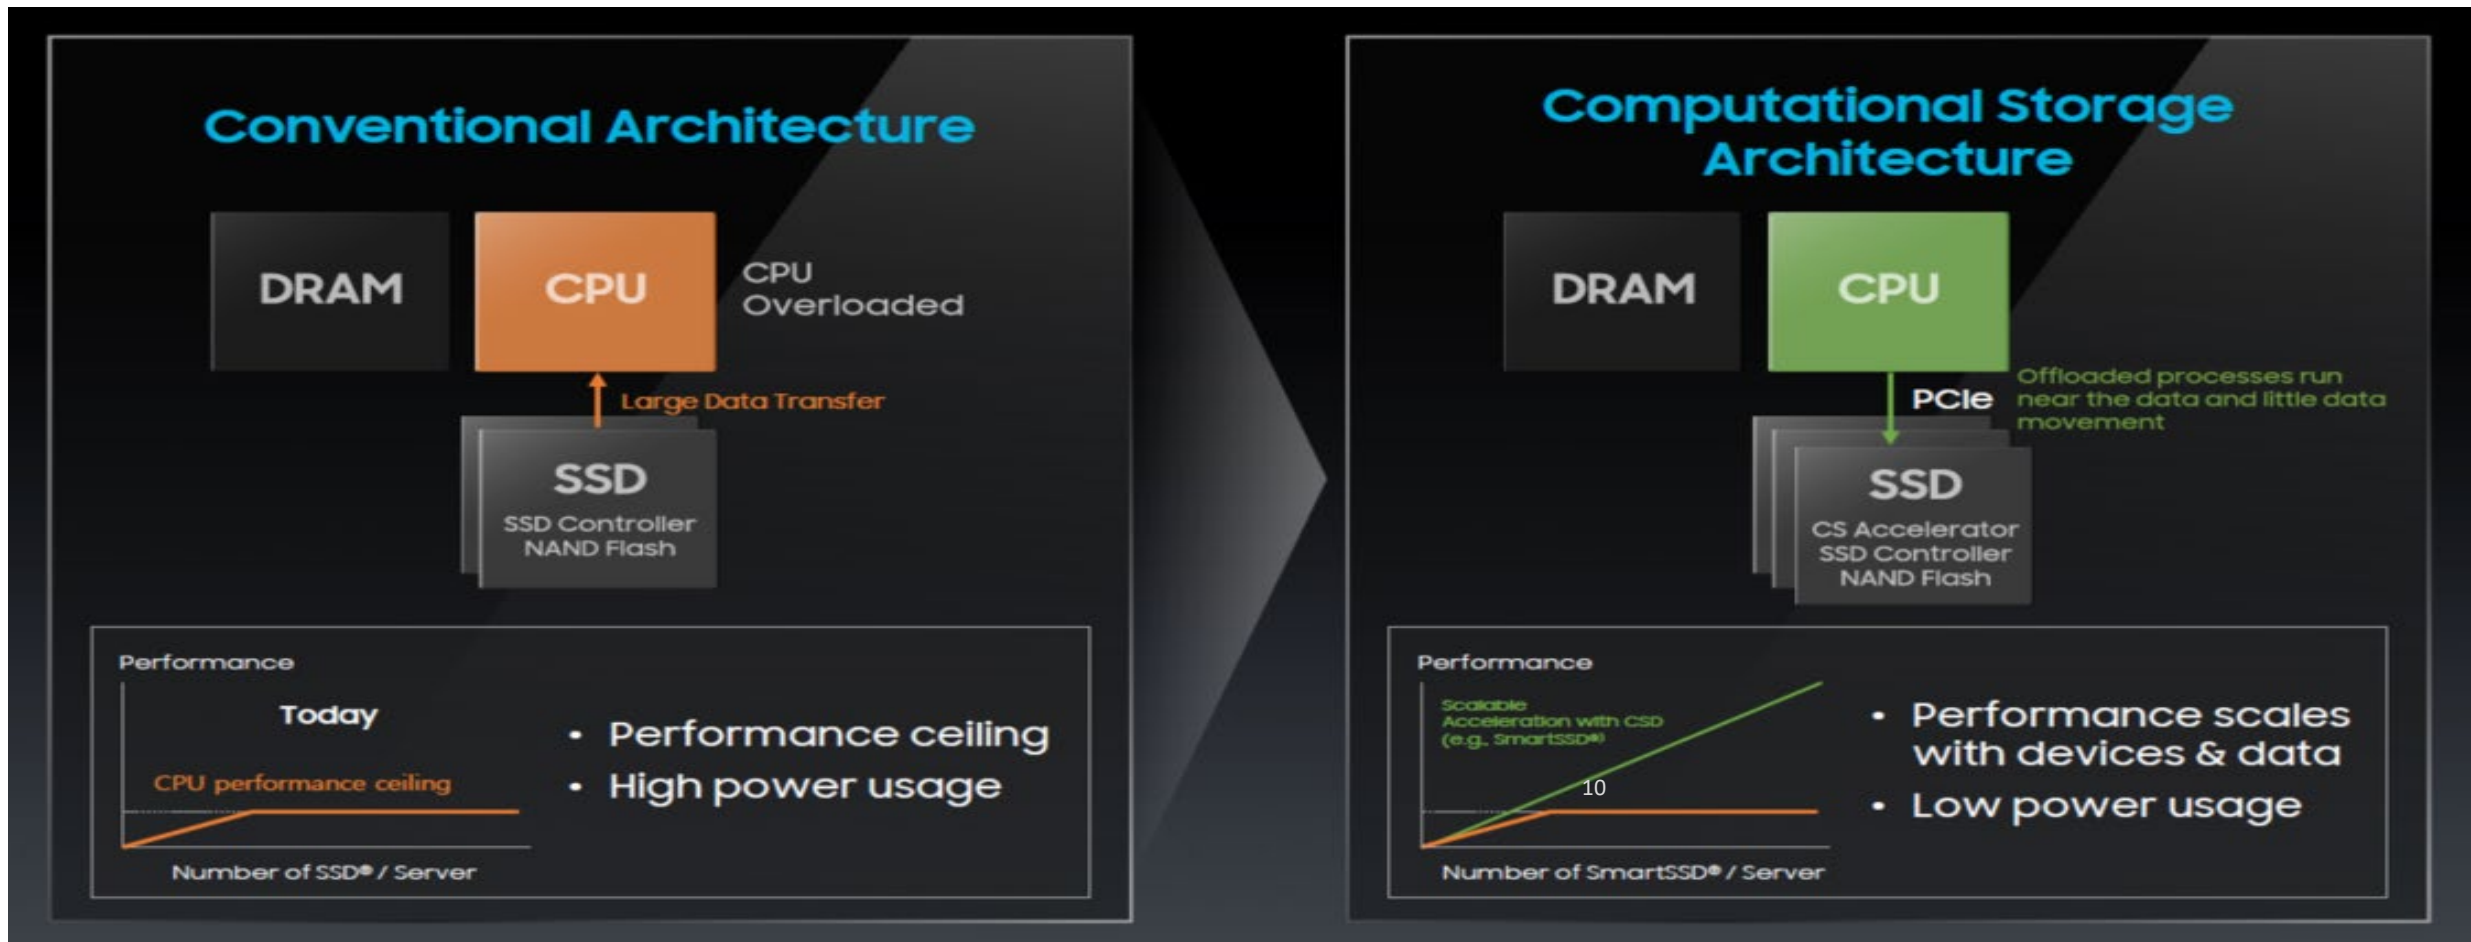


*Coughlin
Associates*

<http://www.tomcoughlin.com/techpapers.htm>

<https://Objective-Analysis.com/reports/#Emerging>

Computational Storage Explained



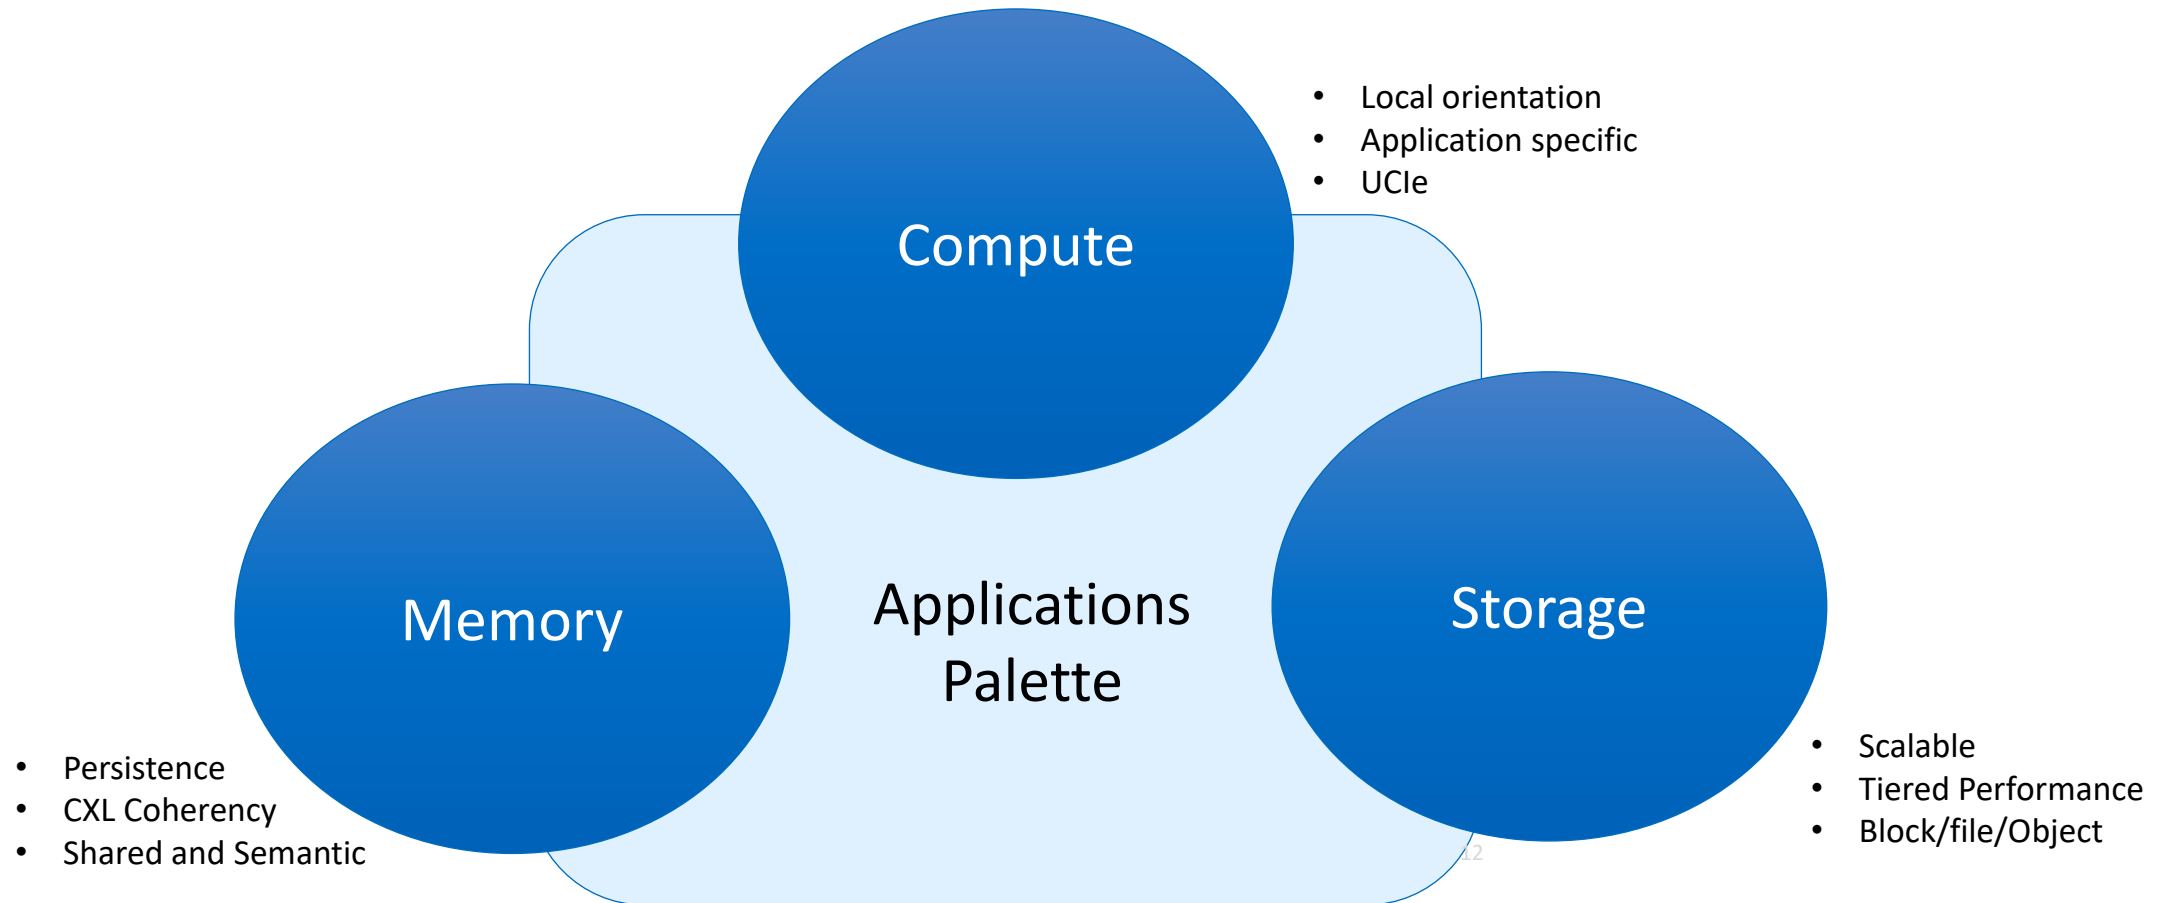
Computational Storage Summary

Benefits	Features	How	Technology Enhancements
TCO Advantage	Host Processor Offload	SmartSSD liberates host processor resources	Accelerator Portfolio- FPGA to SOC
Reduced Latency	Compute where the data resides	Localized accelerator in SSD	<ul style="list-style-type: none"> • CXL • FPGA and SOC/UCle
Scalability	One-to-Many SmartSSDs	Tuned Compute and Storage Resources	Memory: Persistence, Semantic, Pooling, Computational
Green Power Savings	Balanced Compute, Storage and Network resources	<ul style="list-style-type: none"> • Less Processors • Minimal Data Movement 	Energy Saving Form Factors

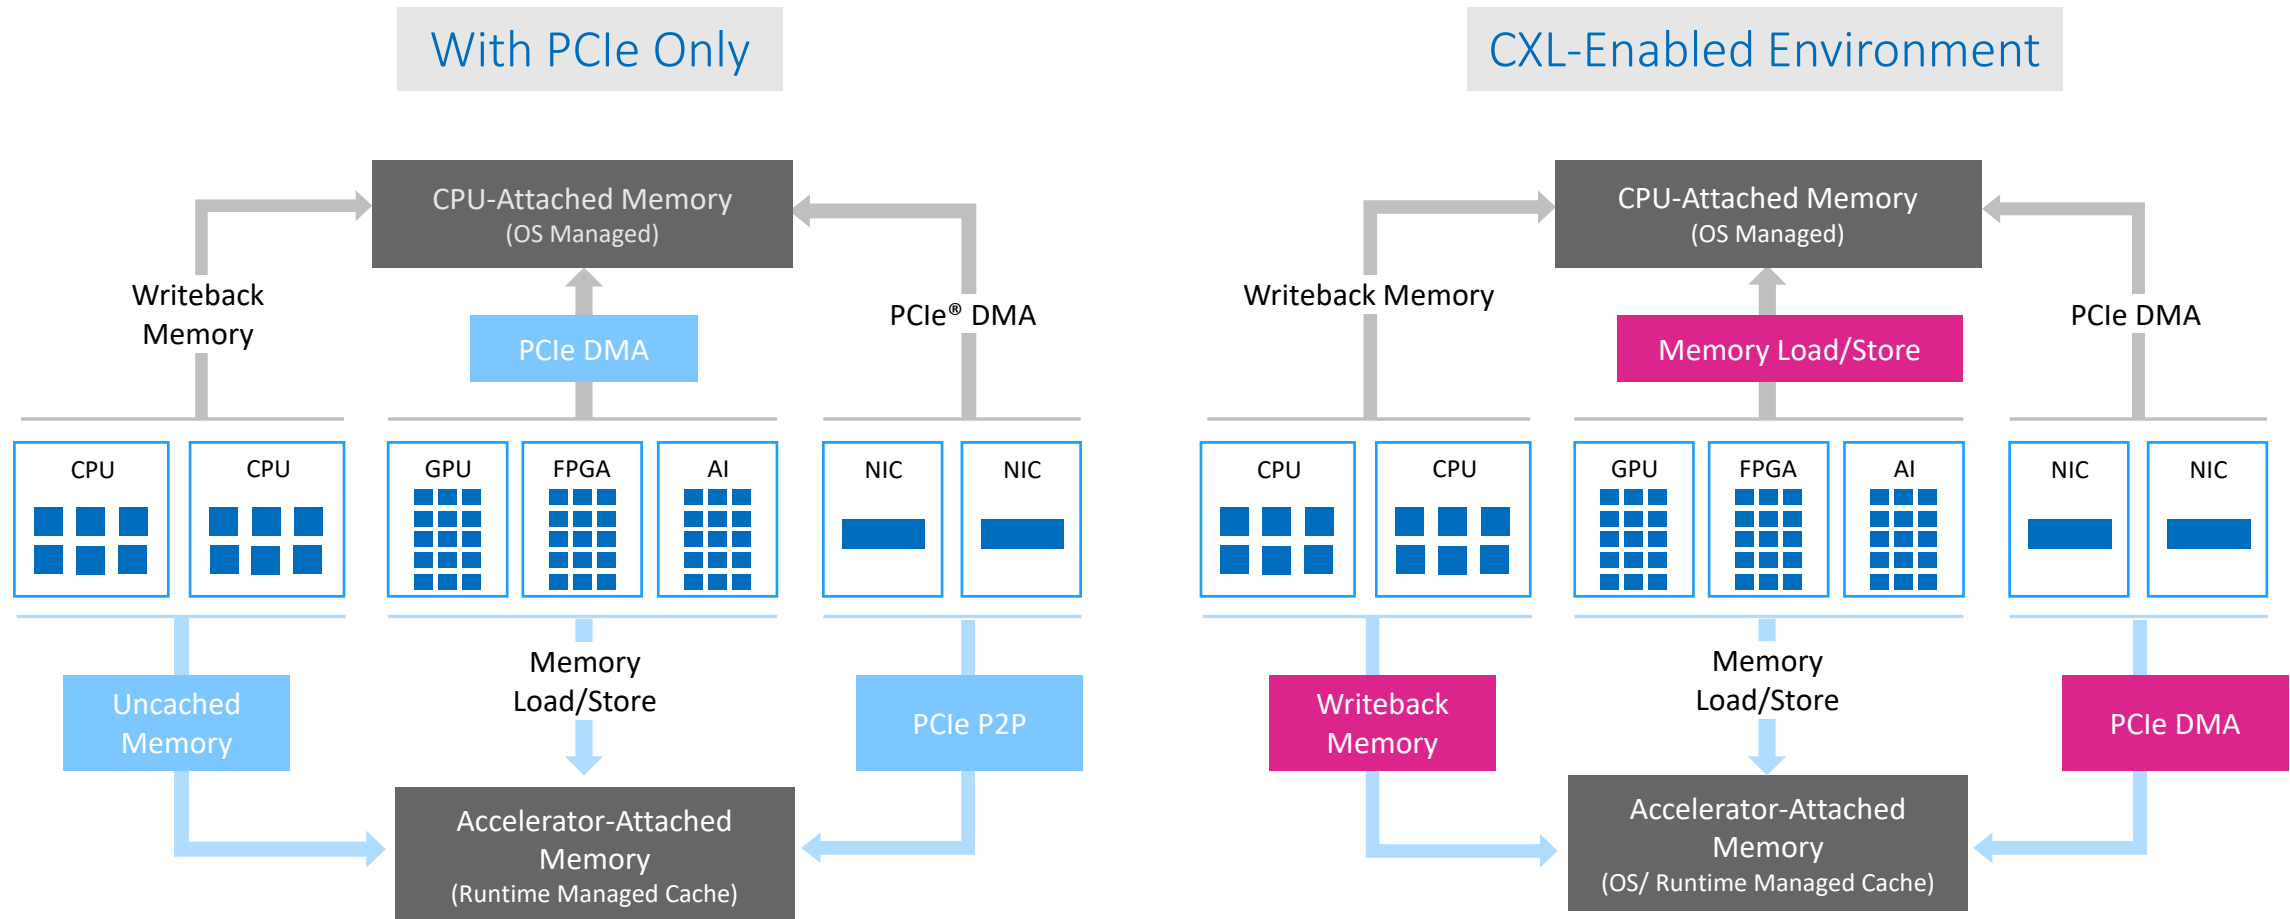
Number of SSD® / Server

Number of SmartSSD® / Server

Balanced Resources with Enhanced Technologies




CXL™: A New Class of Open-Standard Interconnects



The System Memory Challenge


Normalized Growth Rate



Ack: Chris Petersen, Meta

- Increasing core counts drives memory demand
- Increasing bandwidth and capacity
- Memory is not able to keep up → more DDR channels (cost, power and feasibility challenges)

Price per Gb (Log Scale)



- Memory is an increasing % of system power and cost
- Memory price (cost/bit) is flat due to scaling challenges
- Memory power scaling with speed

CXL™ Approach

Coherent Interface

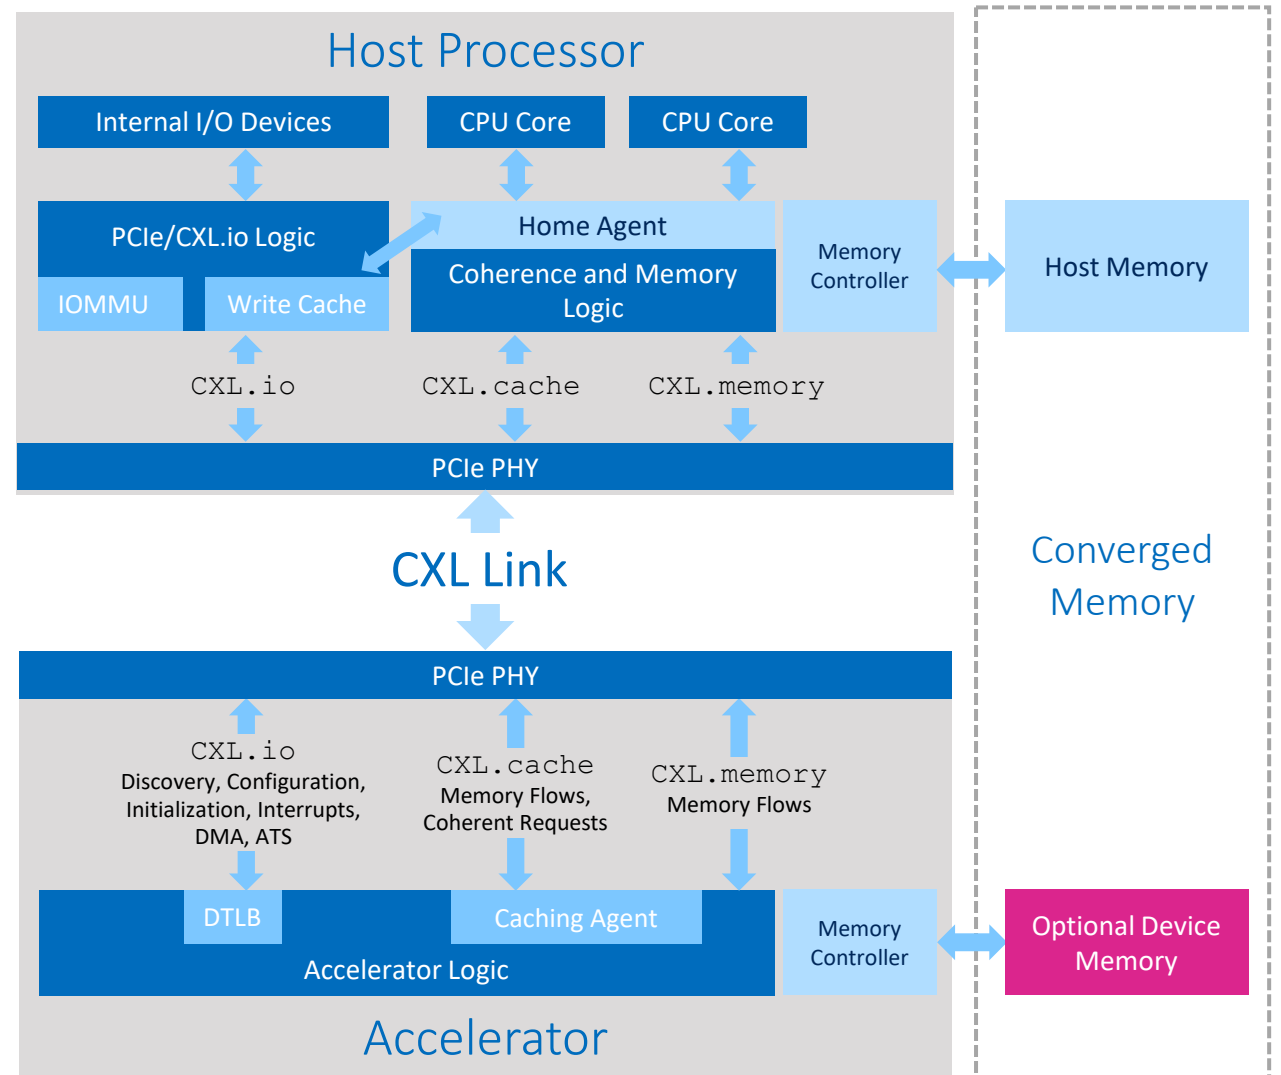
- Leverages PCIe with three multiplexed protocols
- Built on top of **PCIe® infrastructure**

Low Latency

- CXL.Cache/CXL.Memory targets near CPU cache coherent latency (<200ns load to use)

Asymmetric Complexity

- Eases burdens of cache coherence interface designs for devices



CXL™ 1.0/CXL 1.1 Usage Models

Type 1 Device


Caching Devices/Accelerators

Usages:

- PGAS NIC
- NIC atomics

Protocols:

- CXL.io
- CXL.cache



Type 2 Device

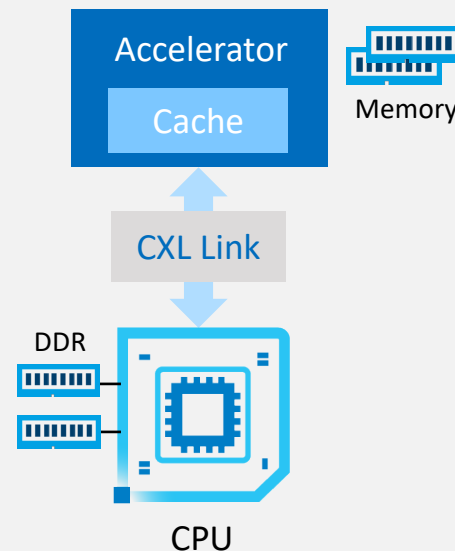
Accelerators with Memory

Usages:

- GPU
- FPGA
- Dense
- Computation

Protocols:

- CXL.io
- CXL.cache
- CXL.memory



Type 3 Device


Memory Buffers

Usages:

- Memory BW expansion
- Memory capacity expansion
- 2LM

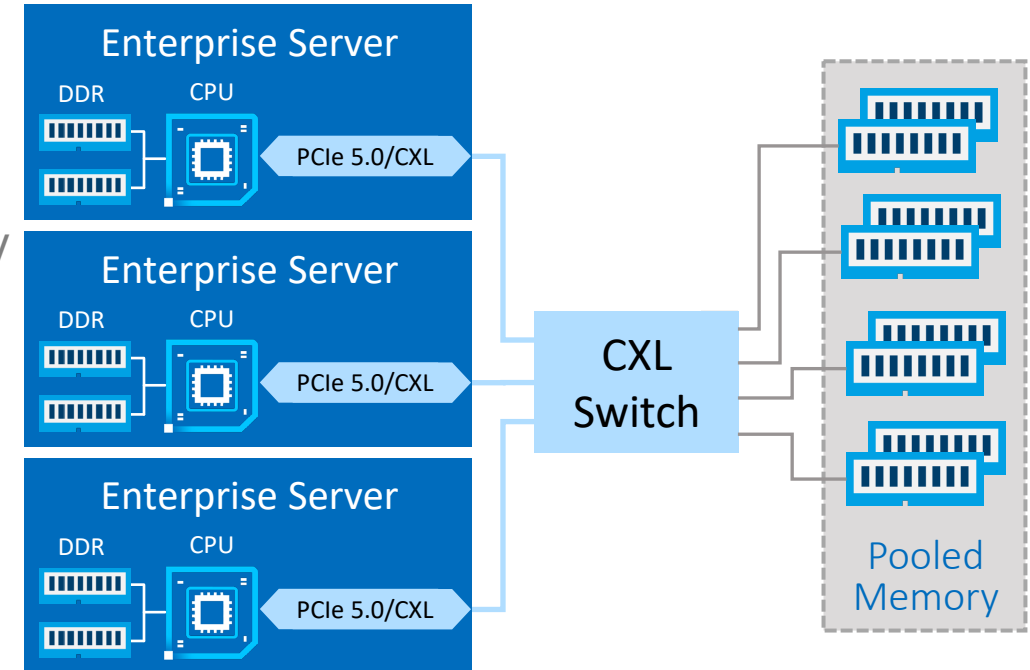
Protocols:

- CXL.io
- CXL.mem



CXL™ 2.0: Resource Pooling at Rack Level, Persistent Memory

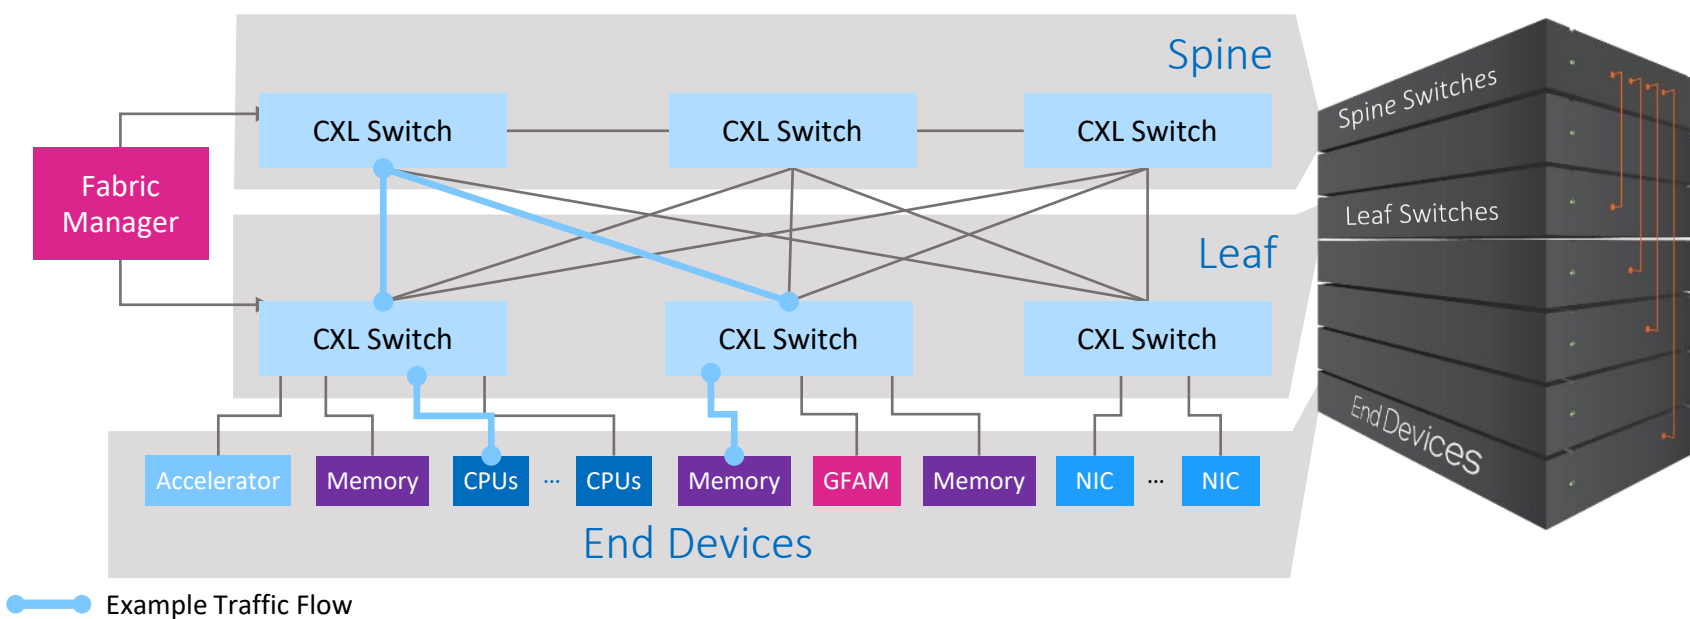
- Resource pooling/disaggregation
 - Managed hot-plug flows to move resources
 - Type-1/Type-2 device assigned to one host
 - Type-3 device (memory) pooling at rack level
 - Direct load-store, low-latency access – similar to memory attached in a neighboring CPU socket (vs. RDMA over network)
- Persistence flows for persistent memory
- Fabric Manager/API for managing resources
- Security: authentication, encryption
- Beyond node to rack-level connectivity!



Disaggregated system with CXL optimizes resource utilization
delivering lower TCO and power efficiency

CXL™ 3.0 Enhancements

- Bandwidth doubling with 64 GT/s at 0-latency add
- Protocol enhancements with direct peer-to-peer to HDM memory
- Composable systems with spine/leaf architecture at rack/pod

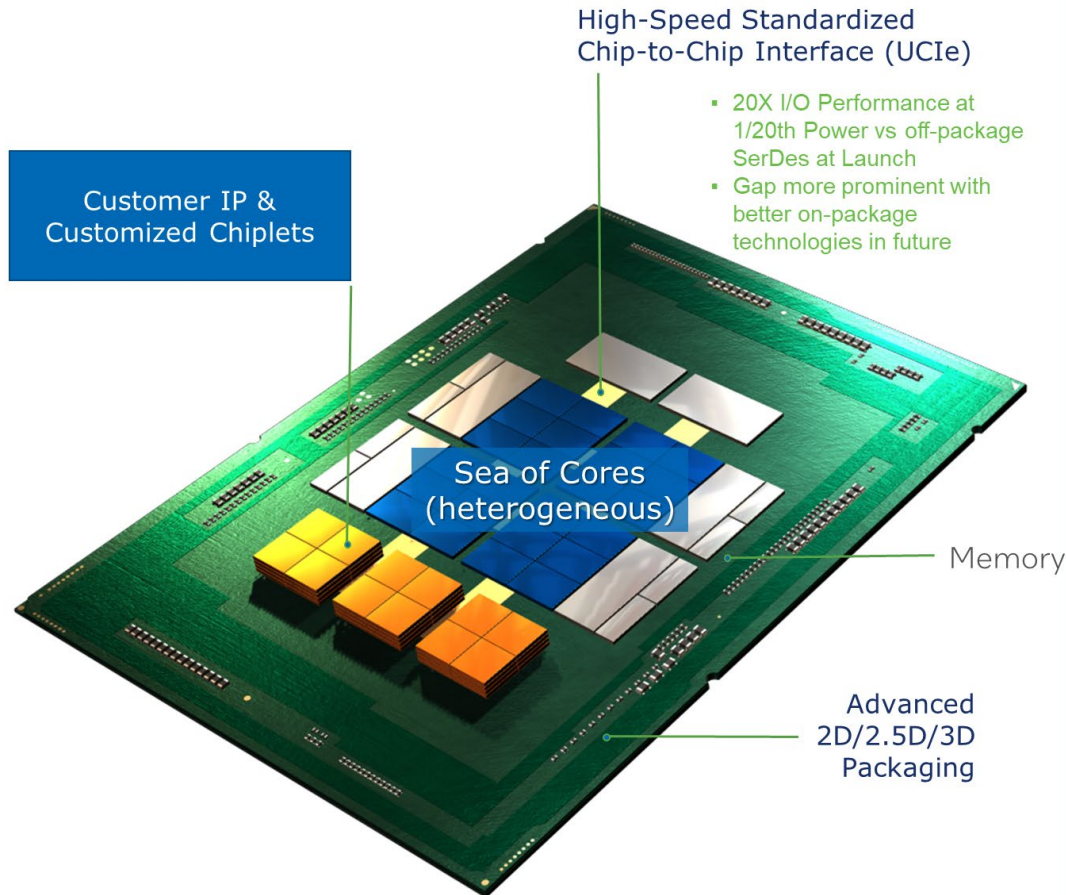


CXL 3.0 Fabric Architecture

- Interconnected spine switch system
- Leaf switch NIC enclosure
- Leaf switch CPU enclosure
- Leaf switch accelerator enclosure
- Leaf switch memory enclosure

Motivation for UCle

OPEN CHIPLET: PLATFORM ON A PACKAGE



Heterogeneous Integration Fueled by an Open Chiplet Ecosystem
(Mix-and-match chiplets from different process nodes / fabs / companies / assembly)

Align Industry around an open platform to enable chiplet based solutions

- Enables SoC construction that exceeds maximum reticle size
 - Package becomes new System-on-a-Chip (SoC) with same dies (Scale Up)
- Reduces time-to-solution (e.g., enables die reuse)
- Lowers portfolio cost (product & project)
 - Enables optimal process technologies
 - Smaller (better yield)
 - Reduces IP porting costs
 - Lowers product SKU cost
- Enables a customizable, standard-based product for specific use cases (bespoke solutions)
- Scales innovation (manufacturing/ process locked IPs)

UCle 1.0: Characteristics and Key Metrics

UCle 1.0 delivers the best KPIs while meeting the projected needs for the next 5-6 years.


Wide industry leader adoption spanning semiconductor, manufacturing, assembly, & cloud segments.

CHARACTERISTICS	STANDARD PACKAGE	ADVANCED PACKAGE	COMMENTS
Data Rate (GT/s)	4, 8, 12, 16, 24, 32		Lower speeds must be supported -interop (e.g., 4, 8, 12 for 12G device)
Width (each cluster)	16	64	Width degradation in Standard, spare lanes in Advanced
Bump Pitch (um)	100 – 130	25 - 55	Interoperate across bump pitches in each package type across nodes
Channel Reach (mm)	<= 25	<= 2	
KPIs / TARGET FOR KEY METRICS	STANDARD PACKAGE	ADVANCED PACKAGE	COMMENTS
B/W Shoreline (GB/s/mm)	28 – 224	165 – 1317	Conservatively estimated: AP: 45u; Standard: 110u; Proportionate to data rate (4G – 32G)
B/W Density (GB/s/mm ²)	22-125	188-1350	
Power Efficiency target (pJ/b)	0.5	0.25	
Low-power entry/exit latency	0.5ns <= 16G, 0.5-1ns >= 24G		Power savings estimated at >= 85%
Latency (Tx + Rx)	< 2ns		Includes D2D Adapter and PHY (FDI to bump and back)
Reliability (FIT)	0 < FIT (Failure In Time) << 1		FIT: # failures in a billion hours (expecting ~1E-10) w/ UCle Flit Mode

Why Persistent Memory Is Needed

- For faster storage arrays

- Metadata and checkpointing data acceleration (or checkpointing elimination)
- Fast write caching improves latency and consistent application response time
- Reducing the time transient data must be replicated to the peer



- For faster AI applications

- Machine learning algorithms stored in PM



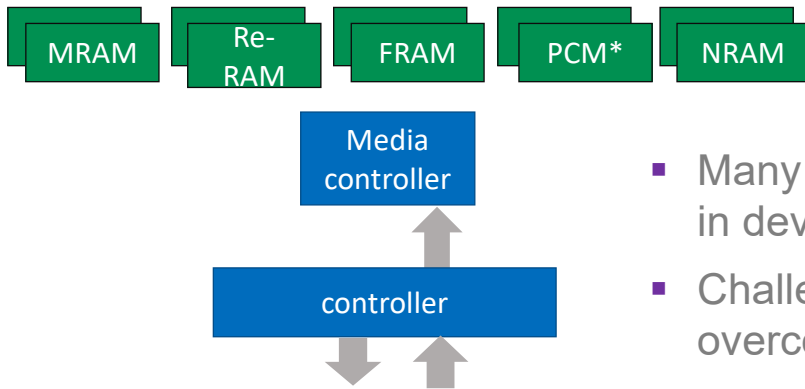
- For faster system recovery from downtime

- Preserve error logs and security events in real time
- Preserve transactions in flight
- Reduce time to retrieve backed up data



Persistent Memory Architecture Trends


Persistent Memory with New or Emerging Non-Volatile Media



- Many promising PM technologies in development
- Challenges still need to be overcome

- CXL™ PM adds tier of memory in-between DRAM and NAND
- CXL allows byte-addressable and cache-coherent memory that is agnostic to the architecture
- CXL-attached memory can be persistent like a traditional NVDIMM, but can also involve micro-tiering technology to take full advantage of onboard DRAM and NAND.

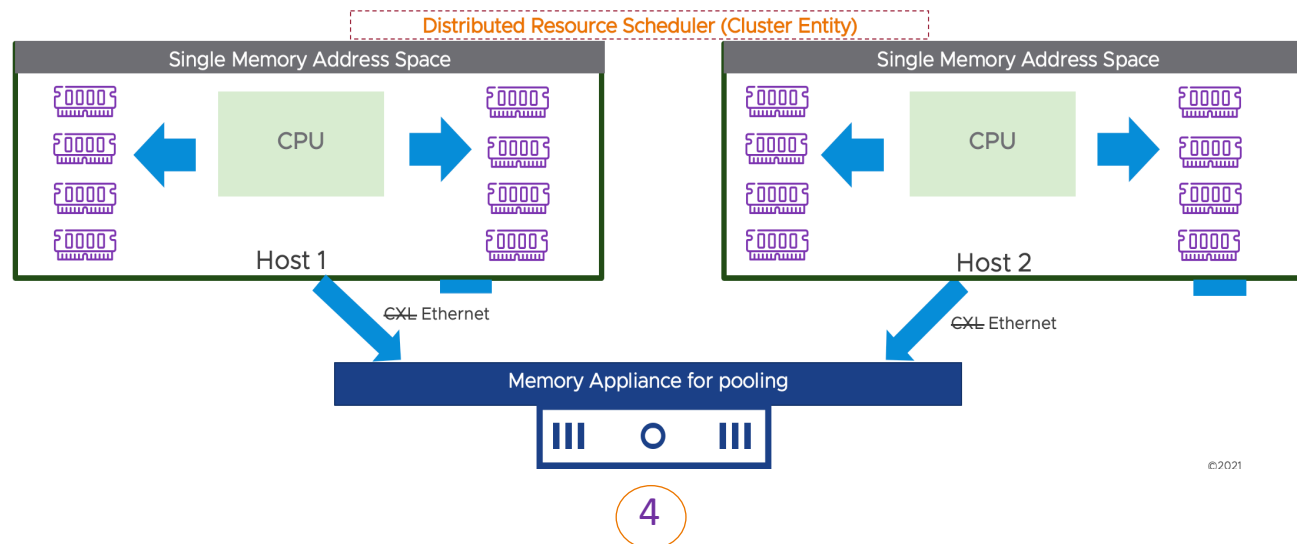
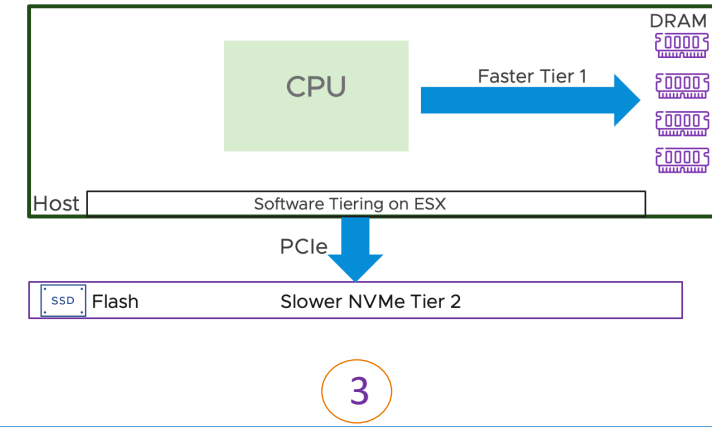
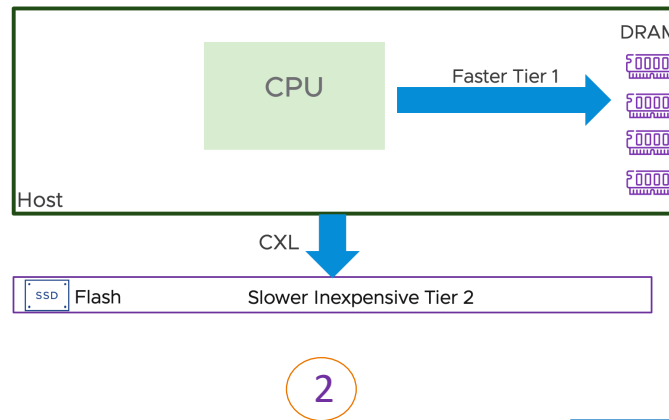
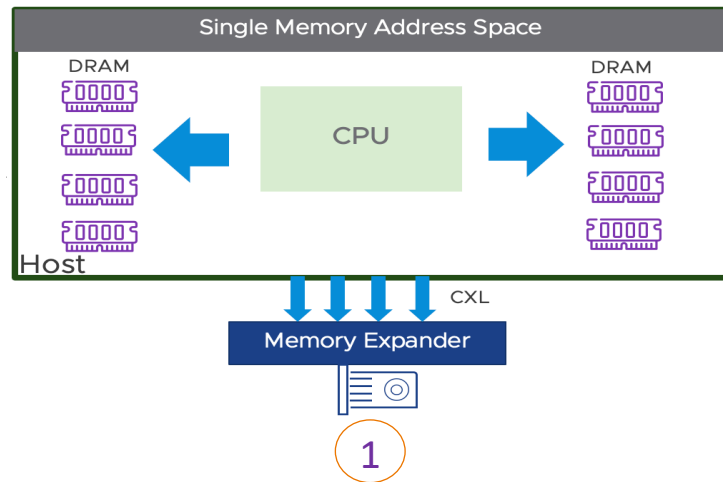
Persistent Memory with DRAM, Flash and Energy Source



- DRAM speed + NAND persistence
- New CXL™ PM solutions
- Leveraging NVDIMM for fast deployment
- Provides a valuable middle ground architecture



CXL/Tiering Deployment Options



Use-cases emerging

- Memory Expansion – capacity, bandwidth
- Core Utilization/Server consolidation
- Lower TCO
- Power/Space/Cooling
- Hardware Offloads
- Consolidate cluster memory usage
- Reduce cluster data movement
- Disaggregation & Composability

CXL (PCIe on steroids!!)

- **New emerging standard**
 - Uses PCIe physical interface
- **Provides memory semantics, more efficient than PCIe**
 - Cache coherency
 - Cache-line granularity as opposed to page-level granularity
 - Coherent Host to Device Or Device to Host memory access
 - Avoids page faults
- **Makes memory access more reliable, cheaper**
- **Allows for composability/disaggregation**

- **Software/Virtualization**
 - Memory tracking
 - Monitoring for latency & bandwidth
 - Min. config changes & min perf.
Degradation
 - Ensuring fairness
 - Mitigating risks through transparent migration
- **Page classification and Tiering**
 - Software or Software-Hardware Co-design
- **Industry CXL ecosystem**
 - Evolving ecosystem activity
 - CPU & memory device vendors, OEMs
- **Number of newer use-cases**
 - Increasing value-prop with Accelerators
 - Device to Host accesses
- **Newer fabric management capabilities**



Questions?

Thank you for attending MemCon!

Visit <https://www.snia.org/memcon23> for more information