G2M RESEARCH

Responsive and Efficient Storage Architectures for CSPs

PLIOPS

NVIDIA

Excelero

Multi-Vendor Webinar

Tuesday May 18, 2021

Webinar Agenda

9:00-9:05	Ground Rules and Webinar Topic Introduction (G2M Research)
9:06-9:37	Sponsoring Vendor presentations on topic (8 minute each)
9:38-9:44	Panel Discussion Question #1
9:45-9:46	Audience Survey #1
9:46-9:52	Panel Discussion Question #2
9:53-9:54	Audience Survey #2
9:54-10:00	Panel Discussion Question #3
10:01-10:08	Audience Q&A (8 minutes)
10:09-10:10	Wrap-Up

G2M Research Introduction and Ground Rules

Mike Heumann
(Managing Partner, G2M Research)

Storage for Cloud Service Providers (CSPs)

- Hyperscalers/CSPs have a variety of challenges when fielding storage for their customers/applications
  - Hyperscalers store roughly half of the world's 1,000 exabytes of digital data
  - A peak bandwidth of >1TB/sec for a CSP is not uncommon
  - Millions of concurrent users (multi-tenant ops on steroids!)
  - All of this is spread over tens/hundreds of datacenters
- Key considerations for CSPs/hyperscalers
  - Ability to rapidly scale # users, bandwidth, and capacity
  - Minimizing CPU utilization
  - Minimizing power footprint
  - Providing predictable end-user performance

Panelists

Excelero

Kirill Shoikhet
Chief Technology Officer
www.excelero.com

PLIOPS

Steve Fingerhut
President
www.pliops.com

NVIDIA

Rob Davis
Vice President,
Storage Technologies
www.nvidia.com

G2M RESEARCH

Mike Heumann
Principal Analyst
www.g2minc.com

Kirill Shoikhet
Chief Technology Officer
www.excelero.com

Characteristics of CSP Operation

Agility: Supporting the widest range of applications

Efficiency: Viable operation requires the most efficient use of the underlying resources

Scale: Success means growth. Growth means scale. Scale exposes and punishes any design flaws

Resiliency: Unavailability costs money and reputation. Failures happen. Embracing failures is hard

What Does it Mean for Storage?

Agility: Don't assume anything about workloads.
Some require low latency; some - bandwidth
Flexibility is the key

Efficiency: Storage needs to expose full capabilities
of the underlying HW: drives and networks

Scale: Scale means storage networking.
Storage is hard but storage networking is harder

Resiliency: No part can have a single point of failure –
neither data plane nor control plane.
You'll be surprised how many products miss
the later

Appliances

Software-Defined Storage

NVMesh Performance Density	PCI EXPRESS 3	40 GB/s /RU 9.5M IO/s /RU
	PCI EXPRESS 4	80 GB/s /RU 16M IO/s /RU

New technology adoption:
- *years for appliances*
- *weeks for SDS*

How Would One Envision CSP Storage?

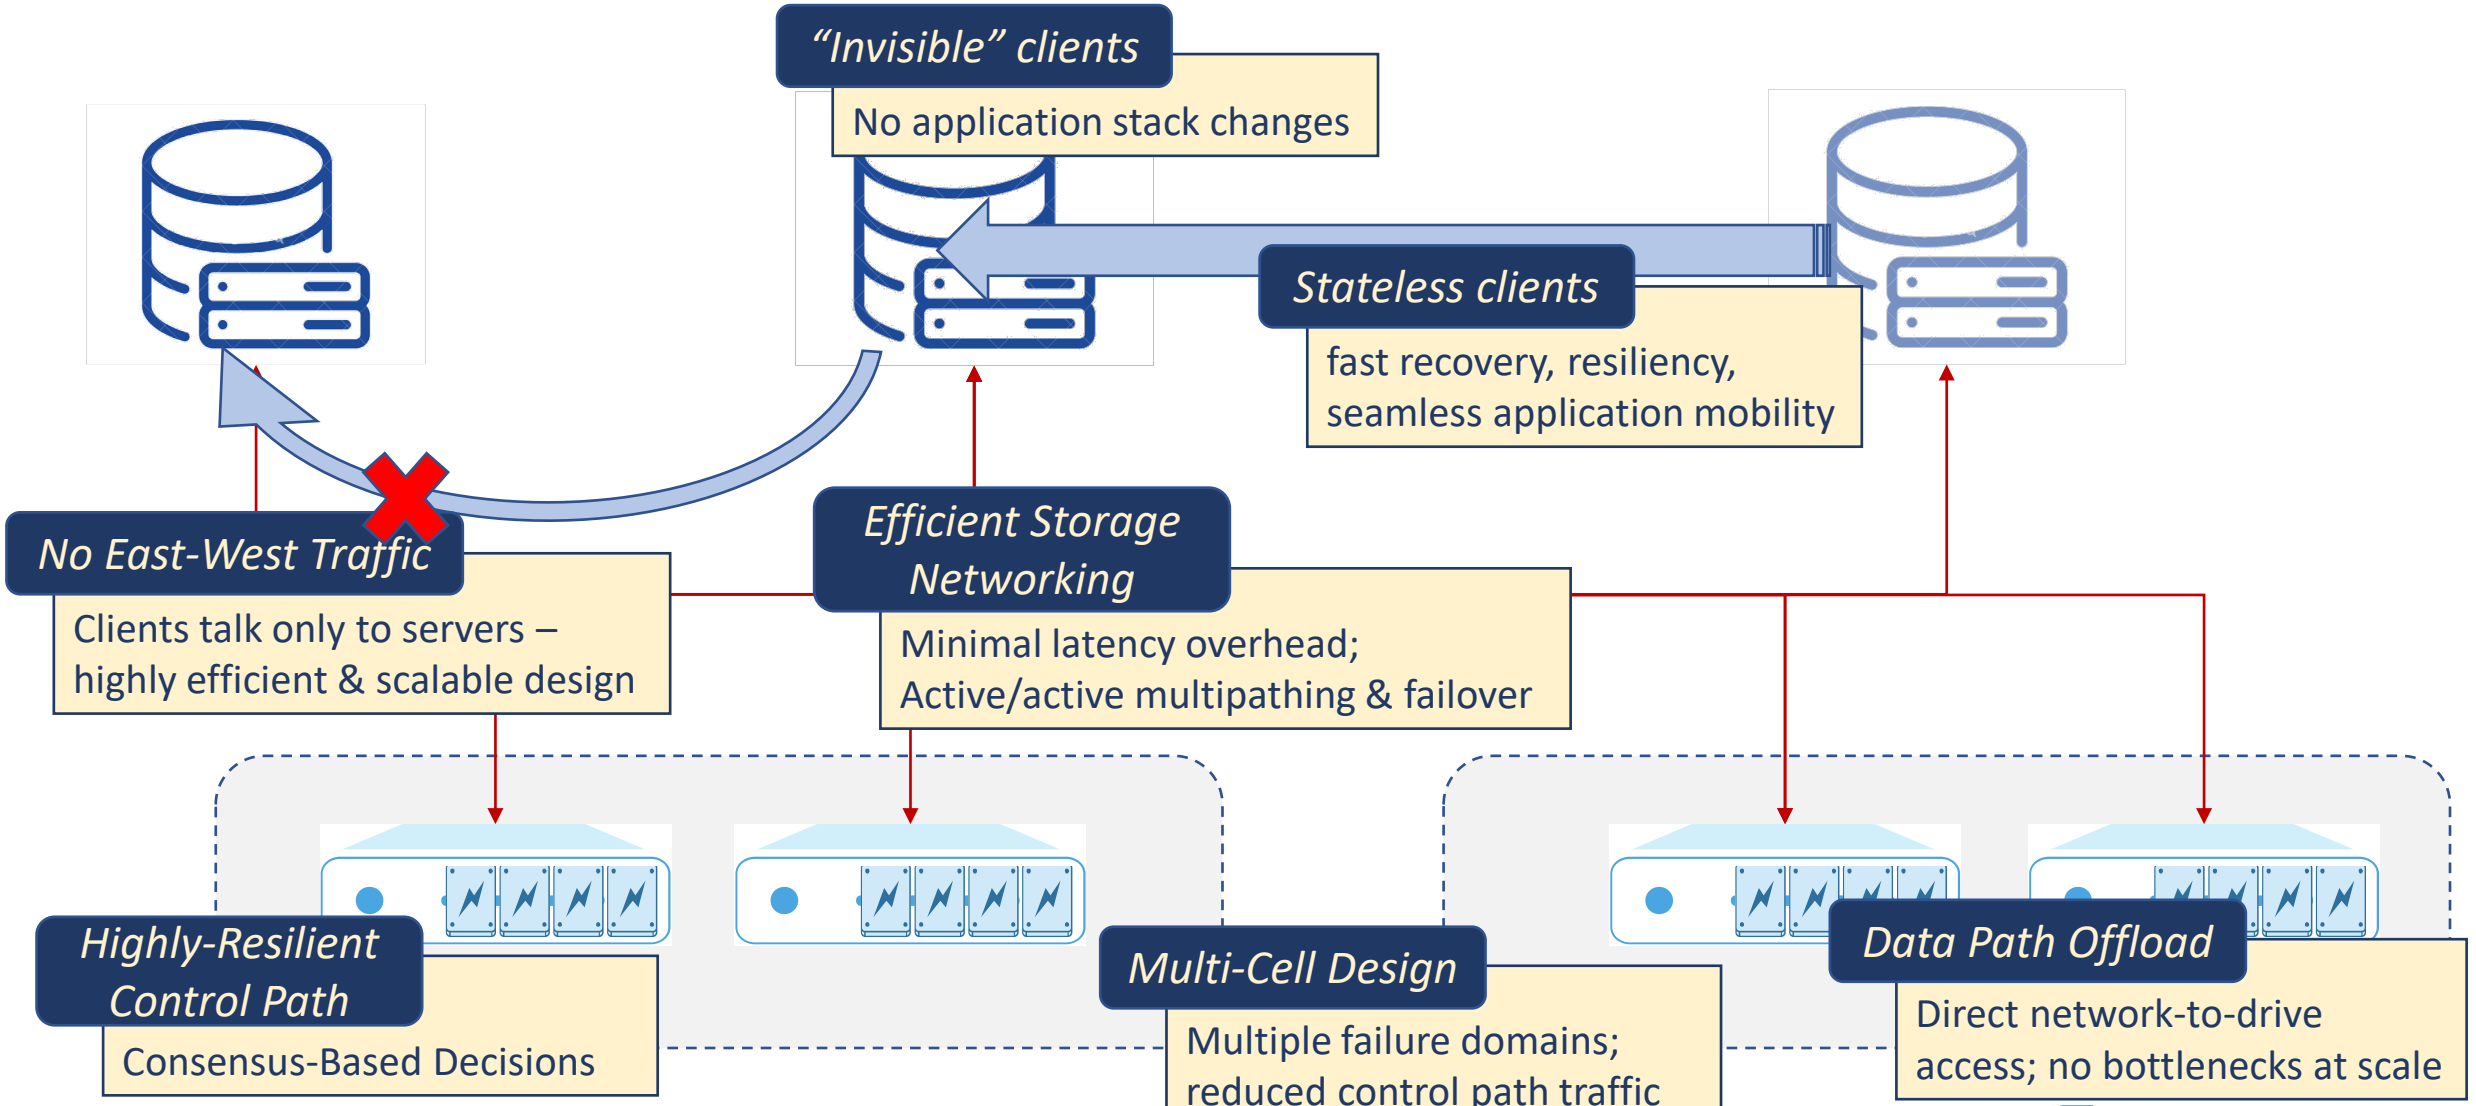

Excelero NVMesh as CSP Storage

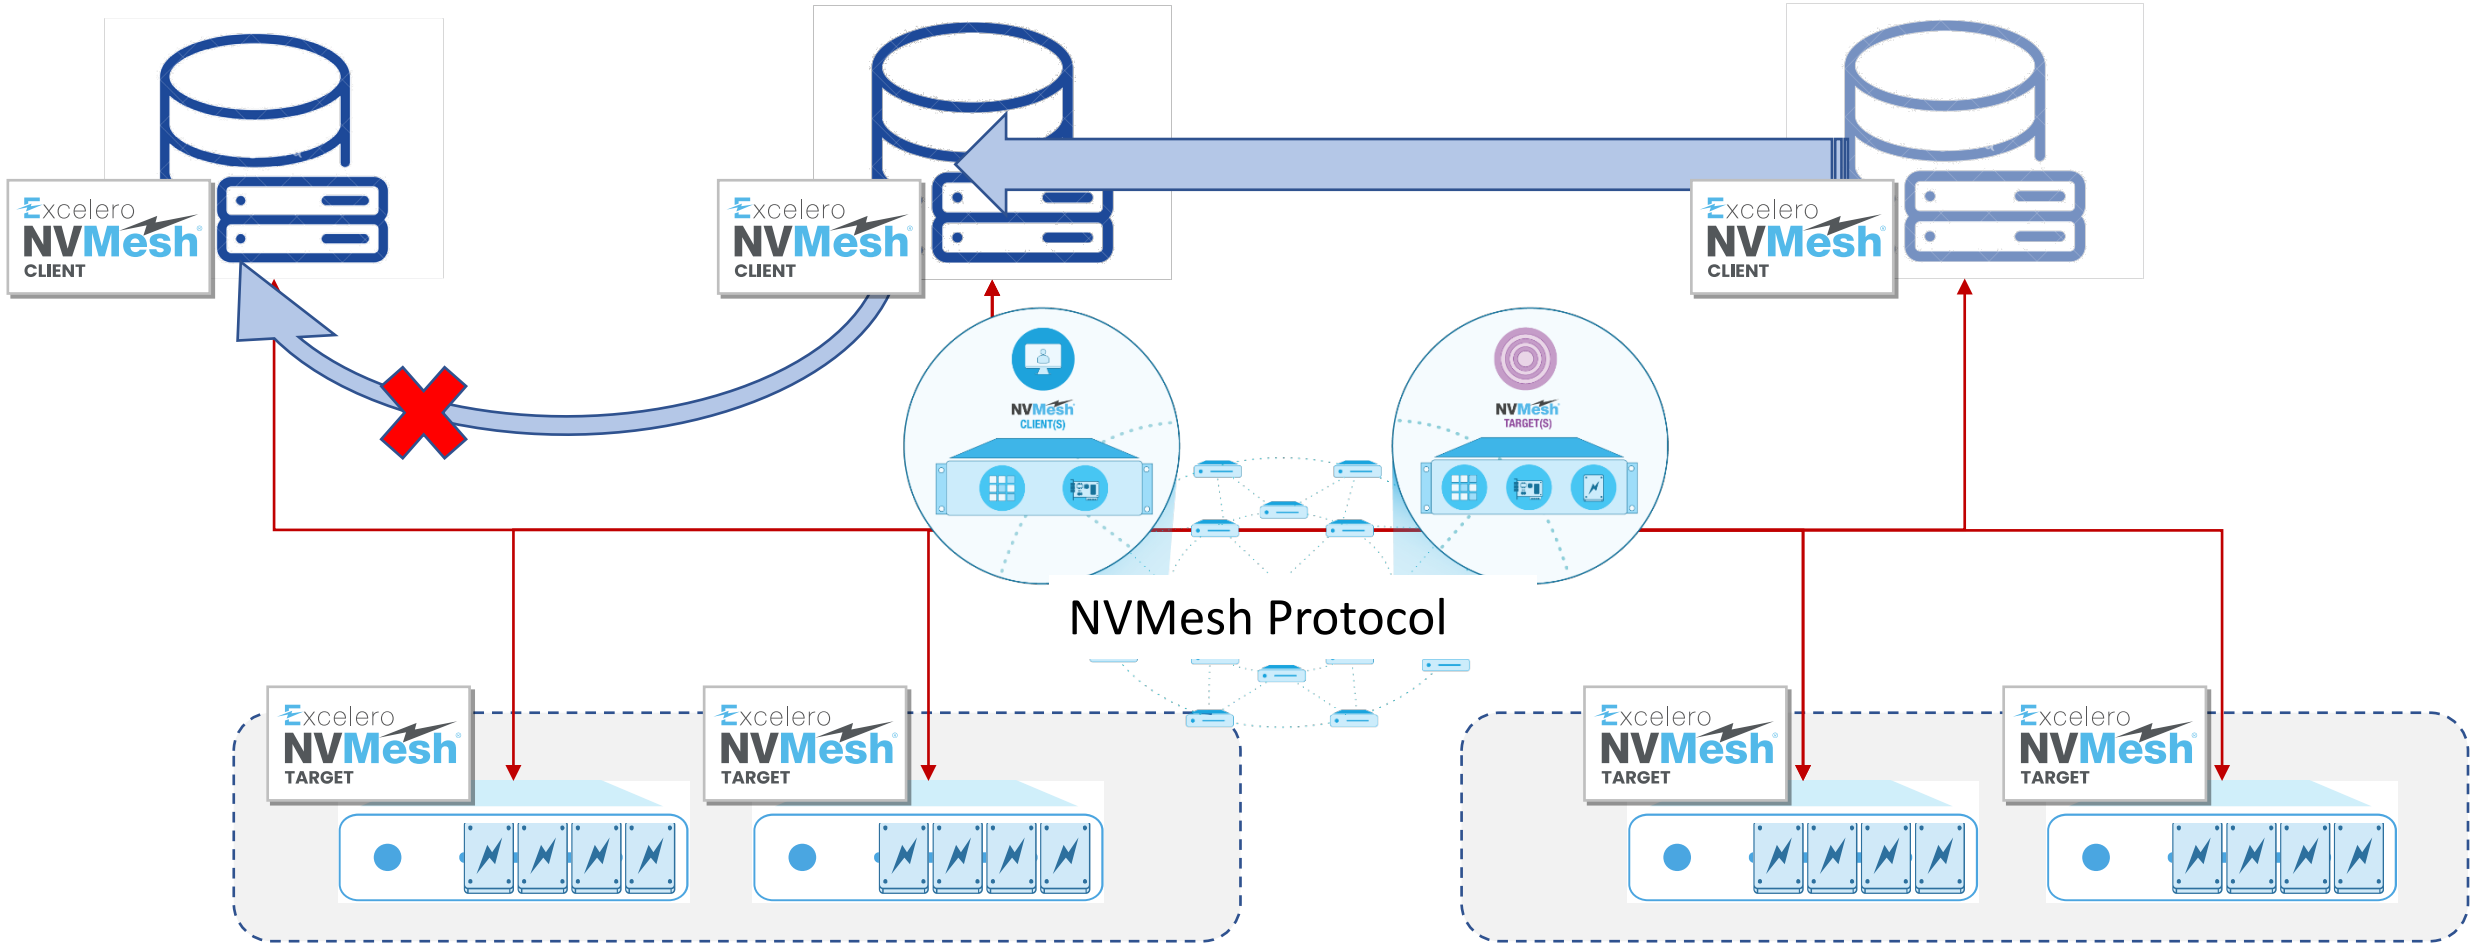

Challenges and Future Directions

- Client-side architectures bring huge benefits to CSPs
- But running storage client on application servers challenges disaggregation

  → Data Processing Units (DPUs) to the rescue!

- Datapath offload reduces CPU dependency
- But memory coherency and bandwidth are still issue

  → Compute-eXpress Link (CXL) protocol!

- Compute-intensive data operations (e.g., compression) require CPU cycles

  → Client-side protocols allows integrating remote HW accelerators with NVMe interface

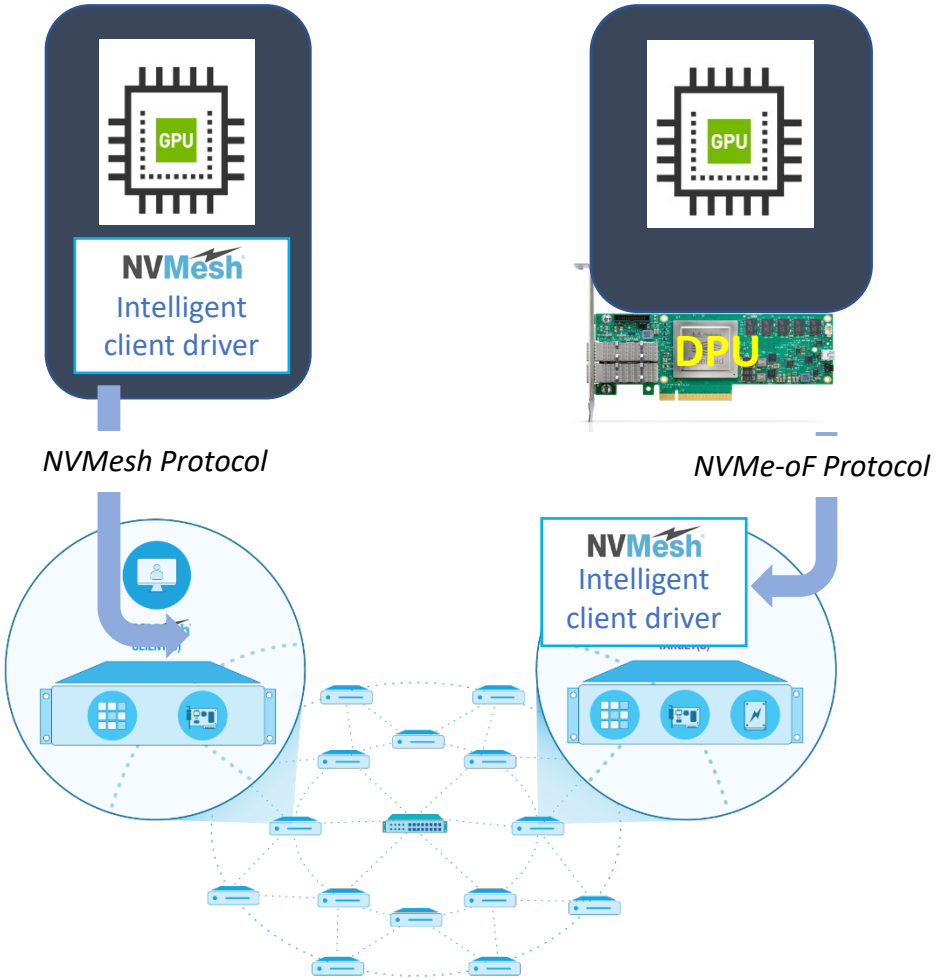

Responsive and Efficient Storage Architectures for Cloud Service Providers

Steve Fingerhut
President & CBO

Challenges with Broad SSD Adoption

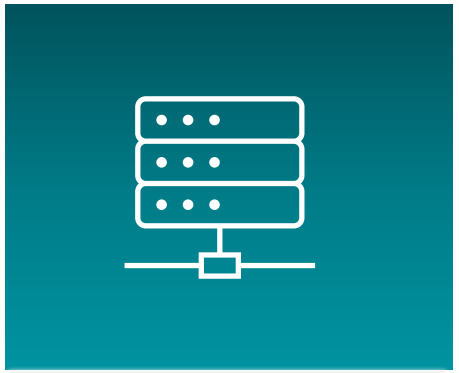

Server Architectures Not Balanced

SSDs' 1000x increase in performance over HDD has not been matched by server advances

Amplified Data

Software that uses SSDs amplifies reads and writes up to 100x, stored data up to 6x, crushing storage and network efficiency

System Reliability Compromised

Traditional RAID is rarely used with NVMe SSDs due to the huge performance penalty, requires costly workarounds

PLIOPS

Pliops Storage Processor

Reliability	Drive Fail Protection 2x > RAID 0
Performance	2-15x
Capacity	Up to 6x
Efficiency	TLC, QLC for any workload

PLIOPS

Accelerates Diverse and Dynamic Workloads

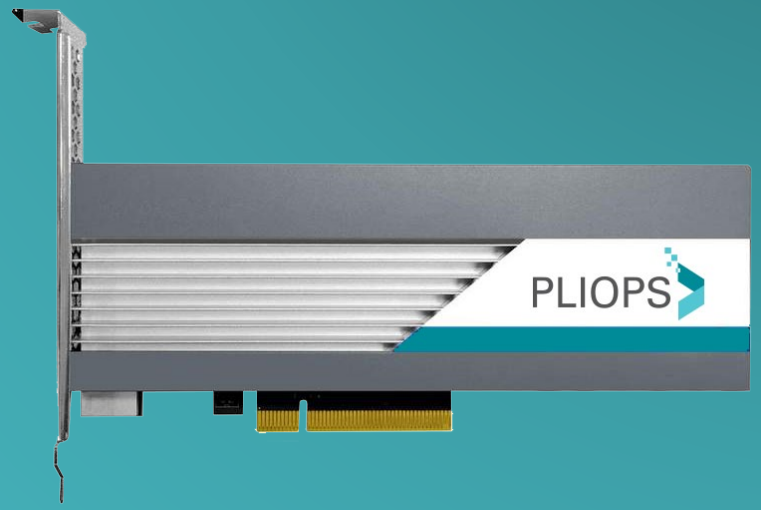

RDBMS

NoSQL

Analytics

Software Defined Storage / Distributed File Systems

PSP accelerates nearly all flash-based workloads

ORACLE
MySQL
MariaDB
Microsoft SQL Server
PostgreSQL
SAP HANA
hadoop
cassandra
elastic
redis
APACHE Spark
mongoDB
GLUSTER
ceph

Yesterday's
COMPUTE CENTRIC

Today's
DATA CENTRIC

Future's
DATA CENTRIC

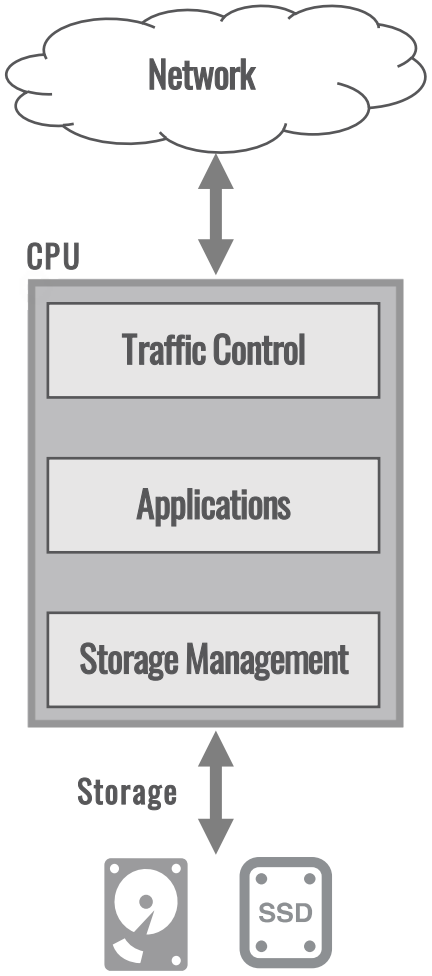

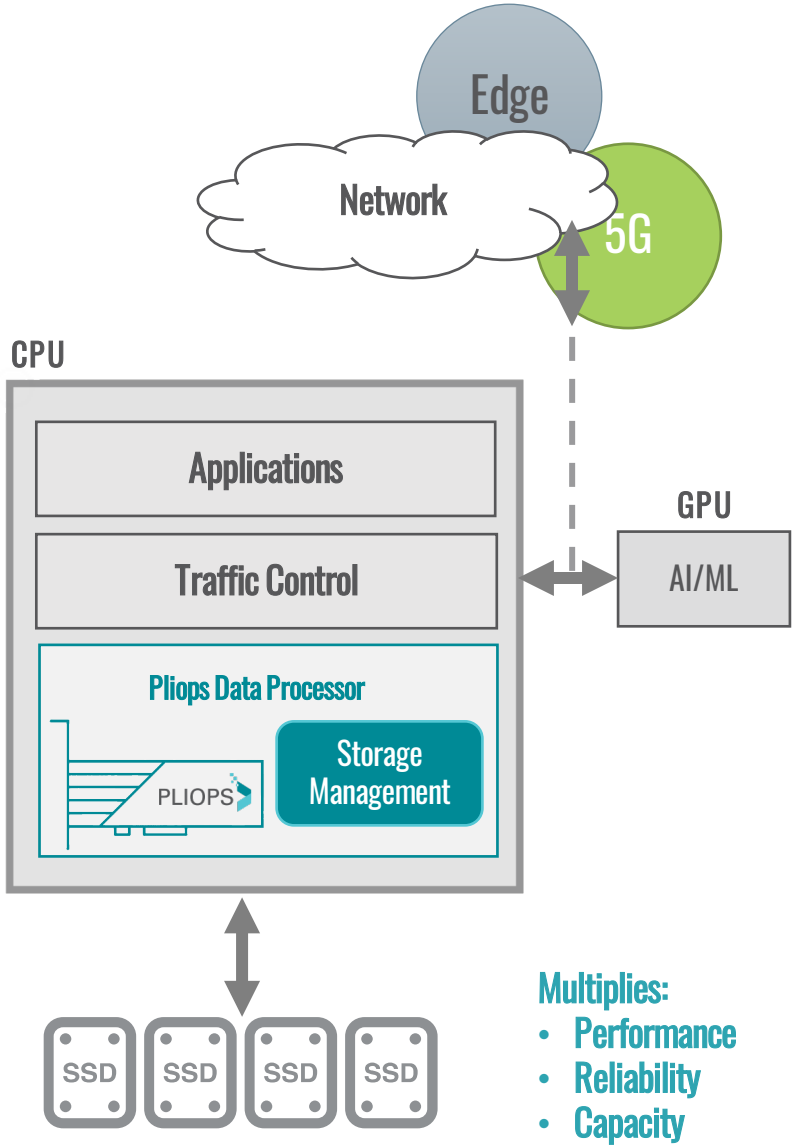

Multiplies:
- Performance
- Reliability
- Capacity
- Efficiency

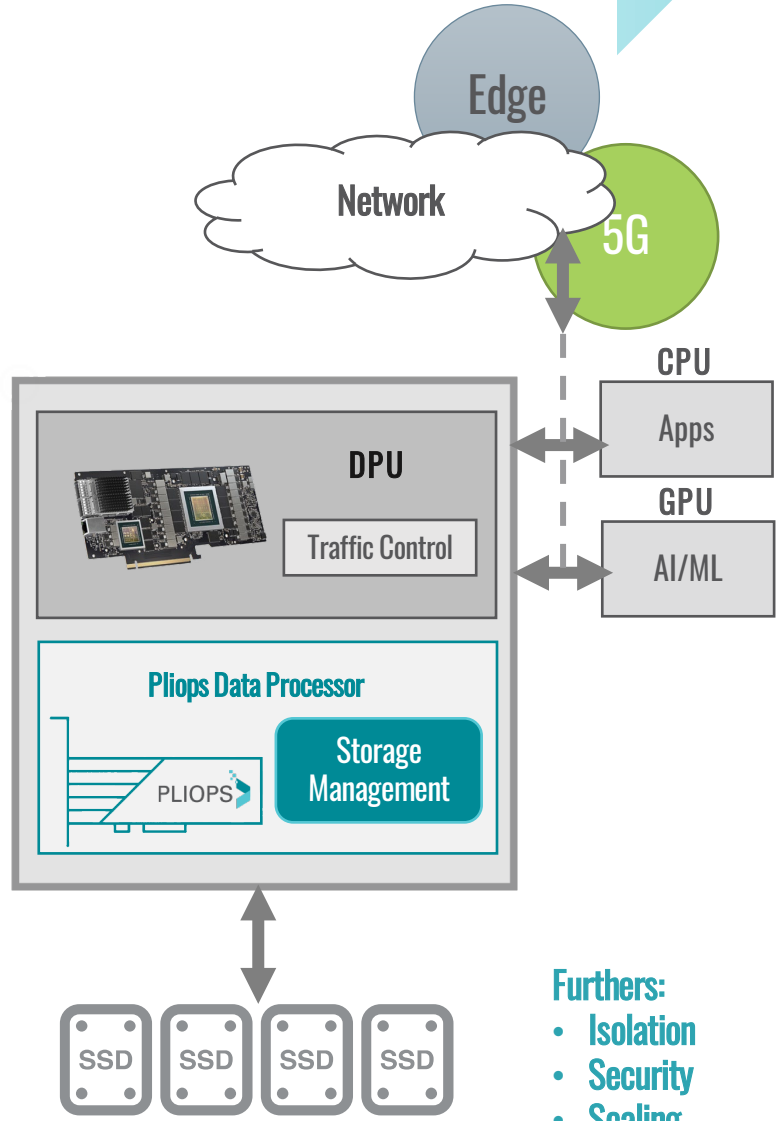

Furthers:
- Isolation
- Security
- Scaling

PLIOPS

Breakthrough Value Across Workloads

Increased Application Scaling

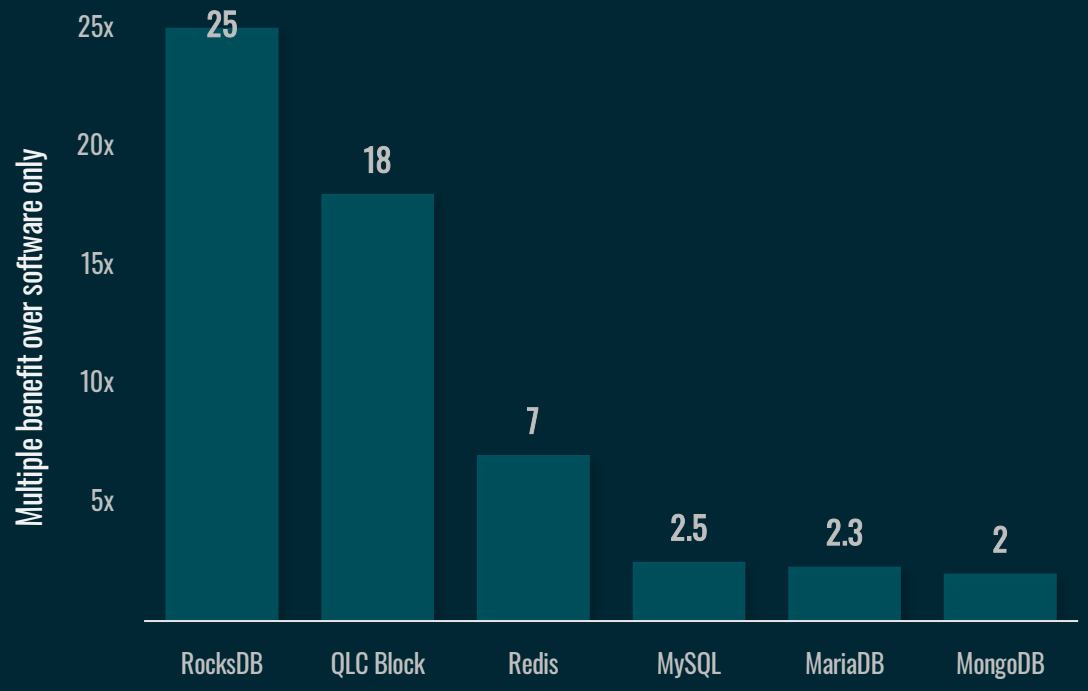

Significant TCO Savings

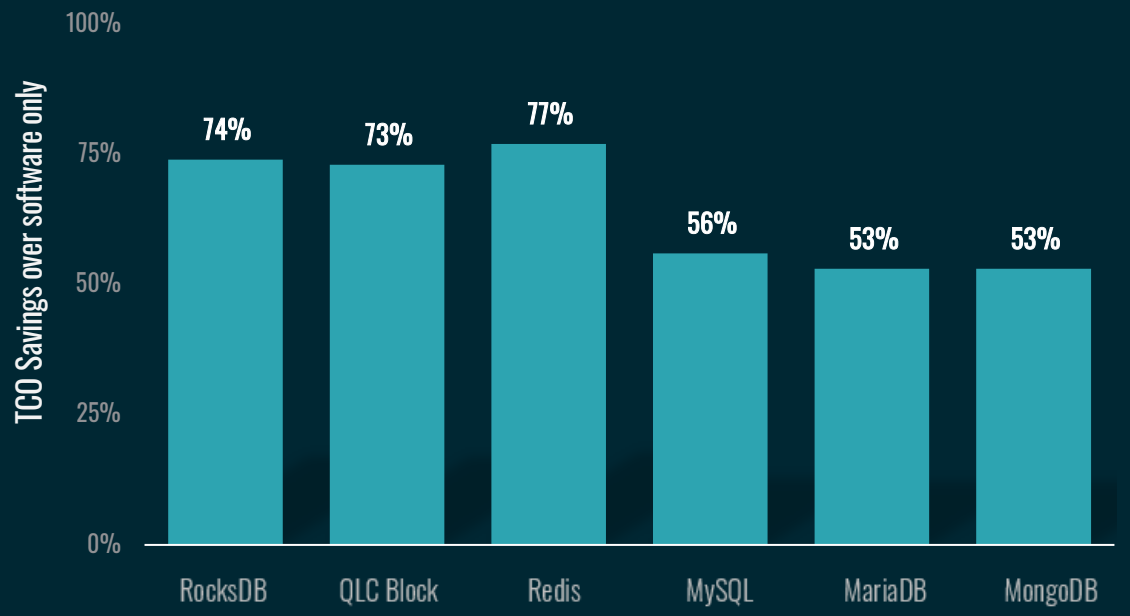

PLIOPS

Large eTailer: ROI with Pliops

Current Software Based Solution

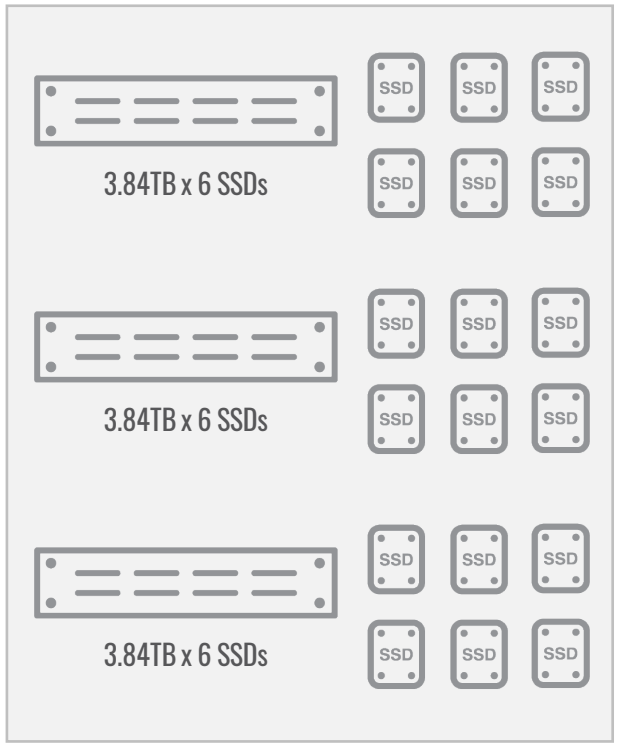

150K Queries/Sec
31 TB Usable
RAID 10

Pliops Accelerated Solution

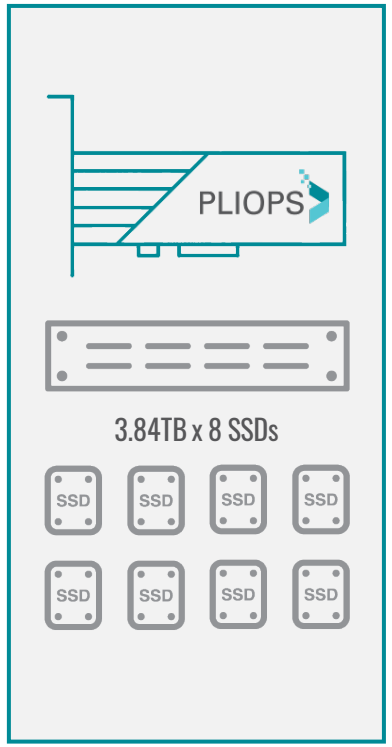

157K Queries/Sec
49 TB Usable
Pliops Drive Fail Protection

Capex Benefit

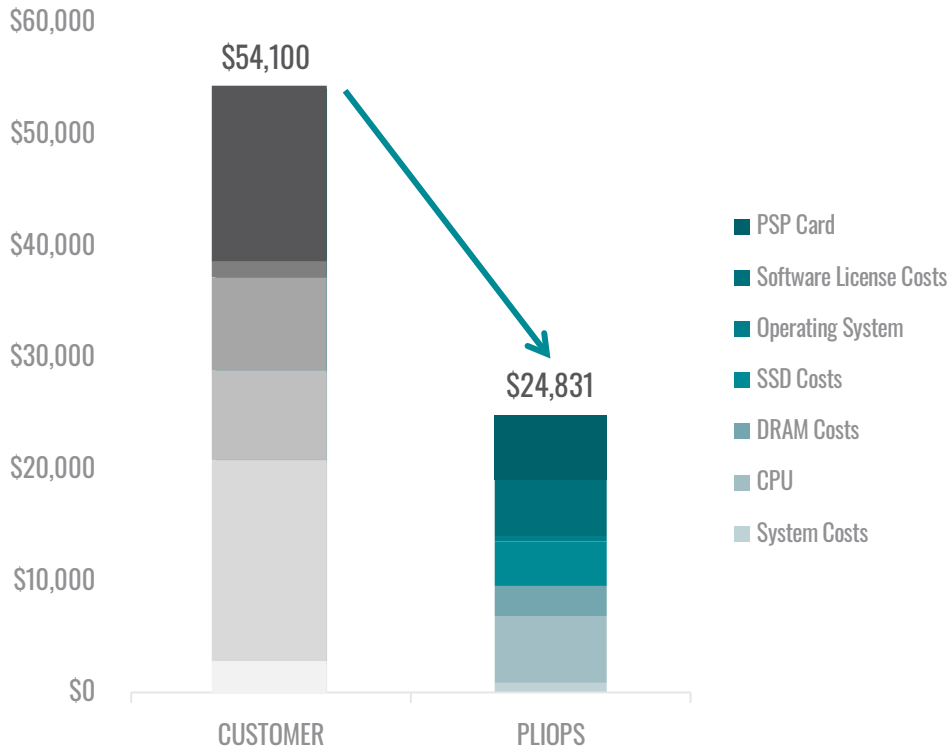

54% lower cost with
58% more usable capacity

PLIOPS

Key Takeaways

1. Data and SSD growth demand a new approach
2. Pliops Storage Processor multiplies TCO
3. Learn more at www.pliops.com

PLIOPS

NVIDIA

Rob Davis

Vice President,
Storage Technology

www.nvidia.com

SSDs Have Changed Networked Storage

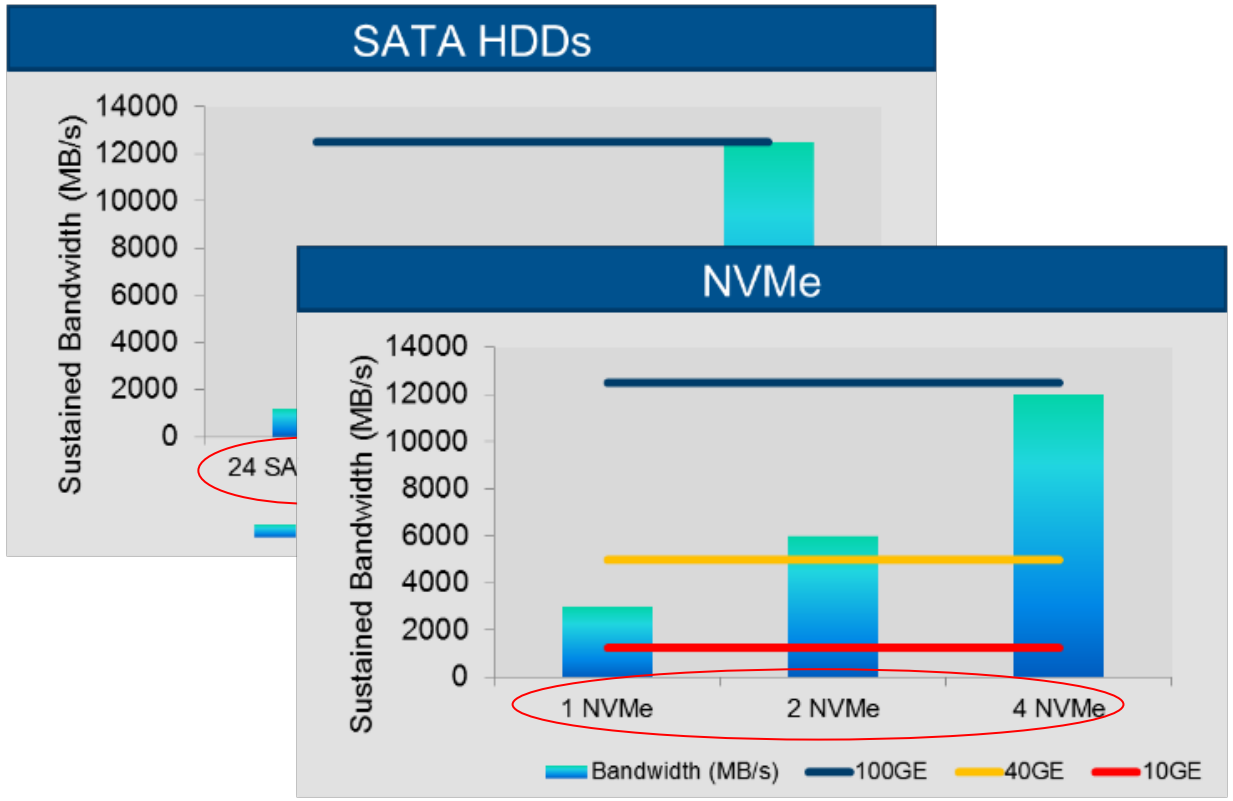

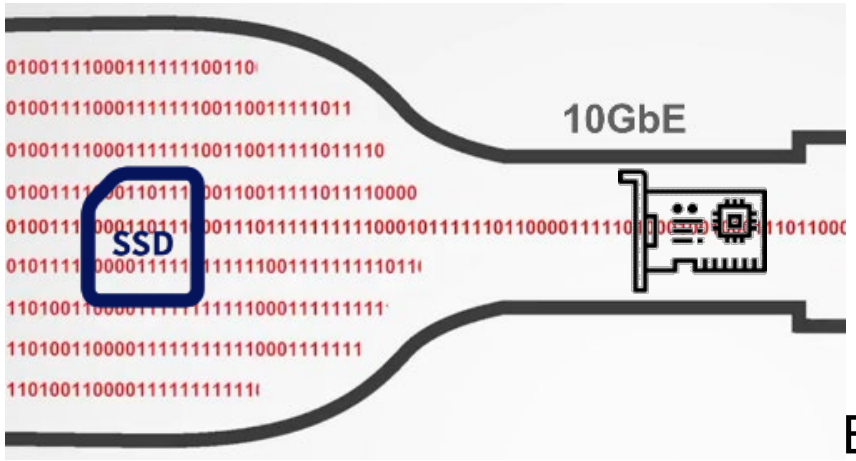

SSDs move the Bottleneck from the Disk to the Network

Faster Network Wire Speeds is Part of the Answer

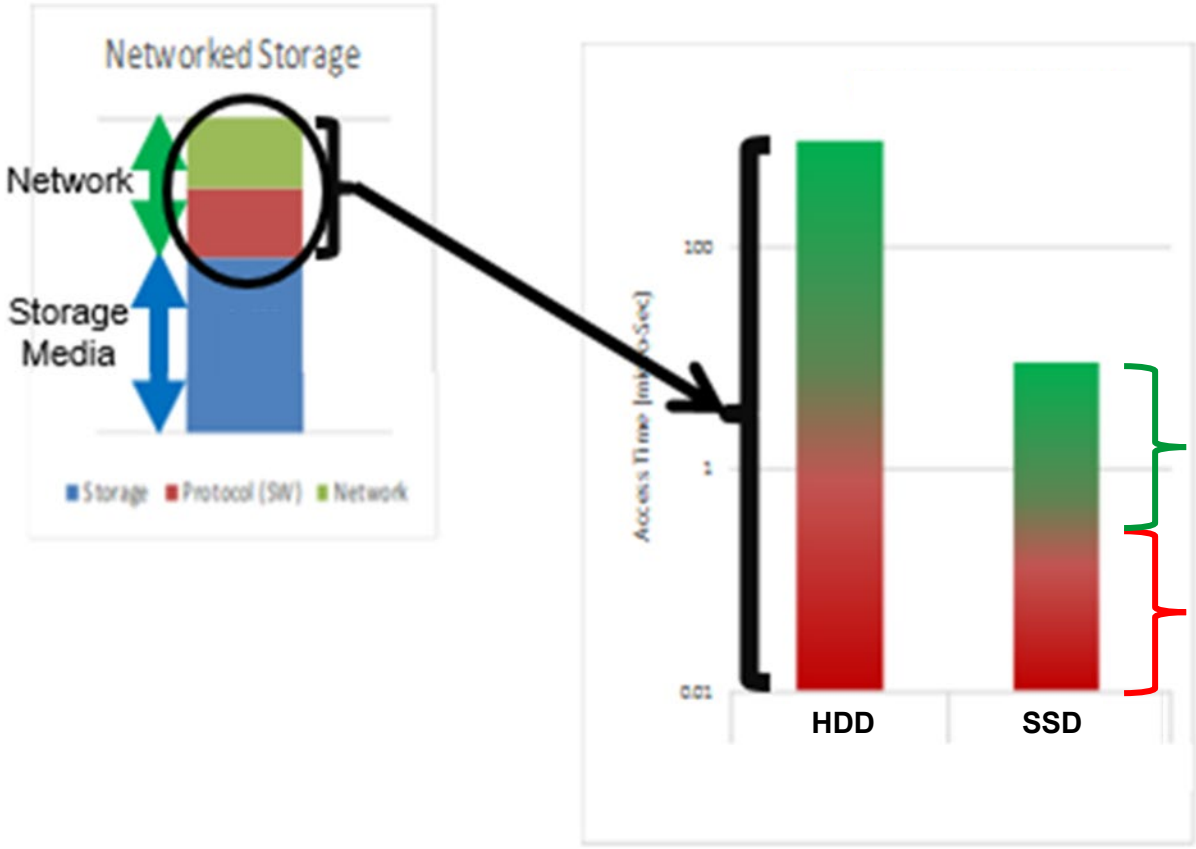

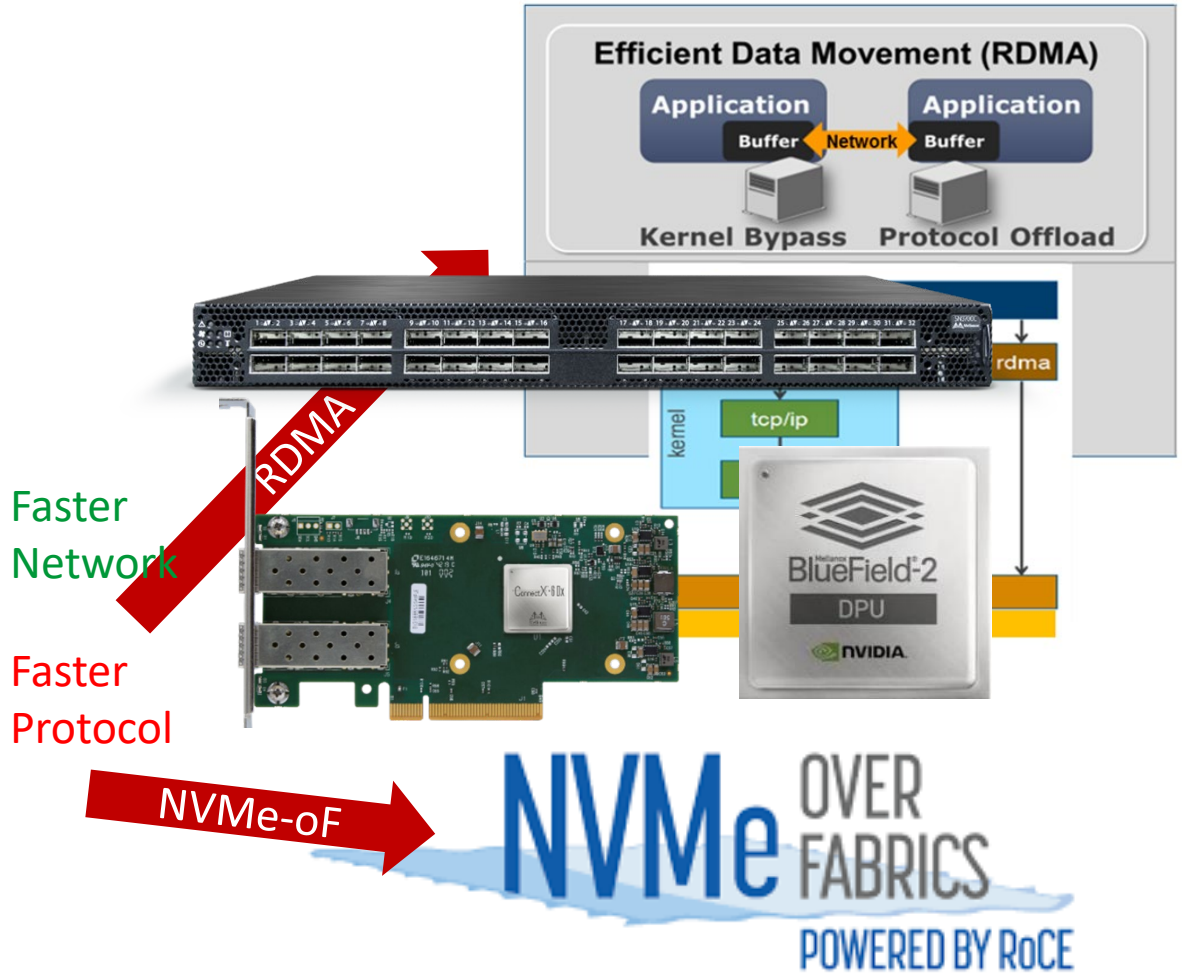

INDUSTRY WIDE RDMA ADOPTION

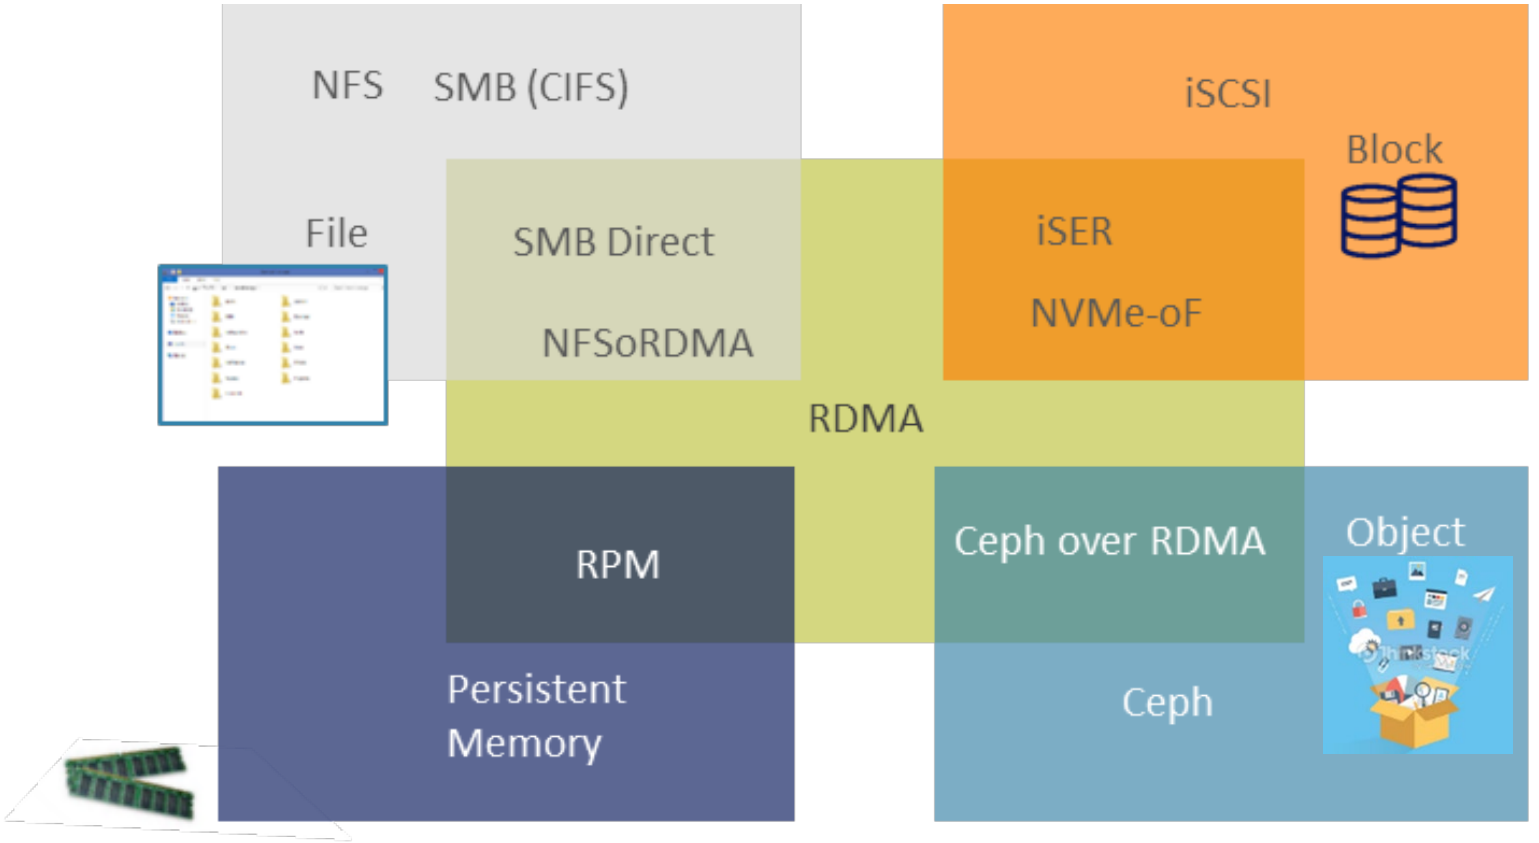

ROCE CAN OPERATE IN ANY ETHERNET FABRIC

Zero Touch RoCE (Lossy)

- No PFC or ECN configured on the fabric/Switches
- For balanced network (no built-in blocking)

ECN Only (Lossy)

- Enable ECN on switch
- Robustness for every network architecture
- Any scale

ECN+PFC (Lossless)

- Supported
- No evident advantage over “ECN only”

CONNECTX-5, ALL-TO-ALL, 8 NODES

Zero touch RoCE provides results like lossless and TCP

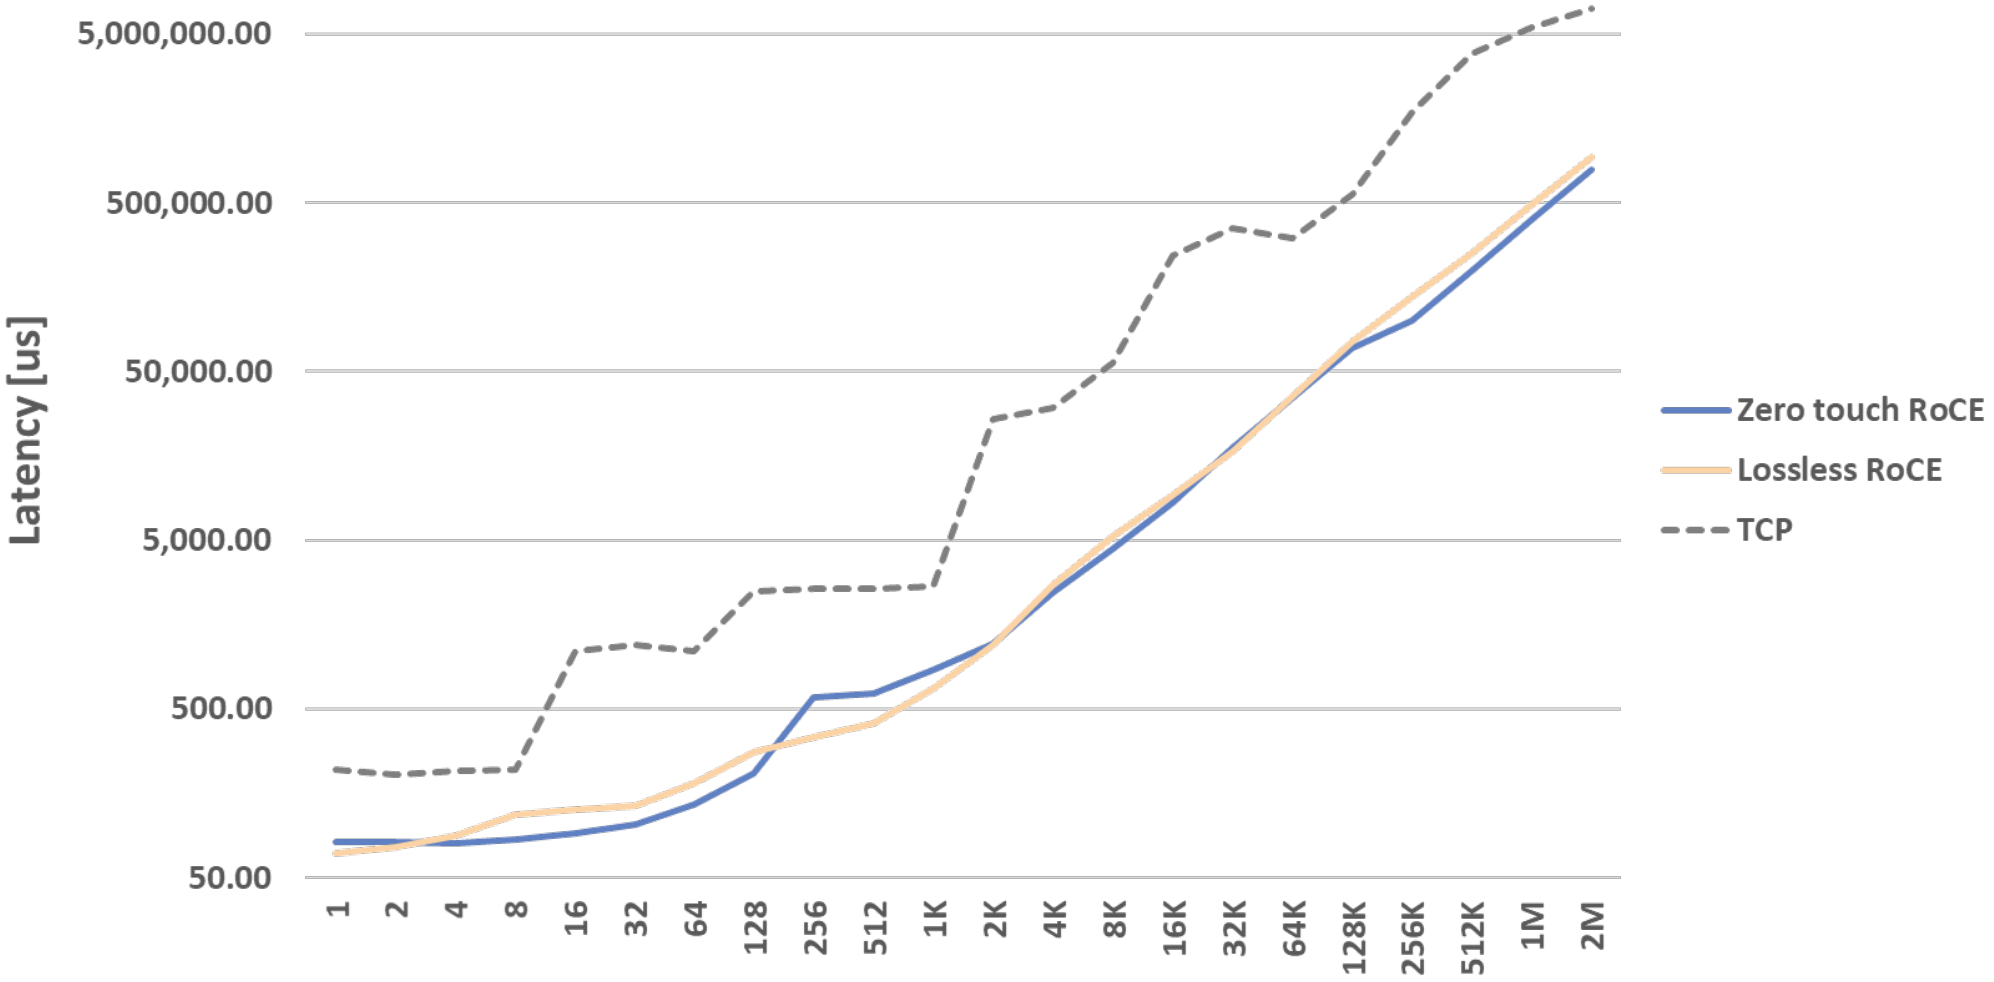

NVIDIA INFINIBAND

In-Network Computing Accelerated Network

METROX-2 LONG-HAUL

SKYWAY GATEWAY

UFM CYBER-AI

CONNECTX ADAPTER

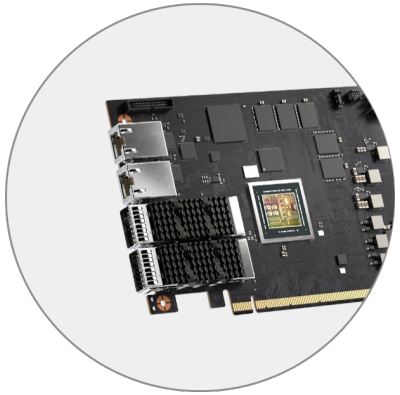

BLUEFIELD DPU

QUANTUM SWITCH

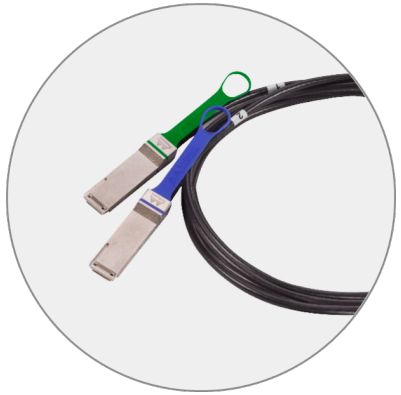

LINKX

IMPORTANCE OF NETWORK LATENCY WHEN NETWORKING NVME BASED STORAGE

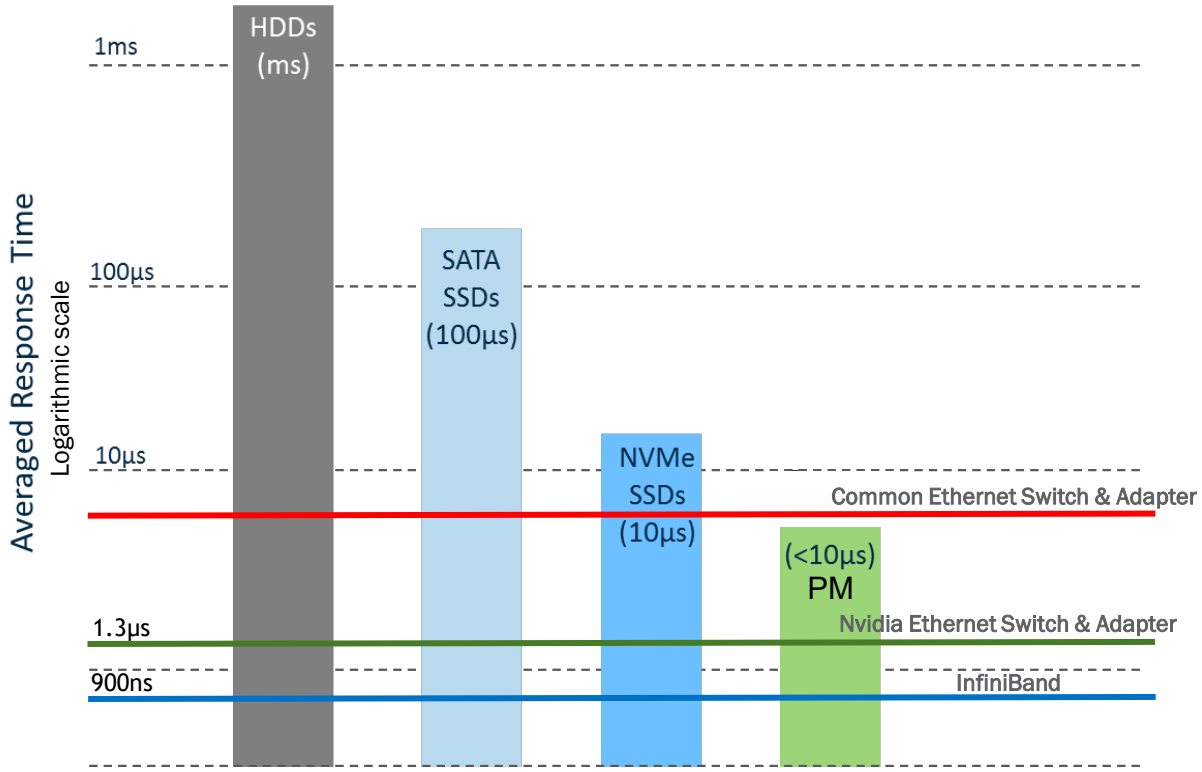

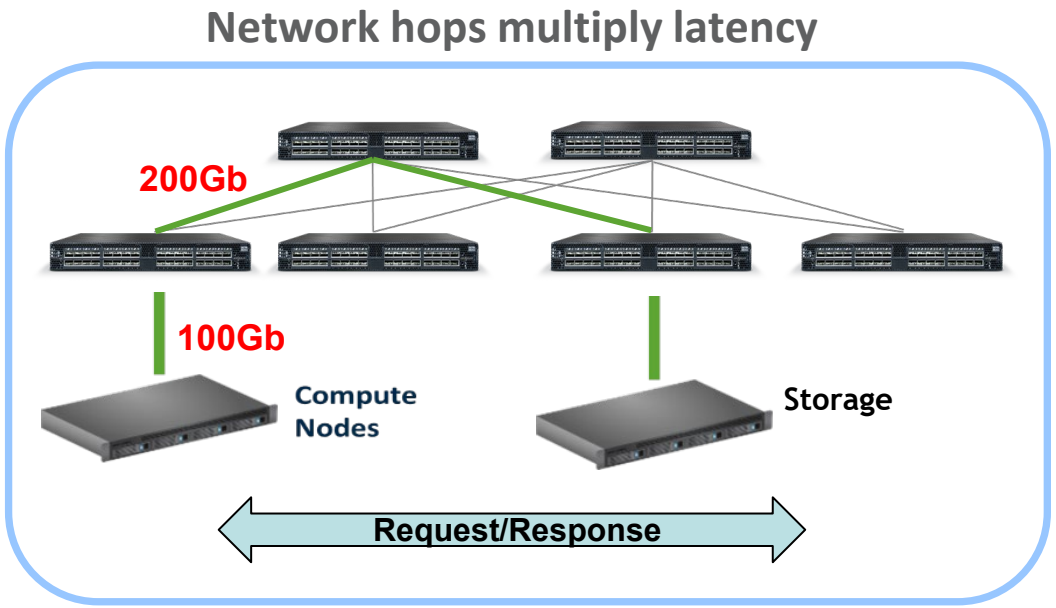

Faster Network Wire Speeds, Protocols, and CPU offload

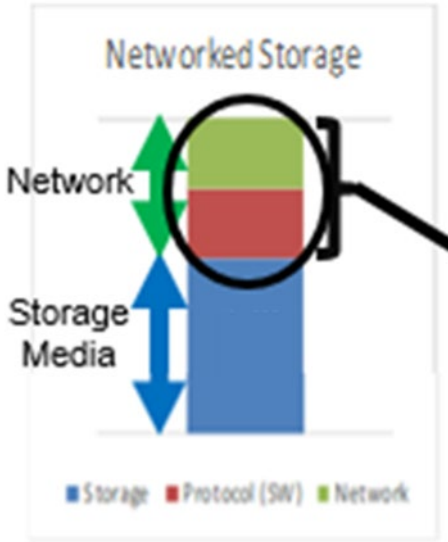

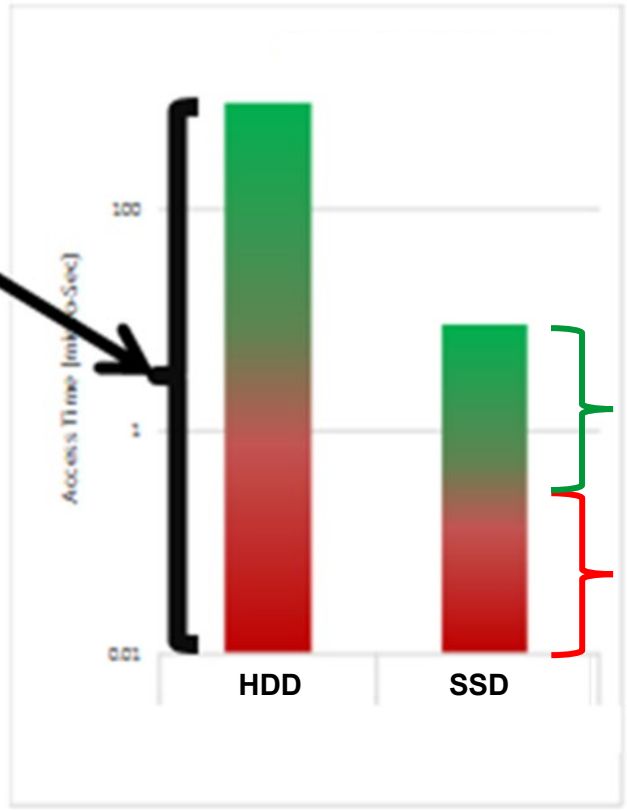

Faster Network

Faster Protocol

Storage and Network Acceleration Offloads

Network Acceleration	RDMA/RoCE
	NVMe-oF Target Offload
	TCP Acceleration - NVMe/TCP
Reliability & High Availability	T10-DIF
	CRC16/CRC32
	CRC32C/CRC64
	Erasure Coding
Security	Encryption at rest - AES-XTS
	Encryption in motion - IPSec/TLS
Lower CAPEX	Compression
	Decompression
	Deduplication (SHA offload)
	Host Chaining
Storage Emulation	SNAP - emulated NVMe / virtio

NVIDIA DPU – THE NEW BUILDING BLOCK OF A DATA CENTER

DPUs Offload, Accelerate, and Isolate Data Movement

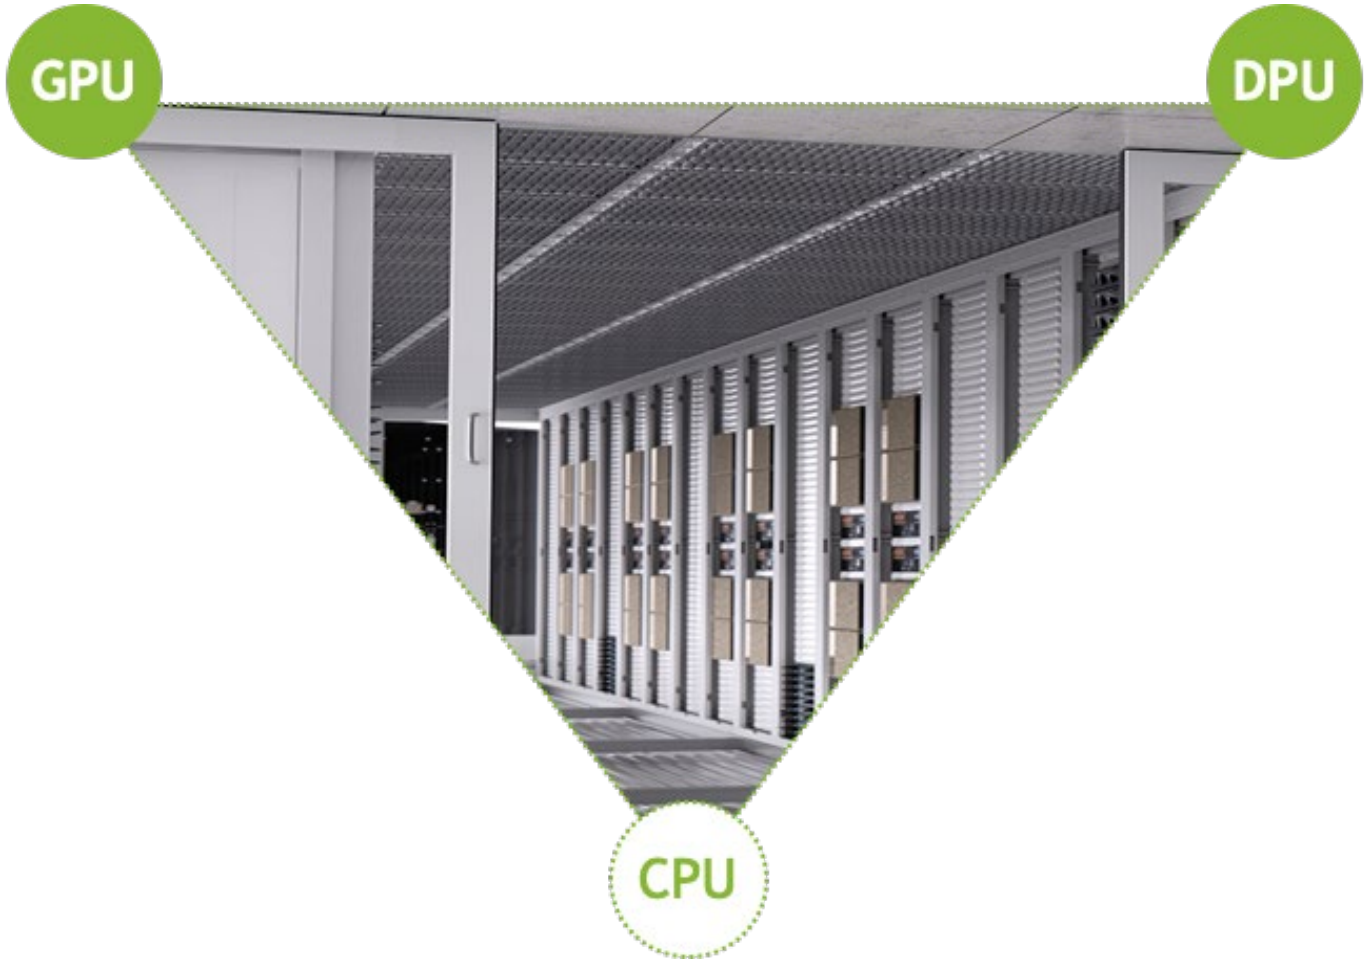

https://blogs.nvidia.com/blog/2020/05/20/whats-a-dpu-data-processing-unit/

BlueField DPU is Everywhere

Private Cloud

Cloud

AI/HPC

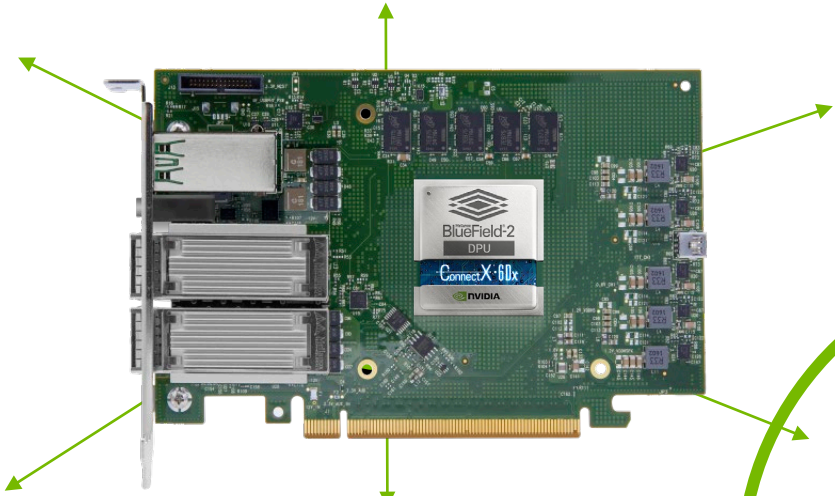

Security

Storage

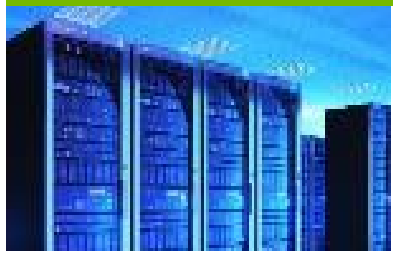

Composable / Disaggregated Storage Security JBOF

Edge/Telco

DPUS ACCELERATE AND SECURE STORAGE

Data Center Infrastructure-on-a-Chip

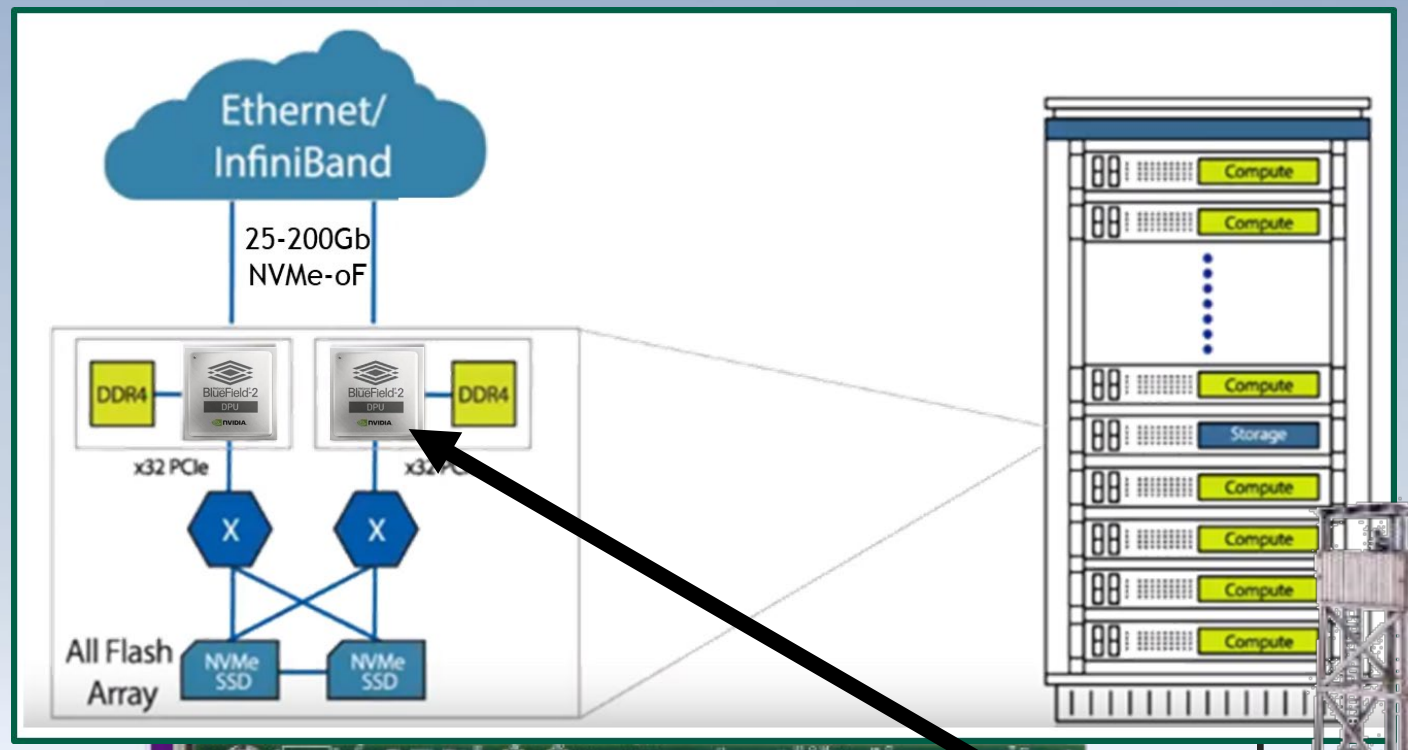

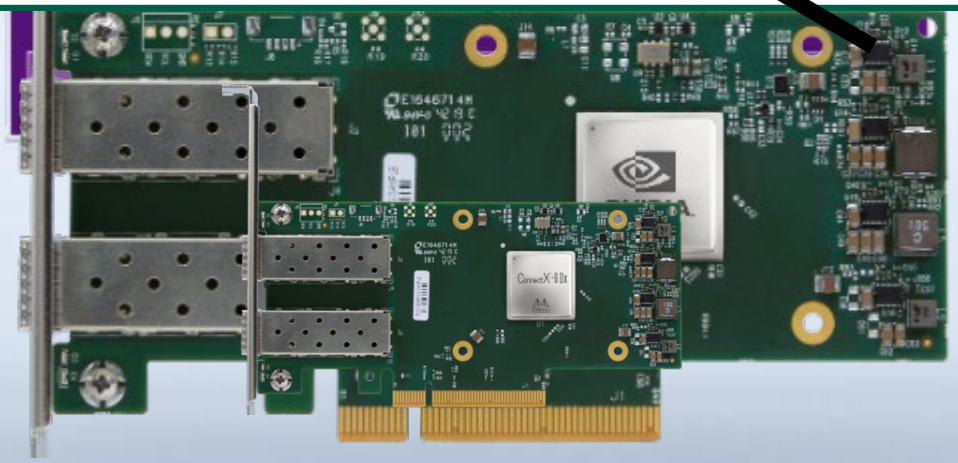

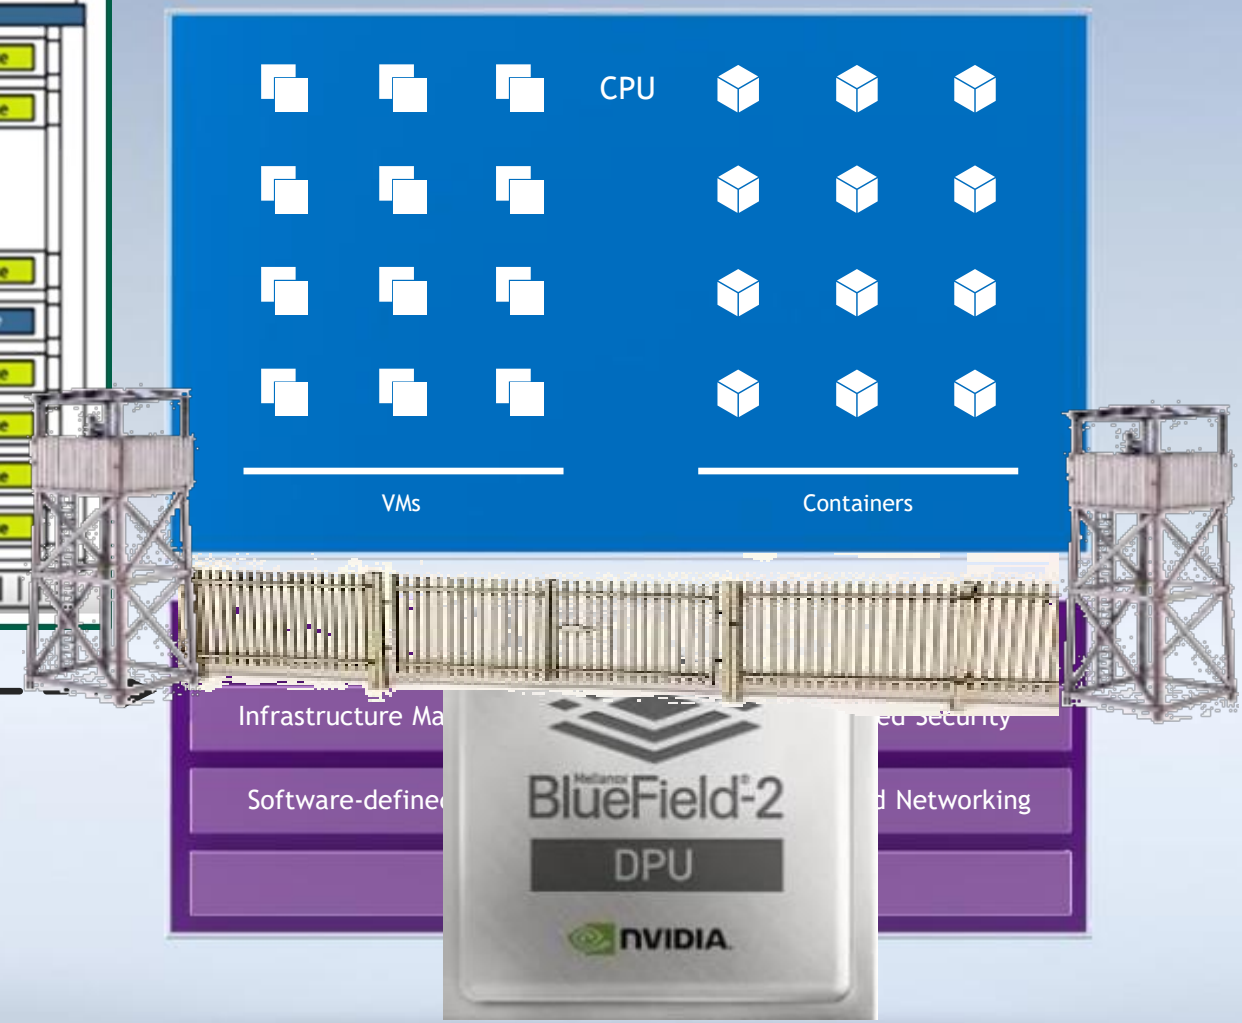

NVIDIA

NVIDIA DOCA

Data Center Infrastructure-on-a-Chip Architecture

SDK for BlueField DPUs

DOCA is for DPUs what CUDA is for GPUs

DOCA preserves App developers' investment as DPUs evolve

Open-source APIs – DPDK, SPDK, P4

Certified reference apps & 3rd party solutions

Support for multiple OS

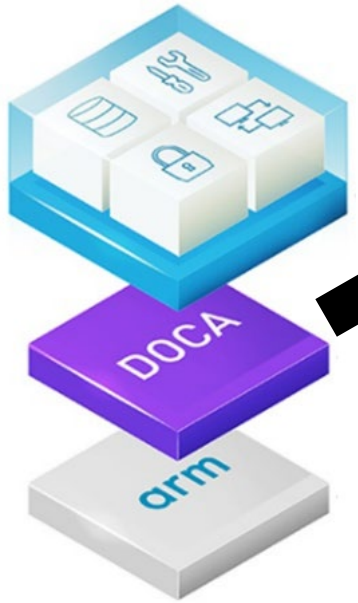

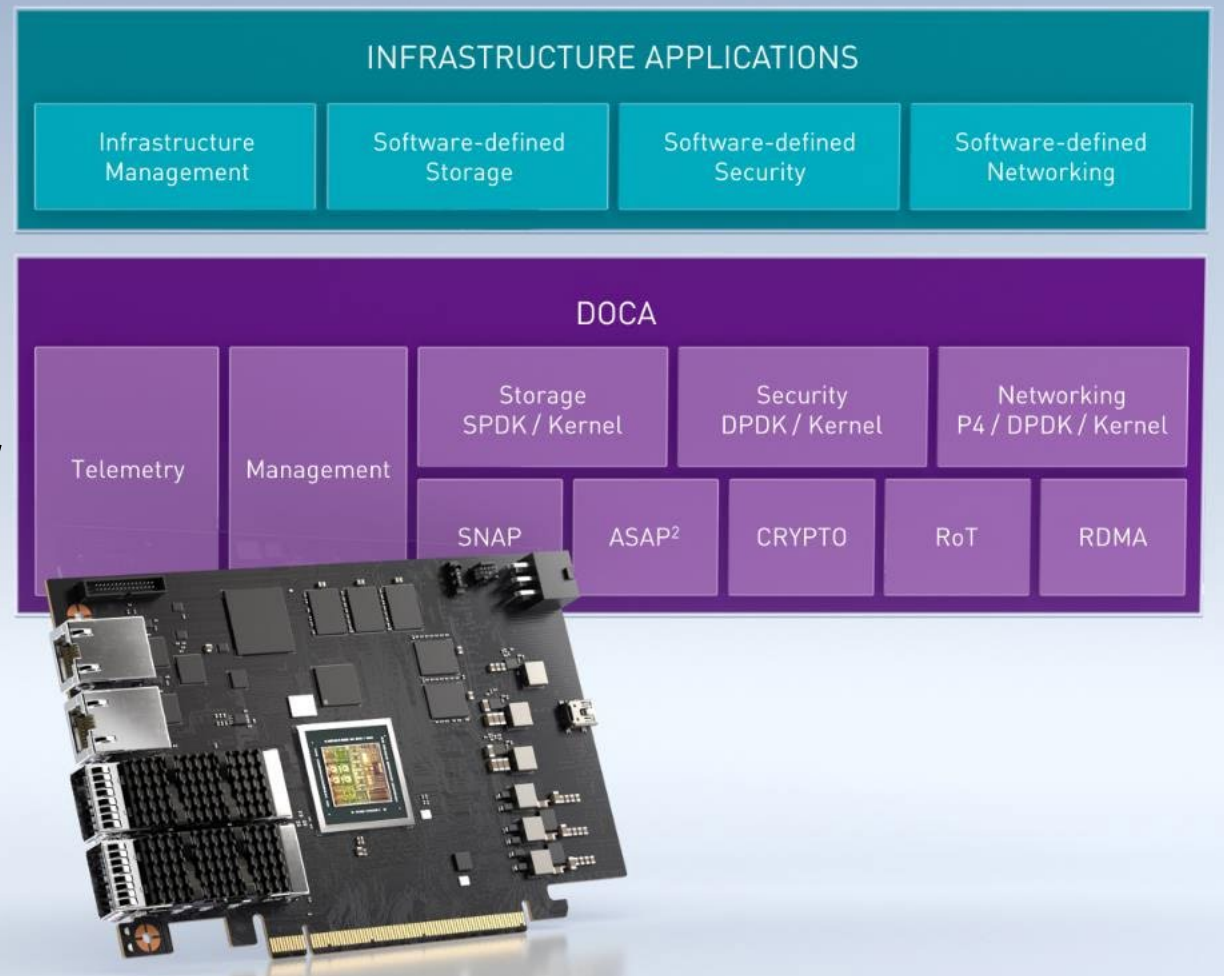

https://developer.nvidia.com/networking/doca

Panel Questions and Audience Surveys

Panel Question #1

- How do you balance scalability, flexibility, performance, and energy footprint in CSP storage and storage networking solutions?
  - Pliops
  - Excelero
  - NVIDIA

Audience Survey Question #1

- What is your organization's greatest concern when using cloud storage? (select one):
  - Overall Cost: 15%
  - Cost Predictability: 12%
  - Performance: 12%
  - Performance Predictability: 15%
  - Security/Data Privacy: 31%
  - Other Concerns: 8%
  - No concerns/no opinion: 8%

Panel Question #2

- Many hyperscalers and CSPs have gone to building their own storage and storage networking hardware over the last decade. What does commercially-available technology have to uniquely offer them?
  - Excelero
  - NVIDIA
  - Pliops

Audience Survey Question #2

- When building a cloud datacenter's storage infrastructure, what is your greatest concern? (select one):
  - Scalability: 17%
  - Storage performance: 25%
  - Minimizing CPU utilization: 4%
  - CapEx/upfront costs: 17%
  - Avoiding vendor lock-in: 25%
  - Other: 0%
  - No opinion: 13%

Panel Question #3

- Many CSPs and hyperscalers now offer cloud-based instances of “open” technologies such as NVMe®, NVMe-oF™, and clustered file systems. What is the best way to integrate these technologies into their storage architectures?
  - NVIDIA
  - Pliops
  - Excelero

Audience Q&A


Thank You For Attending!