STORAGE INDUSTRY SUMMIT

Convergence of Storage and Memory Developing the Needed Ecosystem

JANUARY 20, 2016, SAN JOSE, CA

Rick Coulson

Intel Sr. Fellow

3D XPoint™ Technology Drives System Architecture

SNIA SSSI | SOLID STATE STORAGE

The Journey to Low Latency Storage

- The Latency Journey So Far
- New Media Enables the Next Step
  - A little about 3D XPoint™ Technology
- System and SW Architecture Changes In Process
  - Changes to Block Storage Stack to Minimize Latency
  - Changes to Enable Persistent Memory

Starting Point: The First HDD

1957 IBM RAMAC 350

5 MBytes

$57,000

$15200/Mbyte

~1.5 Random IOPs

Why the Drive for Low Latency?

SNIA™
Solid State Storage Initiative

RAMAC 350
600ms

~100x
Access time reduction

10K RPM
~6ms access

~10,000x
Access time reduction

NAND SSD
~60us access

58 Years

RAMAC 305
100 Hz best
case "clock"

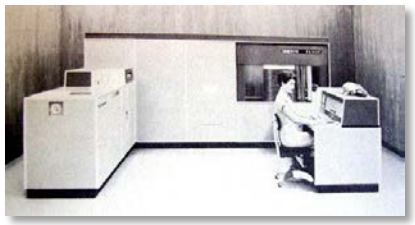

~30,000,000x
Clock speed increase

Core™ i7
~3 Ghz clock

Source: Wikipedia

Continuing the Storage Device Latency Journey

First HDD

Modern HDD

NAND SSD

NVMe SSD

Persistent Memory

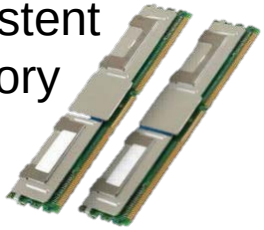

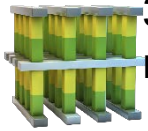

3D XPoint™ memory (SCM)

Ultra fast SSD

Drive for Lower Latency

Media Bottlenecks

Platform HW / SW bottlenecks

Media Enabler: 3D XPoint™ Technology
(Or any SCM)

SNIA™
Solid State Storage Initiative

Crosspoint Structure
Selectors allow dense packing and individual access to bits

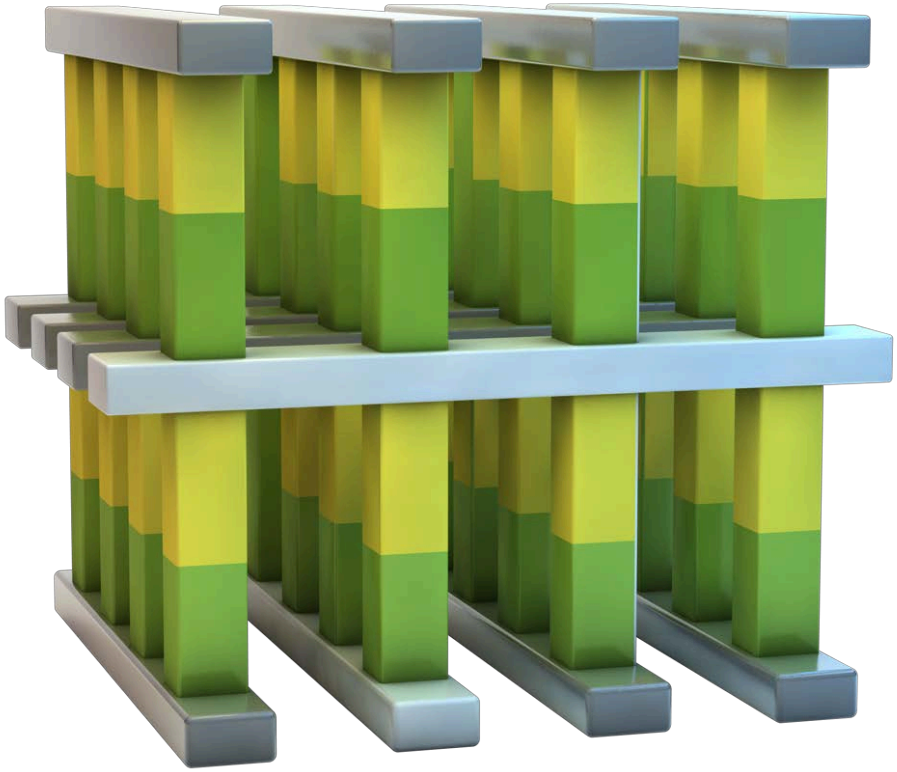

Breakthrough Material Advances
Compatible switch and memory cell materials

Scalable
Memory layers can be stacked in a 3D manner

High Performance
Cell and array architecture that can switch states 1000x faster than NAND

What is 3D Xpoint™ Technology

Video here

3D XPoint™ Technology Instantiation

SNIA™
Solid State Storage Initiative

Intel® Optane™ SSD

DIMMs based on 3D XPoint™

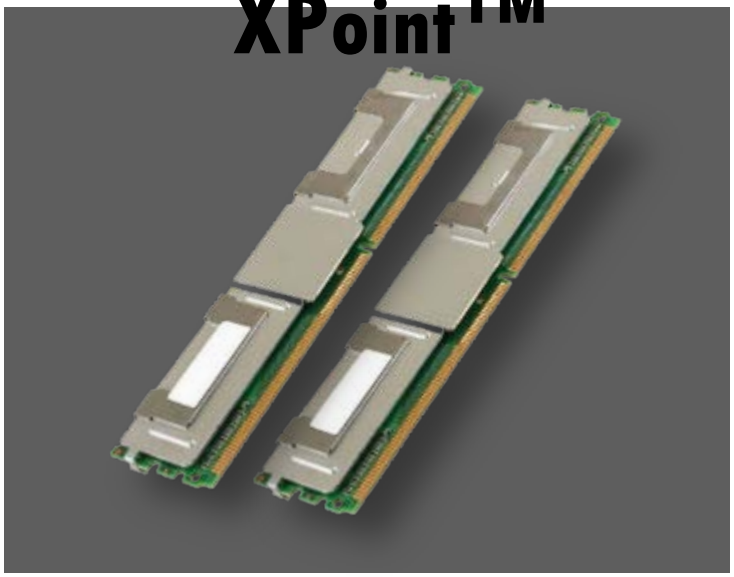

Demonstration of 3D Xpoint™ SSD Prototype

SNIA™
Solid State Storage Initiative

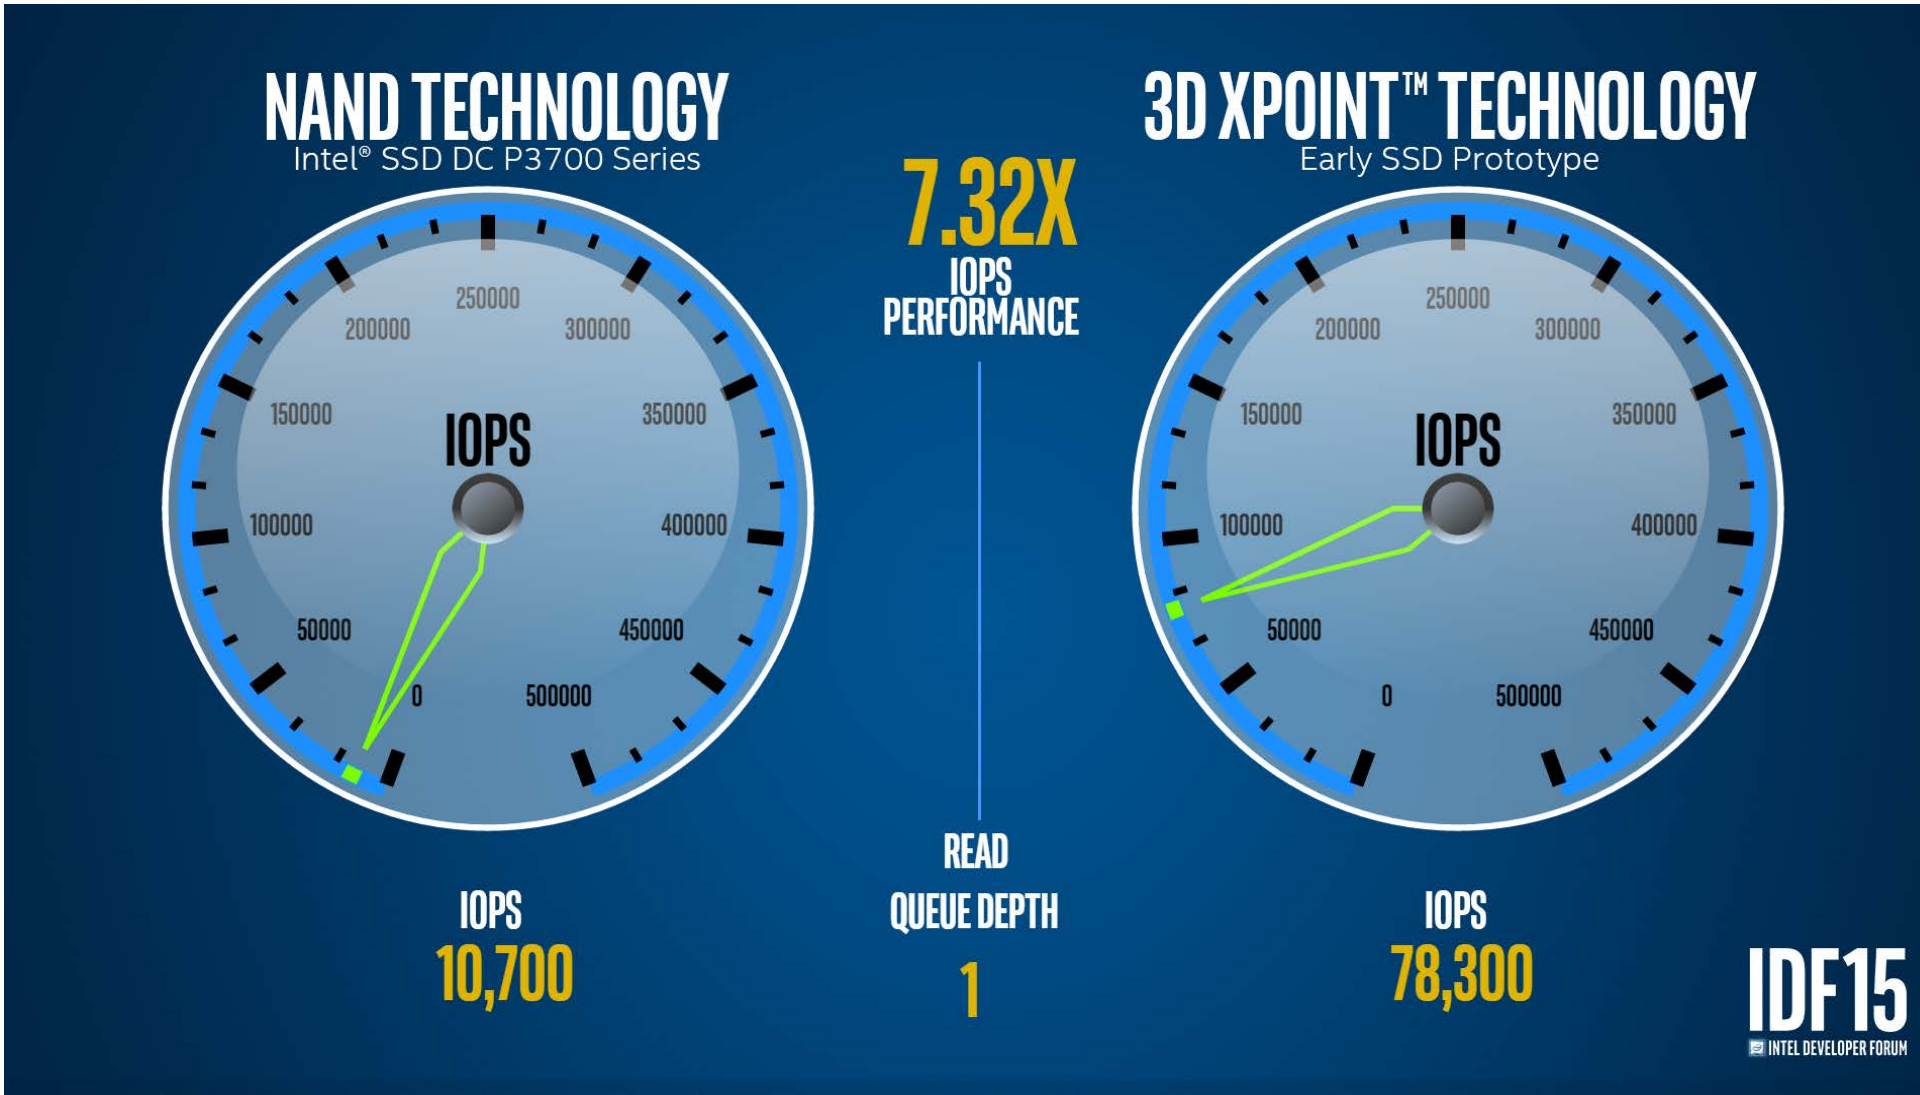

IDF15
INTEL DEVELOPER FORUM

The Need to Address System Architecture

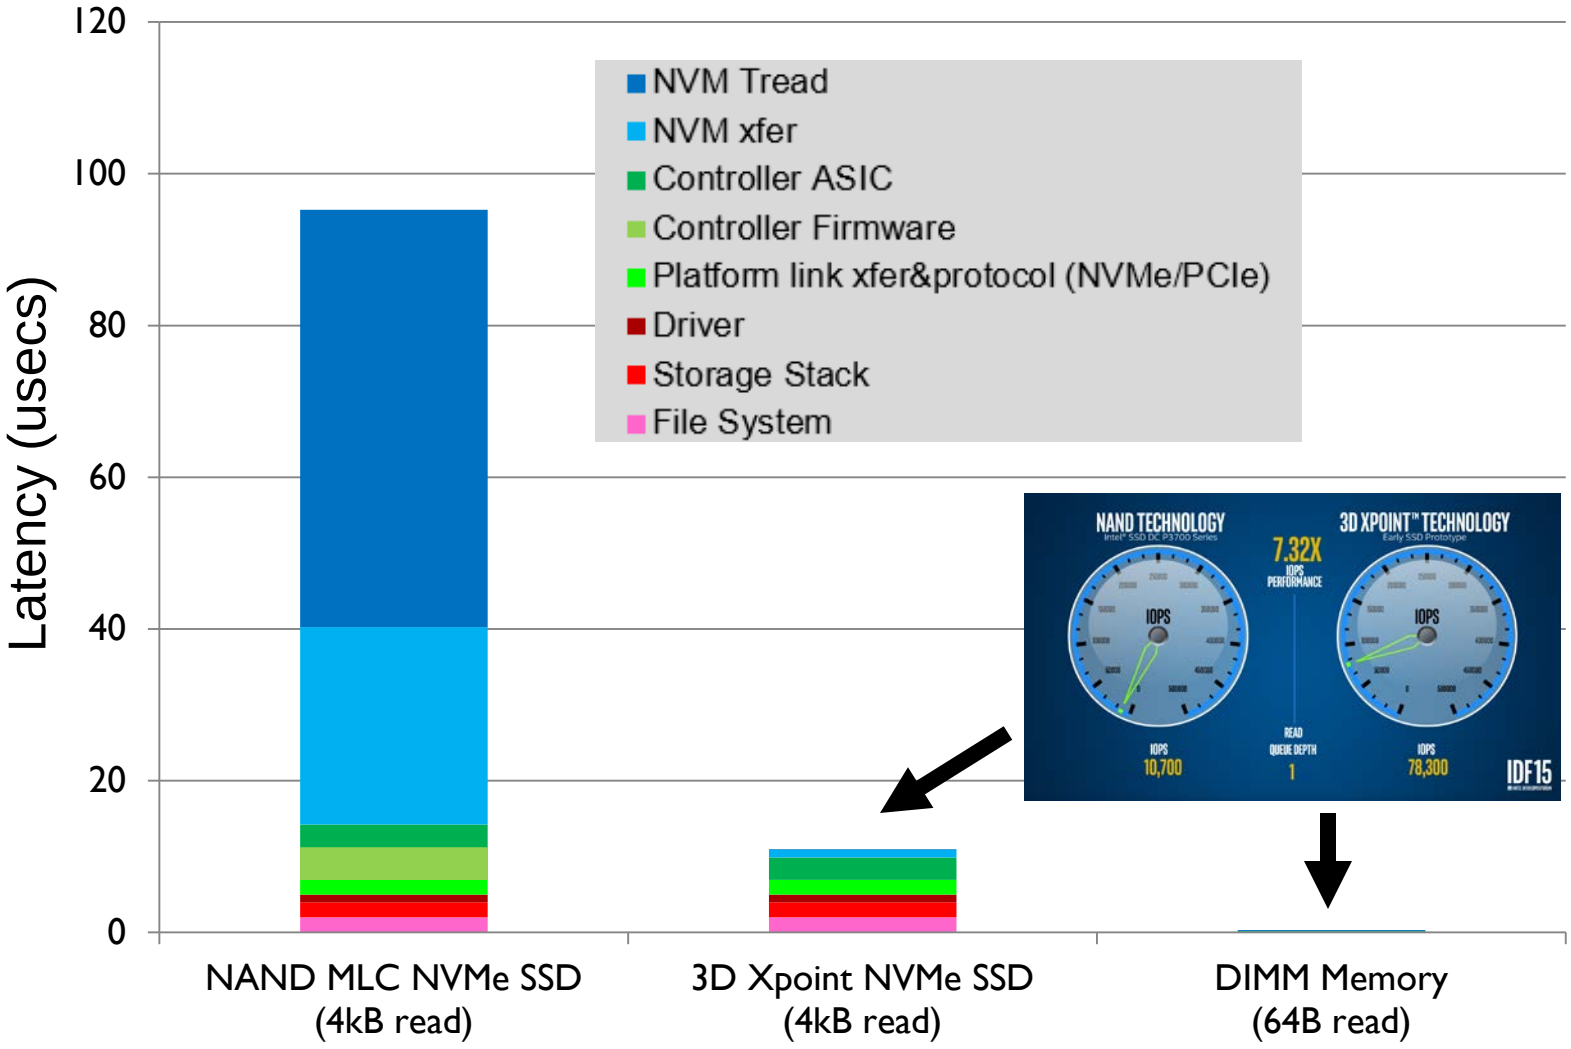

Storage Enabler: NVMe Efficiency Exposes Low 3D XPoint™ Media Latencies

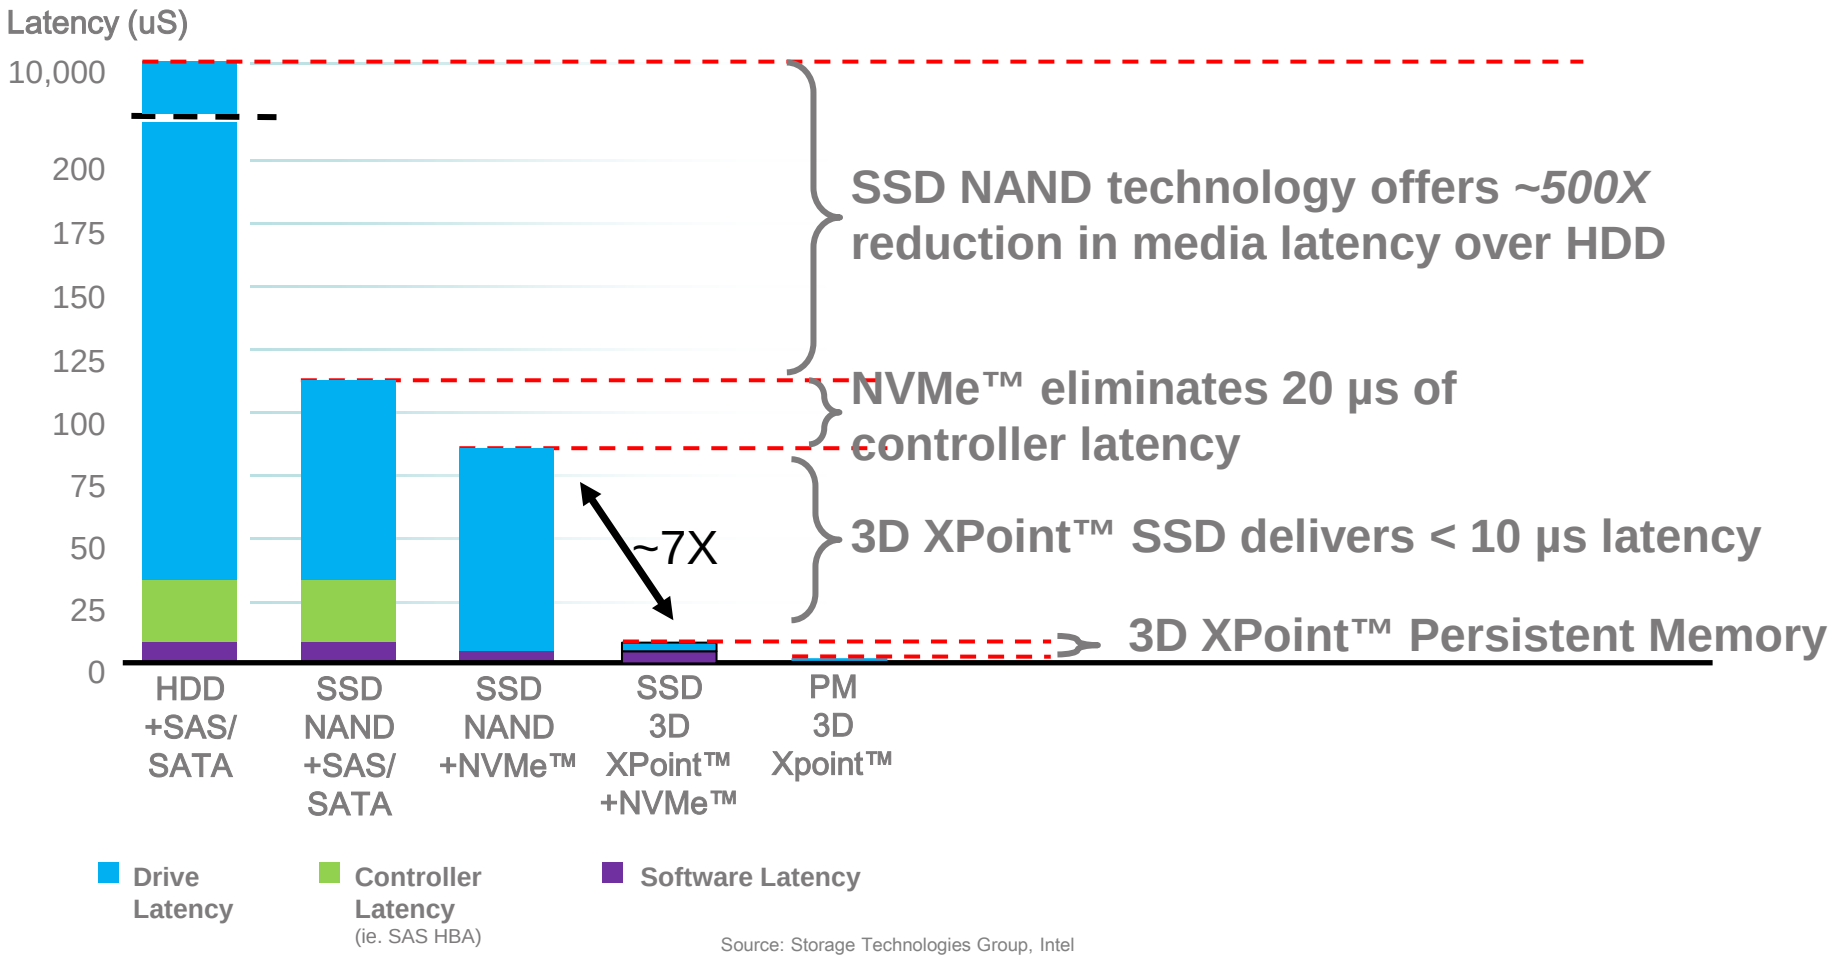

Source: Storage Technologies Group, Intel

NVMe Delivers Superior Latency

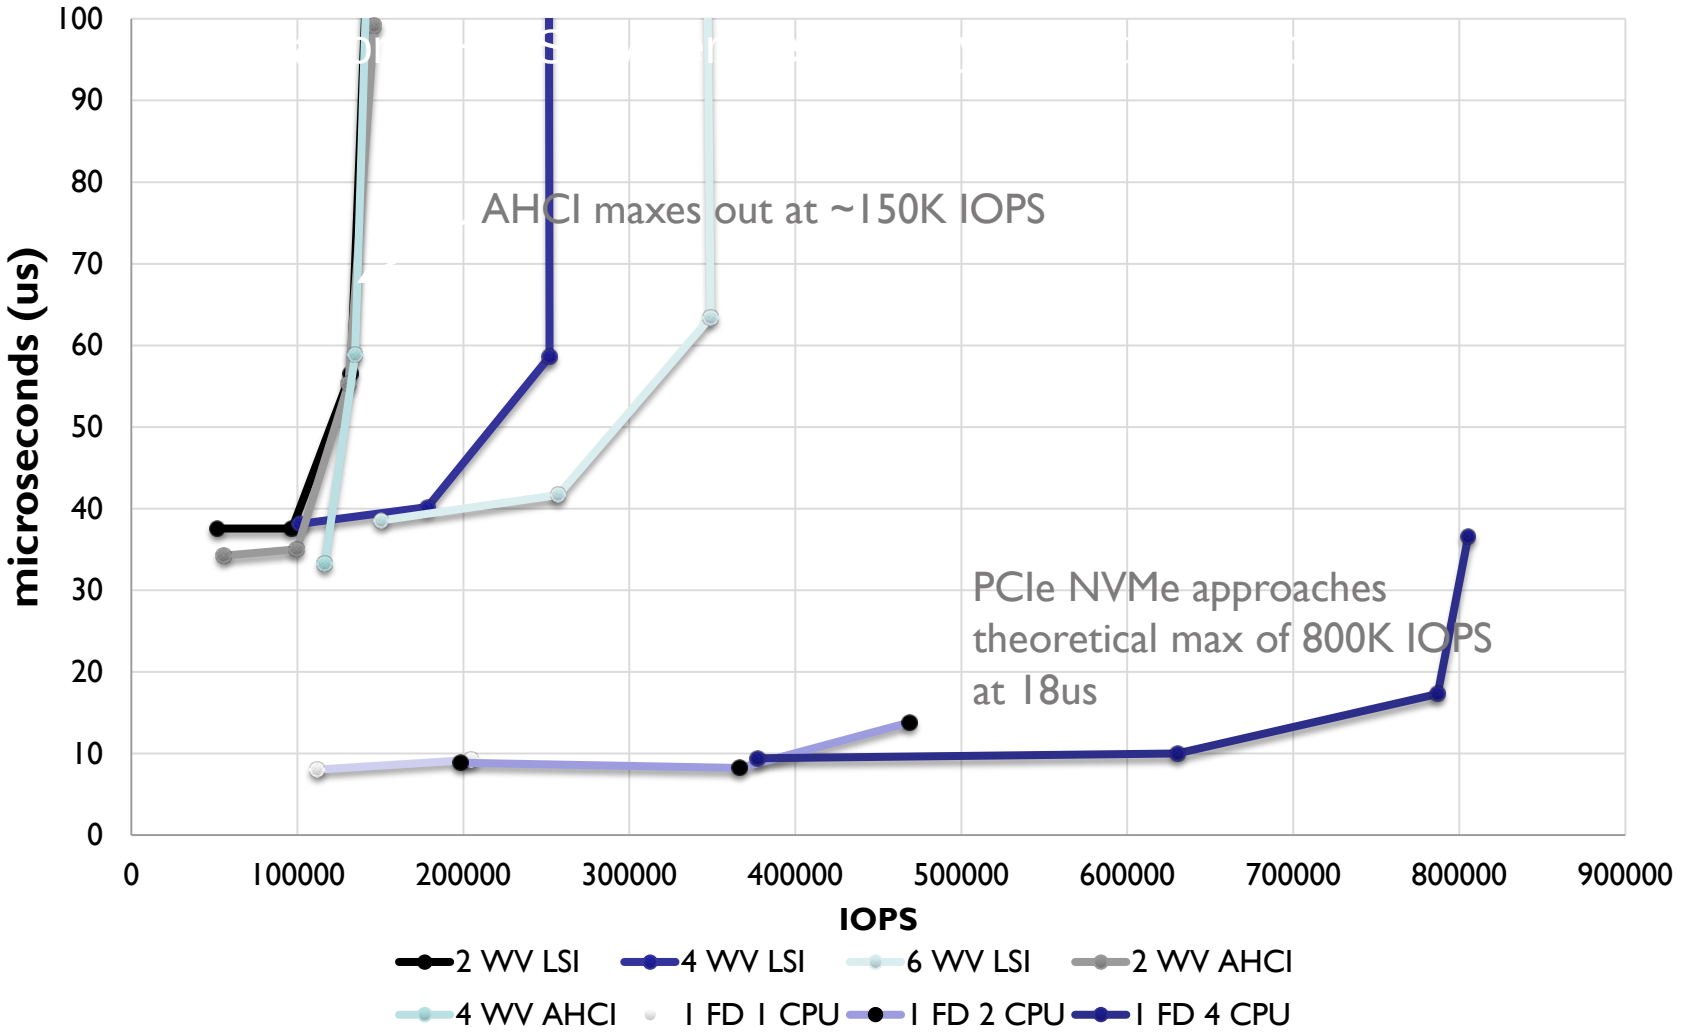

Source: Storage Technologies Group, Intel

NVMe/PCIe Provides More Bandwidth

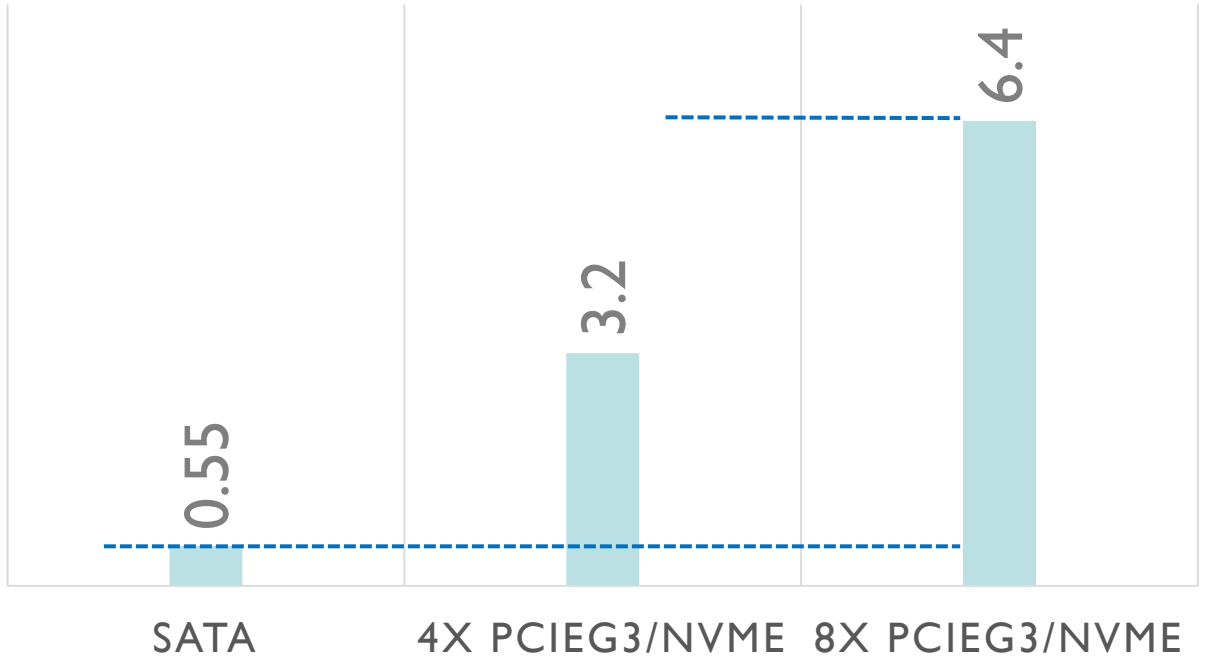

PCIe/NVMe provides more than 10X the Bandwidth of SATA. Even More with Gen 4

Source: Storage Technologies Group, Intel

Enabler: NVMe Over Fabrics

- In most Datacenter usage models, a storage write does not “count” until replicated
- High replication overhead diminishes the performance differentiation of 3D XPoint™ technology
- NVMe over Fabrics is a developing standard for low overhead replication

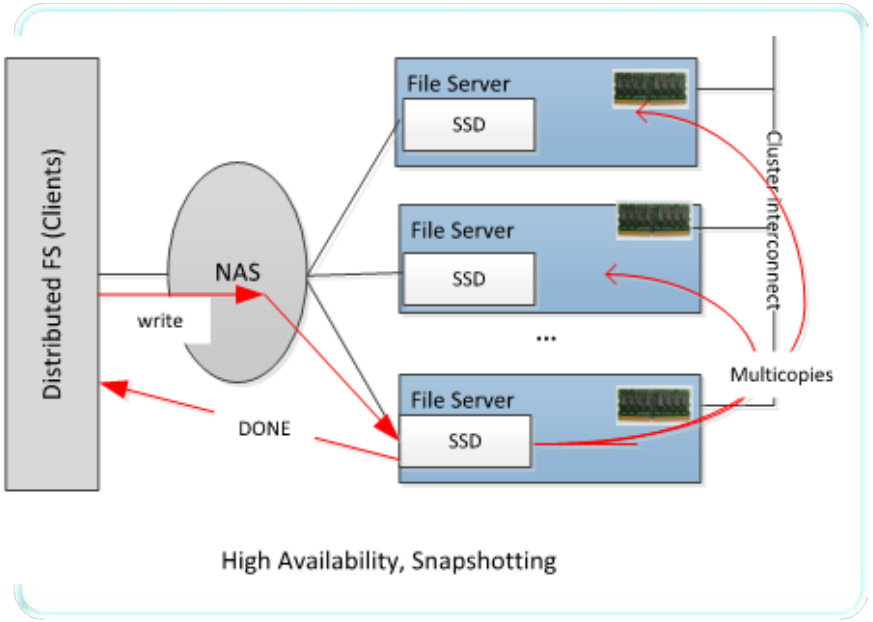

Synchronous Completion for QD1?

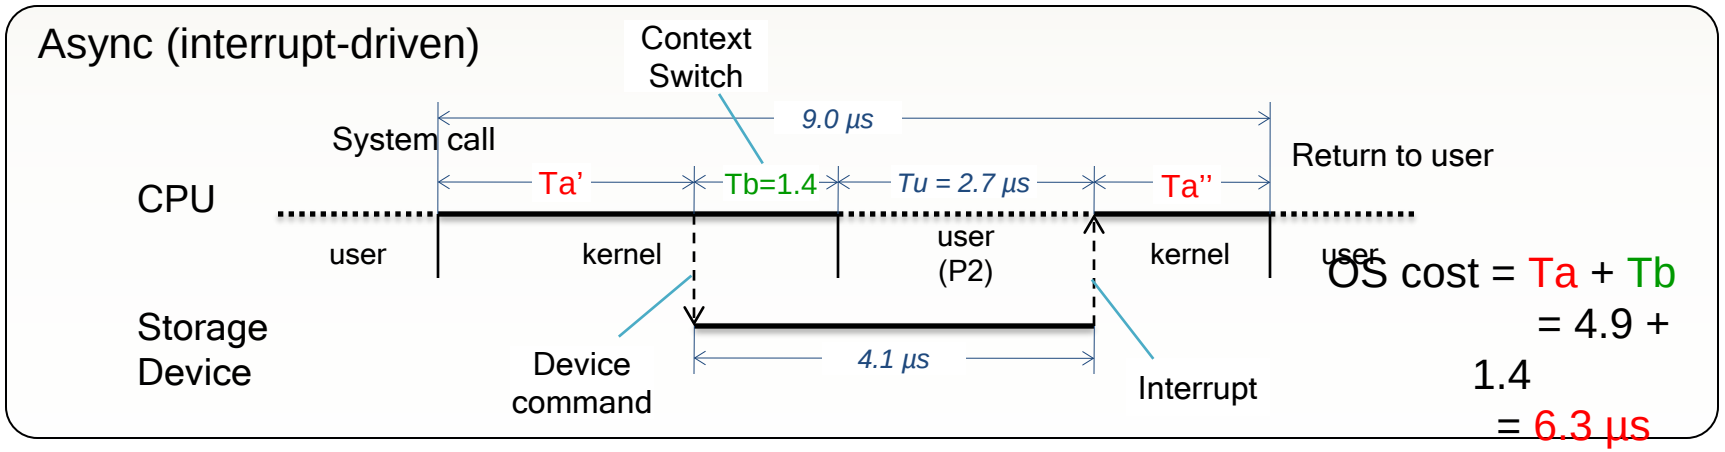

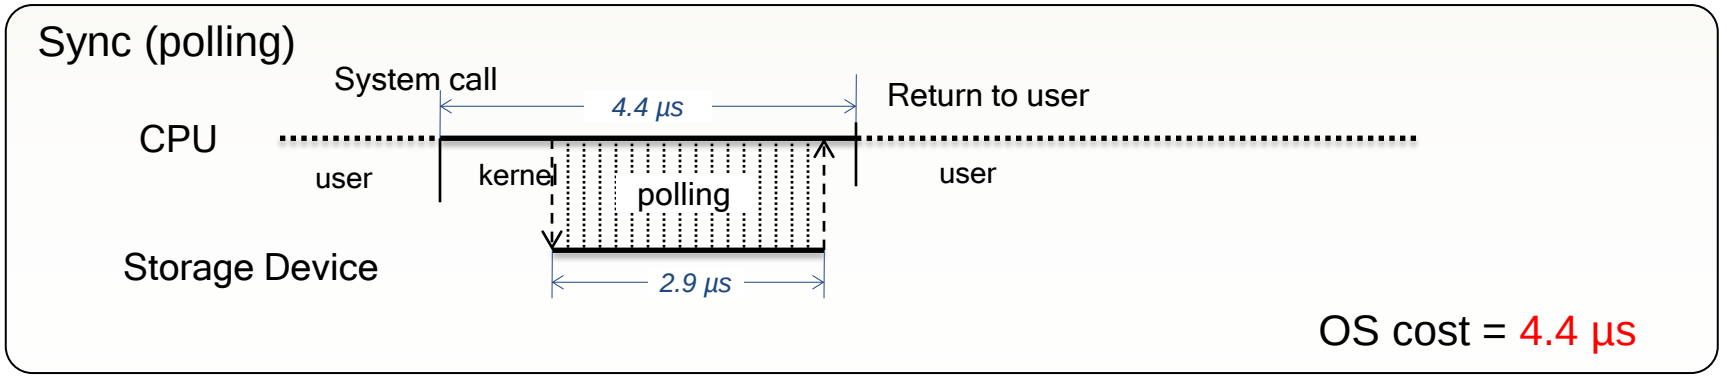

Synchronous completion also costs less OS / CPU time

Other Enablers

- Storage Stack optimizations
- Reduced Paging Overhead
- HW RAID alternatives

Persistent Memory

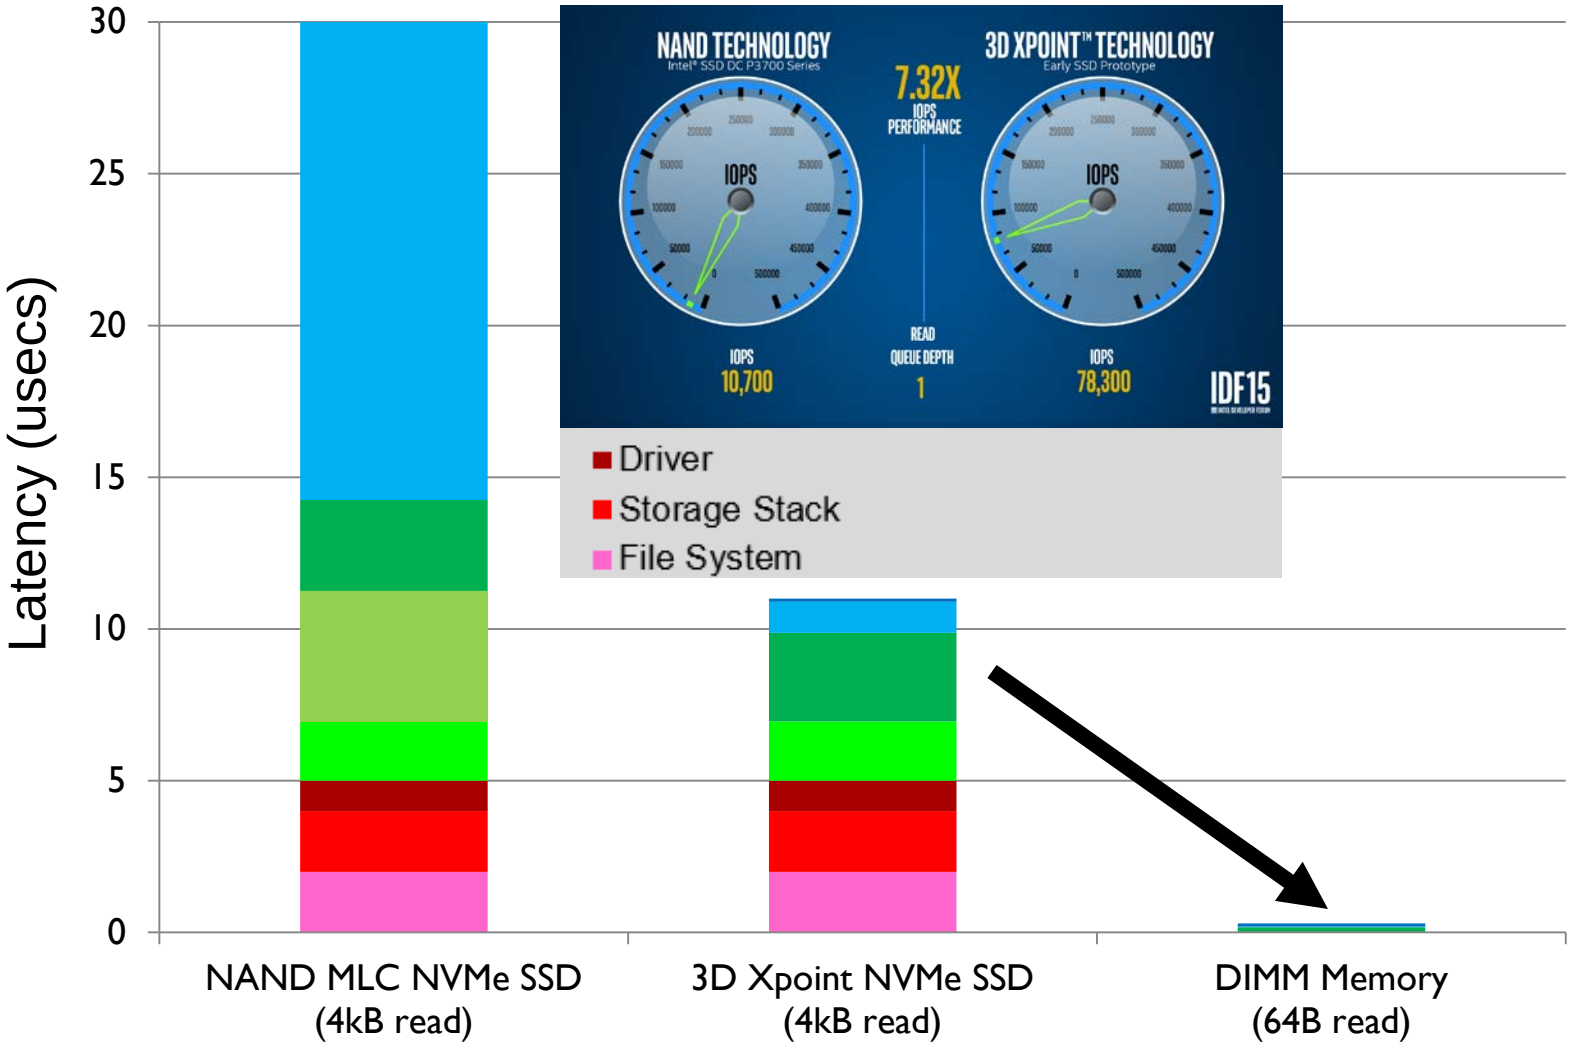

Open NVM Programming Model

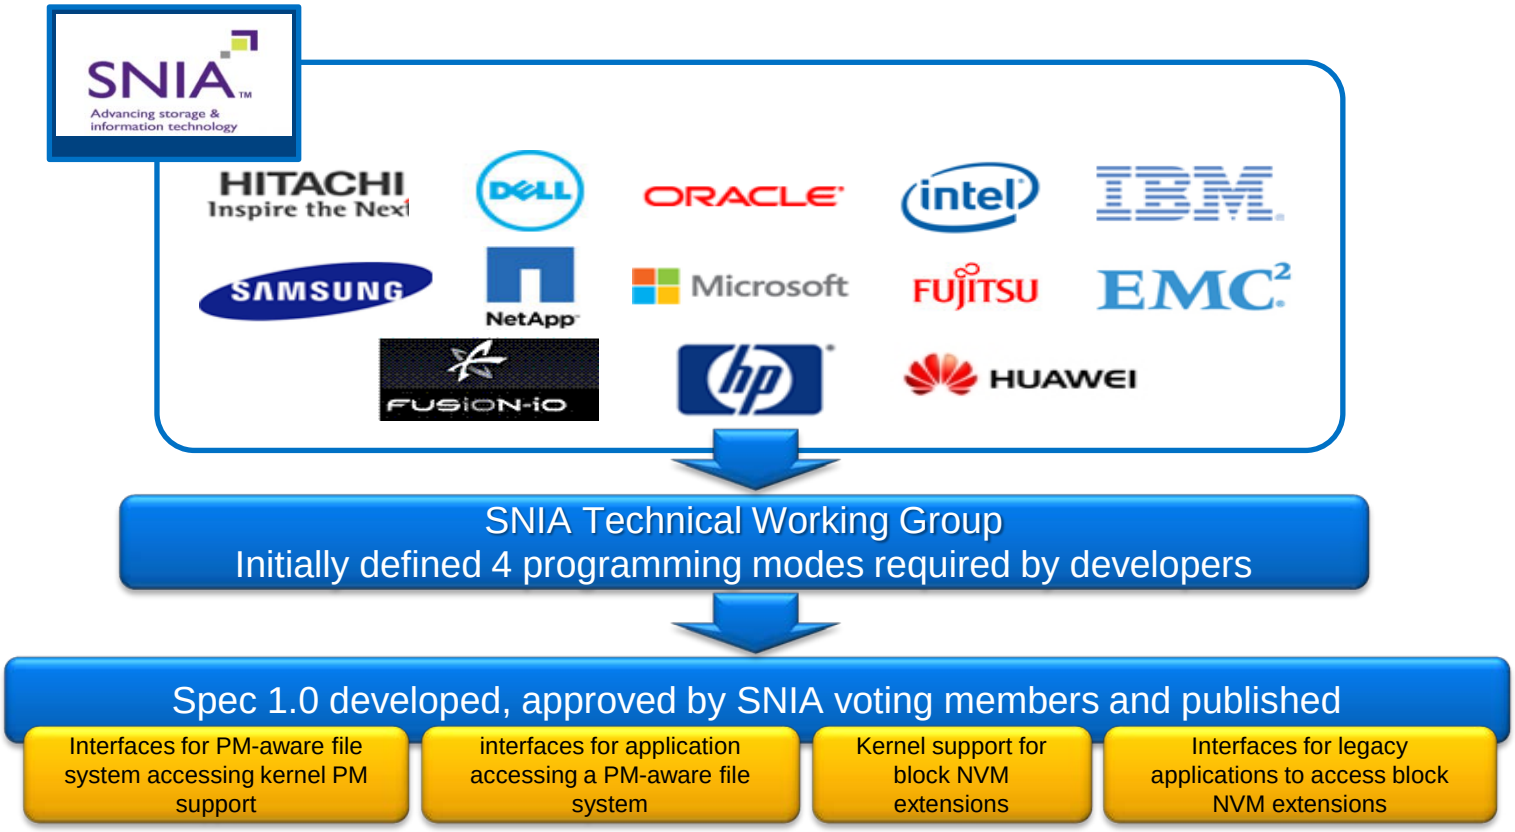

NVM Library: pmem.io

64-bit Linux Initially

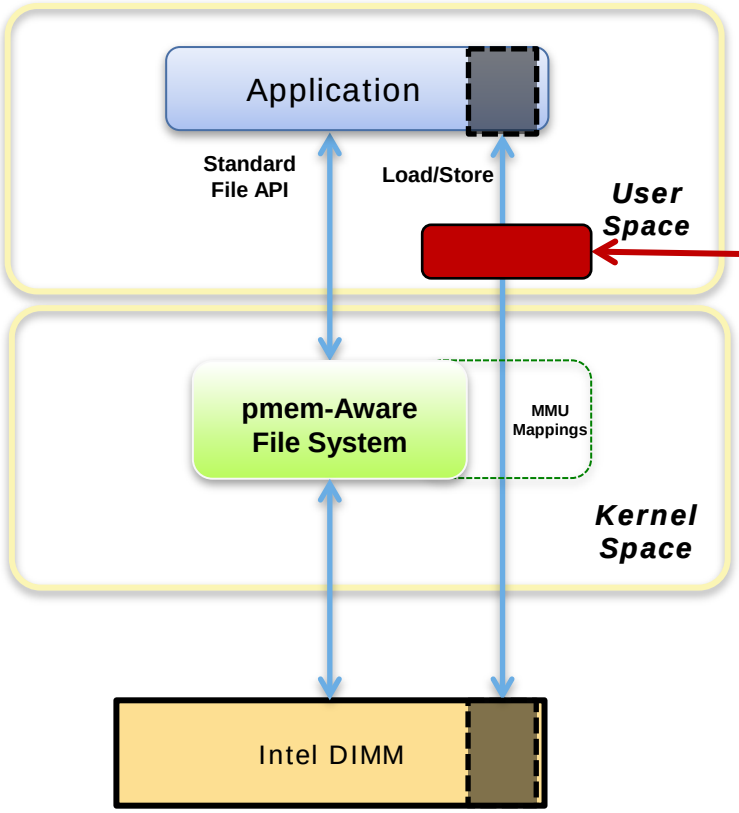

- Open Source
  - http://pmem.io
- libpmem
- libpmemobj
- libpmemblk
- libpmemlog

Transactional

- libvmem
- libvmmalloc

New Instructions

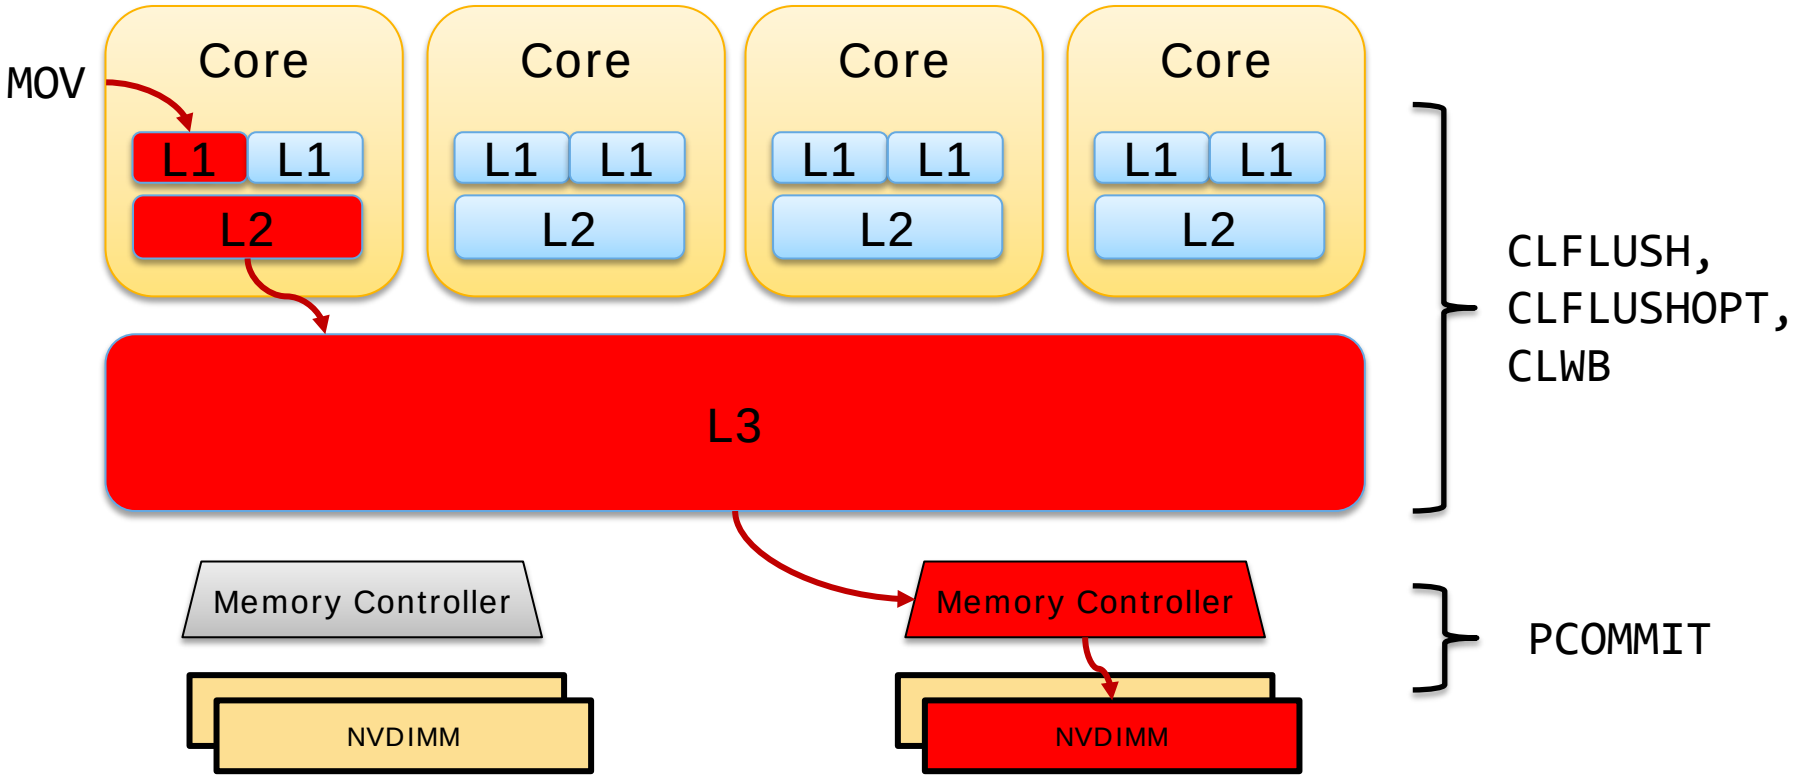

Low Latency Ahead!

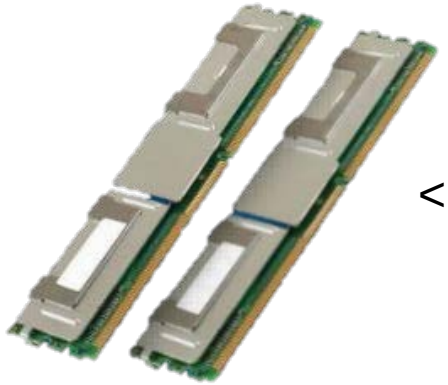

<1 usec

Persistent Memory

3D XPoint™ memory

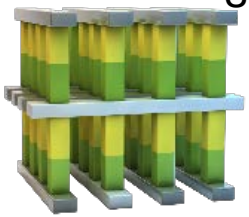

NVMe SSD

Ultra fast SSD

<10 usec

Thank you