

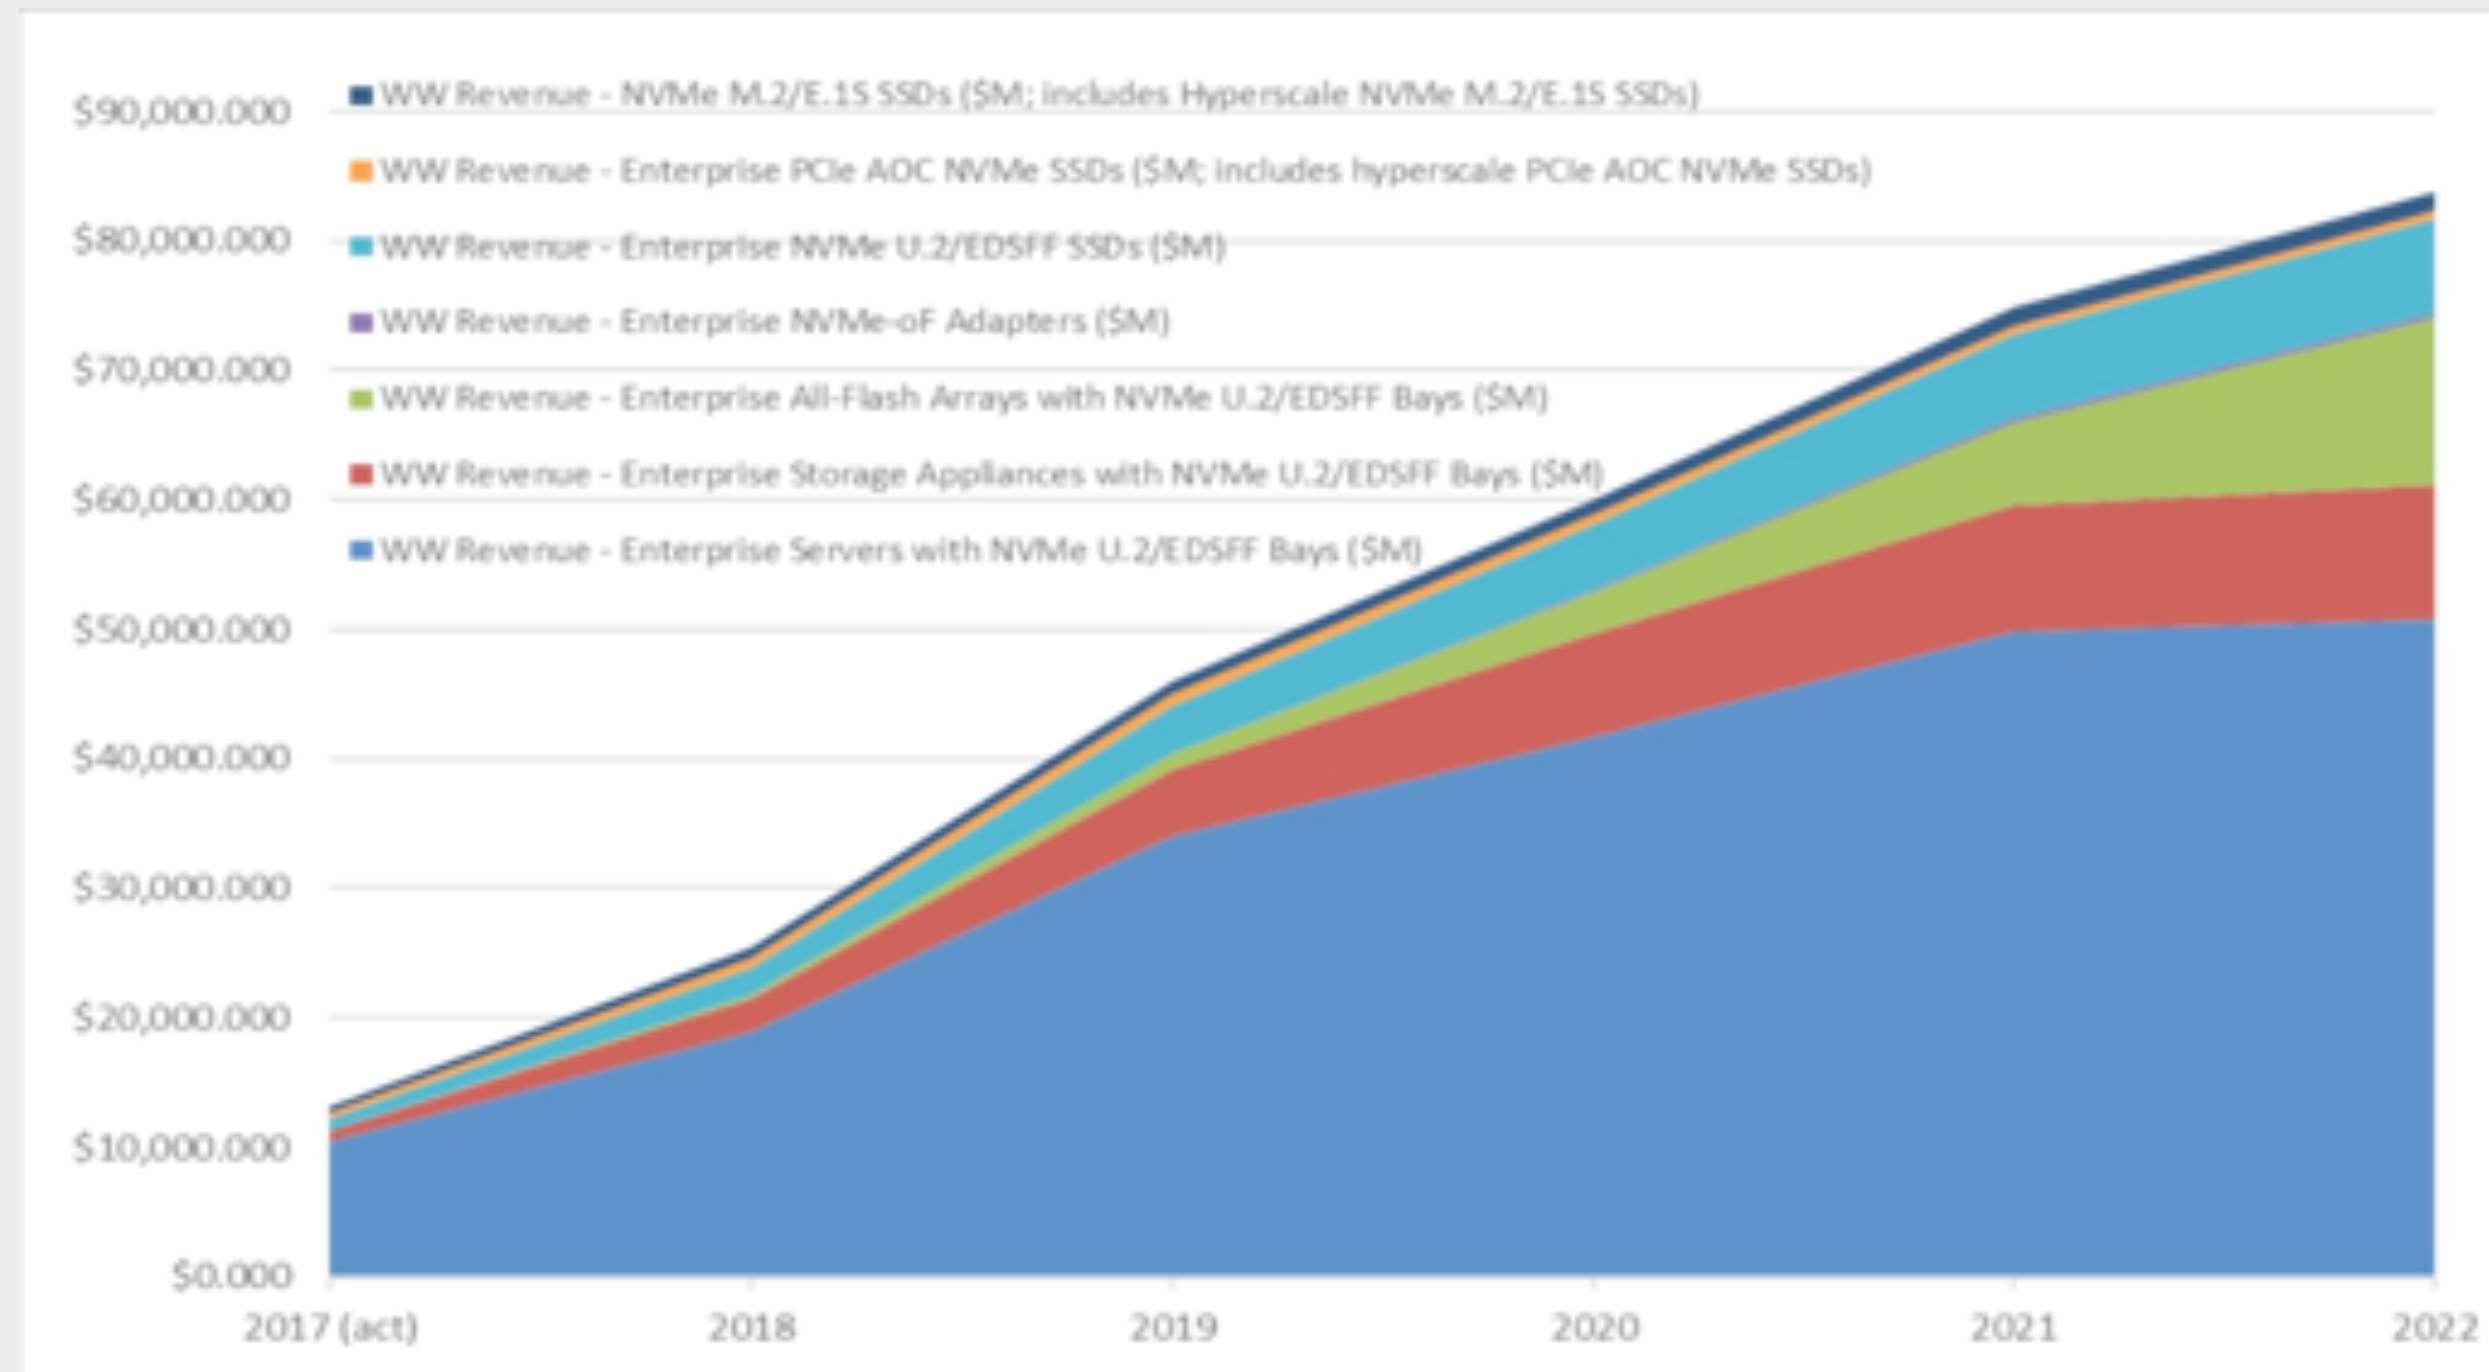
Accelerating Data Intensive Applications with Shared NVMe Storage

E8 storage

When Performance Matters

NVMe Market Adoption

- Adoption shifting over time
 - PC / Servers continue to lead
 - Storage adoption ramping
- NVMe-oF for Storage
 - Standard critical for ramp
 - Major players entering market



Early NVMe Use Cases

Applications that drive business revenue

- Financial Analytics
- Genome Research
- Artificial Intelligence / Machine Learning / Deep Learning
- Fluid Dynamics

The E8 Storage Difference

A new architecture built specifically For high performance NVMe

Direct drive access for near line rate performance

- Separation of data and control paths; no controller bottleneck
- Offloads up to 90% of data path operations to E8 host agents

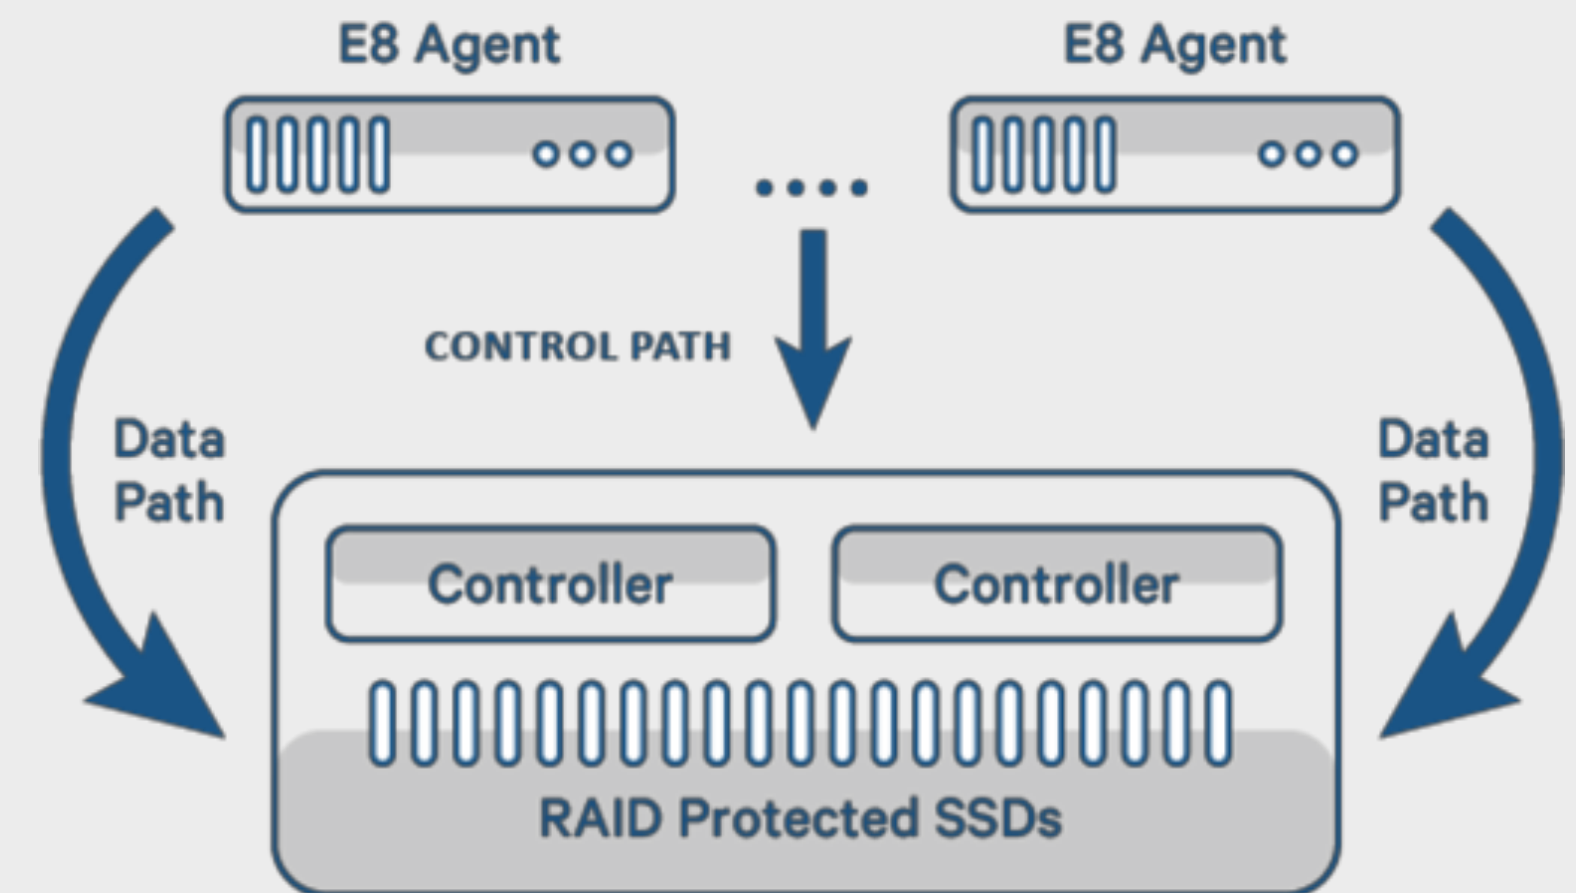
Simple, centralized management

- Intuitive management GUI for host / volume management
- E8 host agents auto-discover assigned LUNs

Unleash the parallelism of NVMe SSDs

Scalable in multiple dimensions

- Scale through host agents
- Scale through target enclosures



Designed for Availability and Reliability

No single point of failure anywhere In the architecture

Host agents operate independently

- Failure of one agent (or more) does not affect other agents
- Access to shared storage is not impacted


RAID data protection (RAID 6/5/10)

Network multi-pathing

- Infiniband, RoCE V1/V2

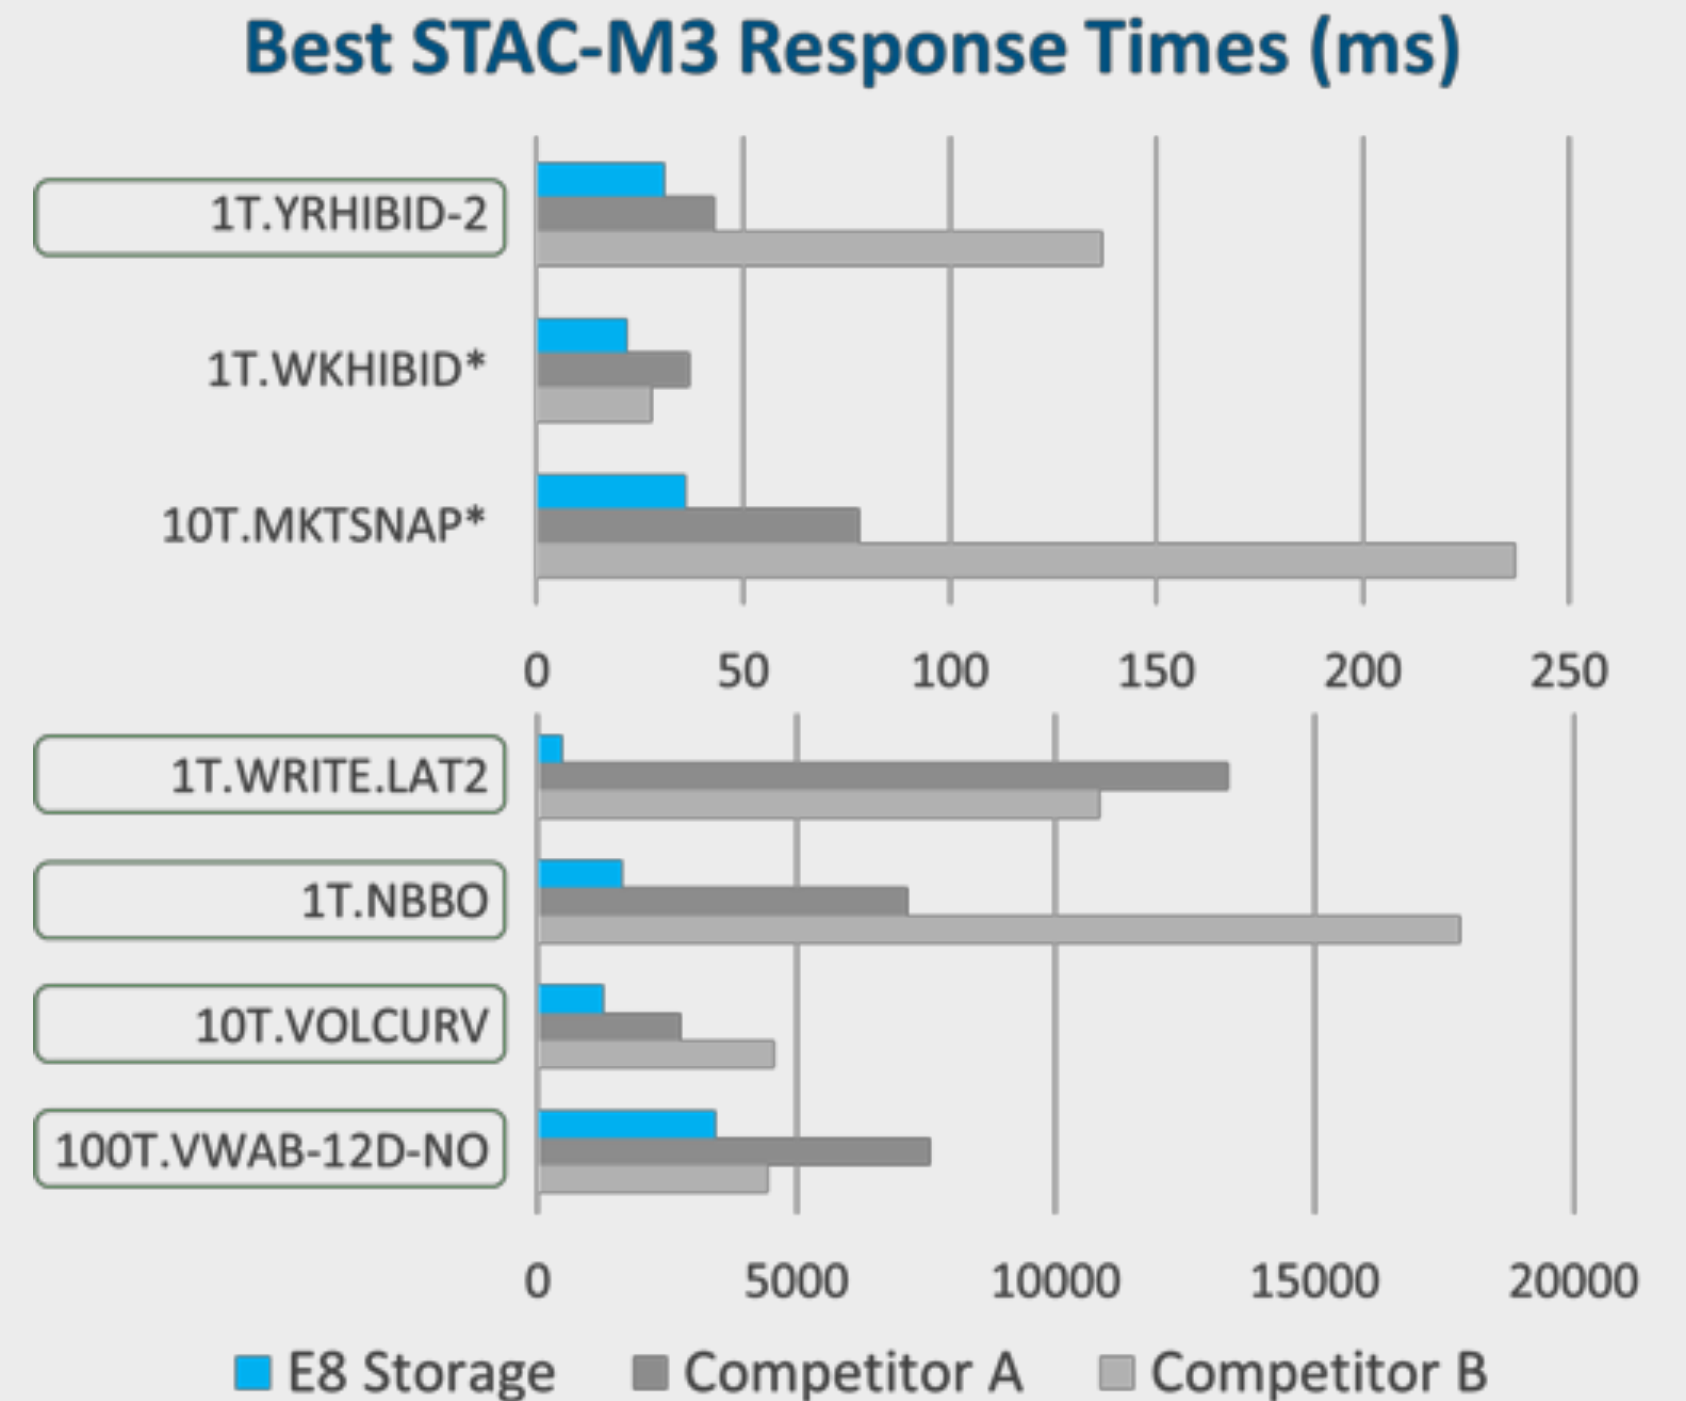
Enclosure high availability

- Option 1: HA enclosure + dual-ported SSDs
- Option 2: Cross-enclosure HA + single-ported SSDs



NVMe for Financial Analytics

- “Latency is King”
- FinTech driving automated trading
- Real-time transaction processing
- Daily trade analytics



Market Data Analytics for Financials

In production with 2 of the world's Top-10 largest hedge funds

Before

- 1152 local SSDs in 72 servers
- Market data copied nightly to all servers
- Restricted to 10TB-20TB



After


- 48 SSDs in 2 E8-D24 appliances
- Market data shared to all 72 servers
- Easily scalable to 300TB



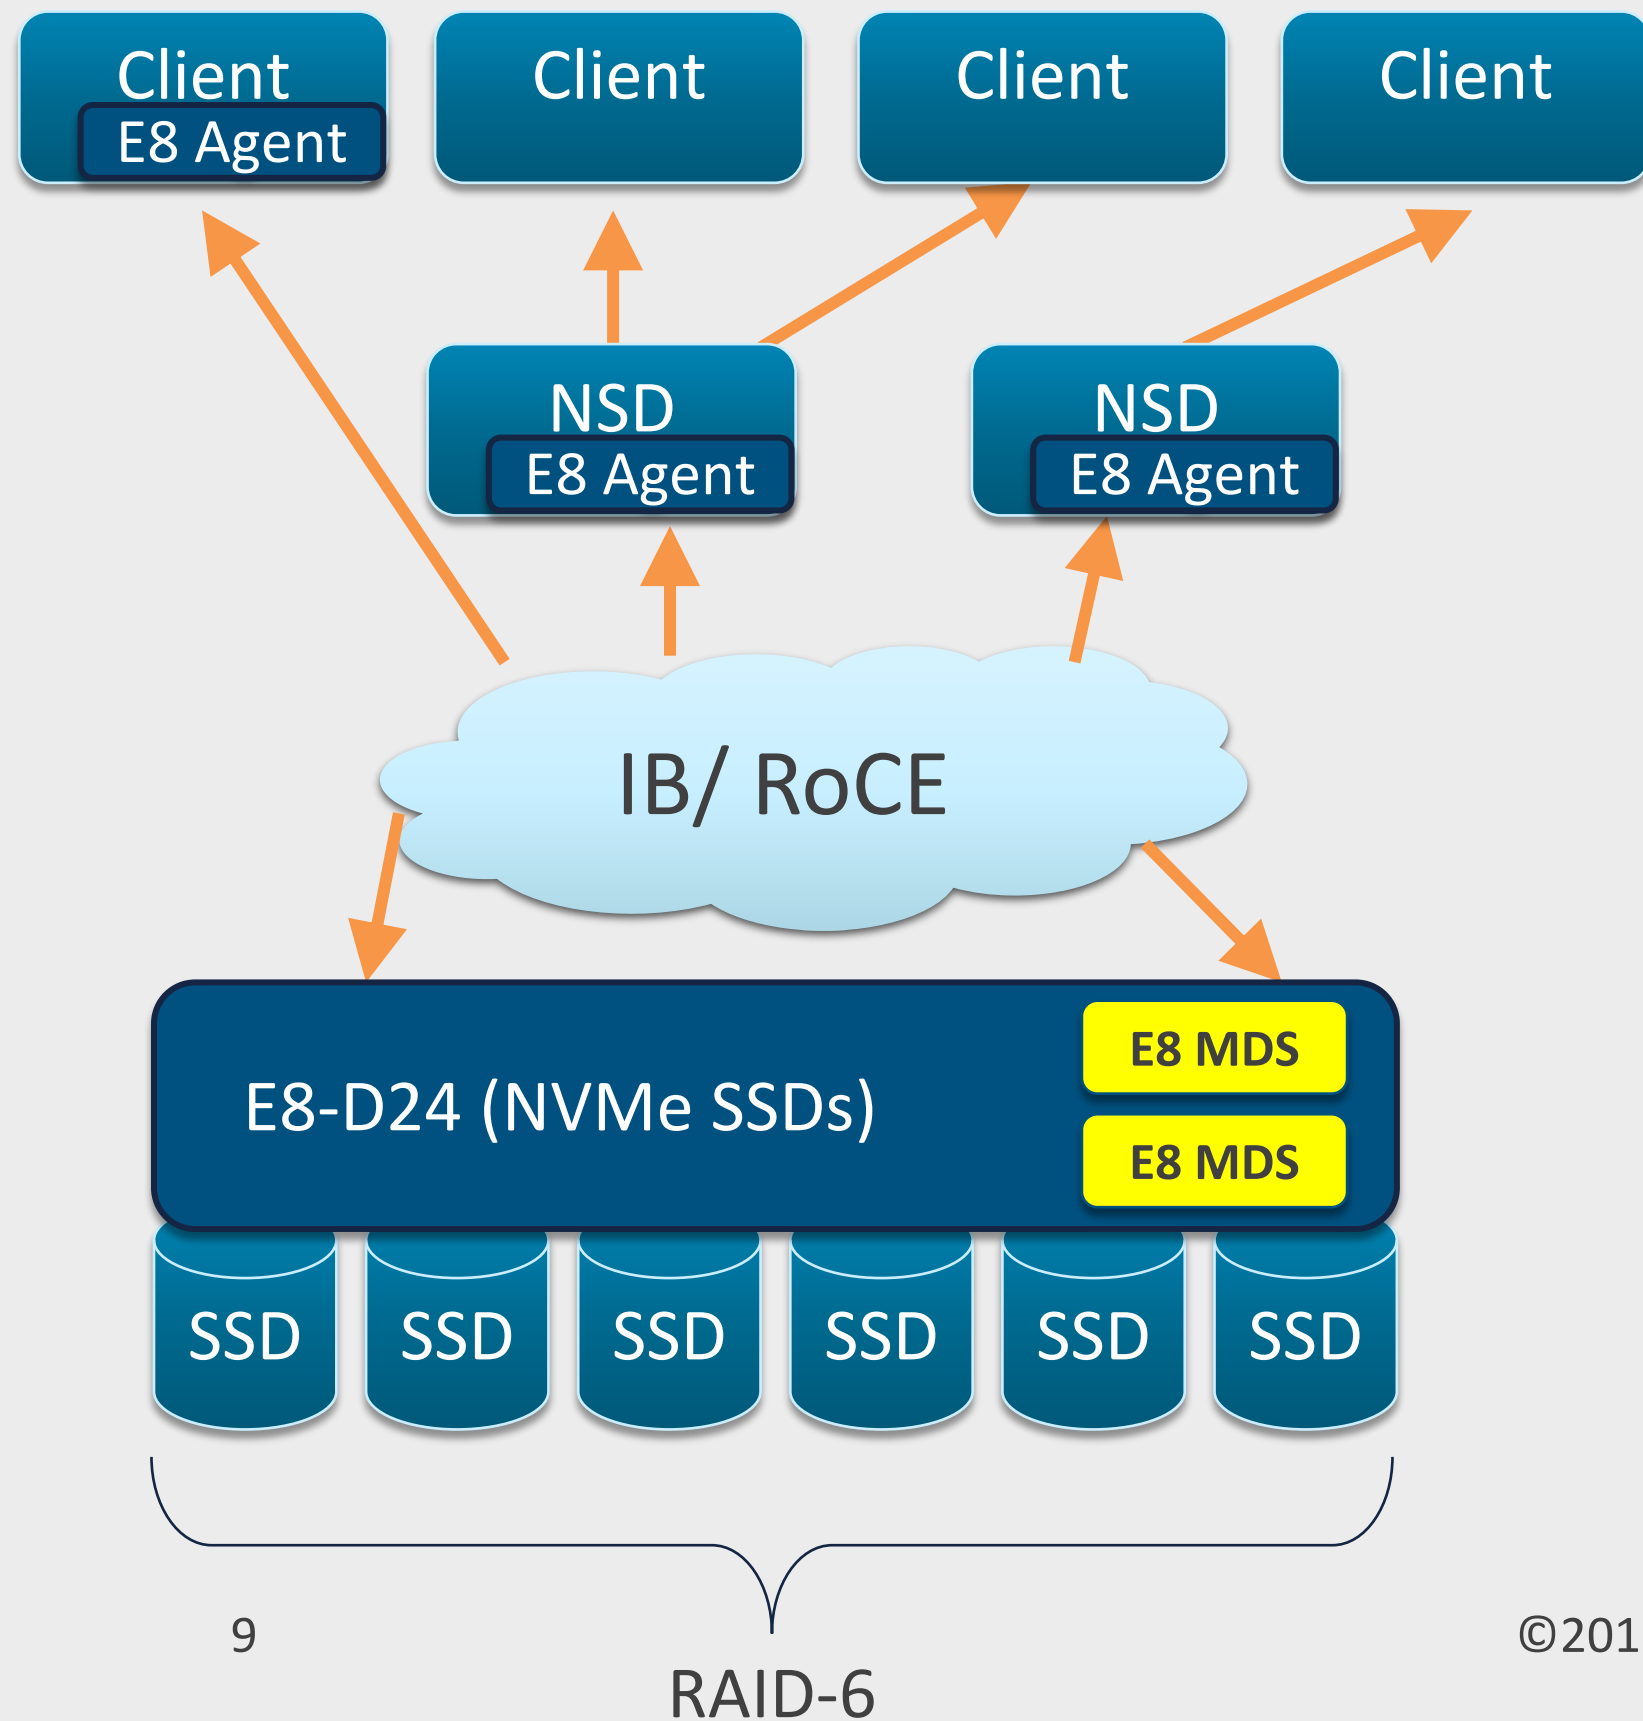
70% Cost
reduction!

Accelerating the Science of Life

- Primary genome sequencing
 - Historically – 10 hours per genome
 - NVMe – 10 genomes per hour!
- Secondary genomic analysis
 - Analyzing genomes per population
 - Requires TB of data



Using E8 with IBM Spectrum Scale



Scalable to larger configurations

- Can mix connectivity depending on requirements

Standalone pool

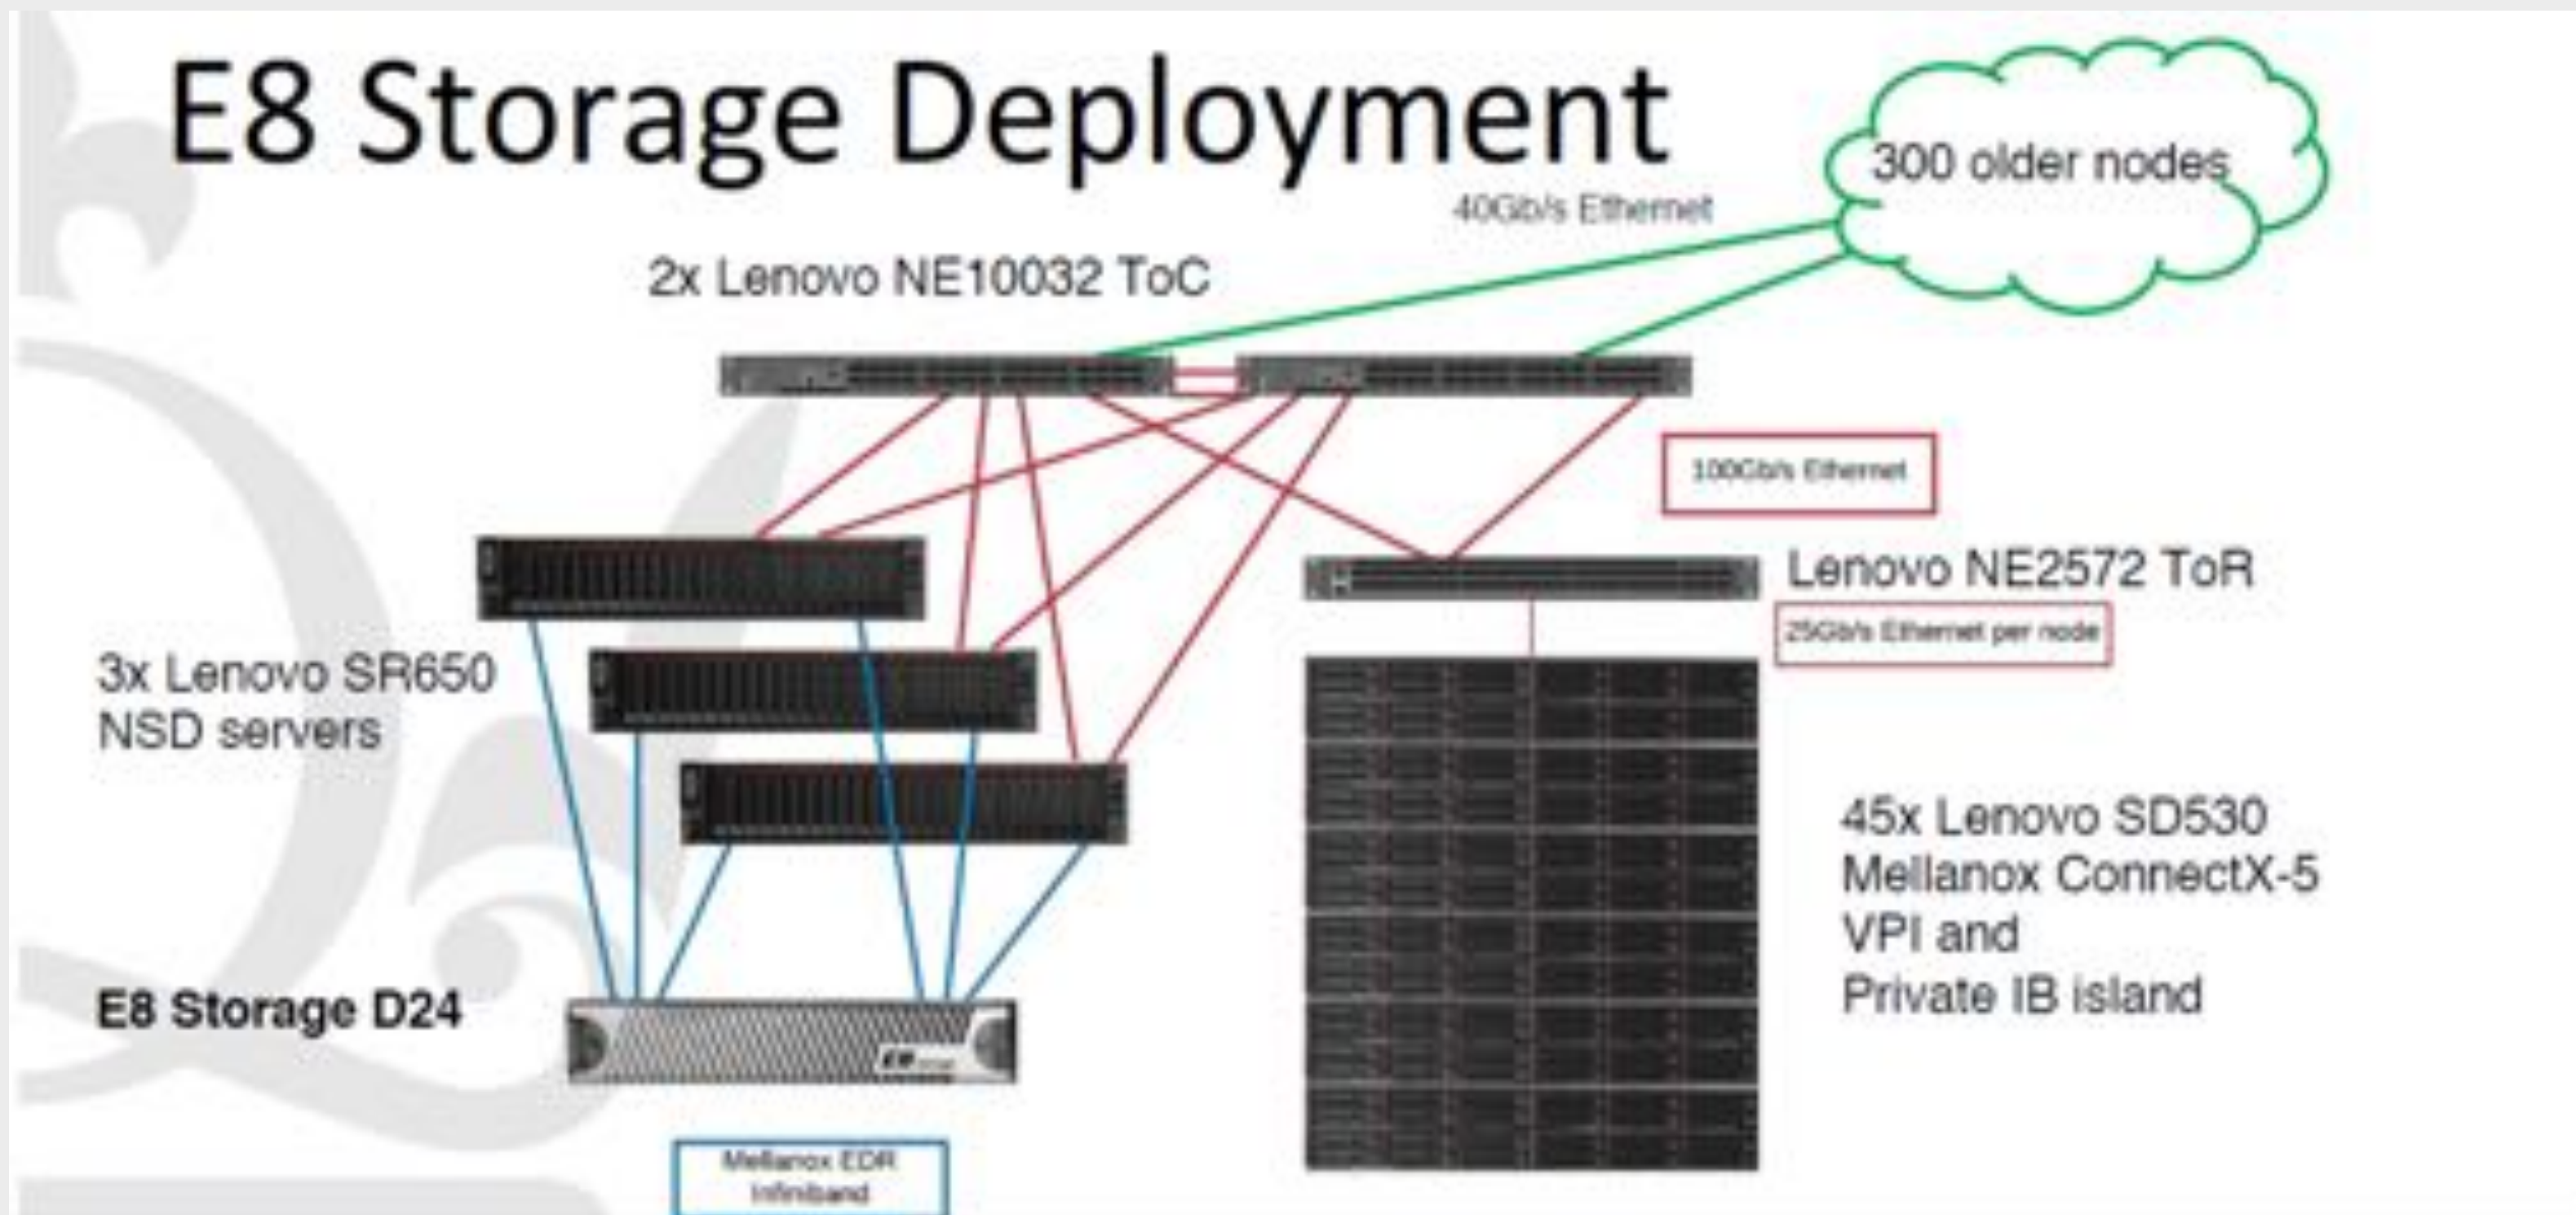
- Shared LUNs

LROC

- Non-shared LUNs (direct connect clients only)

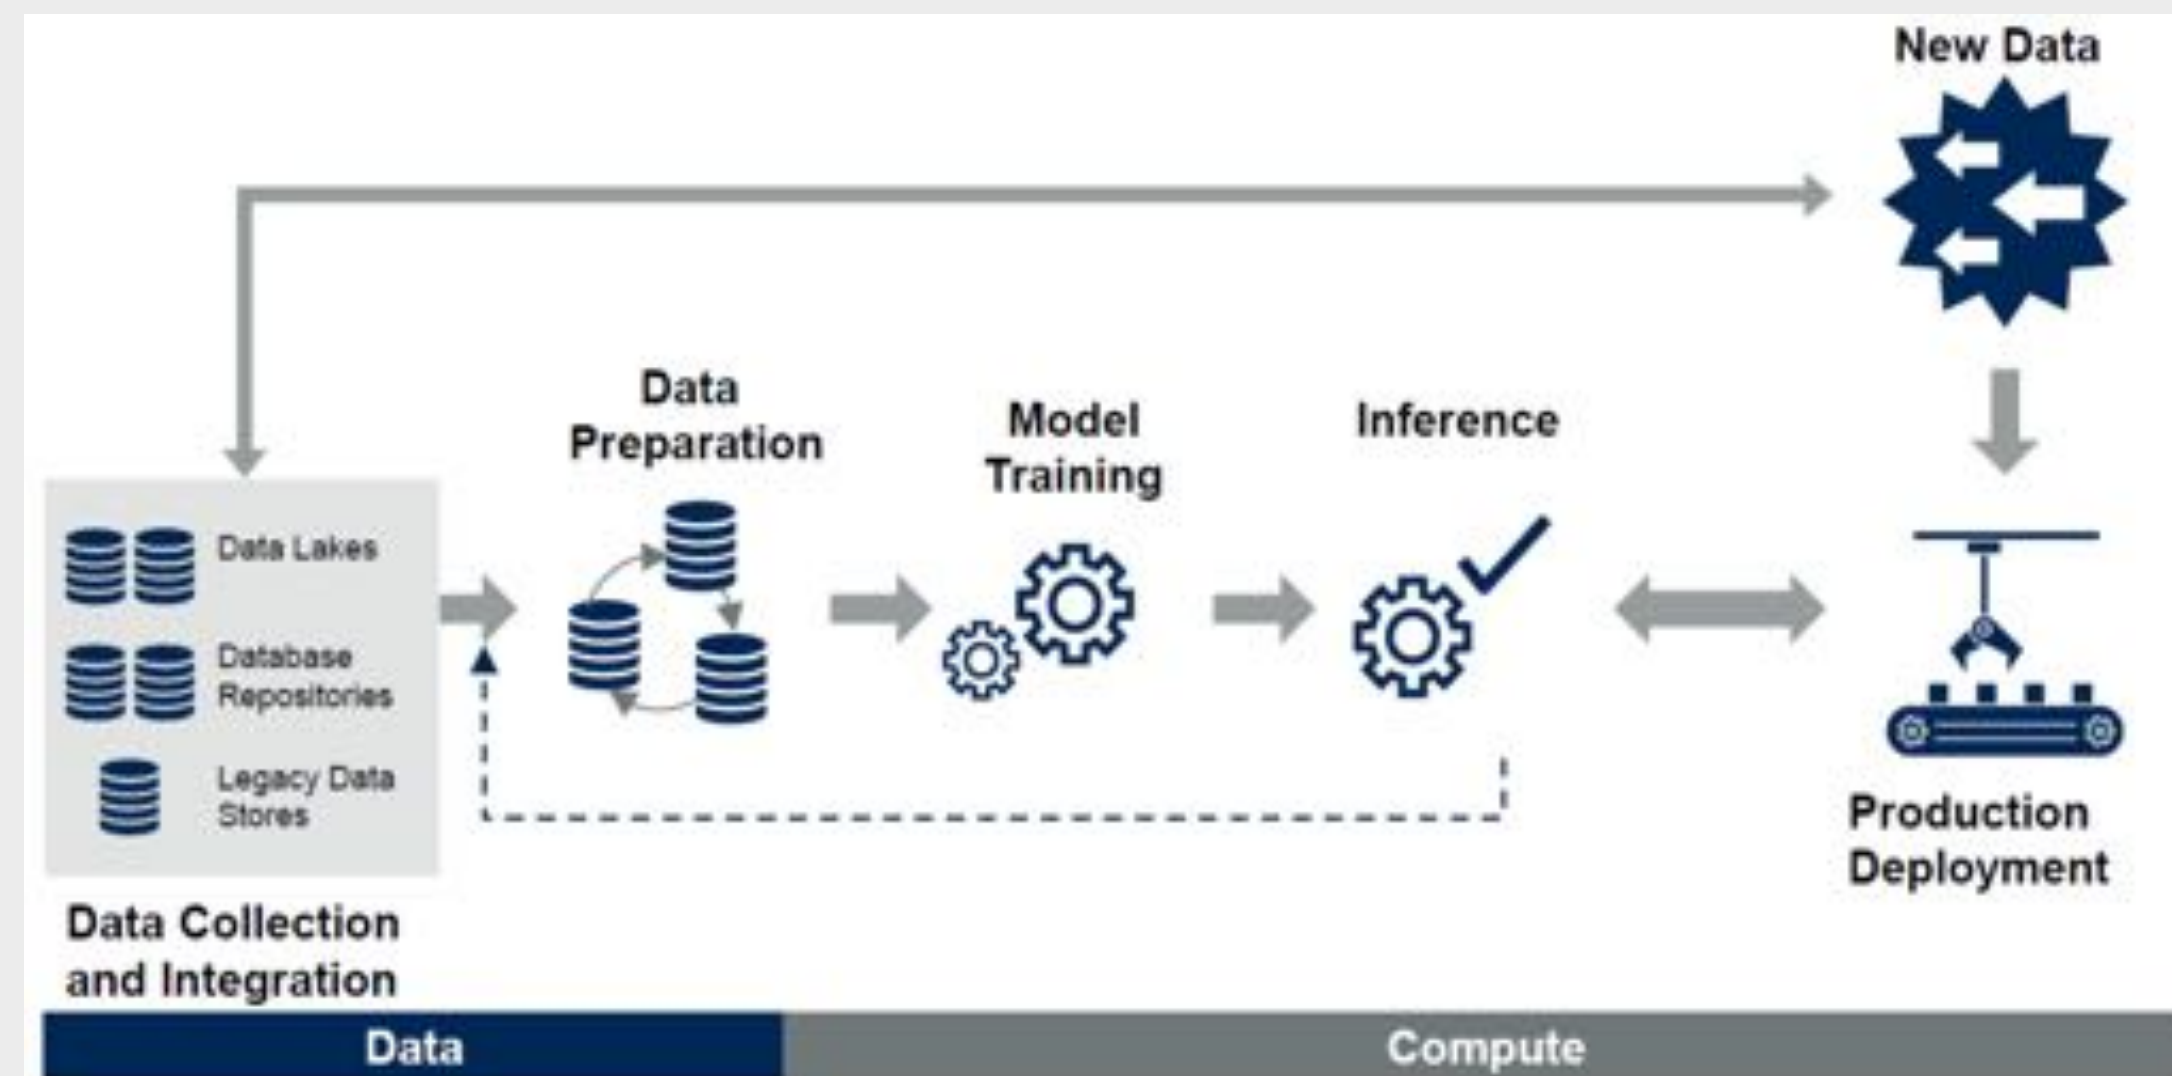
HAWC

Example Architecture



Artificial Intelligence / Machine Learning

- Training phase is critical to AI
 - Massive amounts of data required for deep learning
 - Fast storage required to keep expensive GPUs busy
- Data profile
 - Large and small I/O
 - Millions of files
 - Low latency



© 2018 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates.

AI/ML with IBM GPFS and NVIDIA

Shared NVMe Accelerates Deep Learning

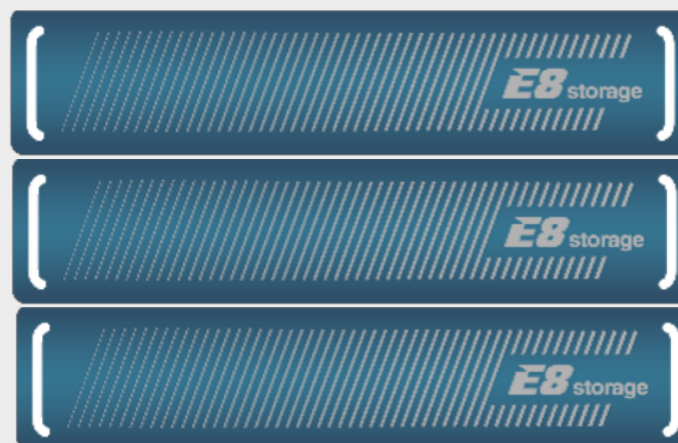


GPU Farm: Nvidia DGX-1

- Up to 8 GPUs per node
- GPFS Client + E8 Agent run on x86 within GPU Server
- Up to 126 GPU nodes in cluster




Mellanox 100G IB



Shared NVMe Storage

- E8-D24 2U24-HA
- Dual-port 2.5" NVMe Drives
- Up to 184TB (raw) per 2U
- Patented Distributed RAID6



High Frequency Trading with Storage Class Memory

Extreme performance for real-time market adjustments

Before

- 20TB cache RAM dedicated to 13 servers
- 20 μ s latency over 40G Infiniband
- 26U @ 9100W storage solution




After

- 24 Intel Optane™ 1TB SSDs in E8-X24
- Data shared to all database nodes
- 2U @ 1200W storage solution



| E8 Storage – Rack Scale Flash. No Compromise.



**PCIe SSD
Performance**



**Centralized Storage
Reliability**



**Hyper-
scalability**



**Affordable
100% COTS**