Computational Storage Use Cases - FMS 2024

JB Baker, VP Products & Marketing

Agenda

- Promises of Computational Storage
- Deployment Examples:
  - HTAP Database
  - Cloud Parallel File System
  - Relational Database

Promises of Computational Storage

Promises of Computational Storage

Moving ***compute functions to the data***

instead of ***the data to a CPU*** to improve:

Efficiency

TCO

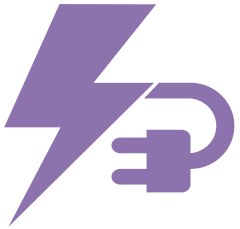

Sustainability

Performance

Security

Deployment Examples

Computational Storage for **HTAP Database**

- Deployment Characteristics
  - Hybrid Transactional Analytical Processing Database
  - Massive scale cloud service
  - 1000’s of customers across 200+ countries and regions
- Desired Outcomes
  - Reduce **Cost** of infrastructure
  - Reduce **Power** consumption
  - Reduce **Maintenance** & **Complexity**
  - All while meeting existing **SLAs**

Computational Storage for **HTAP Database**

Baseline With Ordinary NVMe SSDs

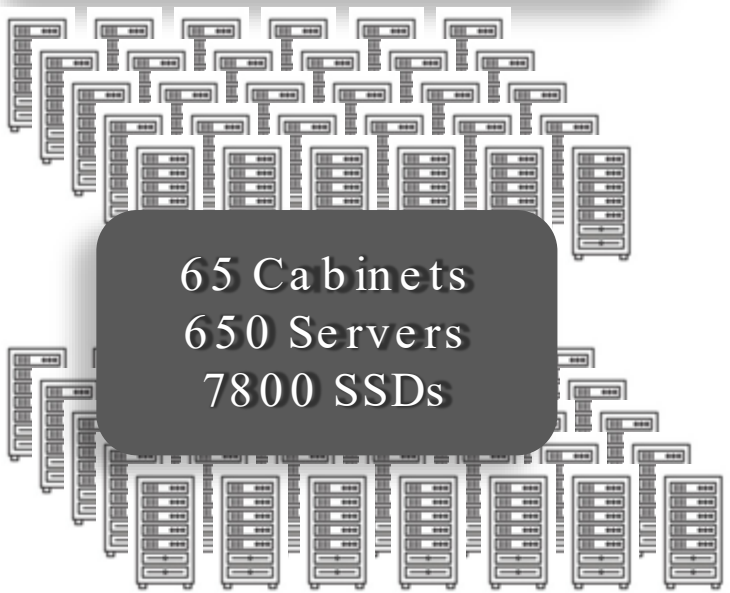

65 Cabinets
650 Servers
7800 SSDs

Using NVMe CSDs for Data Compression

Equipment Needed

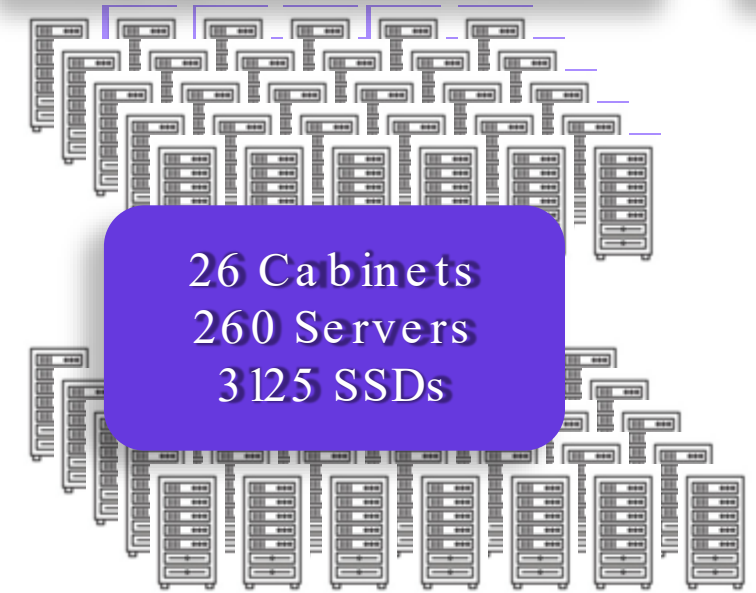

26 Cabinets
260 Servers
3125 SSDs

Improvement

60% lower:
✓ Power
✓ Space
✓ Cost

Computational Storage for **Cloud Parallel File System**

- Deployment Characteristics
  - Massively parallel file system
  - Microsecond access latencies required
  - Serving multiple workloads: AI Training, EDA simulation, CG rendering, and more
  - Triple Replication
- Desired Outcomes
  - Reduce **Cost** of infrastructure
  - Reduce **TCO**
  - Scale **performance** with number of users
  - Keep or improve **Latency** SLAs

Computational Storage for **Cloud Parallel File System**

Baseline With Ordinary NVMe SSDs

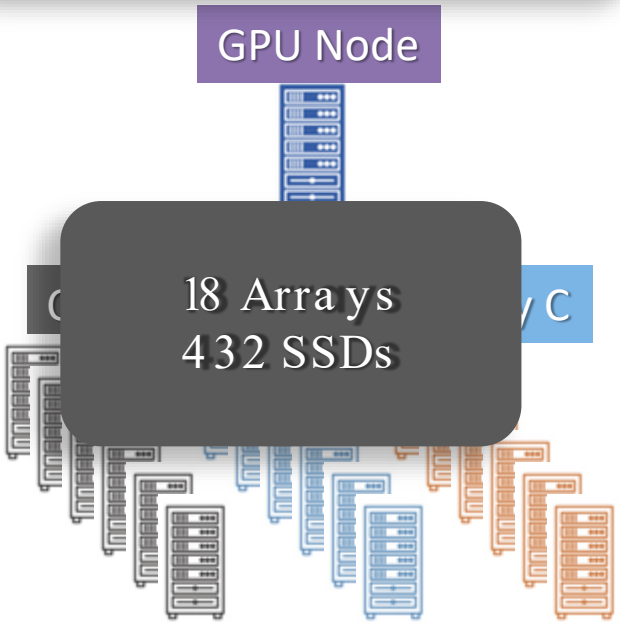

Using NVMe CSDs for Data Compression

Equipment Needed

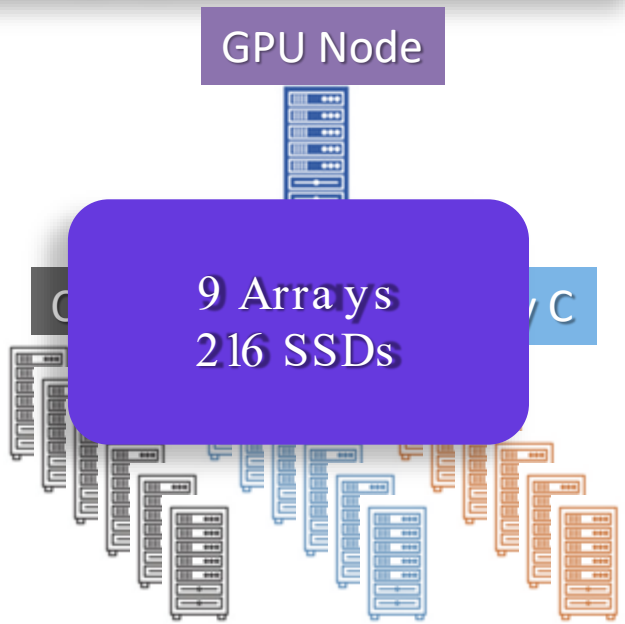

Improvement

50% lower storage Power, Space, Cost
40% lower latency

Computational Storage for **Relational Database**

- Deployment Characteristics
  - Database-as-a-Service provider
  - Tight performance and latency SLAs
  - Redundant systems
- Desired Outcomes
  - Reduce **Cost** of each cluster
  - Keep or improve **Latency** SLAs

Computational Storage for **Relational Database**

Baseline With Ordinary NVMe SSDs

6 App Nodes
12 SSDs
Per Cluster

Using NVMe CSDs for Data Compression

Equipment Needed	Improvement
3 App Node 3 SSDs Per Cluster	50% Fewer Nodes 75% Fewer Drives

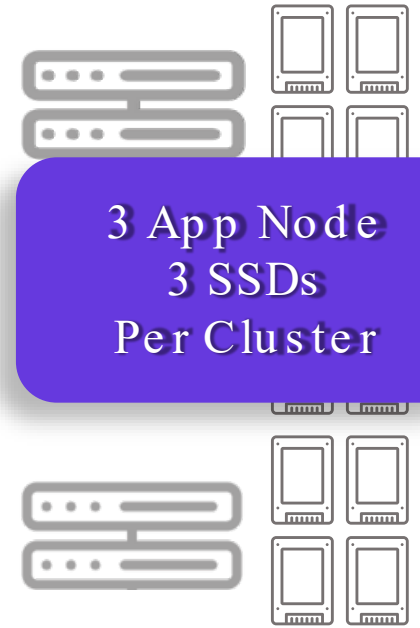

Looking for more info?

info@ScaleFlux.com

www.ScaleFlux.com