



Rack-Scale Flash. No Compromise.

From Rack Scale to Network Scale: NVMe Over Fabrics Enables Exabyte Applications

Zivan Ori, CEO & Co-founder, E8 Storage



NVMe Over Fabrics – Who? Why? How?

Who is using NVMe over fabrics?

Why do they need it?

How do they use it?

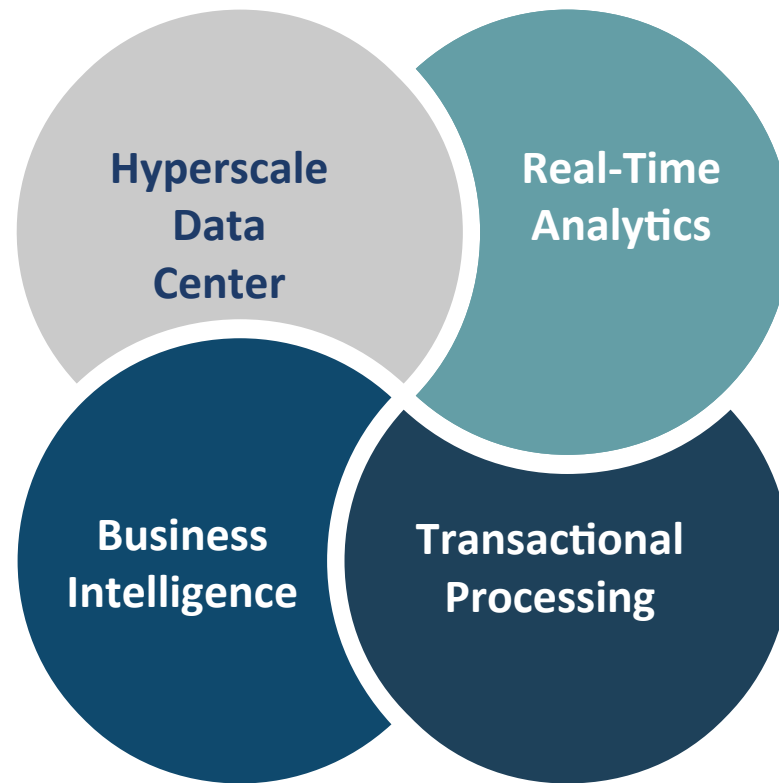
Who needs NVMe over Fabrics



Scalability without sacrificing performance



Parallel processing of shared RW data




Faster time to results at 70% better TCO




Low latency for financial transactions

NVMe Enables New Applications


- **Anyone** using local SSDs would benefit from NVMe
- Ideal for:
 - SQL & NoSQL databases
 - Real-time analytics
 - High performance distributed KV stores
 - HPC
- Enables **new** applications:
 - Machine learning
 - Internet of Things databases and analytics
 - Real-time Application Performance Monitoring
 - Real-time Security Audits



Storage Growth is Data Growth




SSDs are Becoming Faster



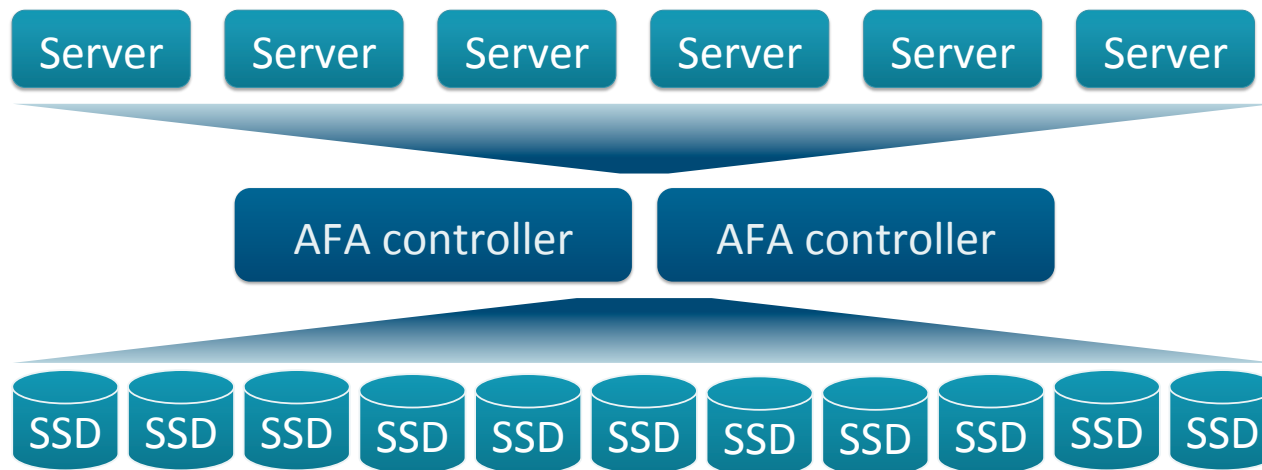
All-Flash Arrays Lagging Behind NVMe

A single NVMe SSD has more performance than an All-flash-array!



Classic Scale-Up Bottleneck

The all-flash array bottleneck is the dual controller itself



Customers are Facing a Tough Choice...


- Use SSDs as local storage in their servers to get high performance
- OR...*
- Put those SSDs in AFAs to gain manageability and provisioning

Inevitably more and more customers choose to use SSDs as local storage

The Problem: Using SSDs as Persistent Local Storage

The provisioning problem


- Islands of storage
- Inefficient capacity utilization
- Cannot share SSD content between servers
- Cannot grow storage capacity



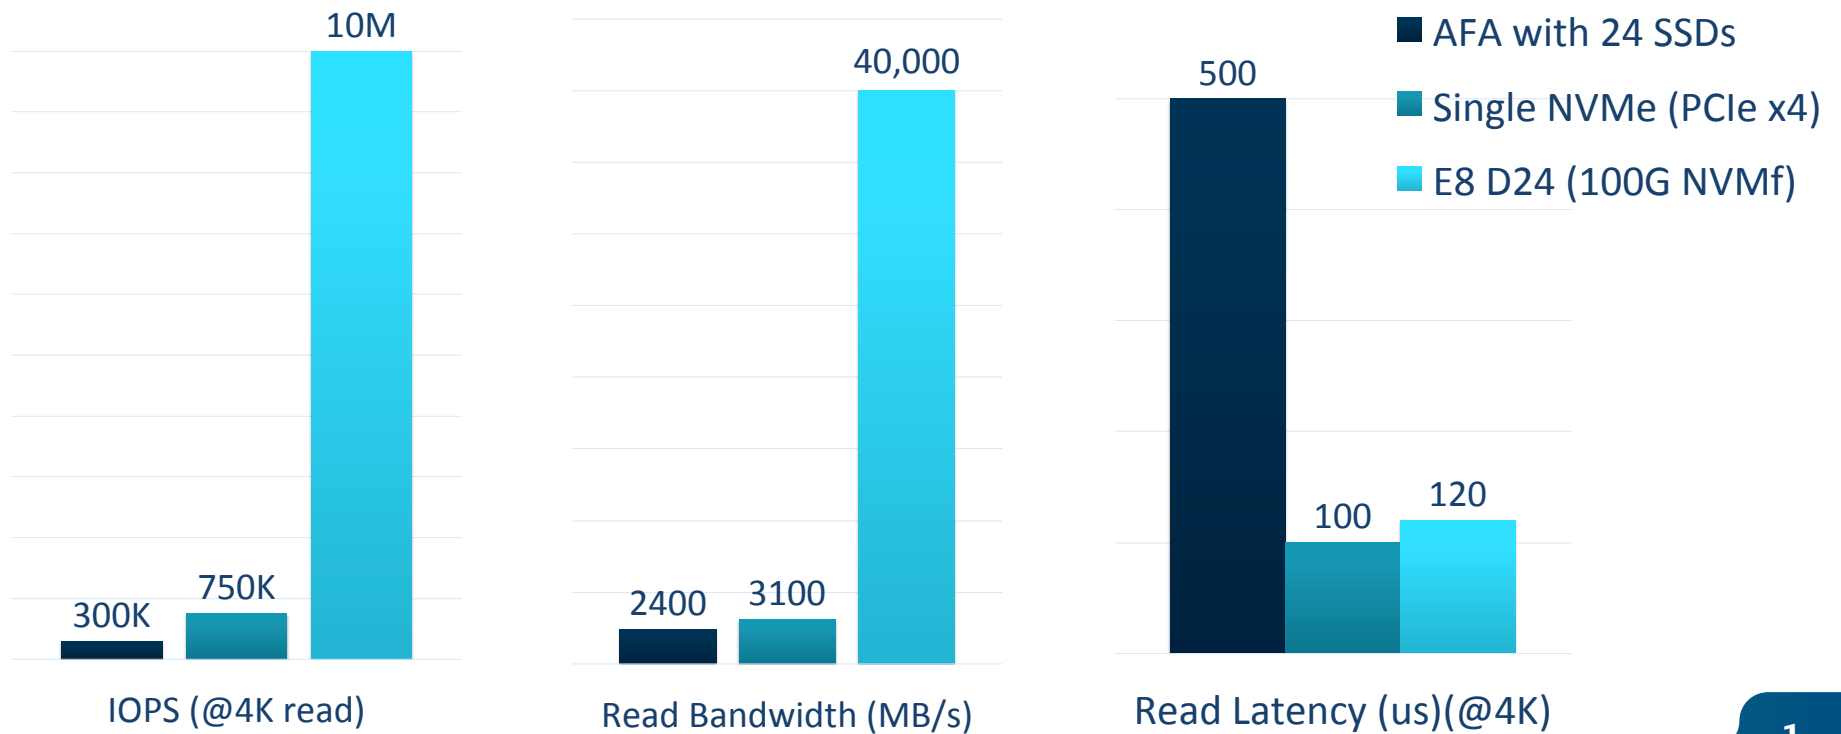
The Problem: Using SSDs as Persistent Local Storage

Coupling of storage and compute

- Up front purchasing decision
- Repair / upgrade challenges
- Lifecycle decisions



End-to-end NVMe Over Fabrics To the Rescue



Real-time Analytics Customer (Financial)

Perform decisions at real-time, e.g. market price adjustments

Before

- 576 SSDs in 72 servers
- Copy nightly the data to all servers



After

- 24 SSDs in E8-D24
- Share the data from E8 to all 72 servers



Shared NVMe reduces the number of replicas needed by 10X - 100X

In-memory DB Replacement using Storage Class Memory

Application example: real-time market price adjustments

Before

- 20TB RAM on 13 servers
- 20us latency over 40G IB



After

- 24 SCM 1TB SSDs in E8
- Share the data from E8 to all database nodes

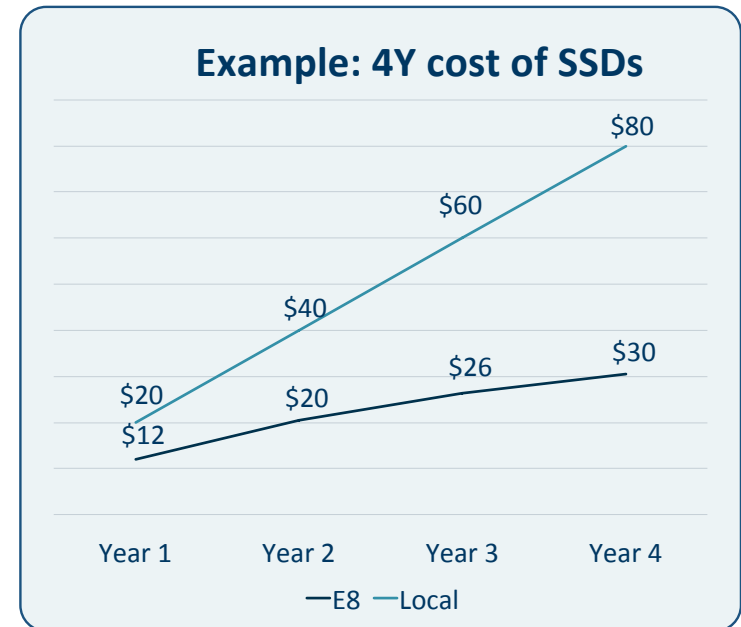


Shared SCM increases storage density and reduces power by over 10:1

Economic Argument for Disaggregated Flash

When used as local persistent storage...

- Need to buy SSD with the server
- SSD gradually fills up: but already paid for on day 1
- Future projected capacity is paid for in today's price




Save \$50M over the 4-year life of a data-center

Hyperscale Customer: Top-of-the-rack Storage

Before

- 96 servers in a rack
- 96 * 2.2TB SSDs in servers



Local SSD

Utilization: 30%

After


- 92 servers in a rack
- 24 * 4TB SSDs in E8 box



Central SSD

Utilization: 88%

Converged Ethernet vs. Segregated Network



Converged Ethernet

- Uses the existing network
- Reduces equipment and cabling
- Allows flash to be accessible by all nodes in the network.


From Rack Scale to Network Scale



200GbE / 400GbE fuels network scale storage



New enhancements for network growth



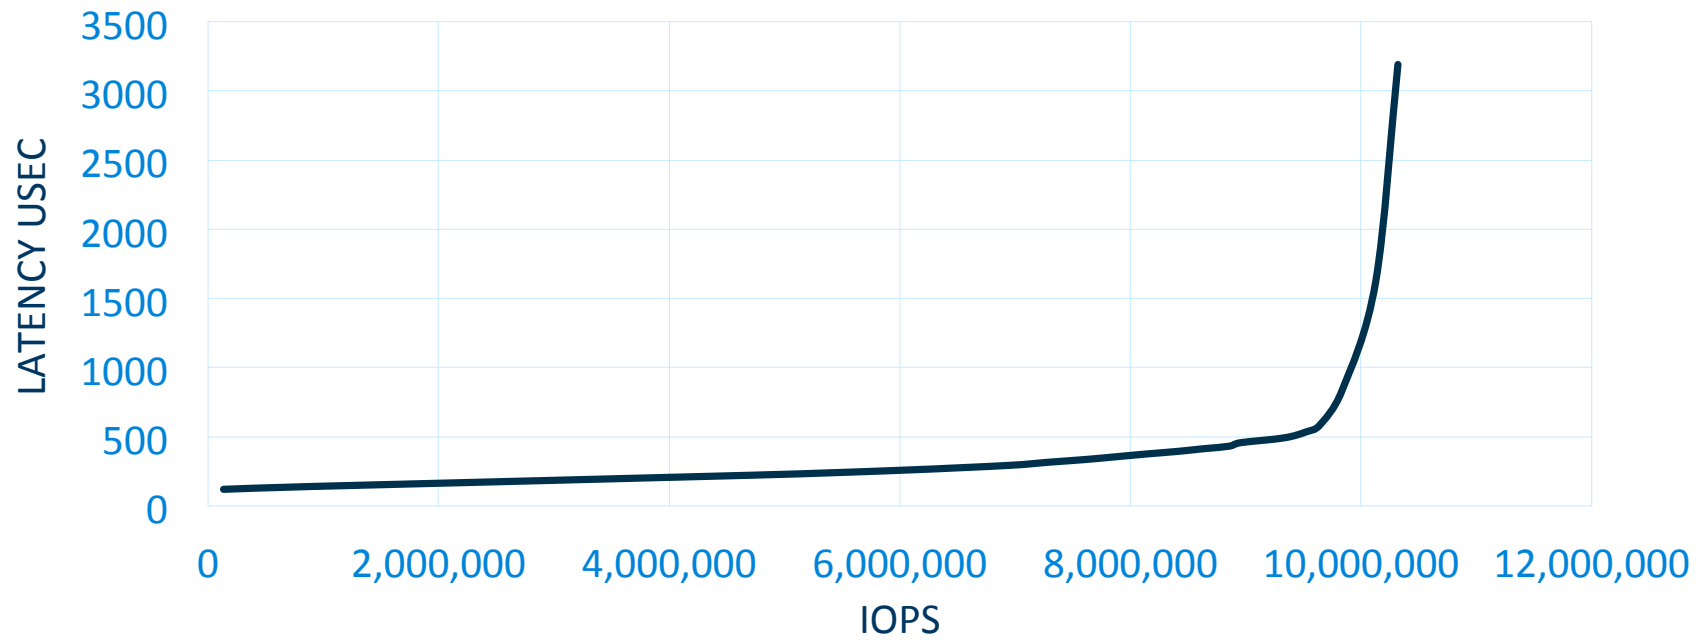
Easily add NVMe capacity to Ethernet Network



Clustered storage with centralized management

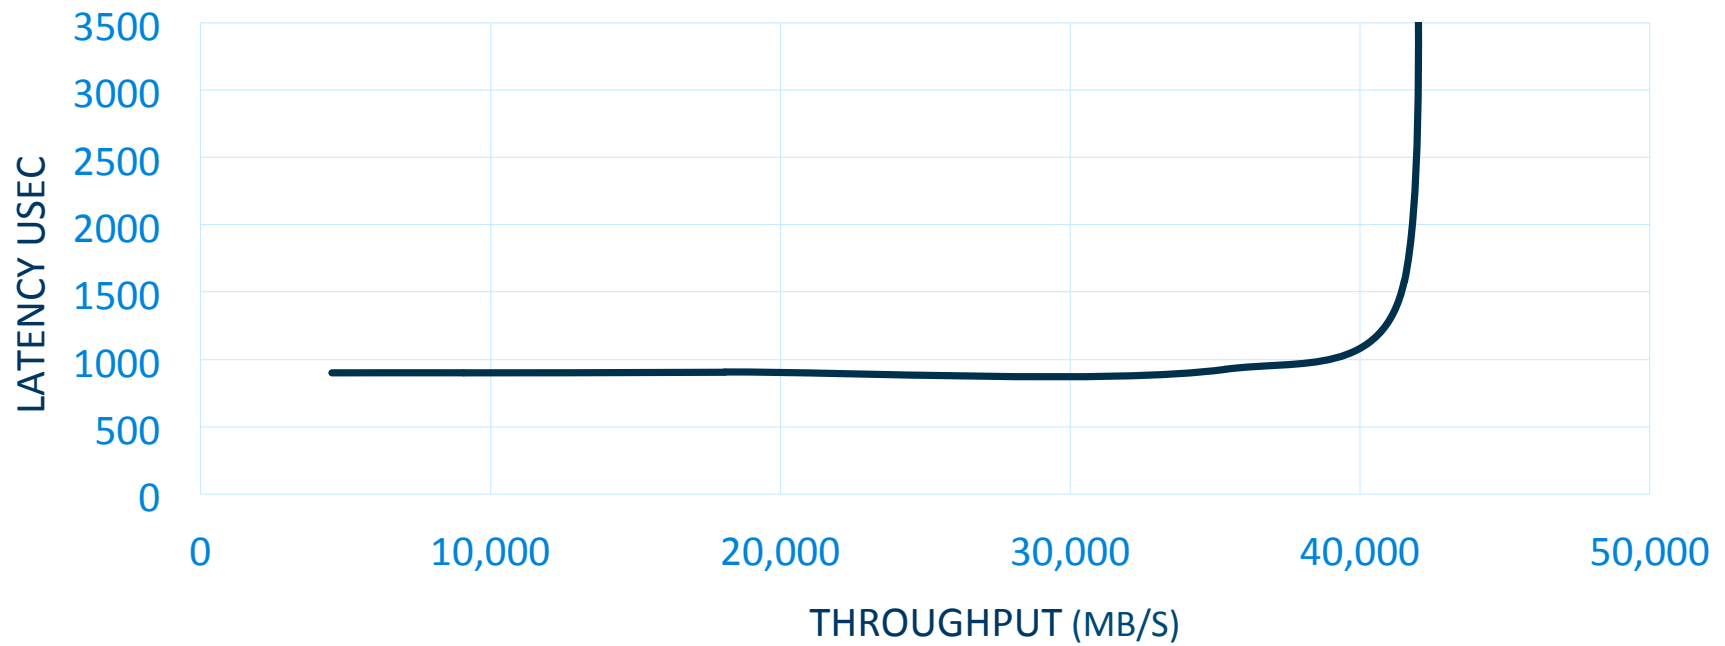
E8 Performance and Latency

4K 100% RR IOPS/LATENCY




E8 Performance and Latency

256K 100% RR BANDWIDTH/LATENCY




E8 Performance and Latency

90% 8K RR IOPS / LATENCY (RAID6)



E8 Storage – Rack Scale Flash. No Compromise.



PCIe SSD
Performance



Centralized
Storage
Reliability



Hyper-
scalability



Affordable
100% COTS