E8 storage
When Performance Matters

Shared NVMe for High Performance Spectrum Scale Clusters

Stuart Campbell
Platform Architect
May 8th 2019

The E8 Storage Difference

A new architecture built specifically for high performance NVMe SSDs + 100GE/IB networks

- **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**
  - Scaling compute scales performance horizontally
  - Scale storage separately from compute by adding SSD enclosures

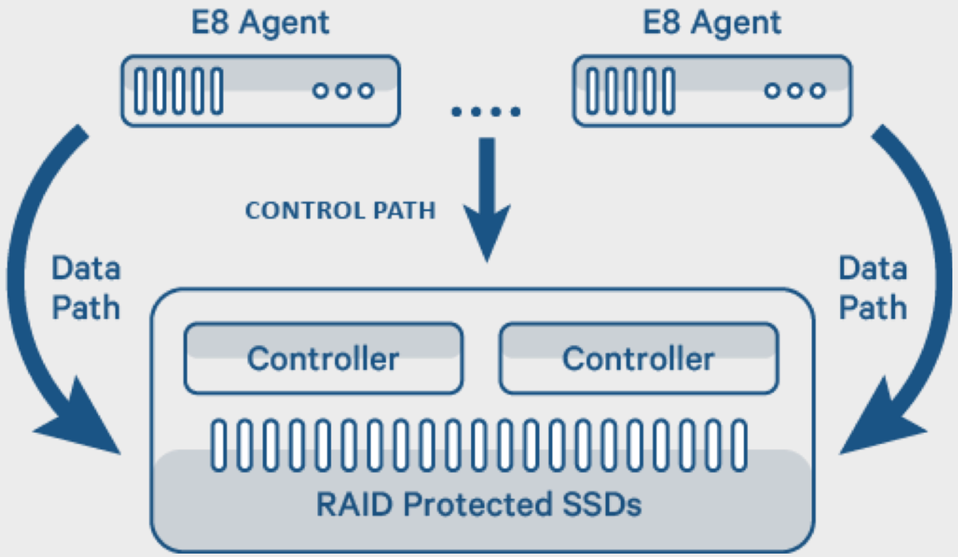

E8 storage
When Performance Matters

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 with virtual spare capacity**
- **Network multi-pathing with fast fail-over**
- **Enclosure high availability**
  - Option 1: HA enclosure + dual-ported SSDs
  - Option 2: Cross-enclosure HA + single-ported SSDs

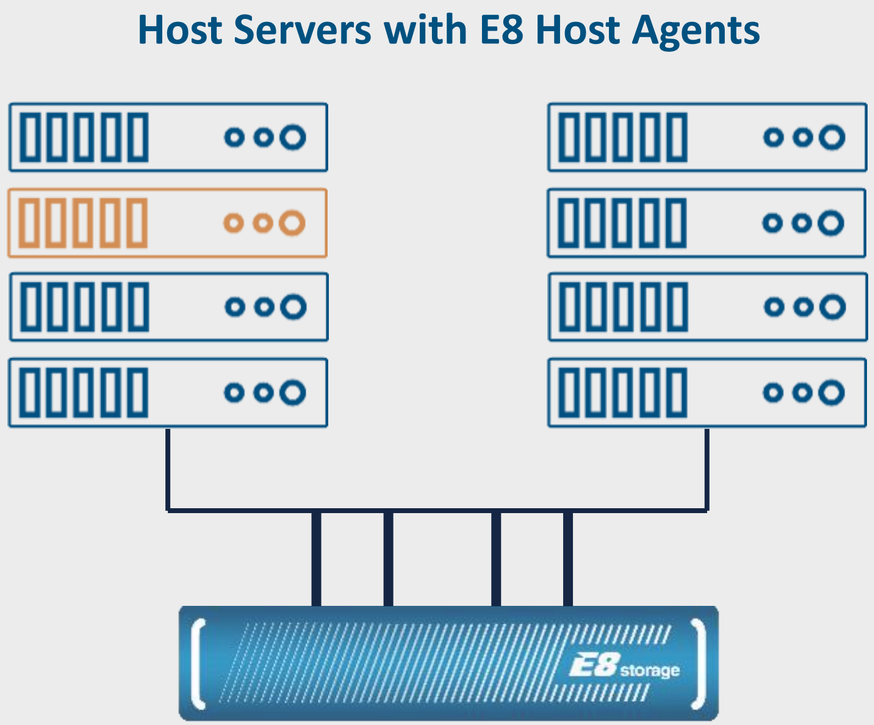

E8 with IBM Spectrum Scale

Multiple Deployment Options

- Standalone pool
- Local Read Only Cache (LROC)
- High Availability Write Cache (HAWC)
- Metadata repository (RAID 10 recommended)

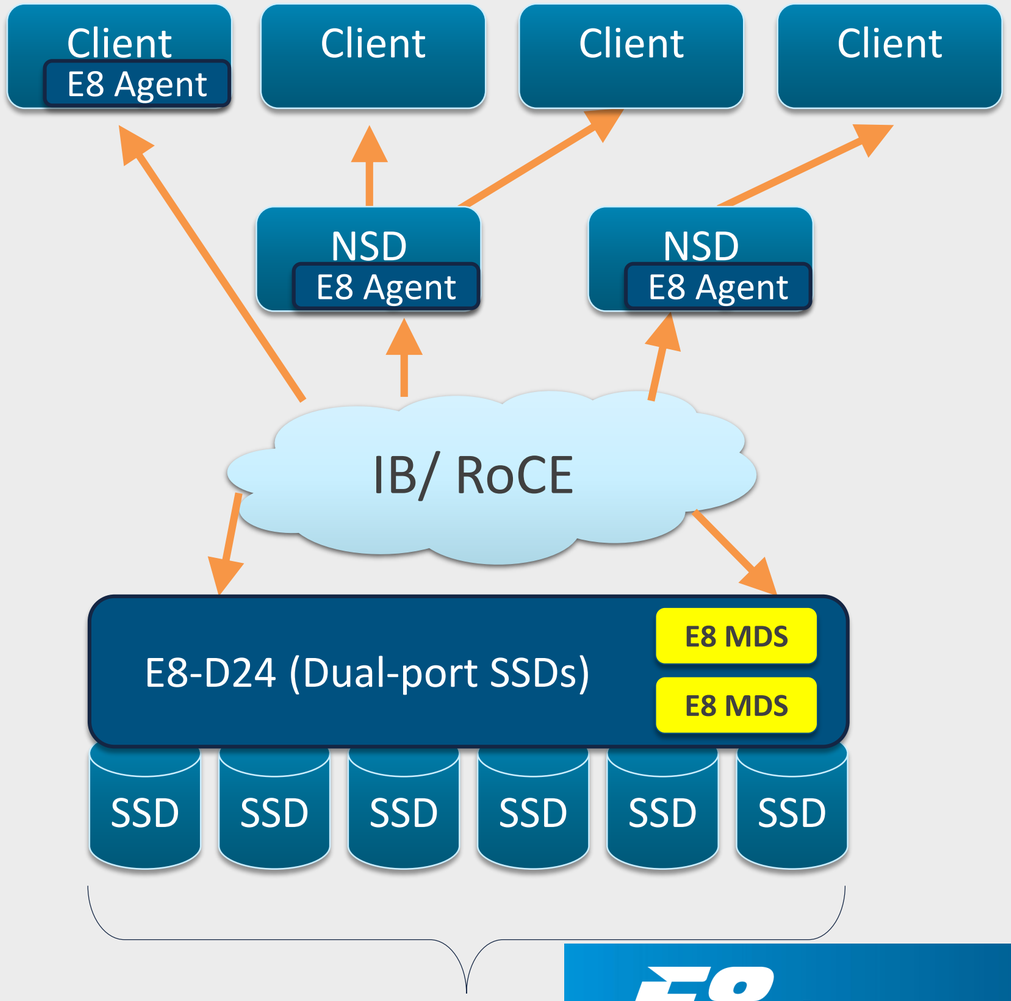

E8 storage
When Performance Matters

E8 storage

When Performance Matters

PCIe SSD Performance

Centralized Storage Reliability

Hyper-scalability

Affordable 100% COTS