SAMSUNG

STORAGE DEVELOPER CONFERENCE SDC 23

BY Developers FOR Developers

Is SSD with CXL Interfaces Brilliantly Stupid or Stupidly Brilliant?

YANG SEOK KI, Ph.D.

Vice President, Memory Solutions Lab., Samsung Electronics

Stupid or Brilliant?

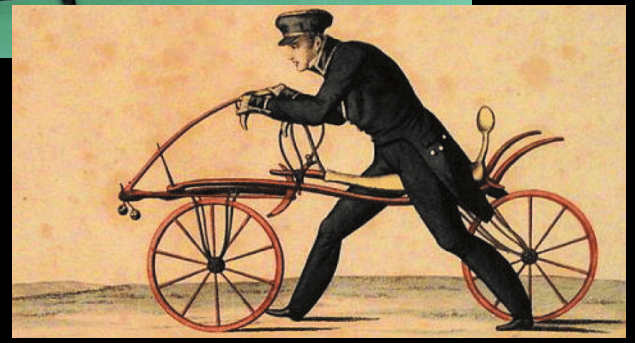

Draisine (Running machine), 1817:
the 1st bicycle in record

The Jazz Singer, 1927: the 1st movie with an audio track

SDC 23

SSD with CXL Interfaces

- Storage with memory and/or storage interfaces

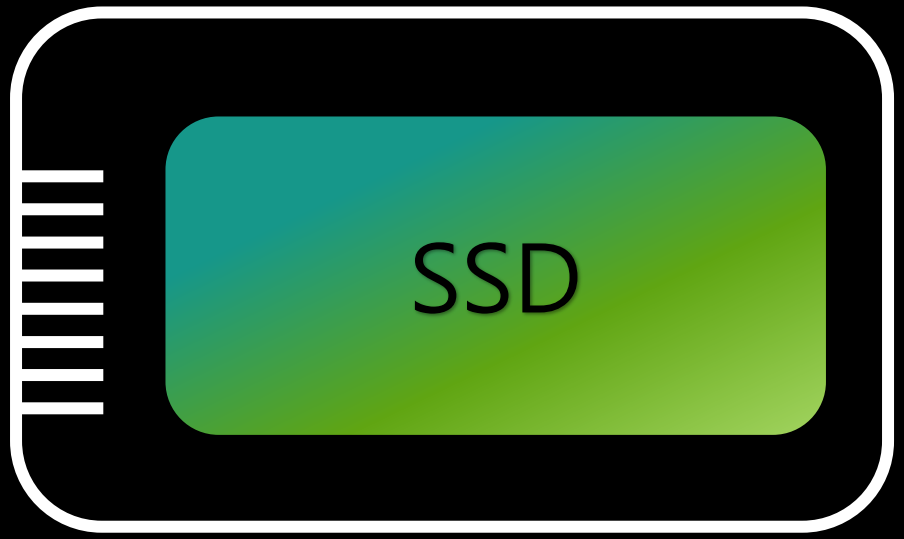

CXL Compute Express Link™

Technical Needs

Memory Hierarchy

- Keep hot data close to CPU using data locality

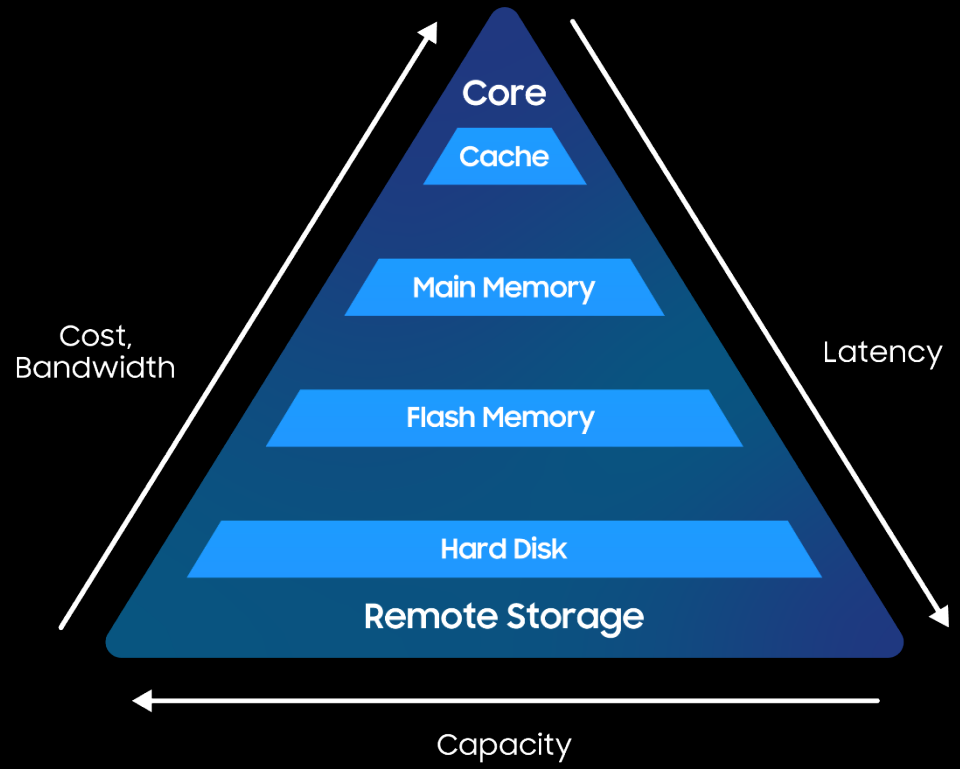

Memory Hierarchy

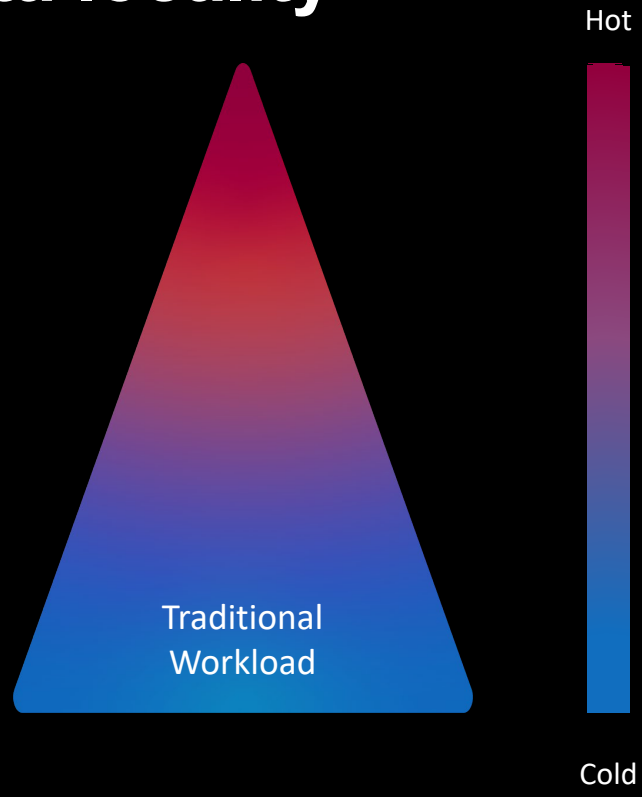

Traditional Workload

SDC 23

Needs (1): Persistent Memory

- Discontinuation of the leading technology

Database Server

Storage Server

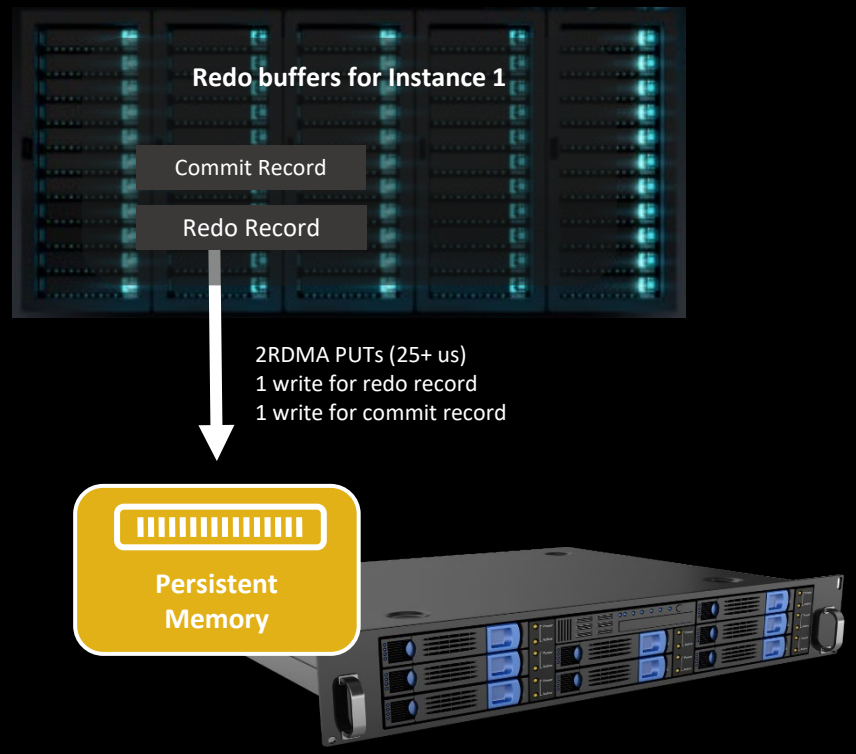

Database: Oracle Exadata Redo Log

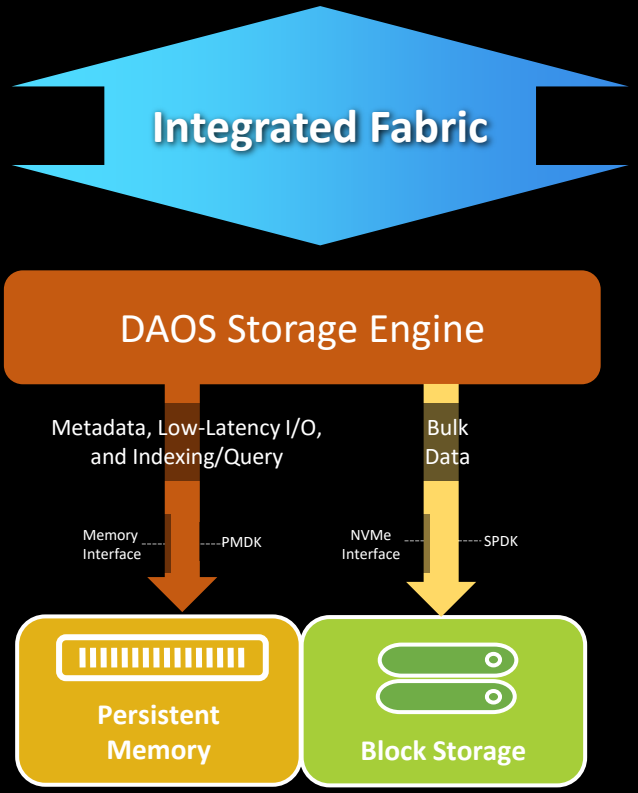

Storage: DAOS
(Distributed Application Object Storage)

Needs (2): Secondary Memory

- **High overhead of virtual memory implementation**

<OS-level>
Swap for memory extension on disk

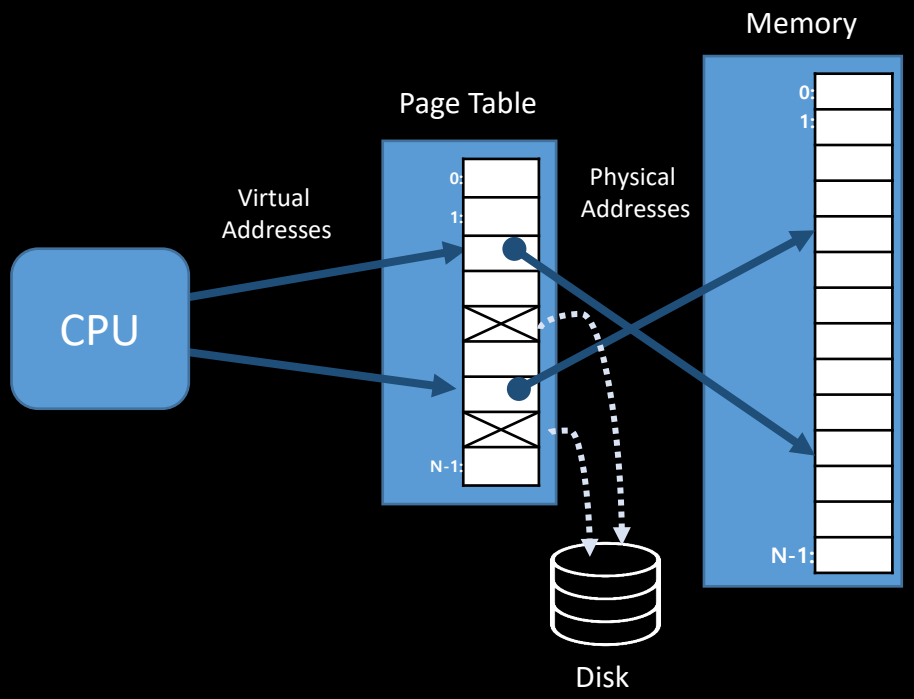

<User-level>
Redis Auto Tiering for memory extension on SSD

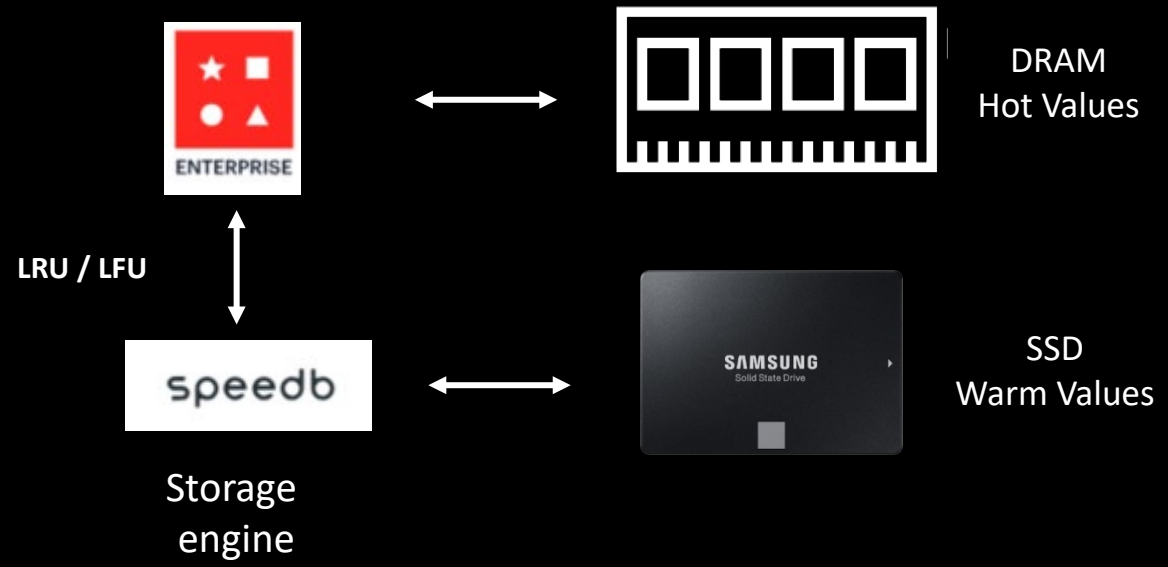

Needs (3): Fast Small IO

- High overhead of IOs smaller than 4KB

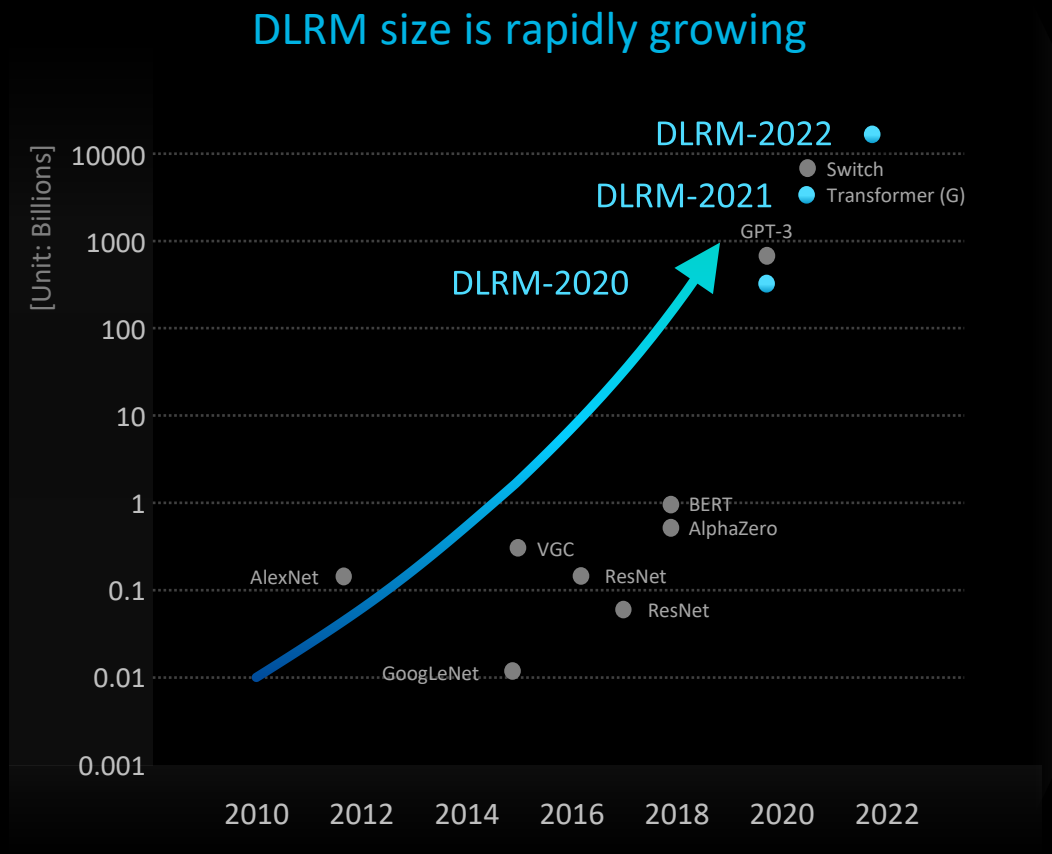

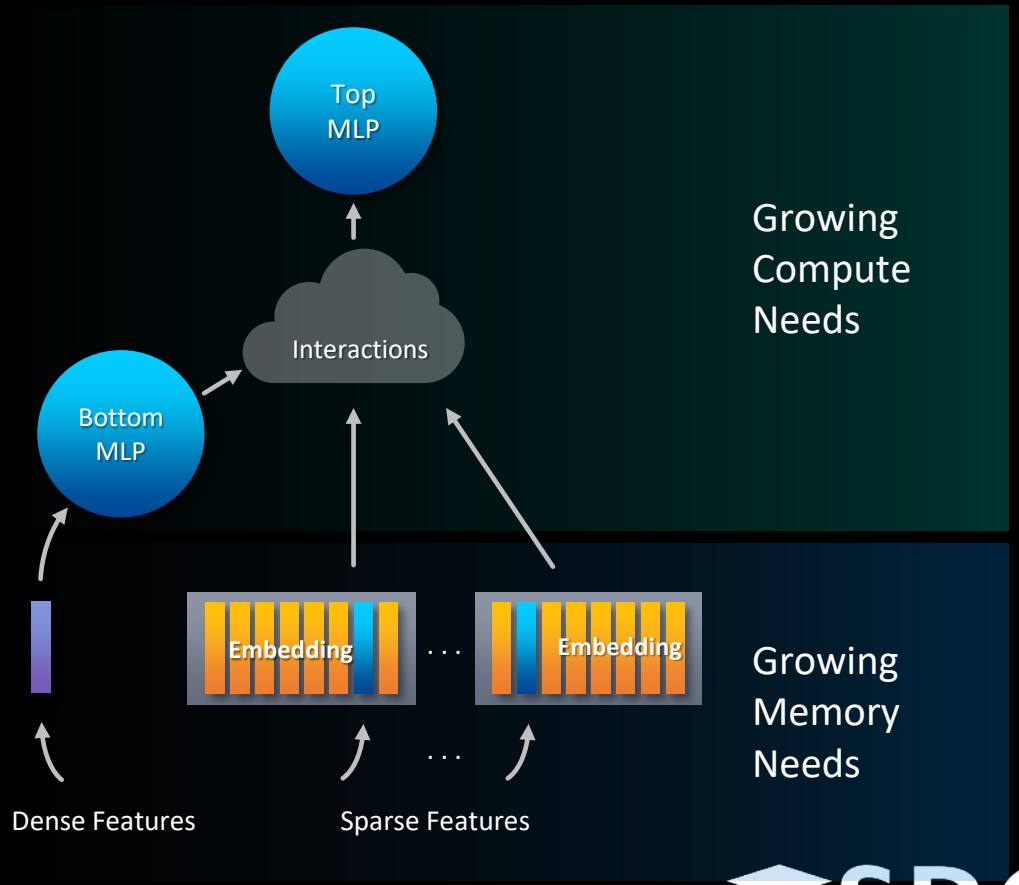

Source: Meta

SDC 23

CXL-based SSD

CXL-based SSD

- A Hybrid device of DRAM and NAND with CXL interfaces

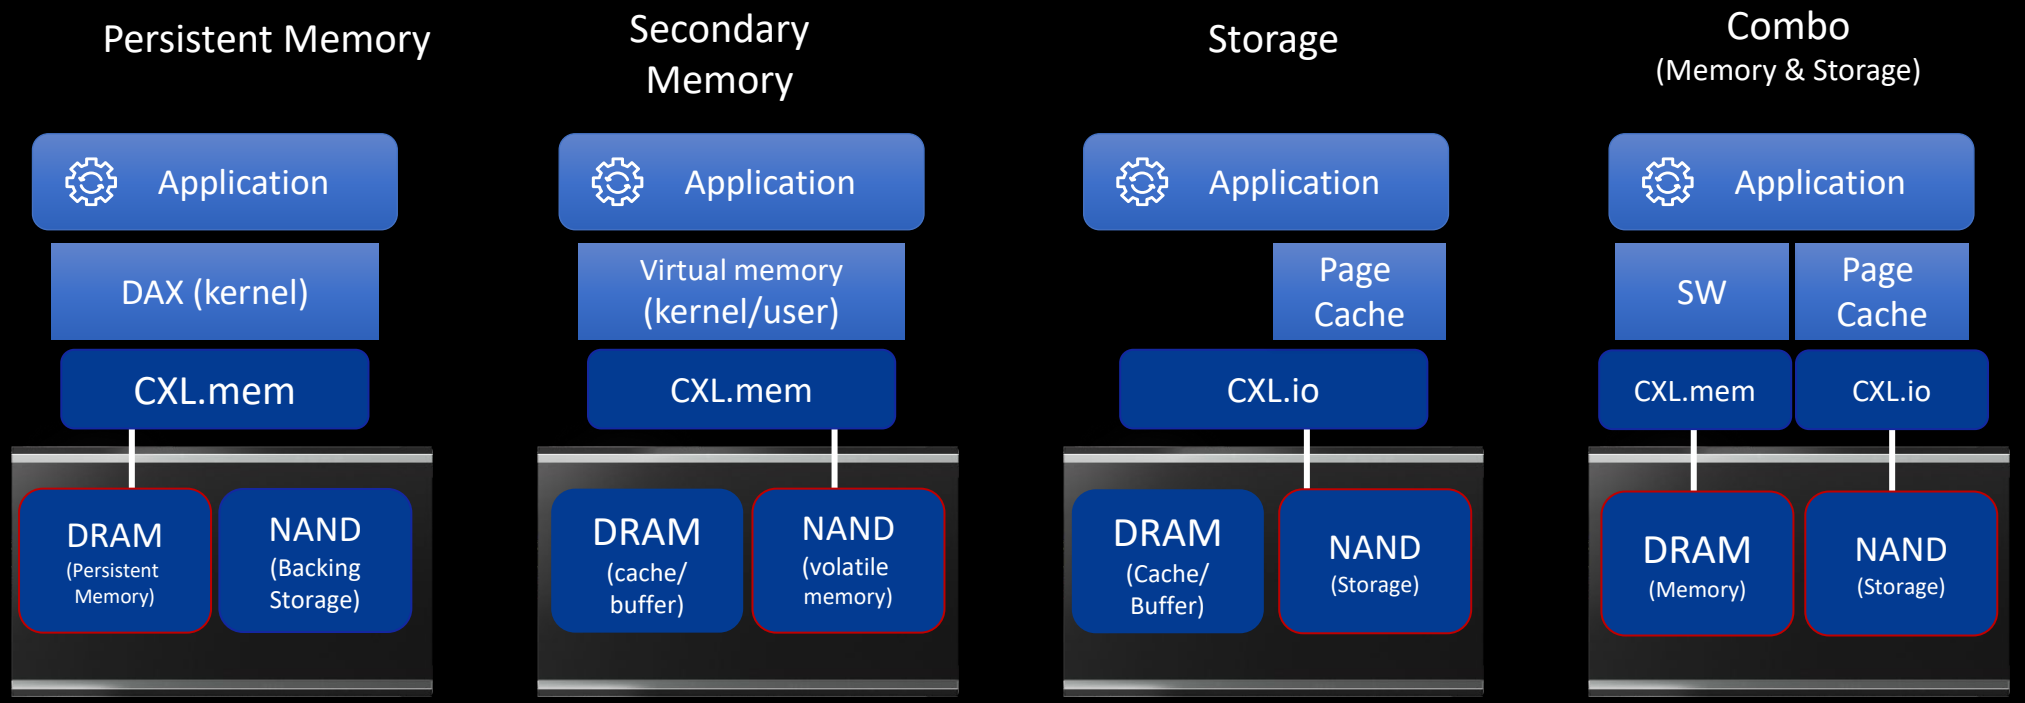

CXL (Compute Express Link)

- Asynchronous blocking memory interface with optional coherency

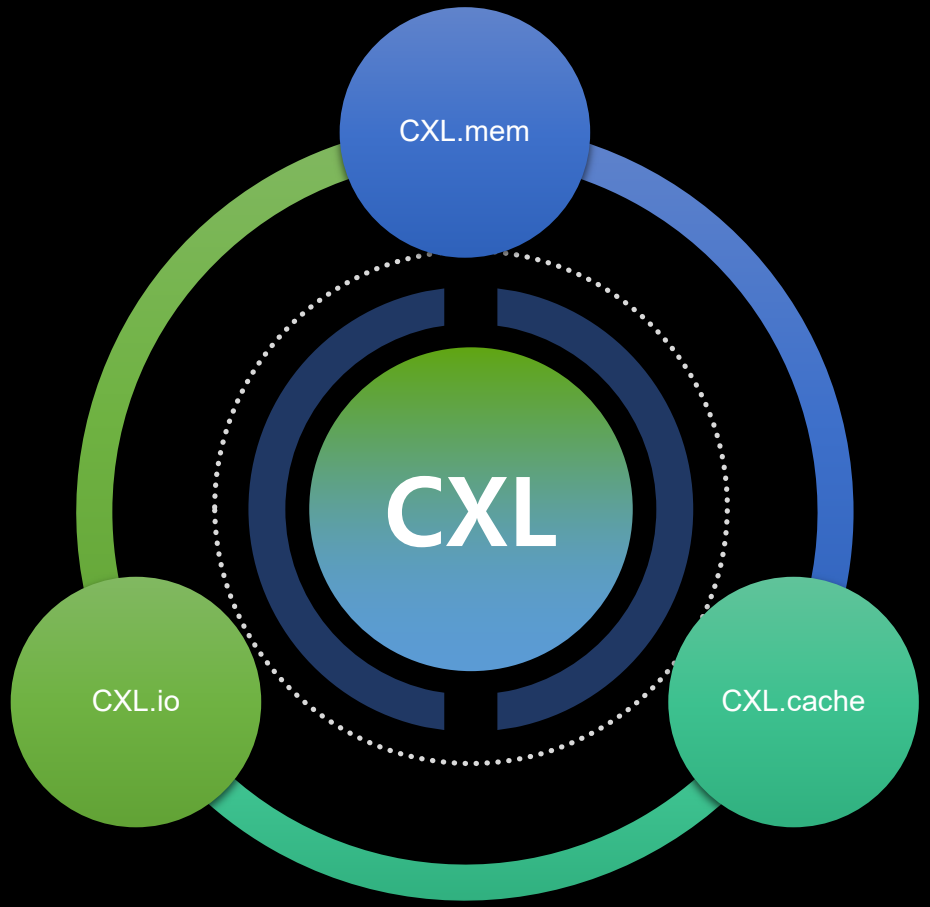

CXL Device Types

- Device types based on protocols, not functions

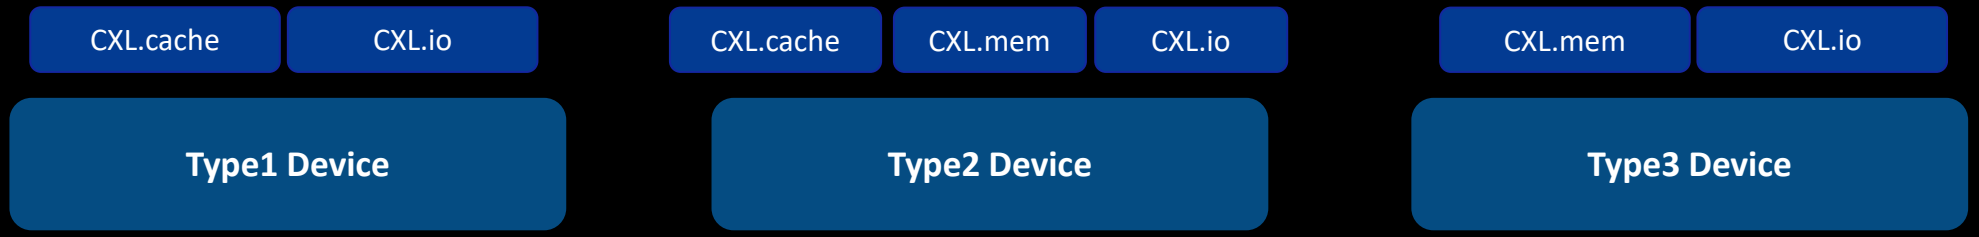

CXL-based SSD as Persistent Memory

- Type-3 device similar to NVDIMM

Dump Size	Dump Time(s)	Dump Energy
16GB	**3.2**	58J
32GB	6.4	115J
64GB	12.8	229J
128GB	25.6	**457J**
256GB	51.2	913J
512GB	102.4	1,825J

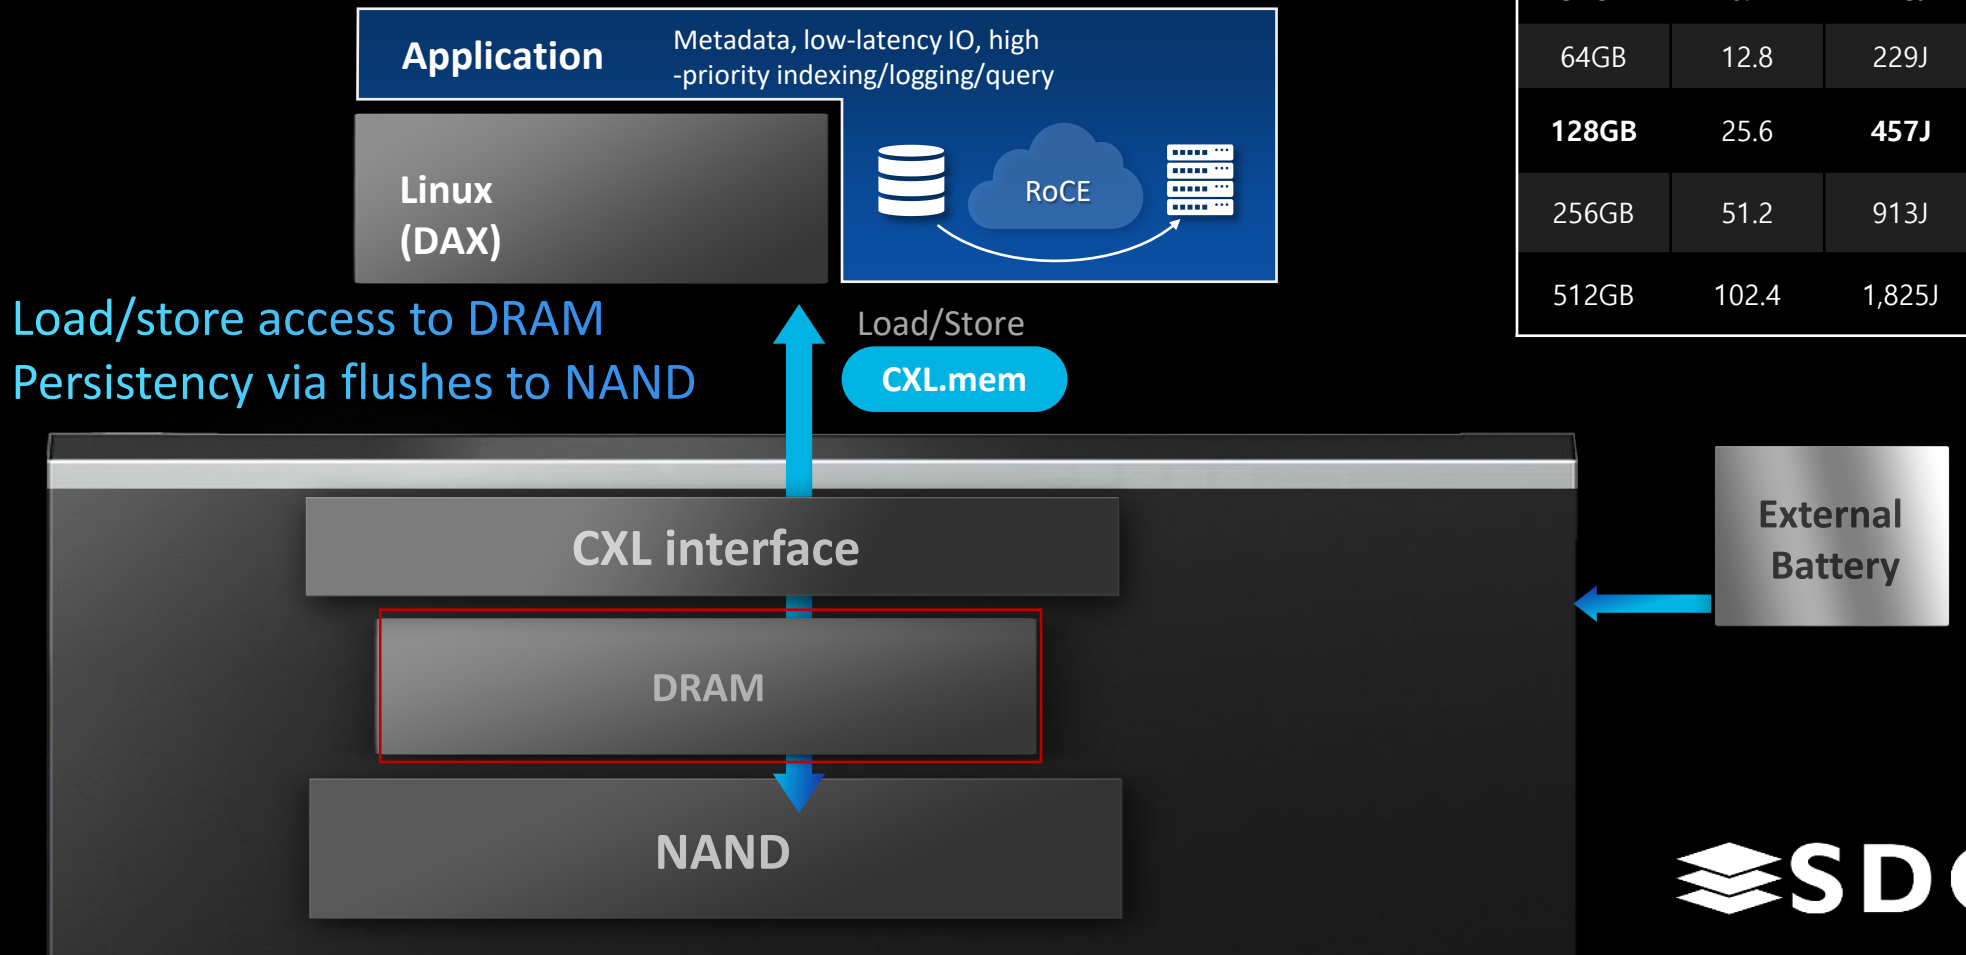

Power Failure Protection

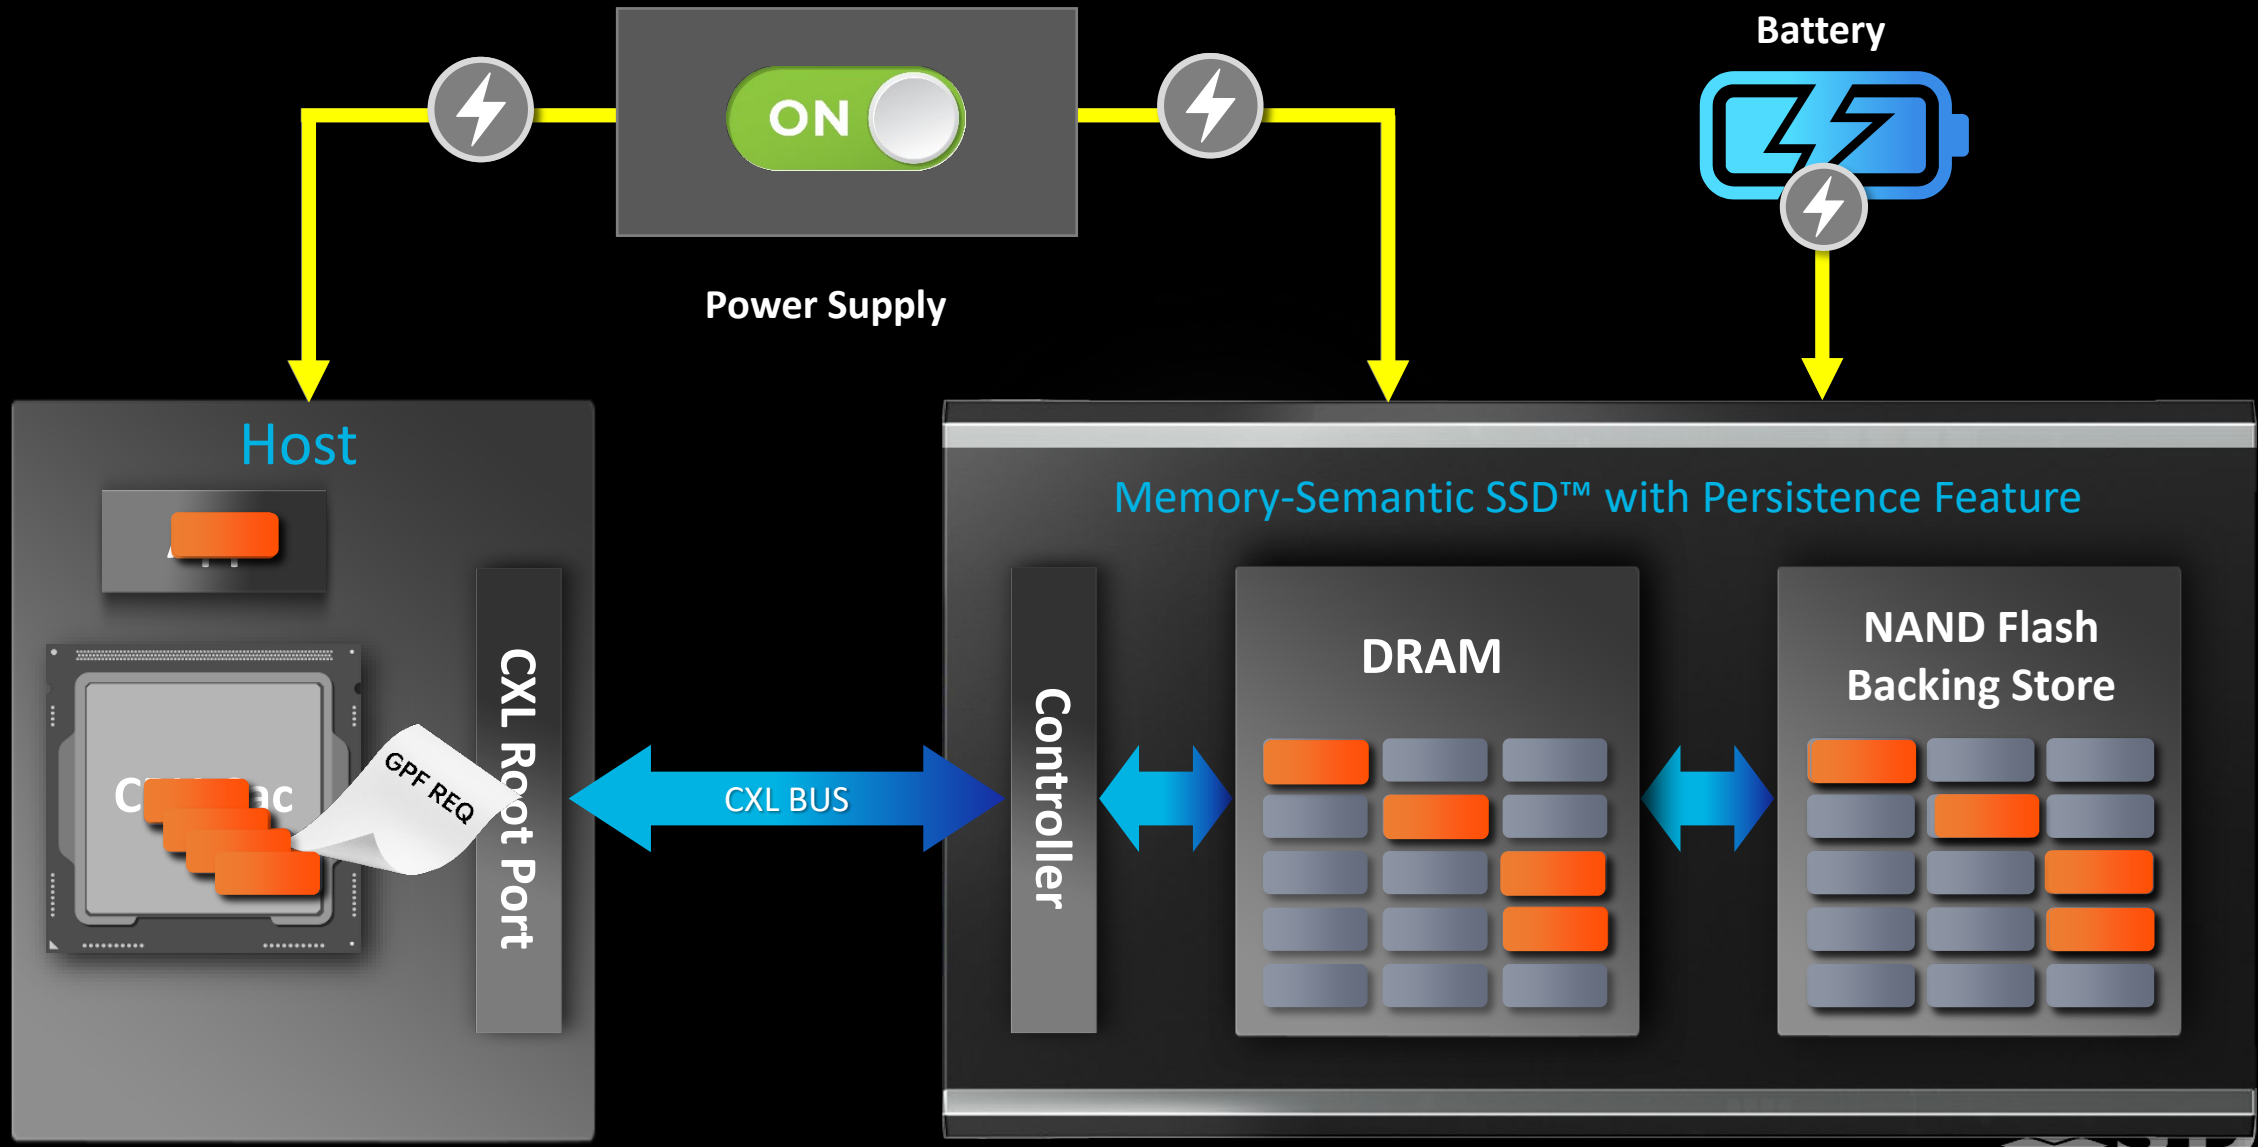

Persistent Memory Performance

Key Features & Benefits

- Battery-backed DRAM with speed comparable to DDR5
- Persistence achieved with data dumps to NAND flash
- Supports flush-on-fail with CXL 2.0 GPF feature

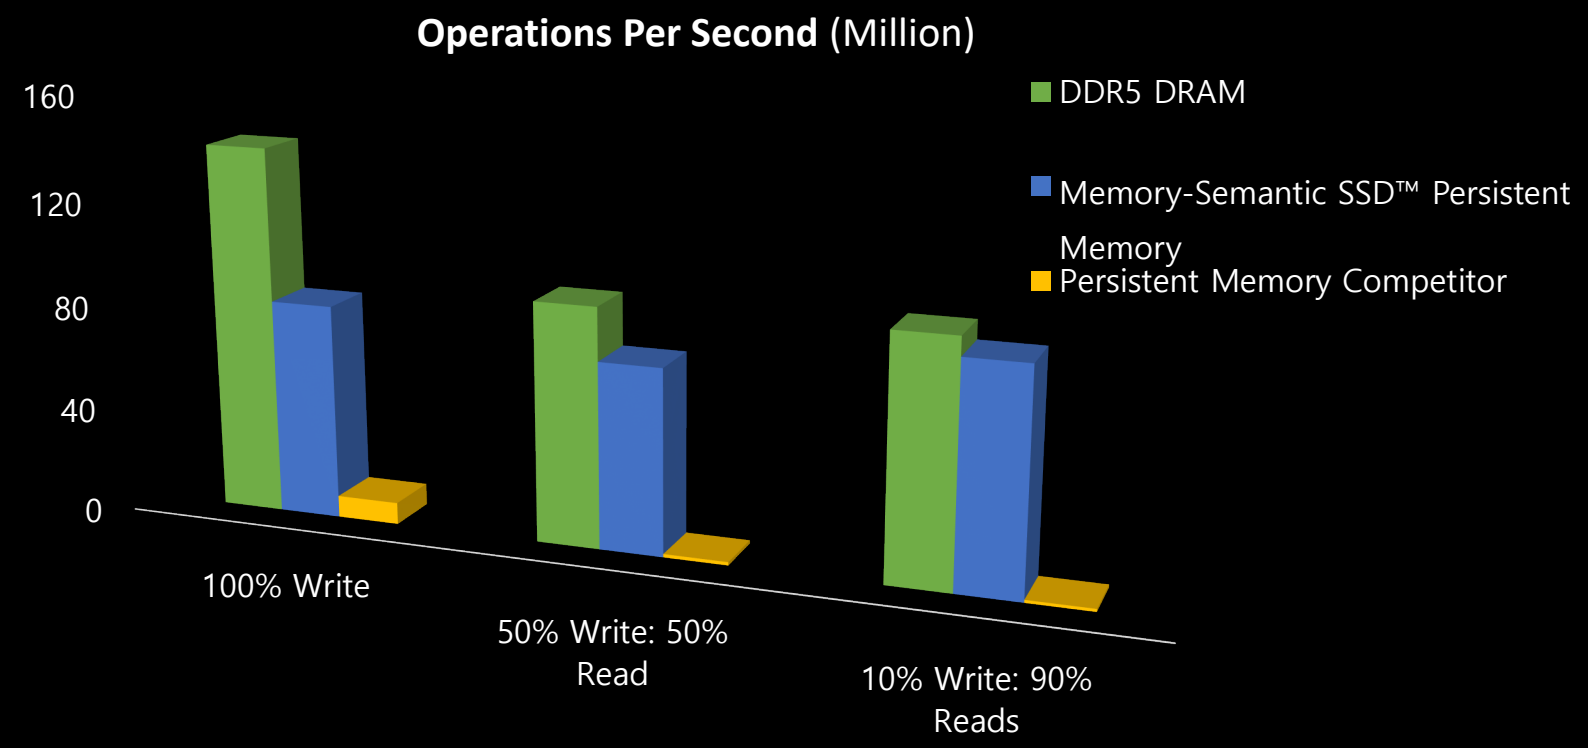

Memory-Semantic SSD™ Persistency Demo

CXL-based SSD as Secondary Memory

CXL-based SSD with built-in DRAM

Built-in DRAM

- Processing AI and ML applications, usually need relatively small-sized data chunks
- Applications can write data to the DRAM cache at DRAM speed

CXL Technology

- Low latency enabled by CXL.memory protocol

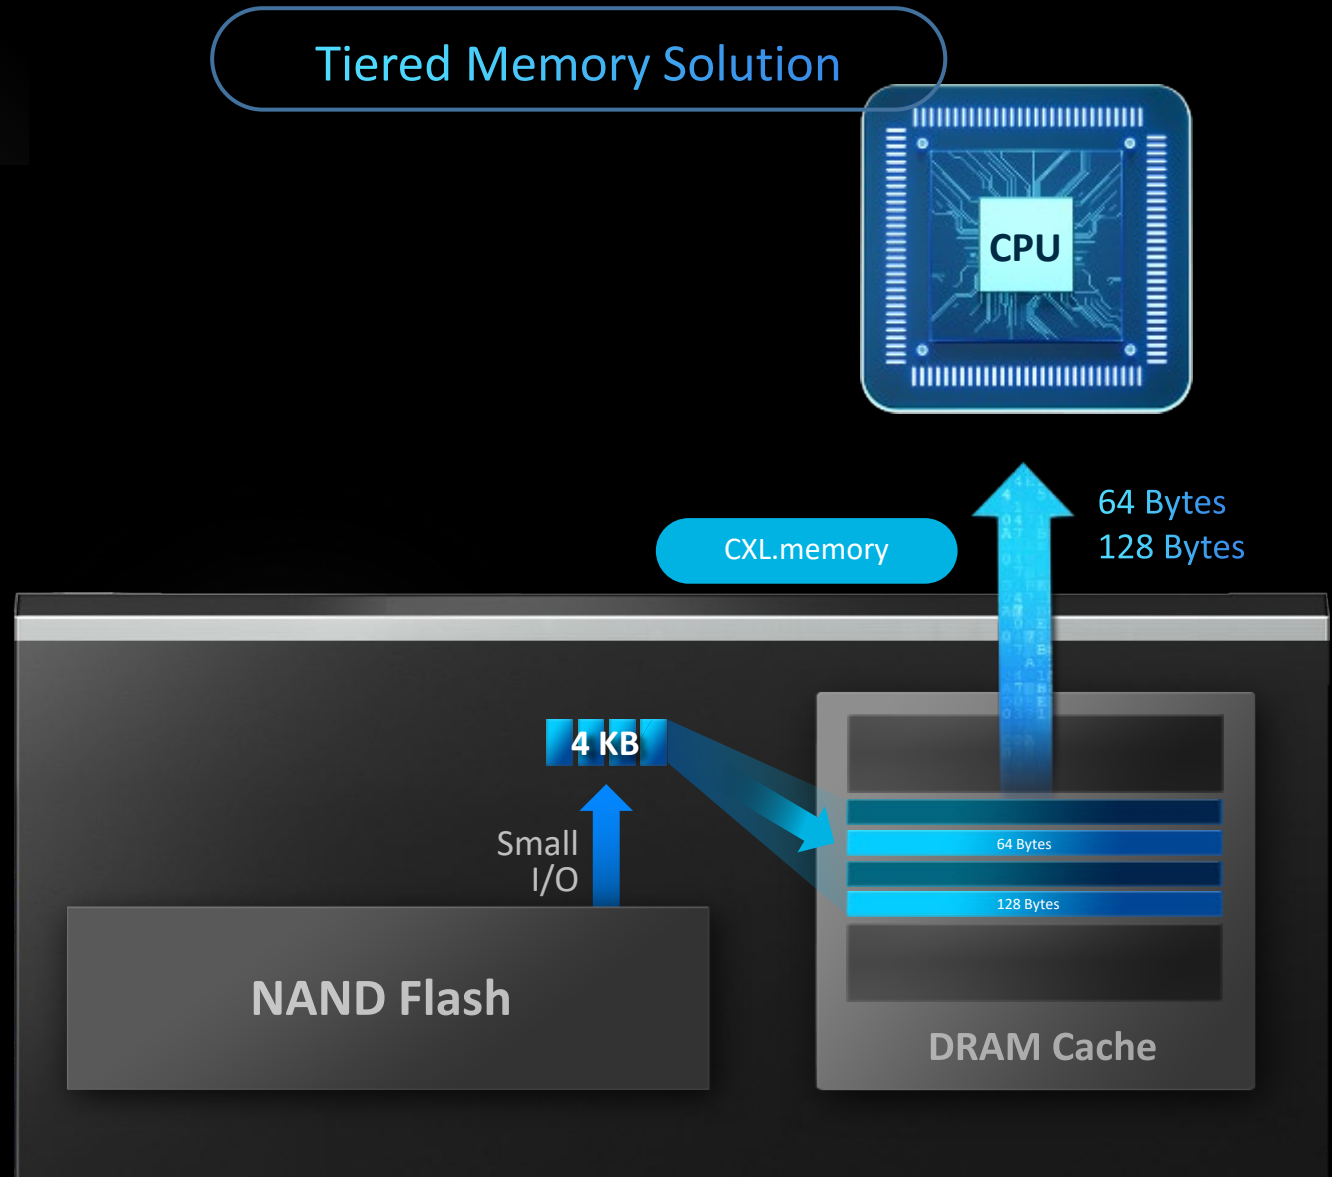

Secondary Memory Options

- Example of Memory Configuration with TM Mode

<2TB Main Memory (Case-1)>

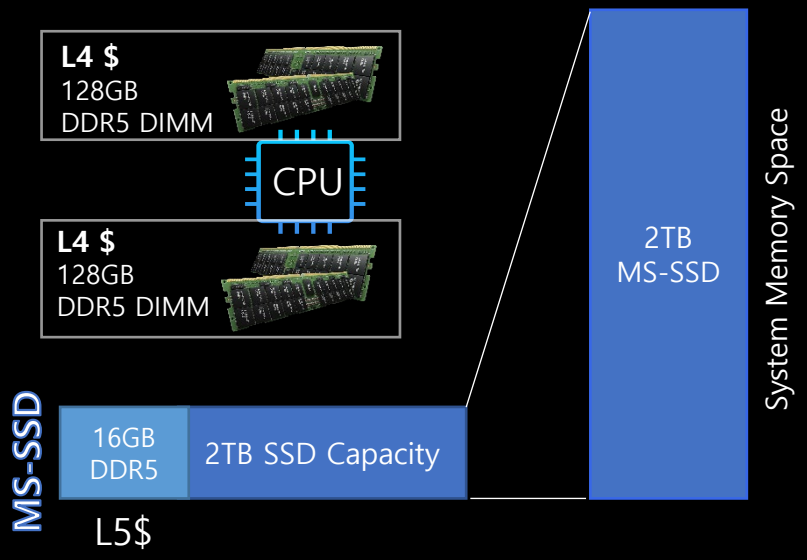

<1-Tier Host-managed Device Memory (Case-2)>

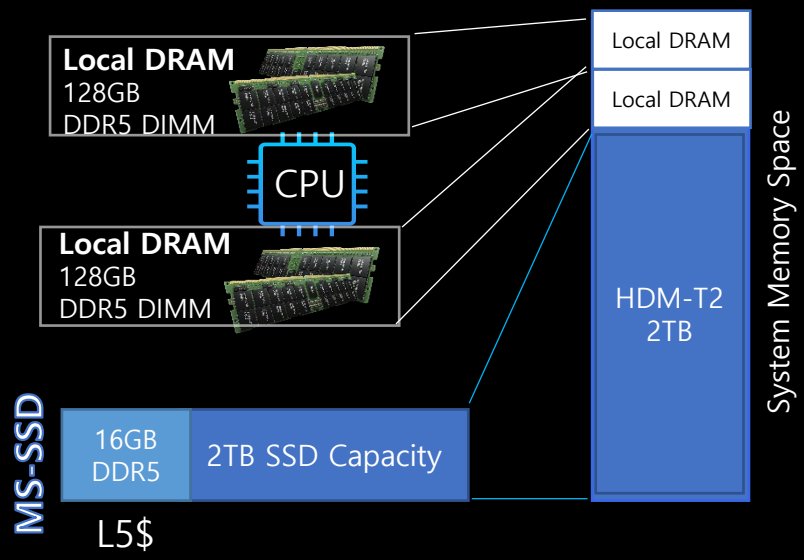

<2-Tier Host-managed Device Memory (Case-3)>

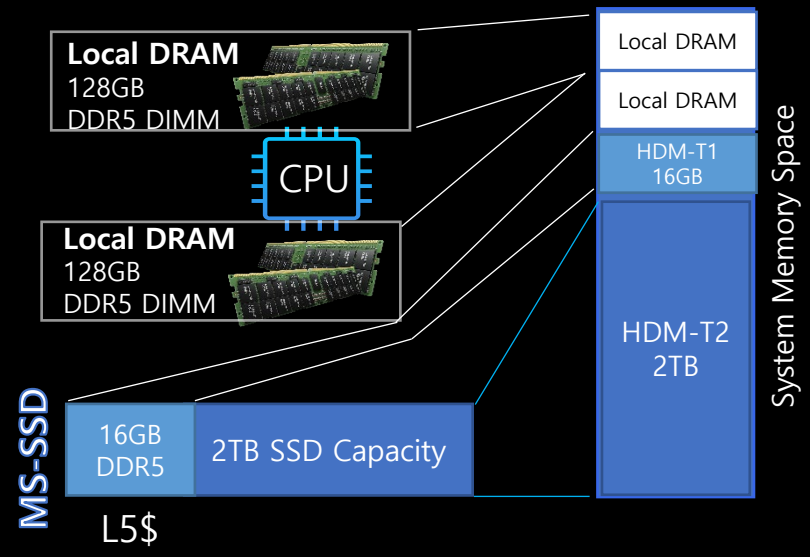

SDC 23

Option 1 Performance

Key Features & Benefits

- Small granularity data access enable performance scales with cache hits
- Direct memory access advantage; no software cache overhead
- Large memory capacity at lower TCO

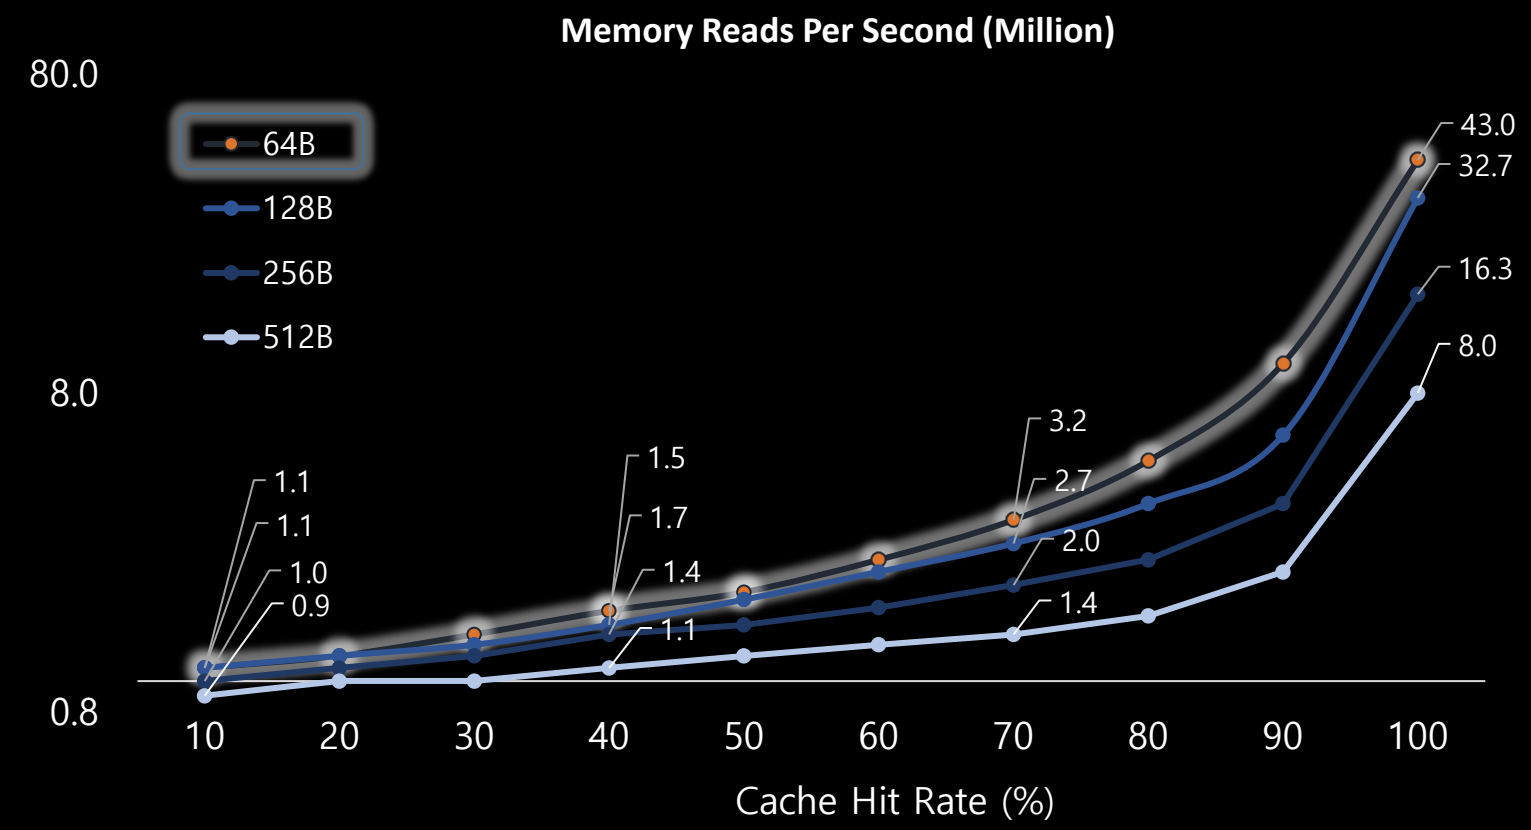

**Compared to PCIe Gen4 NVMe SSD

CXL-based SSD as Fast Small IO Storage

CXL.Mem Read	Random Perf (128B)	Cache hit 0%	0.8 MIOPS
		Cache hit 50%	1.5 MIOPS
		Cache hit 100%	35.0 MIOPS
	Latency	Cache hit/miss	<1us/ 70us
CXL.io	Seq. perf (128KB)	Read: 5,500 MB/s Write: 2,000 MB/s	
	Random Perf (4KB)	Read: 800 KIOPS Write: 85 KIOPS	

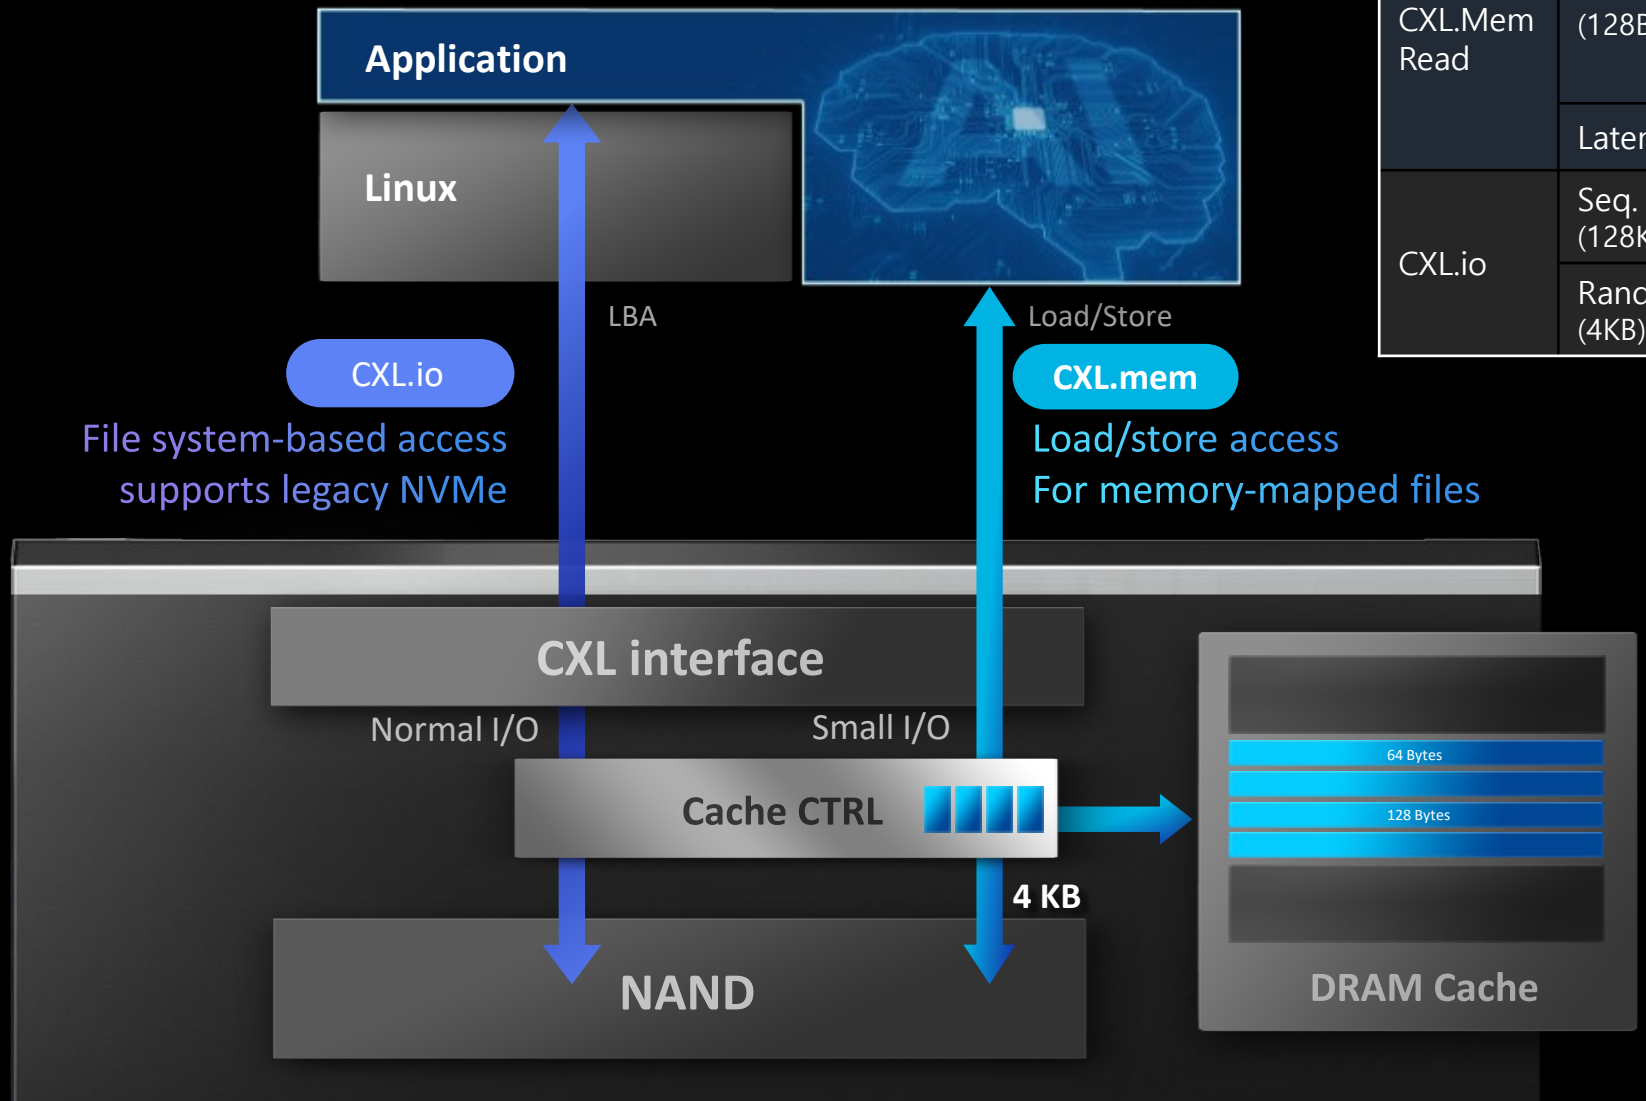

SDC 23

Fine-grain Access to Storage Data

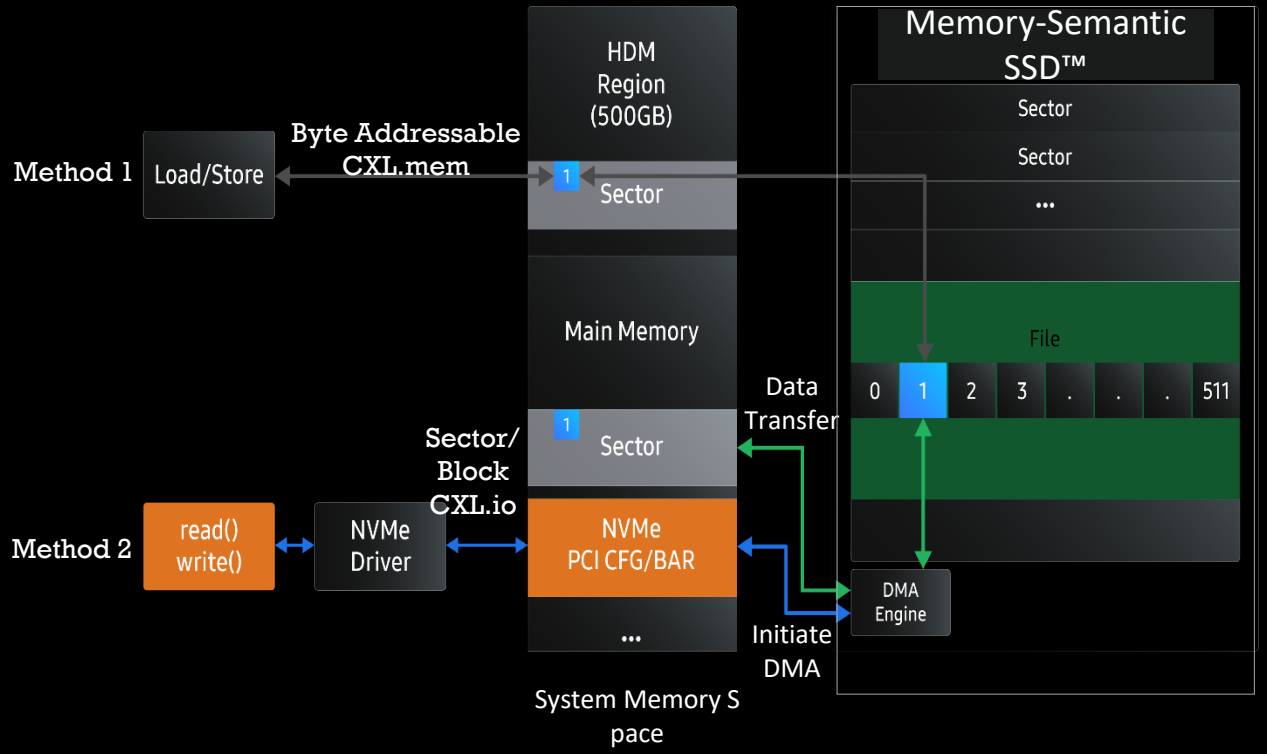

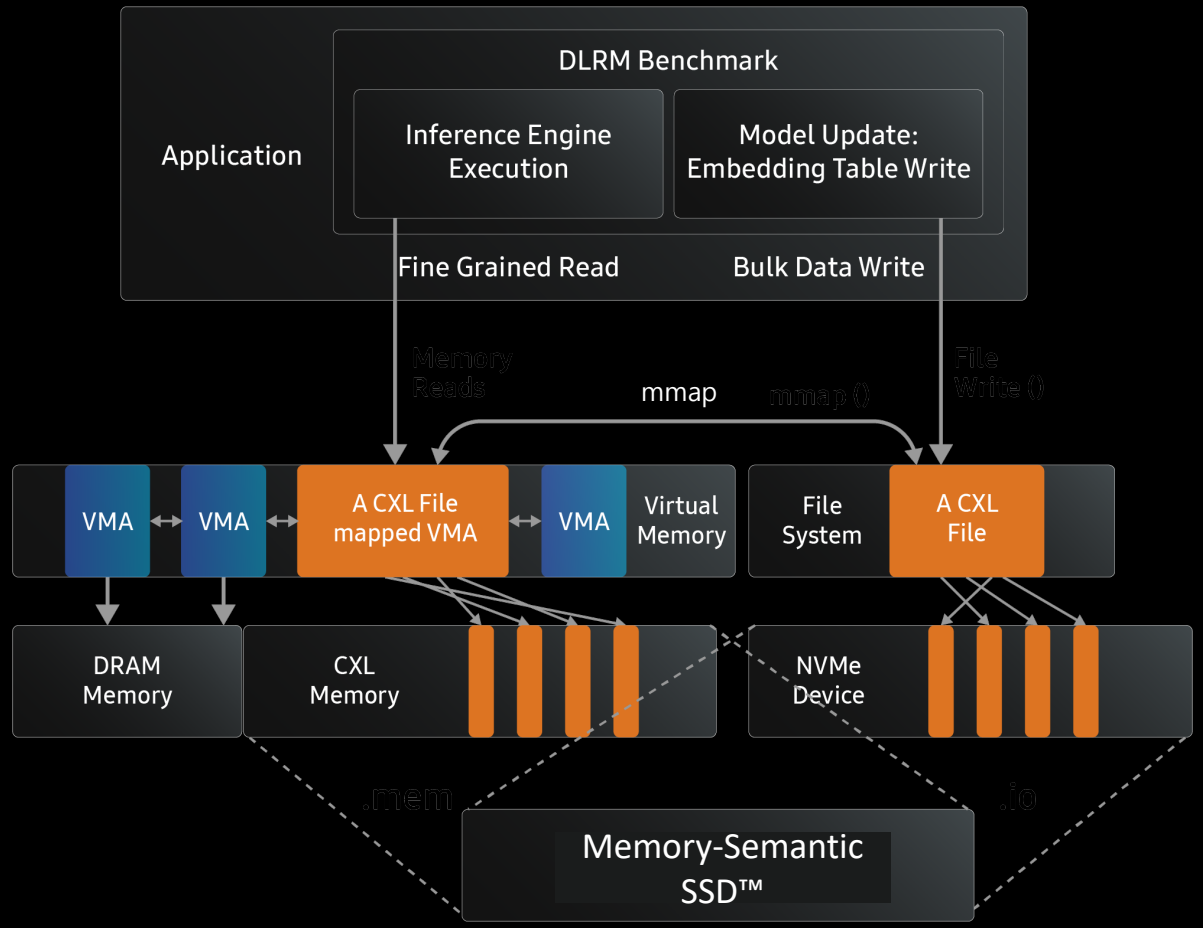

SDC 23

DLRM Performance with Fast Small IOs

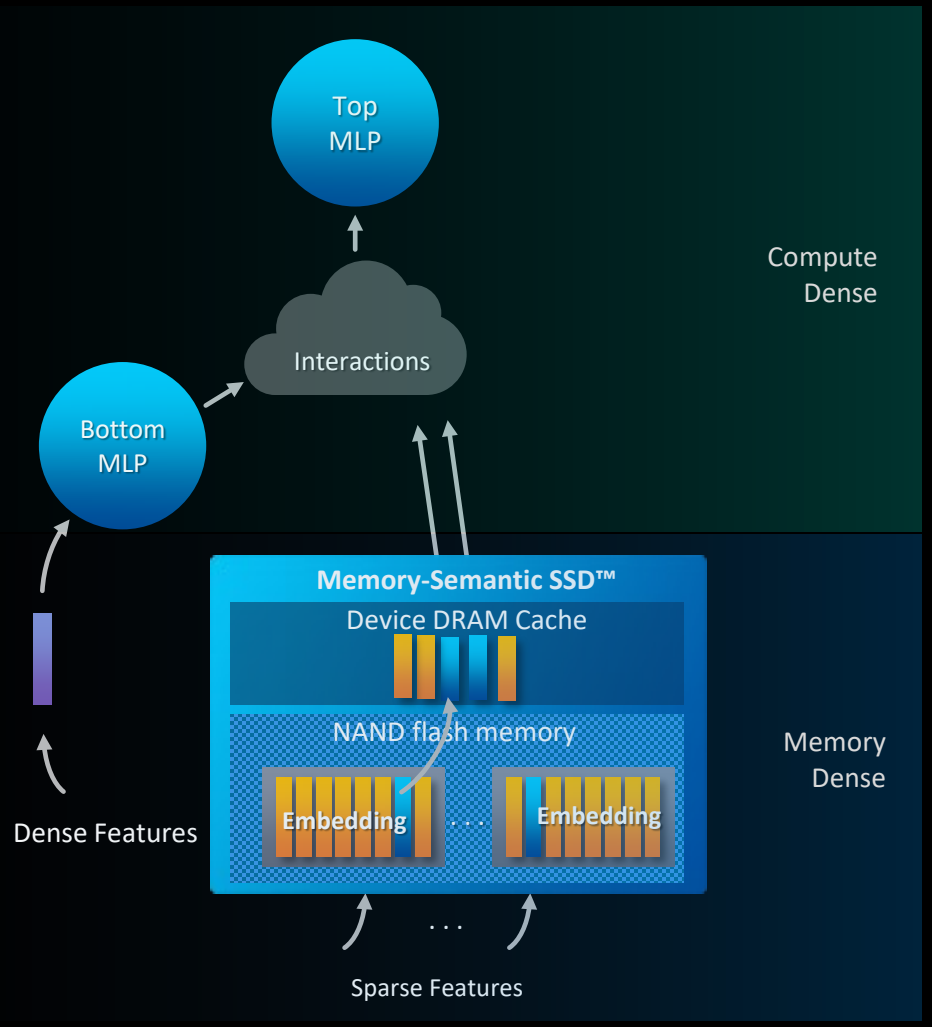

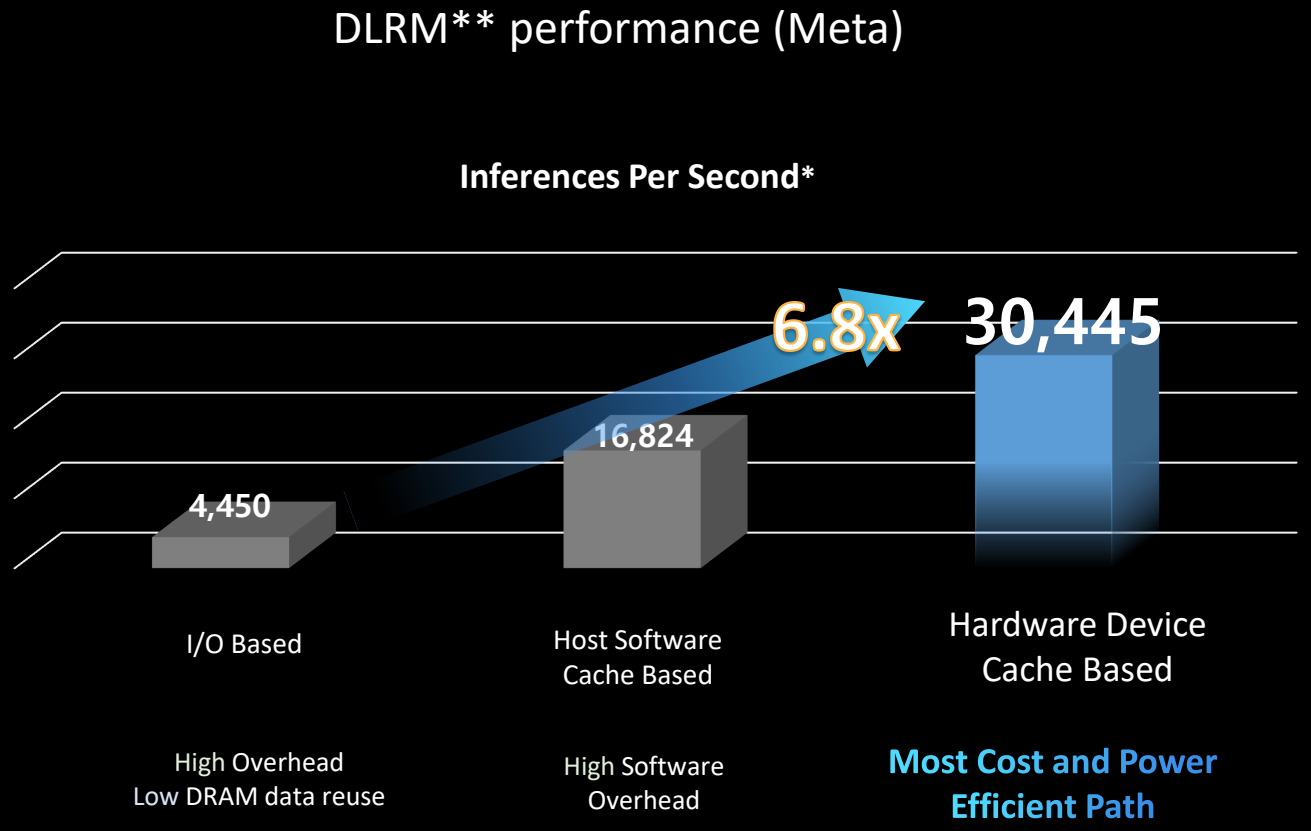

* Results based on publicly available DLRM workload traces from Meta and FPGA based PoC Memory-Semantic SSD™
** DLRM : Deep Learning Recommendation Model

Movie Recommendation System Demo

SAMSUNG

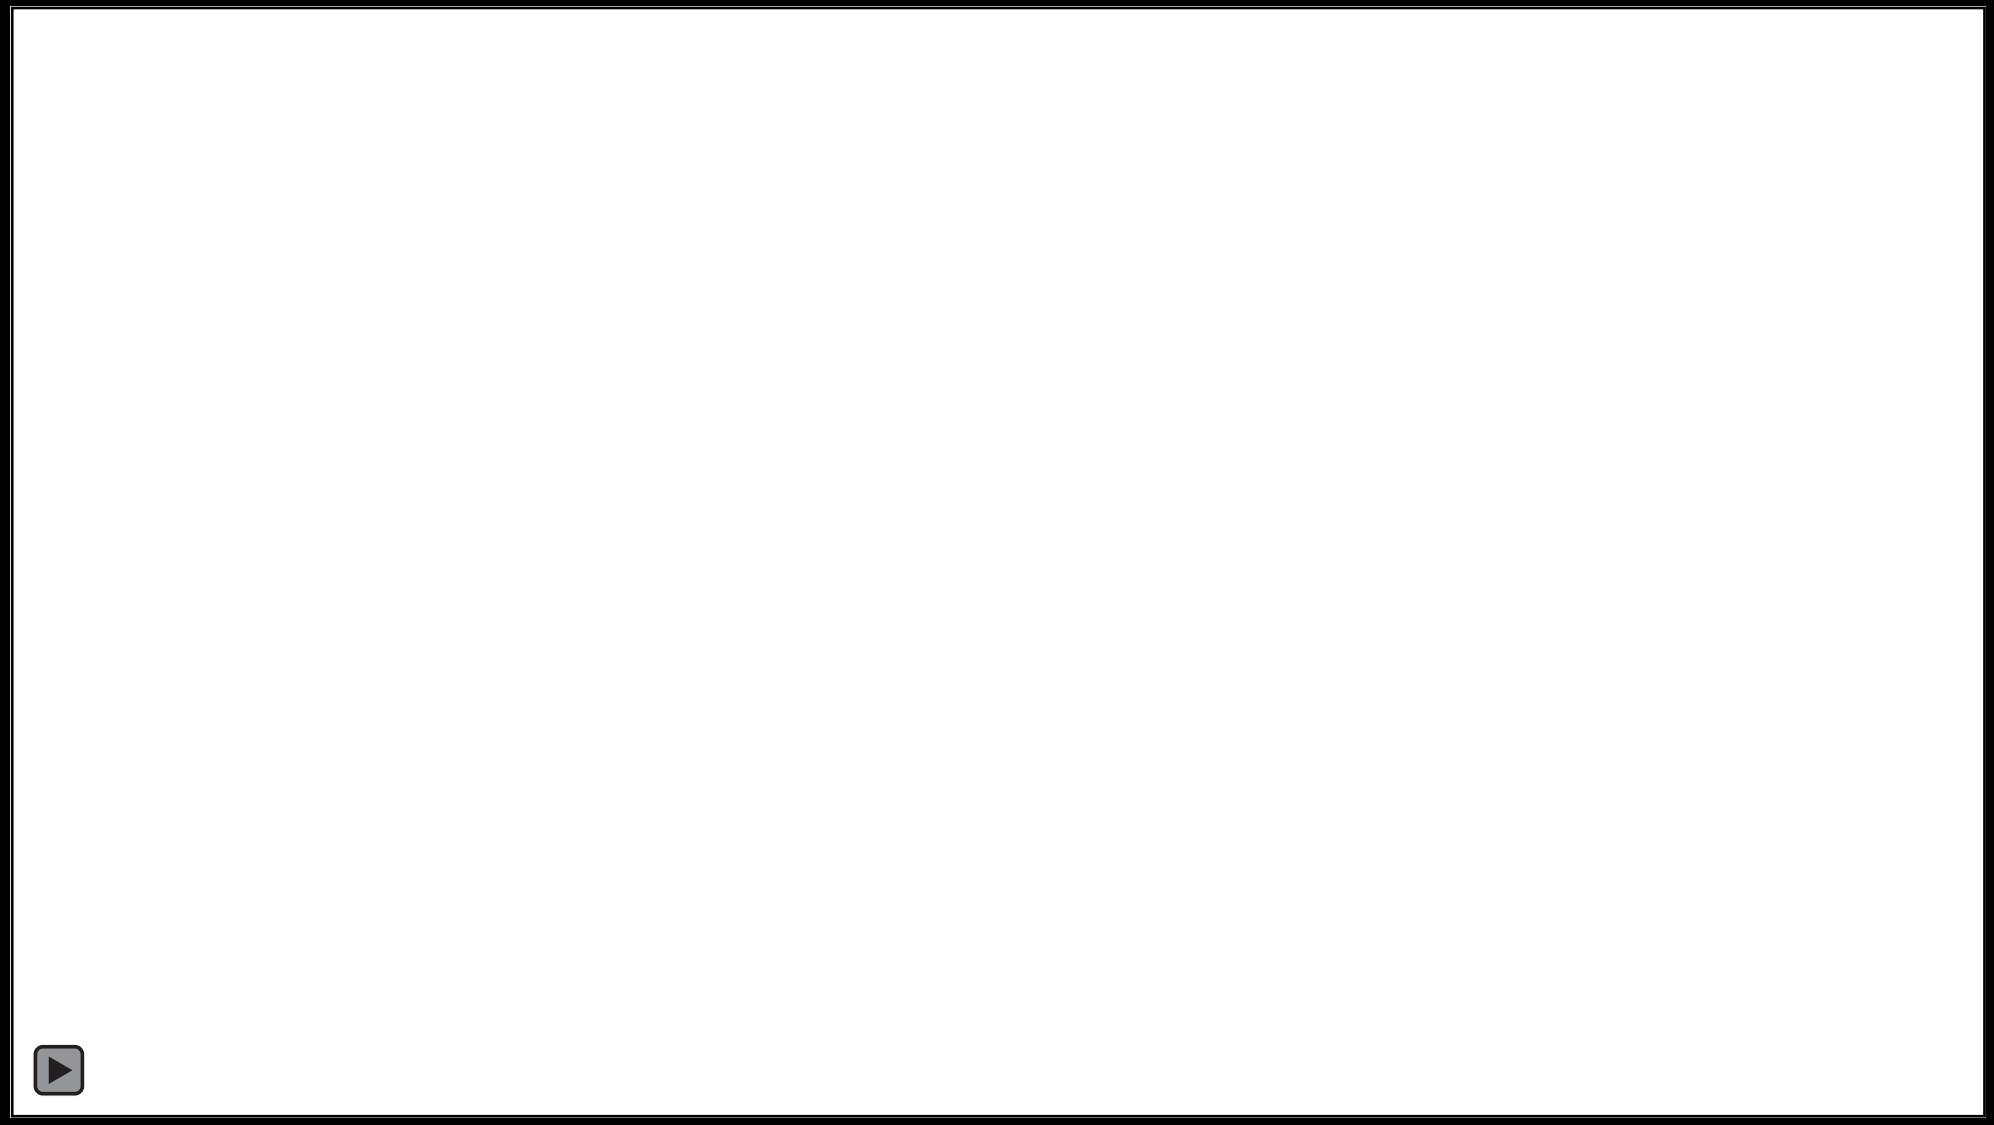

Challenges and Opportunities

Standard and Eco

- No definition and spec

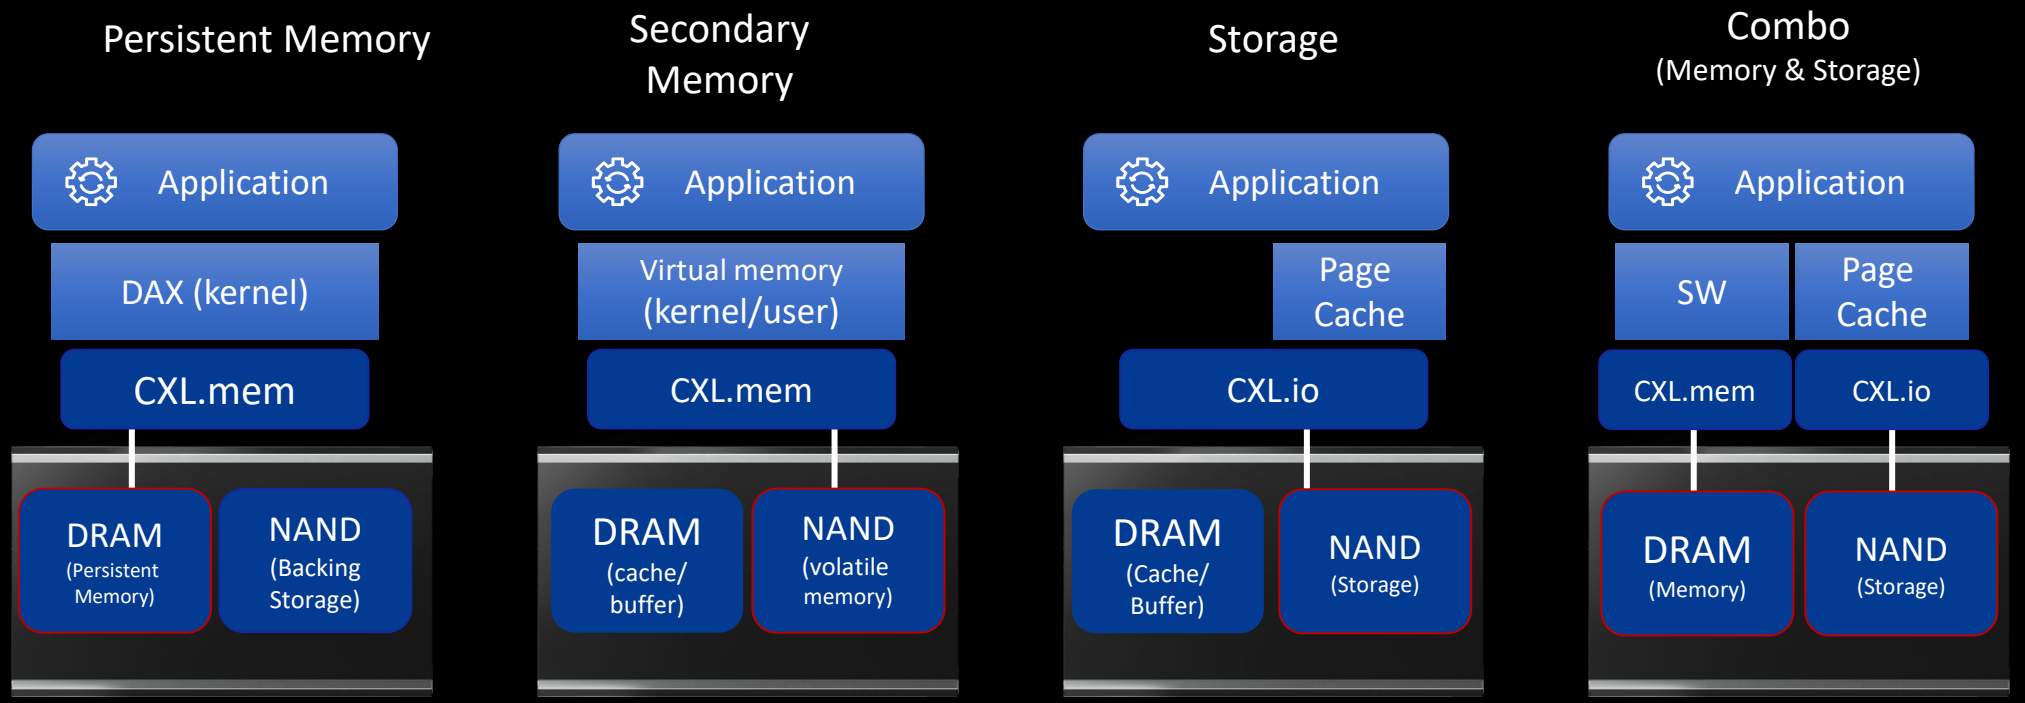

Latency Tolerance

- Impact of long latency on CPU performance

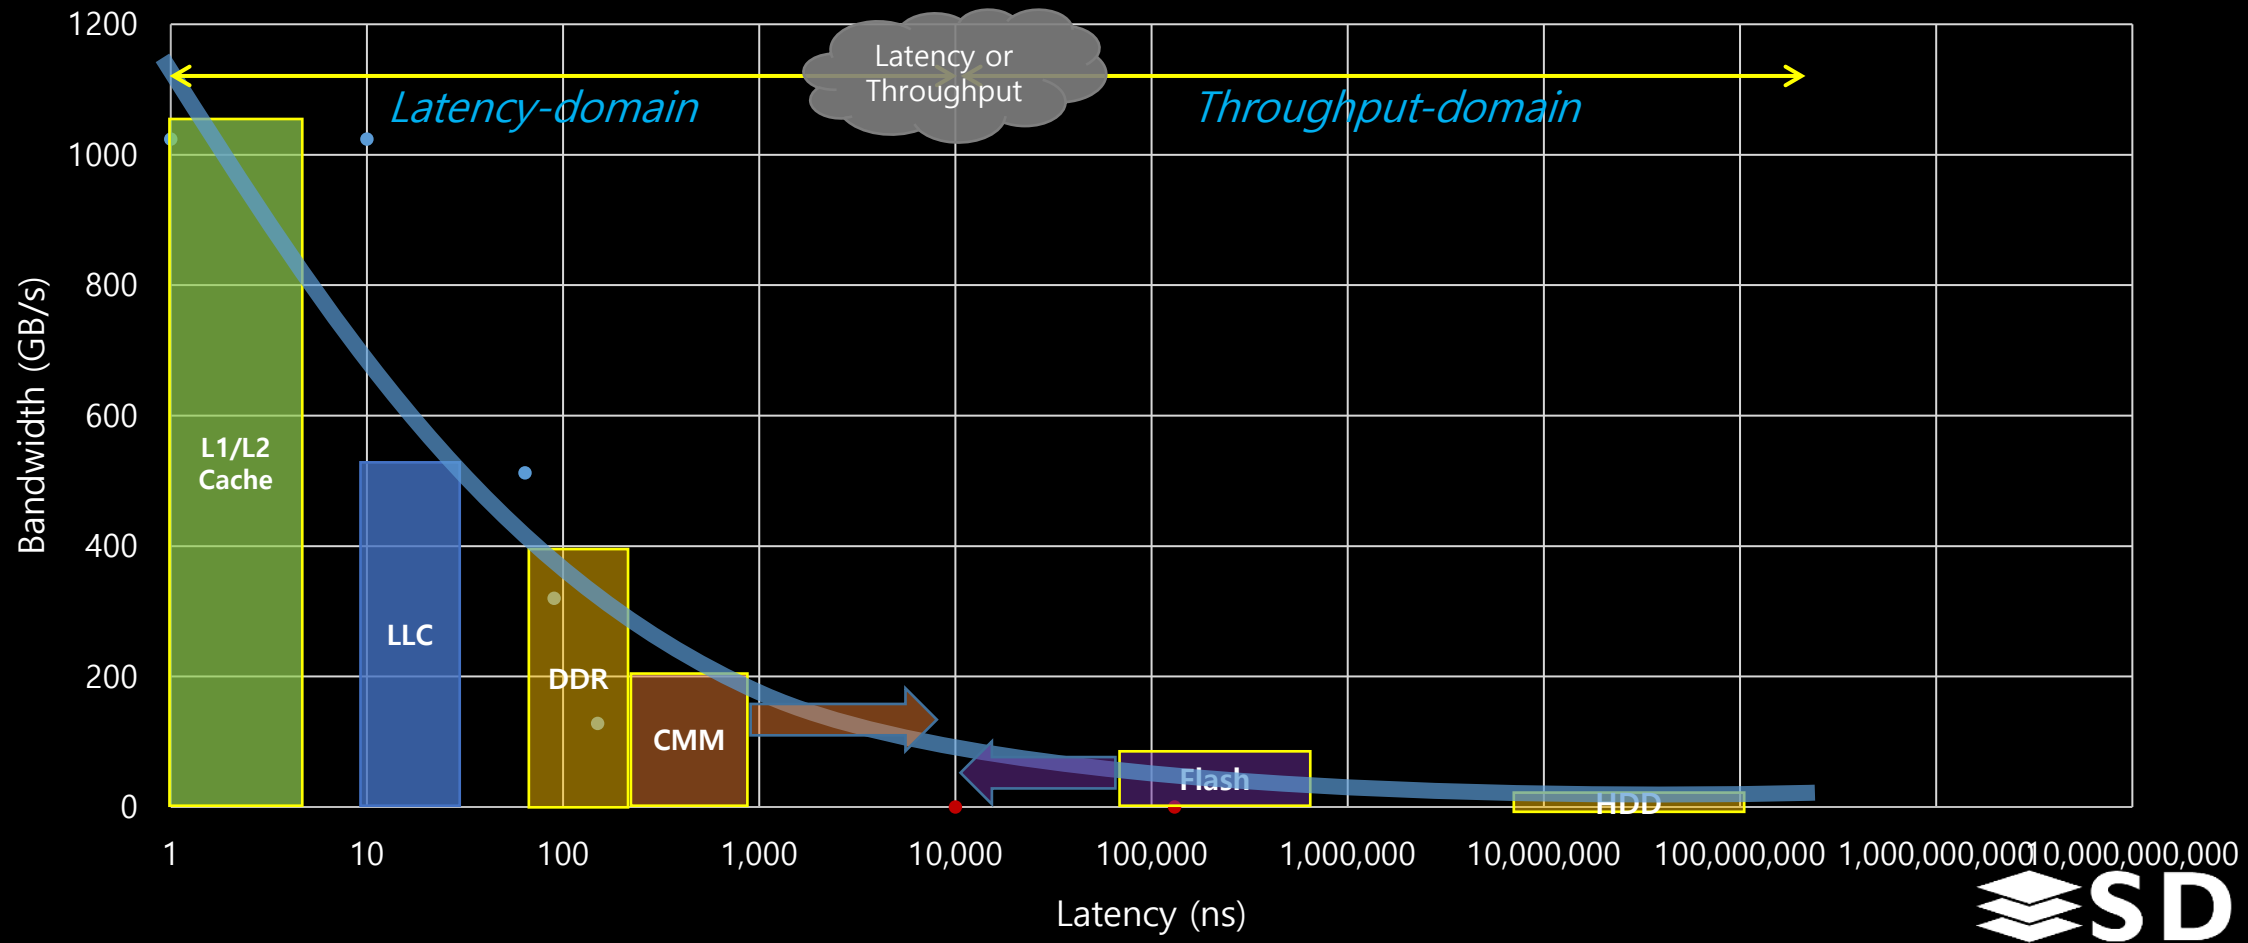

Cache Management

- **Managing in-device DRAM is the key!**

Key Features & Benefits

- Close to DRAM end-to-end performance at a lower TCO*
- Up to ~10x better end-to-end performance with FPGA-based PoC**

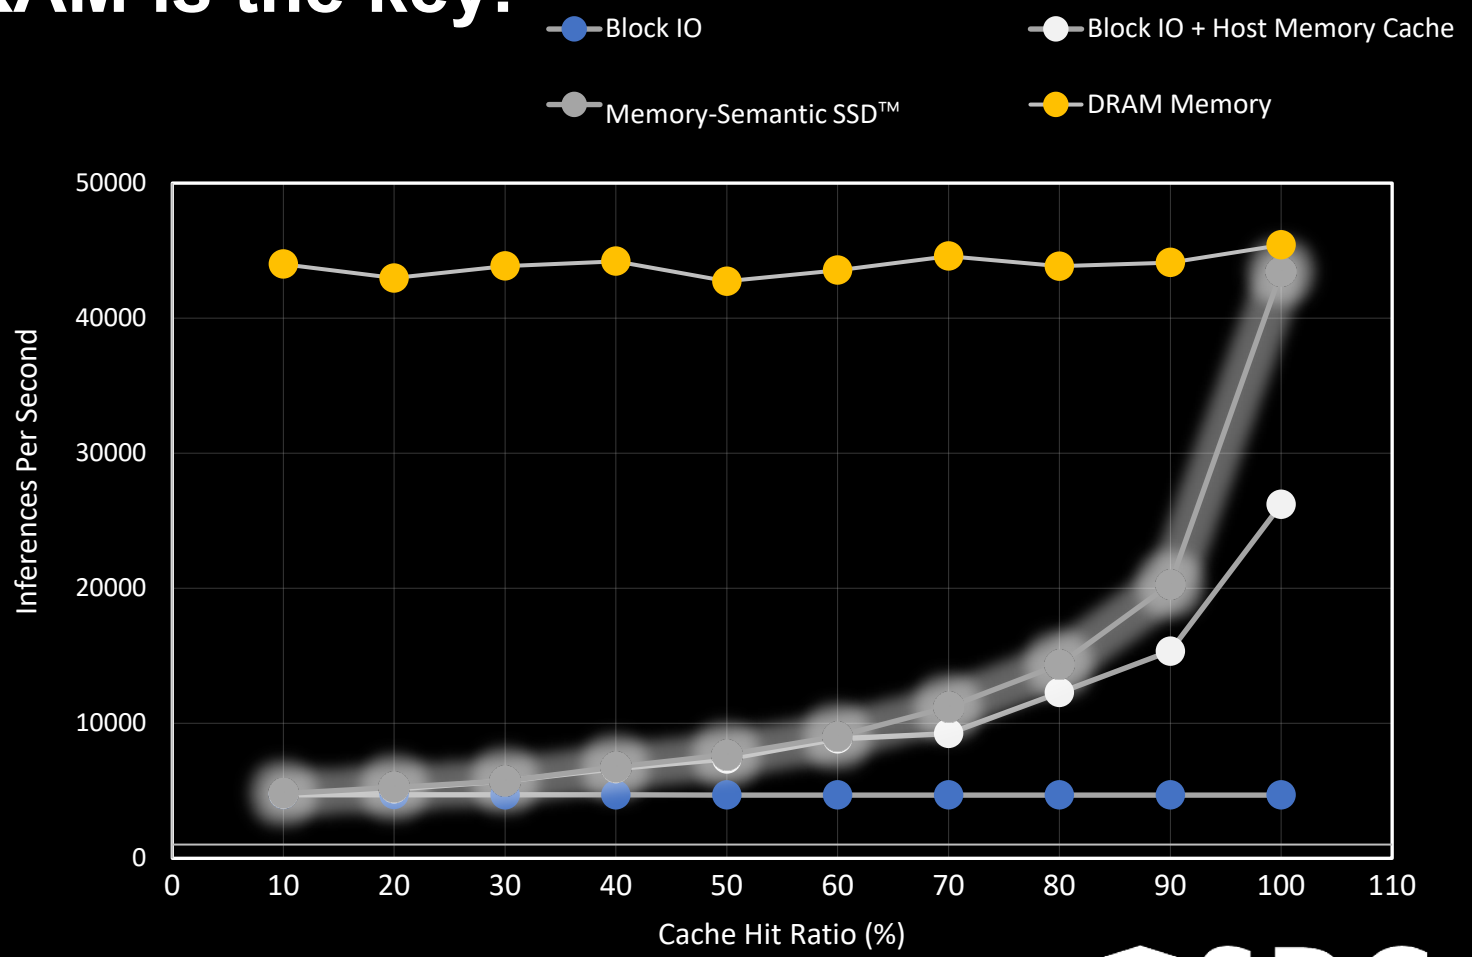

* When 100% hit ratio
**Compared to PCIe Gen4 NVMe SSD

Wrap Up

- **SSD with CXL interfaces for**
  - Persistent memory
  - Performant secondary memory
  - Storage for AI and HPC
  - Near data processing platform
- **Community efforts**
  - Standard for SSD with CXL interfaces (+cache management)
  - Software ecosystem
  - CPU architecture to tolerate long latency

Thank You

STORAGE DEVELOPER CONFERENCE

SDC 23

BY Developers FOR Developers

SAMSUNG