OPENFABRICS ALLIANCE

10TH ANNUAL INTERNATIONAL DEVELOPER WORKSHOP

NVM as a Disruptive Technology

#OFADevWorkshop

Non-Volatile Memory Vision

Fast Like Memory

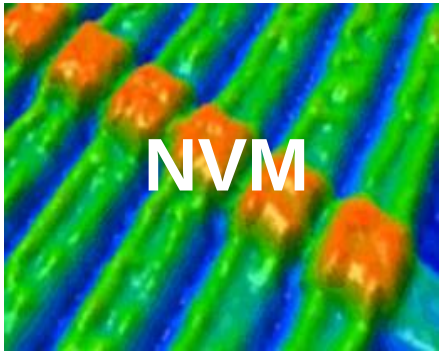

Durable Like Storage

NVM Brings Storage

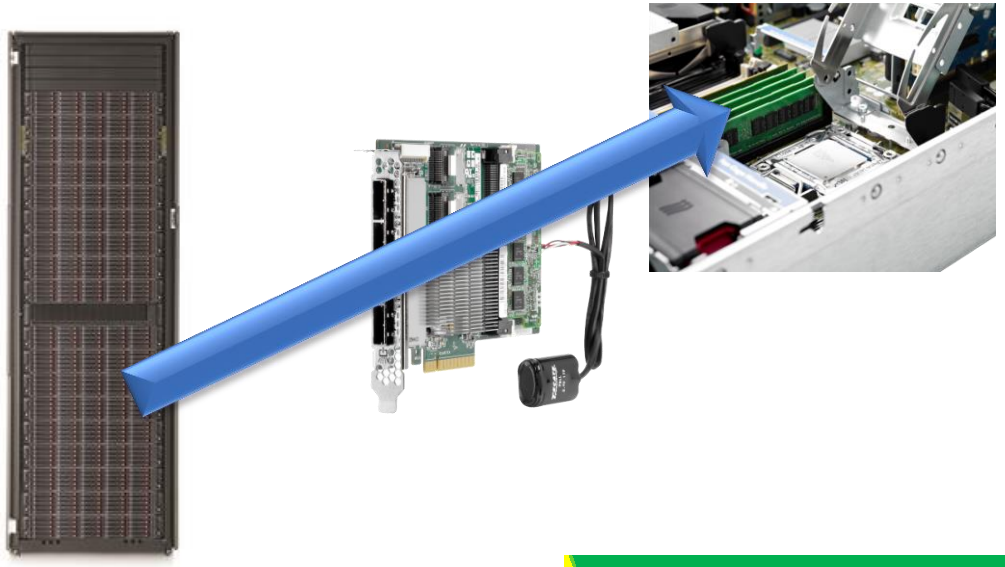

To Memory

Make Data Durable Without Doing IO!

Next Generation Scalable NVM

OPENFABRICS ALLIANCE

Scalable Resistive Memory Element

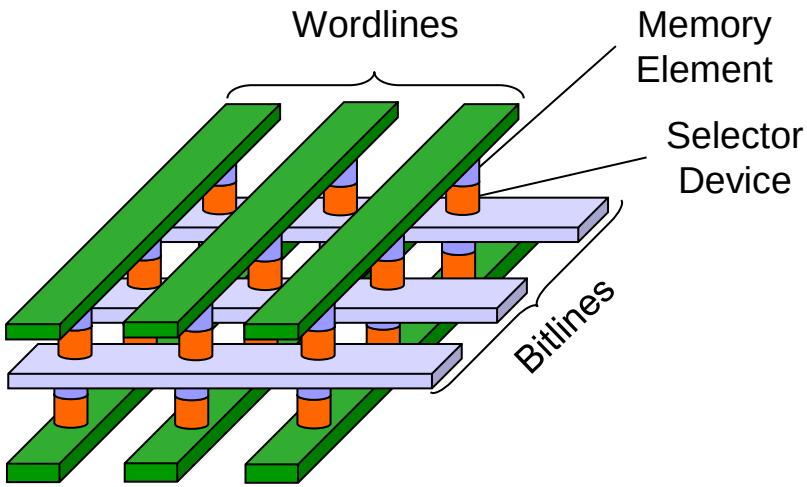

Cross Point Array in Backend Layers

Resistive RAM NVM Options

Family	Defining Switching Characteristics
Phase Change Memory	Energy (heat) converts material between crystalline (conductive) and amorphous (resistive) phases
Magnetic Tunnel Junction (MTJ)	Switching of magnetic resistive layer by spin-polarized electrons
Electrochemical Cells (ECM)	Formation / dissolution of "nano-bridge" by electrochemistry
Binary Oxide Filament Cells	Reversible filament formation by Oxidation-Reduction
Interfacial Switching	Oxygen vacancy drift diffusion induced barrier modulation

~ 1000x speed-up over NAND.

From Jim Pappas, Intel, SNIA NVM Summit

Merging Storage and Memory

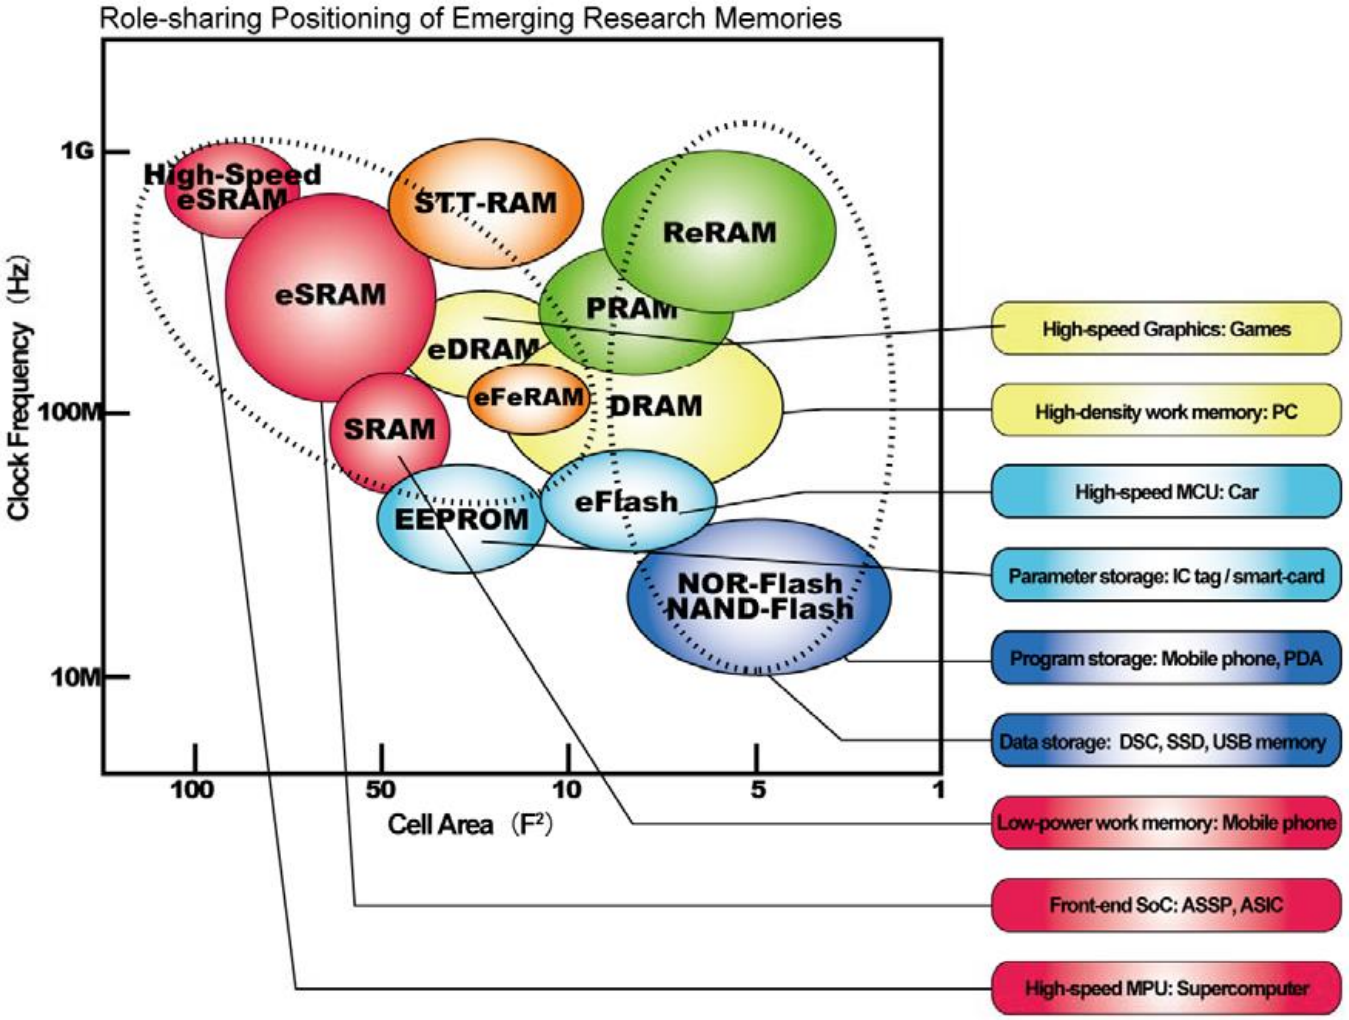

Storage Latency is Changing

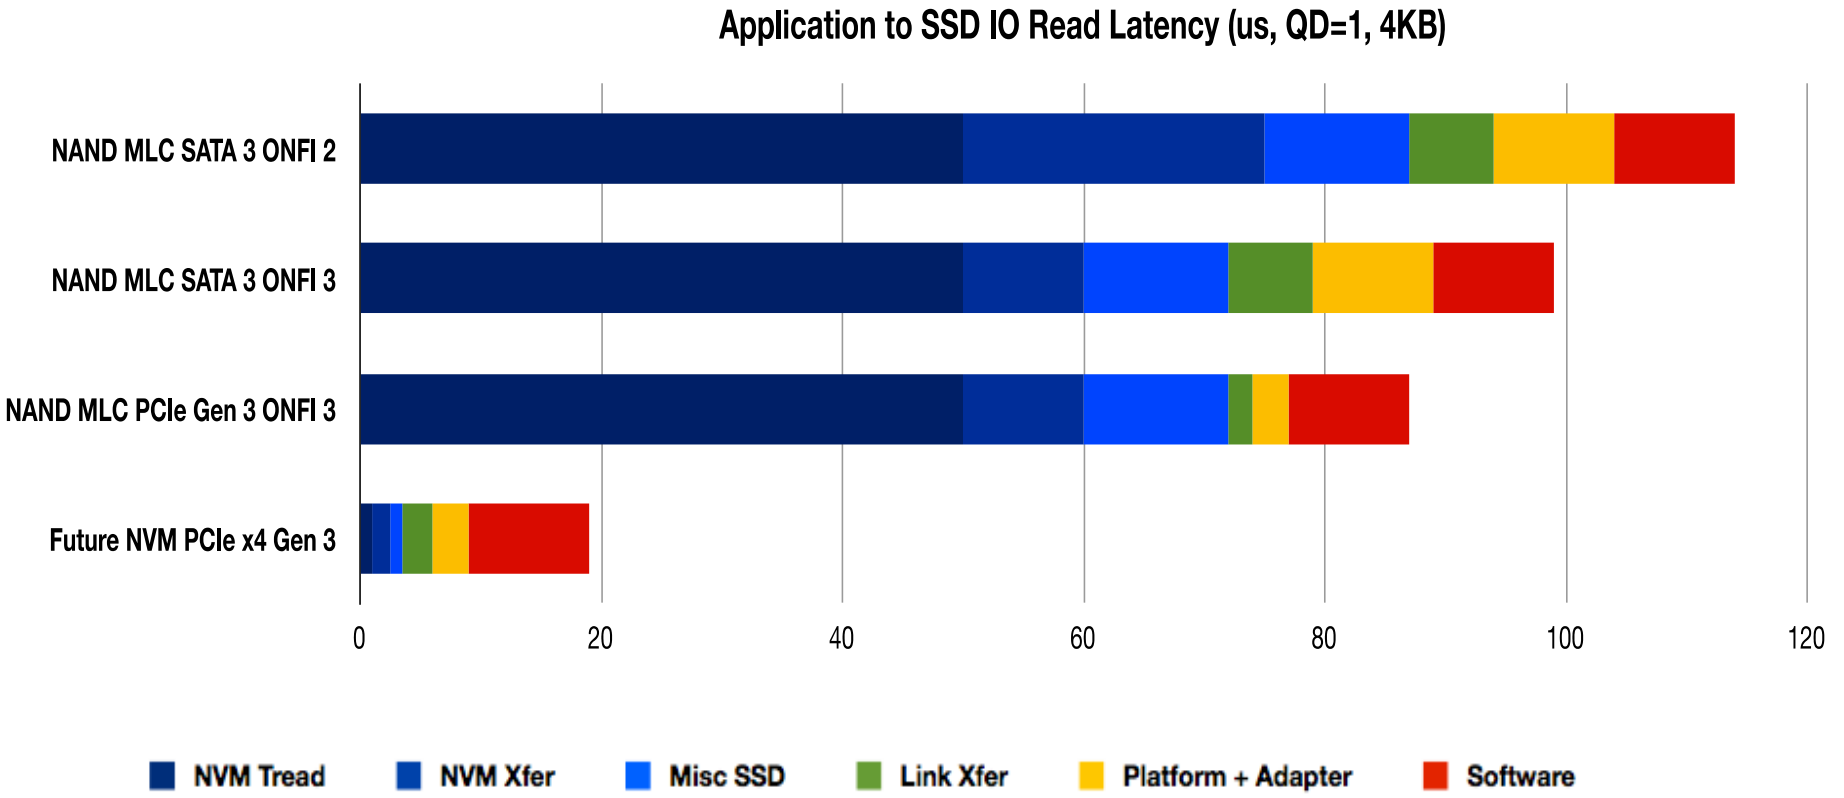

SSD Latency Profile – In the next generation, media is not the bottleneck.

From Jim Pappas, Intel, SNIA NVM Summit

Interconnect and Software Will Soon Dominate Latency

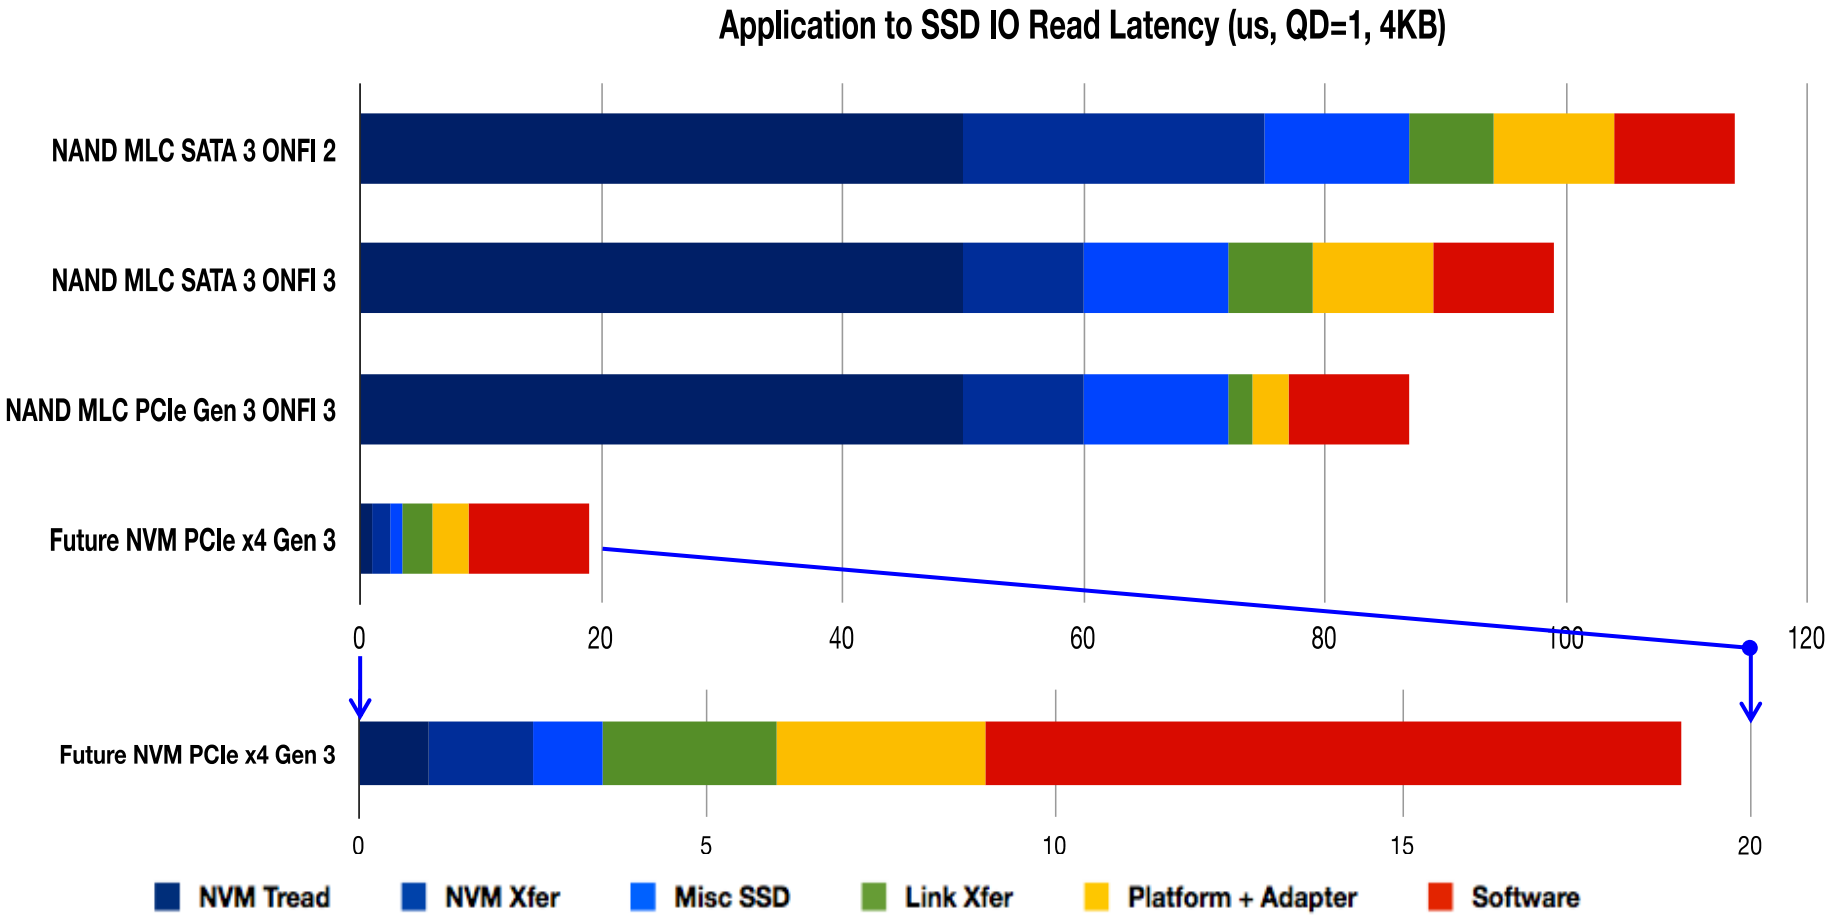

From Jim Pappas, Intel, SNIA NVM Summit

Opportunities with Next Generation NVM

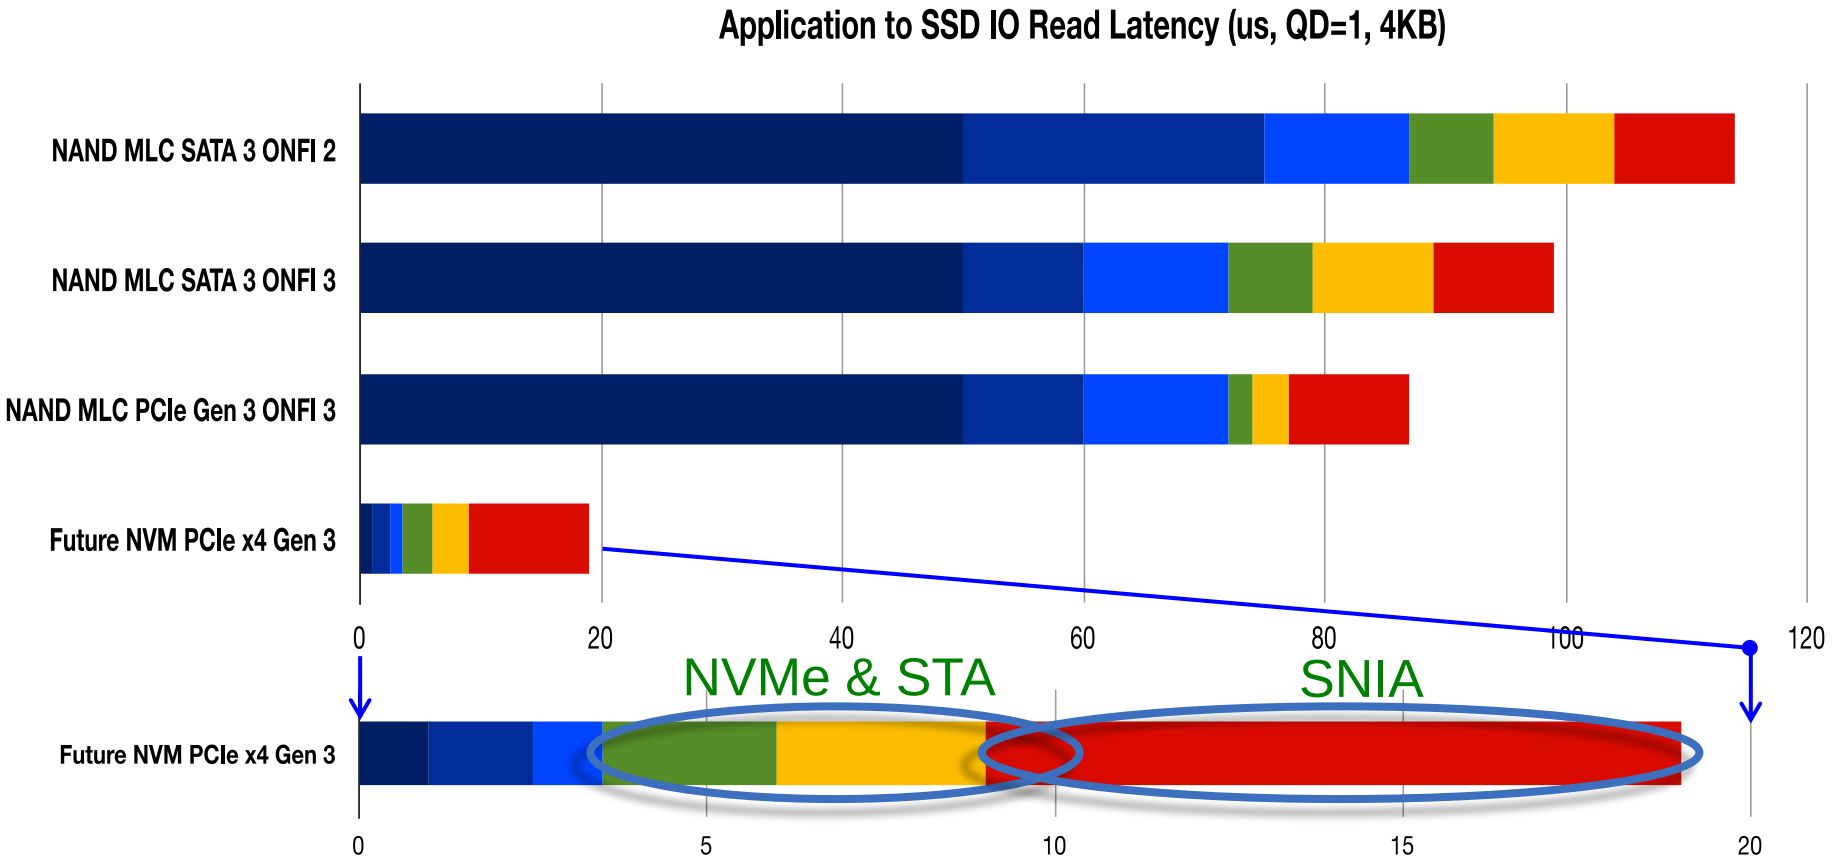

NVM Express/SCSI Express: Optimized storage interconnect & driver
SNIA NVM Programming TWG: Optimized system & application software

From Jim Pappas, Intel, SNIA NVM Summit

Application View of IO Elimination

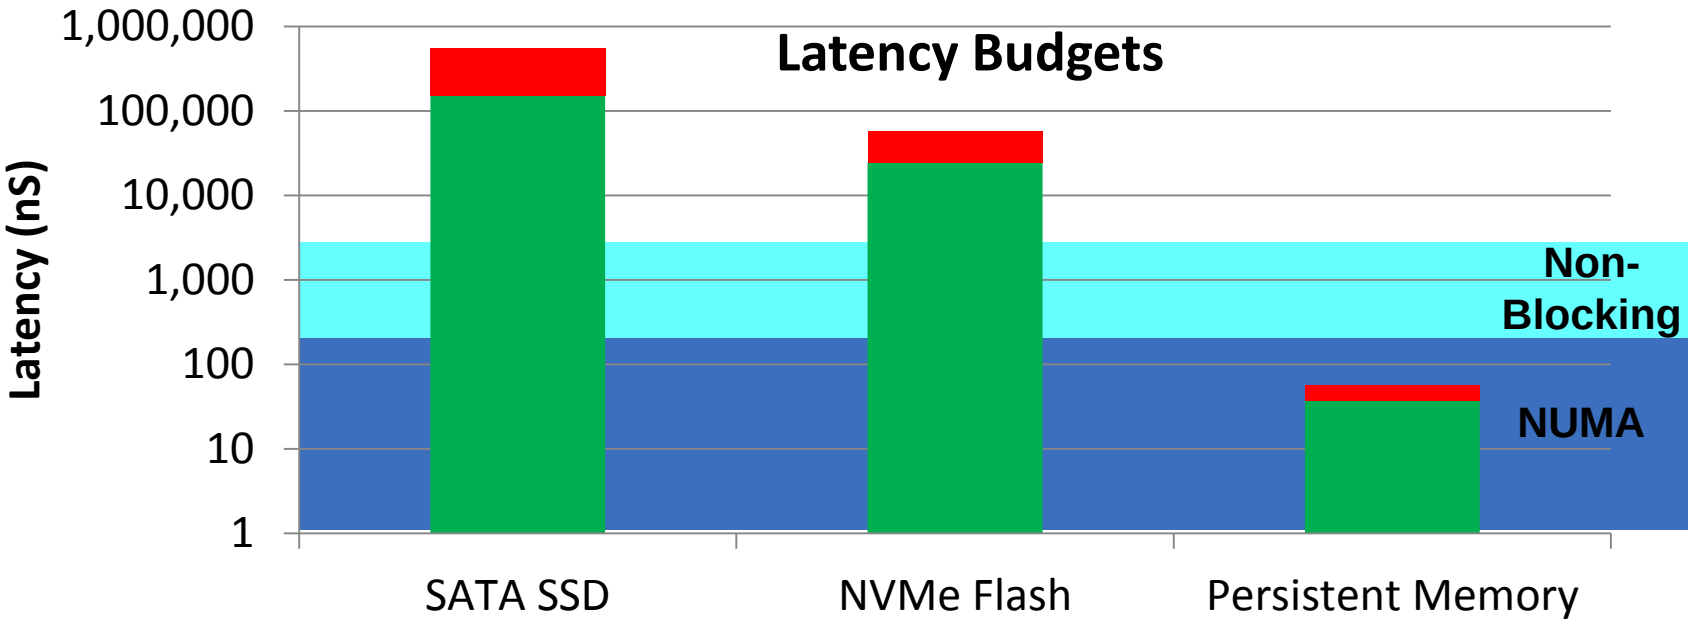

Software overheads are being driven to keep pace with devices
NUMA latencies up to 200 nS have historically been tolerated
Anything above 2-3 uS will probably need to context switch
Latencies below these thresholds cause disruption

Memory Disaggregation

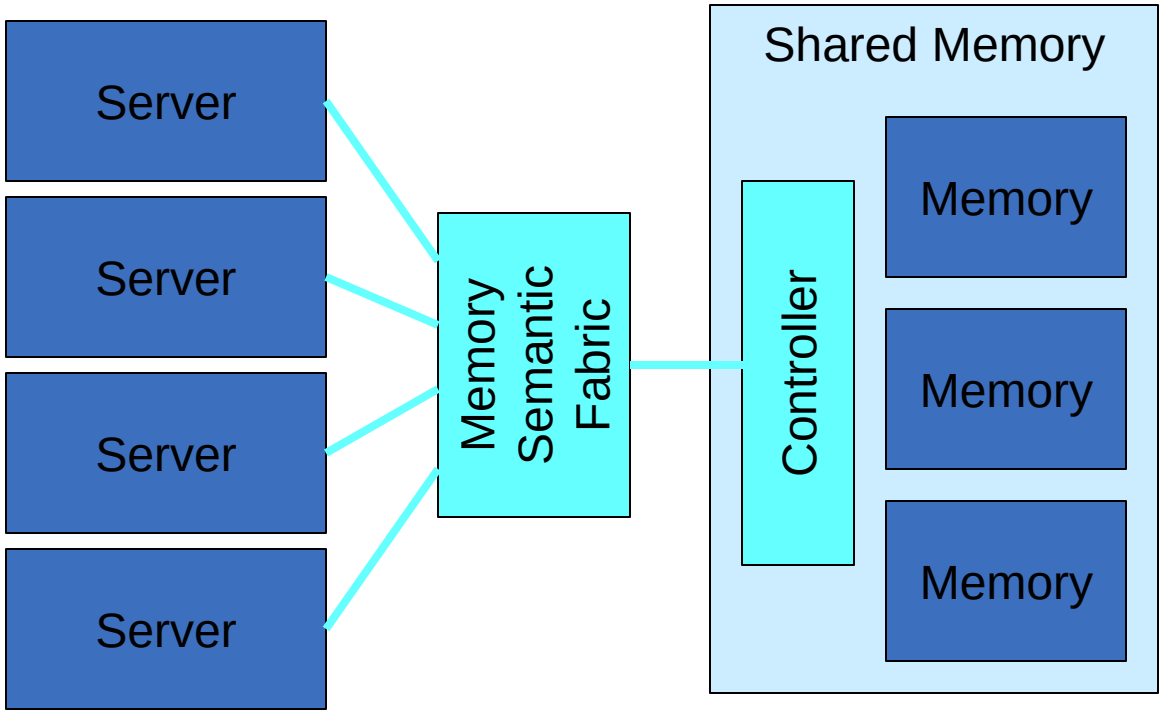

Increased flexibility to provision memory from a shared pool
Pass data or VM's between servers without moving bits
Enabled at rack scale using silicon photonics

Application View of Memory Properties

OPENFABRICS ALLIANCE

Durability

- Applications should know which memory is persistent.

NUMA

- What part of the system manages proximity of processing to memory?
- How much unpredictability is tolerable?

Multi-processing

- SMP boundaries implied by distance and fault domains

Abstractions need to account for…

OPENFABRICS ALLIANCE

Resource Diversity

- Choosing the right memory resources in real time

System Scalability

- Programming model accounts for wide scale

Transactions

- Accelerate recoverability to memory speed

Thriving in Chaos

- Make sure hierarchy levels earn their places
- Think dual stack
- Make middleware take up the slack
- Mind the fault lines

Thank You

OPENFABRICS ALLIANCE

10TH ANNUAL INTERNATIONAL DEVELOPER WORKSHOP